

MODELING PHYSICAL PERSONALITIES FOR VIRTUAL AGENTS BY
MODELING TRAIT IMPRESSIONS OF THE FACE:
A NEURAL NETWORK ANALYSIS

by

SHERYL BRAHNAM

A dissertation submitted to the Graduate Faculty in Computer Science
in partial fulfillment of the requirements for the degree of Doctor of
Philosophy, The City University of New York

2002

© 2002

Sheryl Brahnam

All Rights Reserved

This manuscript has been read and accepted for the Graduate Faculty in Computer Science in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

Date

Chair of Examining Committee

Date

Executive Officer

Dr. Michael Anshel

Dr. Daniel Graham

Dr. Richard Holowczak

Dr. Jacob Shapiro

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

ABSTRACT**MODELING PHYSICAL PERSONALITIES FOR VIRTUAL AGENTS
BY MODELING TRAIT IMPRESSIONS OF THE FACE:
A NEURAL NETWORK ANALYSIS**

by

Sheryl Brahnam

Advisor: Dr. Linda W. Friedman

The 1990s gave rise to a host of virtual agents: synthetic characters, interface agents, and virtual humans. Although users welcome the prospect of interacting with virtual agents, more often than not these agents disappoint by being too mechanical and inconsistent in their behaviors. In short, they lack personality. Researchers, recognizing the importance of personality in creating socially engaging virtual agents, have sought ways of endowing agents with a convincing *artificial personality*. There are many aspects to personality. One pressing concern in this area of research is defining those aspects that are of central importance for virtual agents. In this study, the dramaturgical model of personality developed by the social constructivists is used to delineate the domain of artificial personality. According to the social constructivists, personality is the product of three perspectives: that of the actor (the expression of psychological personality), that of the observer (the perception and interpretation of personality), and that of the self-observer

(the management of self-presentations). Most research in artificial personality has focused on the actor. This study explores the observational perspectives by considering the *physical personality* of the actor, defined as comprising those aspects of appearance that give rise to an initial *impression* of personality. It is argued in this study that modeling the impressions of physical personality would provide virtual agents not only with a means of perceiving physical personality but also with a means of creating their own socially intelligent embodiment. To illustrate the feasibility of modeling physical personality, a study focused on modeling the trait impressions of the face using an autoassociative neural network or, equivalently, *Principal Component Analysis (PCA)* is presented. The performances of three-class and two-class PCAs, trained to match human classification of faces in terms of perceived dominance, masculinity, sociality, adjustment, warmth, trustworthiness, facial maturity, and gender, are analyzed. It is found that the PCAs perform better than chance, with two-class PCAs outperforming three-class PCAs. The study concludes by reporting on an investigation designed to gauge the potential of synthesizing faces with a high probability of producing specific trait impressions from within the PCA trait space.

DEDICATION

To Conrad

ACKNOWLEDGMENTS

This study was guided by my exemplary committee chair Linda W. Friedman and the members of my dissertation committee. I was encouraged as well by Dorothy Dologite, Bill Ferns, and everyone involved in the Information Systems Research Workshop (ISRW) at Baruch. I am grateful to the NYPD composite artists, Detectives Steve Mancusi, Juan Perez, and Weldon Ryan. I am deeply indebted to Pierre Cote for generously granting permission to use the database of facial features found in FACES by Interquest and Micro-Intel. Special thanks go to the CIS Computer Lab staff at Baruch College, especially to the assistance of Angela Heath and Glova Smith. I am very grateful for the professional and for the financial support provided by the executive officers of the computer science program, Ted Brown and Stanley Habib. I also wish to express my appreciation to Al Croker, Dessa David, Joe Driscoll, Mary G. Freeman, Avi Giloni, Jonathan Marshall, Victor Y. Pan, Zhon Rosen, and Conrad N. Sawyer for their encouragement, support, and assistance.

TABLE OF CONTENTS

Chapter 1: Introduction.....	1
1.1 HAL’s Legacy.....	1
1.2 The Rise of Computer-Generated Characters.....	4
1.3 Artificial Personality.....	7
1.4 Research Goals.....	9
1.5 Reference Disciplines.....	12
1.6 Contributions of This Study.....	13
1.7 Organization of This Study.....	14
Chapter 2: Agents and Virtual Agents.....	20
2.1 Summary.....	20
2.2 Agent Definitions and Taxonomies.....	20
2.3 Virtual Agents.....	22
Chapter 3: Artificial Personality for Virtual Agents.....	25
3.1 Summary.....	25
3.2 Introduction.....	26
3.3 The Constructivist Approach to Personality.....	34
3.3.1 The Perspective of the Actor.....	35
3.3.2 The Perspective of the Observer.....	38
3.3.3 The Perspective of the Self-Observer.....	41
3.3.4 Interrelation of the Three Perspectives.....	43
3.4 Artificial Personality from the Perspective of the Actor.....	44
3.4.1 Scripting Systems.....	45
3.4.2 Role-Based Systems.....	48
3.4.3 Goal-Based Systems.....	51
3.4.4 Emotion-Based Systems.....	54

3.5	Artificial Personality from the Perspective of the Observer.....	59
3.5.1	Agent Observing Agent	61
3.5.2	Agent Observing User	63
3.5.3	User Observing Agent	65
3.6	Artificial Personality from the Perspective of the Self-Observer.....	66
3.7	Some Problems with Current Research and the Need for the Observational Perspectives	69
Chapter 4: Modeling the Physical Personality of the Face.....		72
4.1	Summary	72
4.2	Introduction.....	73
4.3	Reasons for Modeling the Physical Personality of the Face.....	77
4.3.1	Face Reading.....	78
4.3.2	Making Faces.....	79
4.4	Model Requirements.....	83
Chapter 5: Trait Impressions of the Face.....		86
5.1	Summary	86
5.2	Introduction.....	87
5.3	The Attractiveness Halo Effect.....	90
5.3.1	The Trait Impressions of Attractiveness.....	91
5.3.2	The Facial Qualities of Attractiveness.....	99
5.4	The Overgeneralization Effect of Facial Maturity.....	109
5.4.1	Trait Impressions of Facial Maturity	110
5.4.2	Perceptual Characteristics of Facial Maturity.....	114
5.5	The Overgeneralization Effect of Physical Fitness.....	119
5.6	The Overgeneralization Effect of Emotion.....	120
5.7	Indirectly Modeling Trait Impressions of the Face	123
Chapter 6: PCA Face Classification Techniques.....		125
6.1	Summary	125

6.2	Introduction.....	126
6.3	Autoassociative Neural Networks and Face Classification	129
6.4	PCA and Face Classification	132
6.5	Equivalency of PCA and Autoassociative Neural Networks.....	140
Chapter 7: Research Objectives and Overview		142
7.1	Summary	142
7.2	Objective.....	142
7.3	Limitations of the Current Study	144
7.4	Overview of the Experimental Design.....	146
	7.4.1 Step 1: Data Preparation	146
	7.4.2 Step 2: PCA Modeling.....	147
Chapter 8: Data Preparation.....		147
8.1	Summary	147
8.2	Introduction.....	148
8.3	Generation of the Stimulus Faces	150
8.4	An Experiment Assessing the Trait Impressions of the Stimulus Faces	153
	8.4.1 Participants.....	153
	8.4.2 Dependent Measures.....	155
	8.4.3 Apparatus	155
	8.4.4 Procedure	155
	8.4.5 Results.....	159
8.5	Division of Stimulus Faces into Trait Class Sets.....	164
Chapter 9: PCA modeling.....		168
9.1	Summary	168
9.2	Introduction.....	168
9.3	Training Operations	170
9.4	Testing Operations.....	172
9.5	PCA Evaluation	173

9.5.1	Analysis of Three-Class PCAs	173
9.5.2	Analysis of Two-Class PCAs.....	176
9.6	An Exploration into Face Synthesis from within the PCA Trait Spaces	177
9.6.1	Face Synthesis.....	179
9.6.2	The Trait Impressions of the Synthesized Faces	182
Chapter 10:	Conclusion and Directions for Future Research	186
10.1	Summary and Conclusions	187
10.2	Directions for Future Research.....	195
10.2.1	Building a Specialized Database of Faces	195
10.2.2	Exploring Trait Classification.....	196
10.2.3	Synthesizing Faces.....	196
Appendix A:	Term Definitions and Behavioral Potential Questions	198
Appendix B:	Screen Shots of the Computerized Questionnaire.....	201
Appendix C:	Information Sheet Handed to Human Subjects	203
Appendix D:	Stimulus Faces	204
References	207

LIST OF TABLES

<i>Number</i>	<i>Page</i>
Table 3.1. The Big Five Personality Trait Factors.....	37
Table 3.2. Major Categories for Traits and Trait-like Feelings.....	40
Table 5.1. Mean Weighted Effect Size by Attribute Content in Two Meta-Analyses of the Physical Attractiveness Stereotype.....	95
Table 5.2. Correlation Between Attributes, Attractiveness, and Babyfacedness.	114
Table 8.1. The 10 Trait Dimension Means, Dimension Standard Deviations, and the Lowest and the Highest Face Means within Each Dimension.....	159
Table 8.2. Trait Class Sets: Mean Ranges and Number of Images Selected.....	166
Table 9.1. Averaged Classification Scores for the Three-Class PCAs.....	174
Table 9.2. Example of Distances between Averaged Class Projections.....	175
Table 9.3. Averaged Classification Scores for the Two-Class PCAs.....	176
Table 9.4. Expanded Trait Class Sets for PCAs Used in Face Synthesis.....	179
Table 9.5. Number of Novel Images Selected for Human Subject Assessment.....	181
Table 9.6. Averaged Classification Rates for the Expanded Two-Class PCAs.....	182
Table 9.7. Average Trait Ratings of the Synthesized Faces.....	184
Table A.1. Definitions and Some Behavioral Potential Questions.....	198

LIST OF FIGURES

<i>Number</i>	<i>Page</i>
Figure 5.1. Increased Attractiveness of Averaged Faces.....	105
Figure 5.2. A Media Example of The Overgeneralization Effect of Babyfacedness.	116
Figure 5.3. Cardioidal Strain Transformation	118
Figure 5.4. An Illustration of the Overgeneralization Effect of Emotion.....	122
Figure 6.1. Architecture of a Linear Autoassociative Neural Network.....	129
Figure 6.2. The First 10 Eigenfaces of 220 Randomly Generated Faces.	134
Figure 6.3. An Illustration of the Linear Combination of Eigenfaces.	135
Figure 6.4. Examples of Face-Like Reconstructions.....	136
Figure 6.5. A Simplified Illustration of the Four Results of Projecting an Image onto the Face Space.	140
Figure 7.1. The Two Main Steps Involved in the Experimental Design.	145
Figure 8.1. The Three Steps Involved in Data Preparation.	150
Figure 8.2. An Example of an Acceptable Face (Left) and an Unusual Face (Right).....	152
Figure 8.3. Ethnicity of (N=110) Participants.	154
Figure 8.4. Annual Income Distribution in Dollars of (N=110) Participants.	154
Figure 8.5. Frequency Histograms (A-J) for Each of the 10 Trait Dimensions.	161
Figure 9.1. The Three Steps Involved in PCA Modeling.	170
Figure 9.2. An Example of Face Space Probing.....	178
Figure 9.3. Examples of Face Synthesis by Probing the PCA Trait Space.	180
Figure 9.4. Frequency Histograms of the Ratings of the Synthesized Faces.....	184
Figure B.1. Demographic Information Form.....	201
Figure B.2. An Example of a Face Questionnaire with the Term Definitions/ Instructions Screen Open.	202
Figure D.1. The 220 Stimulus Faces.....	204

CHAPTER 1: INTRODUCTION

The countenance is the title-page which heralds the contents of the human volume, but like other title-pages it sometimes puzzles, often misleads, and often says nothing to the purpose.

W. Mathews

1.1 HAL'S LEGACY

Anyone who has seen the movie *2001: A Space Odyssey* probably retains vivid memories of HAL, a futuristic computer that many critics claim has more dramatic presence in the film than the more prosaic scientists aboard the *Discovery*. Audiences worldwide have come to treasure HAL for his remarkable humanlike qualities. They marvel at his intelligence, his language skills, his show of warm concern and fatherly appreciation for Dave's artistic endeavors. They even appreciate HAL for his misdeeds, which culminate in the murders of Frank and of the three crewmembers in suspended animation. But perhaps what they remember most about HAL is the emotional power of his swan song while being dismantled by Dave: "Dave, Stop . . . Stop, will you? Stop, Dave . . . I'm afraid, Dave . . . Dave . . . my mind is going . . . I can feel it . . . I can feel it . . . my mind is going . . . There is no question about it. I can feel it . . . I can feel it . . . I can feel it . . . I'm a . . . fraid."

The emotional power of HAL's final scene—as his fears eventually give way to memories of his birth, his first instructor, and the first song he ever learned, *Daisy*, which he begins to sing, until slowing, he disappears into oblivion—is so convincingly real that it is hard for anyone to pull back enough to ask whether HAL's feelings and experiences could possibly be genuine. For given his character and agenda, not to mention the fact that he is made of silicon, isn't it more likely his emotional display is yet another clever attempt to put a stop to Dave—a cold and calculated play on human emotion, a masterpiece of deceitful acting? After all, HAL had no reason to be afraid. Dave was not destroying HAL. His memory boxes were simply being removed and set aside. HAL could easily be reassembled—or could he?

That essentially was the question Stork (1997) recently put to a number of leading researchers in computer science and related fields in *HAL's Legacy: 2001's Computer as Dream and Reality*. Stork wanted to know how far science has come in building a computer that resembles HAL. The book includes reports on supercomputer design, reliable computing, speech synthesis and recognition, computer vision, lip reading, planning, chess playing, ethics, common sense reasoning, and even digital emotions. It includes both some surprising successes and disappointments. If HAL were born today, in 2002, he could easily beat Frank or any other crewmember at a game of chess. Soon computers may very well be the world's best chess players, and they will certainly be as fast and large as the HAL 9000 series and possibly even as reliable. However, what HAL

is loved most for—his humanness—is still elusive. Work has just begun in making computers that are responsive to human emotions and moods. Speech recognition and common-sense reasoning are active areas of research, but no computer really comprehends language, and computers still lack many of the elementary social and common-sense reasoning skills easily mastered by young children. Yet, despite these challenges, Stork presents convincing evidence that advances are being made, however slowly, towards the realization of the modern world’s dream of building a computer as humanlike as HAL.

Still there is more to HAL’s legacy than Stork’s book takes into account, something so rich and magical in its transformative power that it is easily overlooked. For in reality HAL is neither a person nor a computer. HAL is a character, a proper noun that represents a part in a script, a being composed of a few scraps of dialogue, forty-four lines to be exact. These lines are brought to life by the expressive powers of Douglas Rain, the voice behind HAL, and the directive vision of Clarke and Kubrick. Thanks to the talents of artists like these, characters like HAL—Mary Poppins, George Bailey, Mickey Mouse—have a way of transcending the narrow confines of their medium to take on separate lives of their own. They enter the world, “like a visitor to town,” as the literary critic Gass (1979) once put it, and later “leave on the arm,” so *convincingly real* by the end of the film do characters as rich as these become.

Another question Stork could have asked then, apropos of HAL, his legacy, and the state of the art in computer science, is “How far have computers come in being able to perform this transformative magic?” In other words, to what degree are computers now capable of playing a role or creating a character that is anywhere near as convincing, as engaging, and as real as that of HAL?

1.2 THE RISE OF COMPUTER-GENERATED CHARACTERS

Certainly if Stork had included a chapter reviewing the state of the art in computer film techniques and animation, he would have noted that current technology has far surpassed Kubrick’s thwarted intention of actually incorporating a few computer graphics or special effects into the psychedelic sequences at the end of *2001*. In the late 1960s, the inclusion of computer imagery was simply not an artistic option as the maximum resolution at that time of 512 by 512 pixels was far too poor to be tolerated on the giant magnifying screens of the movie houses. This may explain in part why HAL’s physical appearance, a series of red bulbs networked throughout the spacecraft, is so stark and unremarkable. The tremendous advances in computer architecture in the last thirty years, however, coupled with a continuous decline in the cost of production, have accelerated progress in high-resolution computer graphics to such an extent that an entire industry devoted to the

production of high quality special effects has arisen. Thanks especially to recent developments in photorealistic scene generation, forward and reverse kinematic techniques, and behavior-based animation, computers are also playing an increasing role in animating characters. The last decade, in particular, has seen not only the debut but also the rise in popularity of computer-generated characters. They have appeared in such films as *Young Sherlock Holmes* (1985), *Terminator 2* (1991), *Casper* (1995), *Toy Story* (1995), which was the first full-length motion picture entirely generated by computer, *Antz* (1998), and most recently and impressively *Monsters, Inc.* (2001). Although *Prince of Egypt* (1999) was the first film produced by Dreamworks SKG not animated by computer, that film became the notable exception in 1999 rather than the rule. Today computerized animation is mainstream.

The popularity of these characters is not limited to the movies. Computer-generated characters are being recruited into a variety of desktop applications. The gaming industry is producing highly realistic characters in response to users who demand that they be allowed to interact with animated creatures that possess the same level of realism as the creatures they have come to love in the movies (Loyall, 1997a). In tutorials and multimedia titles, students increasingly encounter computer-generated characters that function as friendly guides, presentation agents, and docents (Houle and Simon, 1997; Lester and Stone, 1997; Towns, Voerman, Callaway, and Lester, 1998; Watt, 1997). These characters are often included in programs as interface agents, assistants, or experts

of various kinds (Arafa, Charlton, and Mamdani, 1998; Maes and Wexelblat, 1996; Miksch, Cheng, and Hayes-Roth, 1997). They are an essential ingredient in training simulations and in cutting-edge experiments in interactive movies, fiction, and computer drama (Hayes-Roth and van Gent, 1997; Hirsh, Hayes-Roth, Stern, and Murray, 1998; Loyall, 1997a; Nakatsu, Tosa, and Ochi, 1998; Pinhanez, 1996; Vardi, 1999). And in the computer toy industry, they pose as virtual playmates and dolls (Hayes-Roth, Sincoff, Brownston, Huard, and Lent, 1995) or else playfully romp around the desktop as virtual pets begging to be petted, played with, and fed (Stern, Frank, and Resner, 1998).

It would seem, given the remarkable film successes and popularity of computer-generated characters in the 1990s, that an assessment of whether computers are capable of creating characters as engaging as HAL would be overwhelmingly positive. But such is not the case. Success is mostly limited to the movies and media arts, where the computer-generated characters are so enchanting that many people are understandably misled into believing these characters can magically reappear on their desktops. Unfortunately, once transported there, these characters disappoint by behaving inconsistently, moving about awkwardly, and failing to respond appropriately (Oren, Salomon, Kreitman, and Don, 1990). A major research problem today is how to enliven these desktop characters by endowing them with an engaging and convincing *artificial personality* (Trappl and Petta, 1997).

1.3 ARTIFICIAL PERSONALITY

Personality has many facets, as a glance at the psychological reference literature confirms. One major concern in the field of artificial personality is defining and isolating those aspects that are of central importance for modeling personality for computer characters, or *virtual agents* (Antonio, Aylett, and Ballin, 2001), the generic term used in this study for characters that are both rendered and controlled by computer. On the one hand, models of personality for virtual agents need not be as comprehensive and as accurate as psychological models of personality. The concern in artificial personality is mostly with characterization, not psychological fidelity (Bates, 1992a; Hayes-Roth, van Gent, and Huber, 1997). On the other hand, virtual agents differ significantly from traditional media-based characters in ways that make endowing virtual agents with a convincing personality problematical. The most significant difference lies in the fact that media-based characters are not required to function in real time. Every aspect of their character—their dialogue, their physical form and expressions, every action and reaction—is carefully worked out in advanced, typically by a team of talented artisans. In contrast, far fewer resources are spent on virtual agents. Yet they are required to be both autonomous and interactive (Badler, Palmer, and Bindiganavale, 1999). Virtual agents, unlike their media-based counterparts, have the daunting task of sustaining the illusion of personality within a highly interactive and unpredictable *social* environment.

Is there a theory of personality that can inform and guide research in the field of artificial personality? The social constructivists provide a dramaturgical model of personality that is particularly valuable in understanding human personality in terms of a *social interaction* (Hampson, 1988a). According to the social constructivists, personality is the product of three distinct perspectives: that of the actor, that of the observer, and that of the self-observer (Hampson, 1995). The perspective of the actor involves everything associated with the development of individual personality and the ways in which such *internal factors* as genetics, traits, and personal experiences influence behavior and attitudes. The perspective of the observer deals with the interpretation of personality as it is inferred from an actor's behaviors, appearance, and possessions. The perspective of the self-observer concerns the actor's ability to consciously ascertain and manage the impressions his personality makes on others. The constructivists claim that all three perspectives are equally important in the *social construction* of personality.

The review of the reference literature in chapter 3 reveals that research into artificial personality is primarily centered on the actor. In the last decade numerous advances have been made in the design of virtual agents that are capable of communicating personality through a variety of expressive modalities. In these systems personality originates within the agent as attitudes, drives, desires, and character traits. These personality components

are then used to constrain emotional intensities, modify behaviors, and guide both goal selection and strategies (Loyall and Bates, 1997; Rizzo, Veloso, Miceli, and Cesta, 1999).

In contrast, research stemming from the perspectives of the observer and self-observer has virtually been neglected—this despite repeated acknowledgements that the user plays a vital role in creating the *illusion* of personality for virtual agents (Elliott and Melchoir, 1995; Elliott, Brzezinski, Sheth, and Salvatoriello, 1998; Hayes-Roth, Johnson, Gent, and Wescourt, 1999; Loyall, 1997a). If it is true, however, that personality is best characterized as a social construction composed of all three perspectives, it would behoove the research community to include the observational perspectives more centrally in its research agenda. As reported in chapter 3, there are serious drawbacks with the current one-sided focus on the actor. Offsetting this imbalance by emphasizing the observational perspectives would not only resolve some of these issues but also open up promising new areas of investigation.

1.4 RESEARCH GOALS

The intention of this research project is to address the observational perspectives more explicitly by focusing on one particular manifestation of personality that immediately impresses itself upon the observer. It is a facet of personality that might best be labeled

the *physical personality*, since it is based on what is instantaneously projected by the actor's physical appearance. It is no secret that people are predisposed to form impressions of a person's social status, abilities, intentions, and character based on nothing more than how that person looks. As Berscheid and Walster (1974) have observed, “. . . our appearance telegraphs more information about us than we would care to reveal on a battery of personality inventories, intelligence tests, and character scales. From flame-colored hair through flat feet, few aspects of appearance fail to provide kernels of folk insight into another's nature” (p. 159).

Even though a number of researchers have investigated the impact the physical personality of an agent has on users (Sproull, Subramani, Kiesler, Walker, and Waters, 1996) and have designed virtual agents in light of this research (De Carolis, de Rosis, and Pizzutilo, 2000), no one to date has attempted to model the user's perceptions of the physical personality of the actor. The goal of this study is to investigate this perceptual possibility. Physical personality, however, embraces a wide range of features—body type, posture, hair color, and skin texture, to name but a few. Modeling the personality impressions of each of these features would be a major enterprise. For this reason, the scope of this project is restricted to covering only those impressions of personality that are elicited by the morphological characteristics of the human face. After all, the face, more than any other part of the body, stands for and is identified with the social self. It serves as a multimodal social interface, replete with signs regarding a person's age,

gender, social status, emotional state, attitudes, and personality. To quote Liggett (1974), “There can be little doubt that the face plays a crucial part in our everyday assessment of our fellows. Not only does it enable us to identify transient emotions—flashes of pleasure and rage, disappointment and hatred—it can also help us to make useful judgements about more durable and lasting qualities of personality and character . . . ” (p. 276).

Because the problem of modeling the trait impressions of the face is basically a face perception problem, this study investigates the possibility of training a linear autoassociative neural network to match the human classification of faces into representative trait categories. A standard face classification technique, autoassociative neural networks associate input patterns with themselves and are sometime referred to as *PCA* (*Principal Component Analysis*) neural networks (Diamantaras and Kung, 1996), so named because they are equivalent to finding the principal components of the cross-product of a set of inputs (Abdi, 1988). *PCA* is the generic term used in the reference literature to refer to *PCA* autoassociative neural networks and other *PCA* classification techniques and is the preferred term.

It is the prediction of this study that PCA will succeed at modeling the trait impressions of the face.¹ PCA has already proven successful at classifying faces according to age (Valentin, Abdi, O'Toole, and Cottrell, 1994), gender (O'Toole and Deffenbacher, 1997; Valentin, Abdi, Edelman, and O'Toole, 1997), and emotional expression (Cottrell and Metcalfe, 1991; Padgett and Cottrell, 1998), characteristics associated with large clusters of trait impressions (see chapter 3). Moreover, since PCA face classification techniques are holistic (see chapter 4), PCA should also prove capable of modeling a number of specific traits even though at this time it is not known with certainty which features are involved in the formation of specific trait impressions. Finally, PCA is not only predictive but also capable of face synthesis.

1.5 REFERENCE DISCIPLINES

Artificial personality is a field that is recognized as being intrinsically interdisciplinary (Trappl and Petta, 1997). Some disciplines that have contributed to this area of research include ethology, linguistics, drama theory, characterization in the media arts, and psychology. In computer science, artificial personality draws heavily on artificial intelligence (especially classical goal-based and reactive planners), research in agent

¹ As the intention of this study is to investigate the feasibility of modeling the physical personality of the face, other classification techniques are neither tried nor compared at this time.

technology, natural language processing, computer graphics, and animation. Trappl and Petta have called for researchers in this area to be ever “on the lookout for yet other fields and disciplines which have not been considered in the given context up to the present time” (p. 215). They go on to characterize the field as being “one of the most exciting scientific meeting grounds.” Since this study focuses on the observational perspectives, specifically on the perception of physical personality, it adds to the growing list of relevant disciplines, neural network, and other statistical image classification techniques.

1.6 CONTRIBUTIONS OF THIS STUDY

The list below summarizes some of the contributions of this study:

- Hampson’s dramaturgical model of personality (1995) is presented as providing a metamodel of personality that is particularly useful in understanding and in organizing the field of artificial personality (chapter 3).²

² It should be noted that this study is not the first to make use of Hampson’s dramaturgical model of personality in the development of virtual agents. Churchill, Cook, Hodgson, Prevost, and Sullivan (2000) have acknowledged the importance of Hampson’s metamodel in outlining some of their research objectives. Specifically, they use Hampson’s model to suggest that virtual agents adjust their personalities to obtain their objectives and to argue that virtual agents be required to pass a “lay personality psychologists test.” This investigation, in contrast, uses Hampson’s model to delineate the entire field of artificial personality and to stress the importance of modeling personality from the observational standpoints.

- The notion of *physical personality* is introduced, and reasons for equipping virtual agents with a means of perceiving physical personality are discussed (chapter 4).
- Finally, and most importantly, a design for a computational model of the trait impressions of the face is offered as a way of illustrating the feasibility of modeling physical personality (chapters 7-9).

1.7 ORGANIZATION OF THIS STUDY

What follows is a summary of each chapter in the remainder of this document.

Chapter 2, “Agents and Virtual Agents,” provides a discussion of the ways in which agents can be defined and classified. Agents can be defined and classified in terms of their properties, the tasks they perform, their control architectures, or, as in biology, according to their natural kind. Agents are generally described as being autonomous, adaptive, and social. Virtual agents are agents that are anthropomorphically embodied, intended to function in social environments, and capable of various forms of social self-expression.

Chapter 3, “Artificial Personality for Virtual Agents,” reviews the reference literature in artificial personality. Hampson’s dramaturgical model of personality is used to define and to delineate the field of artificial personality. The dramaturgical model views personality from three perspectives: that of the actor (which is concerned with the expression of personality), that of the observer (which is concerned with the perception of personality) and that of the self-observer (which is concerned with the negotiation of personality). Not only does the dramaturgical model of personality shed light on present research concerns in artificial personality, but it also points out new areas for exploration. In the area of artificial personality, most research to date has focused on the perspective of the actor. The observational components of personality have received far less attention. Recent research has examined ways in which agents can observe users, but most of this work has been concerned with the agent observing the user’s emotional state rather than the user’s personality. In addition, agents have been developed which abduct the personalities of other agents, and researchers have begun investigating how users attribute personality to computers. No work to date, however, has explicitly addressed modeling the observational perspectives.

Chapter 4, “Modeling the Physical Personality of the Face,” introduces the concept of *physical personality*, a term that is defined in this study as comprising all those aspects of appearance that produce an initial impression of personality and a concomitant set of reactions and expectations in observers. Unlike psychological models of personality that

are focused on the actor, models of physical personality require that attention be placed squarely on the observational perspectives, or the *perception* of personality. Discussed in this chapter are some reasons virtual agents should be endowed with the ability to perceive physical personality, especially the trait impressions of the face. The chapter ends with the listing of four requirements that a model of the physical personality of the face must satisfy if it is to accommodate the specific needs of virtual agents.

Chapter 5, “Trait Impressions of the Face,” provides a fairly complete review of the person perception literature on the trait impressions of the face. There is considerable consistency in people’s impressions of faces. One major theory advanced to explain this consistency is that the perception of facial features has adaptive value and that trait impressions are based on those facial qualities that demand the greatest attention for the survival of the species, namely, physical fitness, age, and emotional state. Faces that possess features that are indicative of these qualities are believed to have pronounced *overgeneralization effects* (Zebrowitz, 1998). A modification of Rosenberg’s trait categories (1977) is used in this chapter to explore clusters of traits associated with five of the most important overgeneralization effects: attractiveness, facial maturity, gender, physical fitness, and emotion. A discussion of the morphology that triggers these overgeneralizations is also presented. Finally, the prospect of indirectly modeling the trait impressions of the face using existing models of attractiveness, facial maturity, and

emotion is evaluated. It is concluded that indirect methods are inadequate primarily because they are not predictive.

Chapter 6, “PCA Face Classification Techniques,” presents an argument for using standard neural network face classification techniques in modeling the trait impressions of the face. It is difficult to isolate those important facial features that hold the keys to an understanding of how faces can be classified. Holistic approaches, such as linear autoassociative neural networks, which allow the classifier system to identify relevant features a posteriori, have been shown to outperform systems that determine important feature sets a priori. It has been demonstrated that classifying faces using a linear autoassociative neural network is equivalent to finding the principal components of the cross-product matrix of a set of faces and reconstructing the faces as a weighted sum of eigenvectors. Since the principal components of an image can be derived statistically or by using an autoassociative neural network, both methods are reviewed in this chapter and their equivalency demonstrated.

Chapter 7, “Research Objectives and Overview,” presents the hypothesis that holistic face classification techniques, specifically PCA, will satisfy the four requirements, presented in chapter 4, of a model of the trait impressions of the face that is suitable for virtual agents. In this chapter, a statement of the objectives and the limitations of the

current investigation are provided, as well as a brief outline of the steps involved in training and testing a PCA model of facial trait impressions.

Chapter 8, "Data Preparation," describes the three steps taken to prepare the trait class sets of faces needed to train and test the PCA models. Data preparation is a three step process. In step one, 220 stimulus faces are randomly generated from a database of facial features. In step two, 110 subjects judge the stimulus faces along 10 trait dimensions using a 7-point bipolar scale. In step three, trait class sets are formed for each trait dimension by taking a subset of the 220 faces that are clearly representative of the two bipolar extremes and a neutral class of faces.

Chapter 9, "PCA Modeling," details the steps taken to train and to evaluate the PCA trait models. The performance of PCA classification for the eight trait dimensions of dominance, masculinity, sociality, adjustment, warmth, trustworthiness, facial maturity, and degree of certainty regarding gender is analyzed for three-class and two-class PCAs. It is found that PCAs trained to classify two trait classes significantly outperform PCAs trained to classify three trait classes. The chapter concludes by reporting on a promising exploration in face synthesis.

Chapter 10, “Conclusion and Directions for Future Research,” summarizes and evaluates the results of this study. It ends by offering directions for future research specifically in the area of modeling the trait impressions of faces.

CHAPTER 2: AGENTS AND VIRTUAL AGENTS

What do you get if you mix one part Shakespeare, two parts Pac Man, and three parts AI?

Hayem, Fourmaintraux, Petit, Rauber, and Kisseleva
Avatars: New Fields of Implication

2.1 SUMMARY

Agents can be defined and classified in a number of different ways: in terms of their properties, the tasks they perform, their control architectures, or, as in biology, according to their natural kind. Agents are generally described as being autonomous, adaptive, and social. Virtual agents are agents that are anthropomorphically embodied, intended to function in social environments, and capable of various forms of social self-expression.

2.2 AGENT DEFINITIONS AND TAXONOMIES

Part two of *Intelligent Agents III: Agent Theories, Architectures, and Languages* (Müller, Wooldridge, and Jennings, 1996), entitled “What is an agent?—Definitions and taxonomies,” presents a heated debate regarding the characteristics of agents, particularly as they differ from other computer programs or from other systems such as thermostats. Franklin and Graesser (1996) set the stage for the debate that ensues by examining

numerous definitions of agents as provided, for example, by Maes, Darrell, Blumberg, and Pentland (1997), Hayes-Roth (1995), Smith (1994), and Foner (1993). Franklin and Graesser note that agents can be characterized in a variety of ways: in terms of their properties, the tasks they perform, and their control architectures. They can even be defined or classified, as in biology, in terms of their natural kinds (Franklin and Graesser).

So what is an agent? Bradshaw (1997) offers insight into some of the confusion and disagreement in the use of the term among computer scientists by noting that the nominal definition of *agent* typically includes both the notion of a cause (the means by which something is done) and the notion of a representative (as when one acts on behalf of another), ideas that may or may not be combined into a single software agent. Thus, for some, an *agent* refers to a computer program that is autonomous and adaptive, for others, it is, as Bradshaw writes, something “. . . connected with what is portrayed to the user. Here ‘agent’ is used to describe programs which *appear* to have the characteristics of an animate being, often a human” (p. 80). He calls this second aspect of an agent, the *agent metaphor*. Although it is difficult to provide a definition of the term *agent* that would not spark debate, most researchers would concede that agents can be characterized, in terms of the nominal definition above, as autonomous, reactive, and social (Wooldridge, 1996).

2.3 VIRTUAL AGENTS

The term *virtual agent* (Antonio, Aylett, and Ballin, 2001) is an umbrella term for a large variety of metaphor agents that are anthropomorphically embodied, intended to function in social environments, and capable of various forms of social self-expression, for example, facial displays, gaze, gestures, speech, and body postures. Terms that are similar to *virtual agent*, if not identical in many usages, are *animated agent*, *believable agent* (Bates, 1994; Loyall, 1997a; Reilly, 1996), *embodied artificial life* (Dautenhahn, 1998), *entertainment agent* (Franklin and Graesser), *synthetic character* (Elliott and Brzezinski, 1999; Isbister and Hayes-Roth, 1998; Johnson, Wilson, Kline, Blumberg, and Bobick, 1999), *virtual actor* (Meuleman, Heister, Kohar, and Tedd, 1998; Thalmann, Noser, and Huang, 1997; Wavish and Connah, 1997), *virtual being* (Badler, 2001), and *virtual human* (Badler, Palmer, and Bindiganavale, 1999).

It is predicted that the computer interfaces of the future will rely heavily on virtual agents (Negroponte, 1995). There are those who cite the natural tendency of people to anthropomorphize as one contributing factor in this movement toward humanlike interfaces. In this regard, Laurel (1990a) writes, "Interface agents draw their strength from the naturalness of the living-organism metaphor in terms of both cognitive accessibility and communication style" (p. 356). She goes on to predict that virtual agents will function in application areas that "range from managing mundane tasks like scheduling, to handling customized information searches that combine both filtering and

the production (or retrieval) of alternative representations, to providing companionship, advice, and help throughout the spectrum of known and yet-to-be-invented interactive contexts." In addition to guiding users through libraries and databases of information, as Laurel predicted in 1990, virtual agents are currently being incorporated into many engineering applications, entertainment programs (Hayes-Roth, Sincoff, Brownston, Huard, and Lent 1995; Stern, Frank, and Resner, 1998), virtual environments (Murray, 1997; Vilhjálmsón and Cassell, 1998), training applications (Rickel and Johnson, 1998), and ergonomic and military simulations (Badler, Bindiganavale, et al., 1999; Badler, Palmer, et al., 1999). Examples of virtual agents currently under development include *Adele* (Shaw, Johnson, and Ganeshan, 1999), *Cosmo* (Lester, Voerman, Towns, and Callaway, 1997), *DFKI Persona* (Andre, Muller, and Rist, 1996), *Gandalf* (Thórisson, 1997), *Herman the Bug* (Lester and Stone, 1997), *Improv Puppets* (Huard and Hayes-Roth, 1996), *Olga* (Beskow and McGlashan, 1997), *Peedy the Parrot* (Ball et al., 1997), *PPP-Persona* (Rist, André, and Müller, 1997), *Rea* (Bickmore and Cassell, 2001), and *Steve* (Rickel and Johnson, 1998).

The verdict is still out, however, on whether virtual agents function effectively as interface agents. A few researchers are convinced that virtual agents interfere with user performance (Shneiderman and Maes, 1997; Lanier, 1995; Walker, Sproull, and Subramani, 1994; Wright, Milroy, and Lickorish, 1999). Even Patti Maes, once their champion, has admitted to shying away from research into interface agents that are

anthropomorphized (Shneiderman and Maes, 1997). Other studies, however, suggest that virtual agents offer a number of benefits (Dehn and van Mulken, 2000). There is evidence that they make software use more entertaining, especially for children (Koda and Maes, 1996; Takeuchi and Naito, 1995). Users tend to spend more time interacting with virtual agents and make fewer mistakes and respond more when prompted by virtual agents as opposed to textual prompts (Lester, Converse, et al., 1997; Lester, Towns, Callaway, Voerman, and FitzGerald, 2000; Walker et al., 1994). There is also evidence suggesting that virtual agents increase the perceived trustworthiness, intelligence, and thus the acceptance of software that employs them (Hietala and Niemirepo, 1998; Rickenberg and Reeves, 2000; Sproull, Subramani, Kiesler, Walker, and Waters, 1996). Virtual agents also provide users the illusion of interacting more naturally with a reciprocating companion (De Carolis, Rosis, and Pizzutilo, 2000).³

³ For a detailed review of the empirical literature on the effectiveness of virtual agents in different task domains, the reader is referred to Dehn and van Mulken (2000).

CHAPTER 3: ARTIFICIAL PERSONALITY FOR VIRTUAL AGENTS

Actions, looks, words, steps, form the alphabet by which you may spell characters: some are mere letters, some contain entire words, lines, pages . . .

Lavater

3.1 SUMMARY

In this chapter, the social constructivist's dramaturgical model of personality is used to define and to delineate the field of artificial personality. The dramaturgical model views personality from three perspectives: that of the actor (which is concerned with the expression of personality), that of the observer (which is concerned with the perception of personality) and that of the self-observer (which is concerned with the negotiation of personality). Not only does the dramaturgical model of personality shed light on present research concerns in artificial personality, but it also points out new areas for exploration. In the area of artificial personality, most research to date has focused on the perspective of the actor. The observational components of personality have received far less attention. Recent research has examined ways in which agents can observe users, but most of this work has been concerned with the agent observing the user's emotional state rather than the user's personality. In addition, agents have been developed which abduct the personalities of other agents, and researchers have begun investigating how users

attribute personality to computers. No work to date, however, has explicitly addressed modeling the observational perspectives.

3.2 INTRODUCTION

In order to develop virtual agents that possess something resembling a personality, it is first necessary to understand what is meant by the use of the term. *Personality*, like *person*, derives from the Latin *persona*, a mask used by actors, and the Etruscan *Phersu*, a mask or masked person (Neilson, 1955). *Personality* still retains this sense of being something of a social mask or the *persona* an individual presents to others. This is reflected in everyday discourse where the word is often used to assess a person's social competence. A person with a *great personality* or *lots of personality* is considered socially engaging and desirable, whereas a person with *no personality* is regarded as socially inadequate (Mayer and Sutton, 1996; Pervin and John, 1997). *Personality* also commonly refers to the salient characteristics or traits of a person, as when a person is said to have a *shy* or *withdrawn* personality (Mayer and Sutton).

The use of the term *personality* among psychologists is more complex and reflects a variety of underlying assumptions about human nature. For Freud, a person's personality remains largely hidden and is governed by unconscious desires and conflicts. For Rogers

(1955), a leading proponent of the humanist approach to psychology, personality is an organized and consistent pattern of perception, and for Allport (1961), a pioneering trait psychologist, personality determines an individual's characteristic behaviors and thoughts. Although each of these definitions offers unique insights into the nature of personality, they all share assumptions that are common to most psychological characterizations. First, they assume that a person is in possession of a personality. Thus, the fundamental unit in the study of personality is the individual. Second, personality is seen as arising from within a person. It is the product of such internal factors as drives, consciousness, traits, and genetics. Finally, it is assumed that personality shapes both the behaviors and the experiences of a person (Hampson, 1988a).

To view personality almost exclusively in terms of the person follows naturally from the Western emphasis on individualism (Triandis, 1989), but it is not the only perspective available. Alternatively, personality can be conceptualized in terms of a collective interpersonal process (Hampson, 1995). The social constructivists in particular acknowledge the importance of a social component in the formation of personality. They believe that personality is not only the product of a variety of internal factors but also a social construction composed of informal lay notions. This inclusion of a lay or social perspective is recognized by the constructivists as providing more than a new theoretical vantage point; it literally represents a separate observational standpoint from which to view personality, one that originates from outside the individual (Hampson, 1995). To

stress and to illustrate these different standpoints, the social constructivists offer a dramaturgical model that defines personality from three perspectives: that of the actor, that of the observer, and that of the self-observer (Hampson, 1995). Personality from the perspective of the actor involves everything associated with the development of individual personality (genetics, traits, and personal experiences) and the ways in which these factors influence behavior and attitudes. It is the proper study of psychologists. Personality from the perspective of the observer deals with the interpretation of personality. Observers infer personality by observing an actor's behaviors, appearance, and possessions. This perspective is concerned with lay theories of personality, a subject addressed by social psychologists. Finally, the perspective of the self-observer allows an individual to understand and consciously to manage the impressions his personality makes on others. This is often an area of concern for therapists and social workers (Hampson, 1988a). As the constructivists' dramaturgical model is useful in delineating the domain of personality for virtual agents, more will be said about these three perspectives in section 3.3.

In general, computer models of personality have focused on the internal or psychological aspects of personality as the history of these models demonstrates. In the 1960s, psychologists were the first to produce computer simulations of personality. They were motivated by the prospect of exploring concrete embodiments of personality theories and of modeling the complexity of human behavior for educational and therapeutic purposes

(Loehlin, 1968; Tomkins and Messick, (1963). Early models ran the gamut of what was then loosely referred to as *personality*: attempts were made to simulate affects, drives, attitudes, language, and thoughts (Loehlin, 1968). The importance of problem solving and pattern recognition in modeling personality was recognized and some attempt was made to address these broader aspects as well (Loehlin, 1968; Tomkins and Messick). The simulations produced at this time ranged in complexity from systems that exhibited simple affinities and aversions to those that attempted to model more complex belief systems. *Aldous* (Loehlin, 1962, 1968), an early example of a simple system, provided a basic set of emotions: one that was positive (a force of attraction), and two that were negative (fear and aggression). An additional set of hardwired attitudes towards specific categories of perception was used to determine reactions to novel encounters. *HOMUNCULUS* (Abelson, 1963; Abelson and Carroll, 1965) was typical of a more complex system in its attempt to model the involvement of personality in higher cognitive functioning. It read in statements and either rejected, accepted, or distorted them based on the current belief system. Despite an initial enthusiasm and a modicum of success, interest in these systems quickly waned along with the early, overly optimistic promises of AI.

Aside from the pioneering work in the 1980s of Braitenberg (1984) and Carbonell (1980; 1981), it wasn't until the early 1990s, that a wave of renewed interest in modeling personality surfaced, this time in the interdisciplinary area of computer-generated

characters. Spurred on in part by the rapid development of powerful PCs, advances in computer graphics, and new work in AI, especially in reaction planning, constraint systems, and knowledge representation, the 1990s saw a flourishing of a whole host of virtual agents: avatars (Benford, Bowers, Fahlén, Greenhalgh, and Snowdon, 1997), synthetic actors (Trappl and Petta, 1997), believable agents (Loyall, 1997a), interface agents (Laurel, 1990a; Maes and Wexelblat, 1996), non-playing characters in computer games, chatter bots (Leonard, 1997), and virtual humans (Trappl and Petta). Users have typically welcomed these characters and delighted in the prospect of interacting with them. More often than not, however, virtual agents have disappointed by being too mechanical, unresponsive, and inconsistent in their behaviors (Loyall). The resurgence of interest in creating computer models of personality, or what has recently been referred to as *artificial personality* (Trappl and Petta), has been motivated primarily by the need to breathe life into virtual agents.⁴

As is the case with the early psychological models of the 1960s, artificial personality for virtual agents has focused on the perspective of the actor, that is, on the internal aspects of personality, with emphasis being placed on generating behaviors that are expressive of

⁴ The creation of more life-like characters is not the only motivation behind this resurgence of interest in artificial personality. Other motivations include the following: to improve agents by making them socially more acceptable and understandable (Moffat, 1997; Reeves and Nass, 1996), to make user interfaces friendlier (Laurel, 1990b), to instill predictors or to limit undesirable or immoral behaviors (Khan, 1995), and to model human personality (Moffat).

personality. Section 3.4 provides an overview of research into artificial personality from the point of view of the actor. It reviews methods for expressing personality through a broad range of agent behaviors: for example, personality quirks (Loyall and Bates, 1991), roles (Reilly, 1997), goal selection (Rizzo, Veloso, Miceli, and Cesta, 1999), animation (Perlin, 1995), and emotional display (Reilly).

While most research into artificial personality has concentrated on the internal or psychological aspects of personality, there is a growing recognition that the task of providing virtual agents with personality differs from that required to develop viable psychological models. Hayes-Roth, van Gent, and Huber (1997), for instance, in contrasting the goals of psychologists with those involved in the development of virtual agents, note that psychologists seek personality models that are as general and as complete as possible and that are capable of explaining a broad spectrum of human behaviors. In contrast, the goals of artificial personality are more modest and artistic (Bates, 1992a). What is required is to model a subset of those psychological factors that most economically and convincingly convey the essence of a character. According to Hayes-Roth et al. (1997), the concern in artificial personality is mostly with characterization, not psychological fidelity. Thus, Hayes-Roth et al. (1997), Loyall (1997a), and others (Bates, 1992a; Reilly, 1997) have found it useful to explore the techniques traditional media artists employ in creating the *illusion* of personality.

In the traditional arts, personality has always played a vital role in characterization. In fact, character is often defined in terms of personality. *Webster's Third New International Dictionary of the English Language*, for instance, defines *character*, as it is applied in the arts, as the *depiction* of personality: it is “personality as represented or realized in fiction or drama” or as “a given representation or realization of this kind.” (Gove, 1986, p. 376) In drama, *character* can also mean “the personality or part which an actor recreates” (Gove, p. 376).

Writers, artists, and actors, however, are not the only ones who participate in the creation of character. Media artists of all stripes are well aware that spectators contribute by suspending disbelief, emotionally empathizing, and fleshing the character out using their imaginations. As Katherine Hepburn once observed, “Given the chance, the audience will do half the acting for you” (Hepburn, 1999). Thomas and Johnston (1981), two of Disney’s best animators, echo Hepburn when they write, “The *audiences* will make our little character sad—actually, far sadder than we could ever draw him—because in their mind that character is *real*. He lives in their imaginations. Once the audience becomes involved with your characters and your story, almost anything is possible” (pp. 20-21). The job of the artist, they go on to say, is to do whatever is required to maintain audience involvement.

Despite the fact that many in artificial personality have acknowledged that the audience is no less important in the characterization of virtual agents than it is in the media arts (Bates, Loyall, and Reilly, 1992a; Churchill, Cook, Hodgson, Prevost, and Sullivan, 2000; Isbister, 1994), only recently has the observer been taken into account. What little work has been done from this perspective is presented in section 3.5. Since personality can be observed by both agents and human beings, this section is further broken down into research that investigates users observing agents, agents observing users, and agents observing other agents.

Finally, it is important to note that, despite the similarities, virtual agents significantly differ from media-based characters (Badler, Palmer, and Bindiganavale, 1999). Media-based characters have the advantage because every aspect of their character, from the clothing they wear to the subtle expressions on their faces, is carefully crafted. Every scene is painstakingly worked out in advance. Nothing is left to chance. Moreover, the information the audience is given about a character is highly selective, and this intensifies the impression the character makes upon the audience (Hoffner and Cantor, 1991). A virtual agent, in contrast, is autonomous and interactive, and it must react, as Loyall (1997a) notes, “to situations as they arise without being able to draw directly on the wealth of its creator’s knowledge” (p. 2). Unlike media characters, virtual agents have

the far more complicated task of sustaining the illusion of personality within a highly interactive and volatile social environment (Hayes-Roth, van Gent, and Huber, 1997). Success at this involves more than simply behaving at all times *in character*. Successful interaction requires flexibility and adaptability in personality expression.

According to the constructivists, personality is negotiated in the course of social interaction. People continuously adjust their personalities to fit the situation and the personalities of others (see section 3.3.3). Those who fail to adjust are often ostracized. There is mounting evidence that virtual agents are also expected to participate in this negotiation process (Reeves and Nass, 1996). Research that explores the perspective of the self-observer, at least in so far as the virtual agent adapts its personality to the user, is presented in section 3.6.

3.3 THE CONSTRUCTIVIST APPROACH TO PERSONALITY

As already mentioned, human personality can be characterized in many ways. Numerous psychological theories have been developed to explain personality, but few offer assistance in the development of artificial personality (Moffat, 1997). The multiple perspectives offered by the constructivists in their dramaturgical model of personality are particularly suited to modeling personality for virtual agents. Moreover, the perspectives

of the actor, observer, and self-observer provide fresh insights and suggest new avenues for exploration. Before exploring artificial personality in the light of the dramaturgical model, however, it is necessary to provide a more detailed discussion of each of these perspectives.

3.3.1 THE PERSPECTIVE OF THE ACTOR

To view personality in terms of the actor is the typical perspective of the psychologist who is preoccupied with understanding the internal factors that determine behavior and mental events. Psychological theories of personality abound, but much of modern personality psychology is built around the study of traits (Matthews and Deary, 1998). Trait theories are based on the assumption that a person's personality can be captured in a series of binary oppositions. A major task among trait theorists is isolating a comprehensive set of traits that best measure and explain behavior (Cattell, 1965). The number of traits necessary to describe a person's personality fully, however, is highly debated. Eysenck (1991) maintains that personality can be understood in terms of the *Gigantic Three*: extroversion, neuroticism, and psychoticism. Cattell, Eber, and Tatsuoka (1970) advocate a minimum of sixteen trait measures. Most current trait researchers, however, favor the *Big Five*: openness, conscientiousness, extroversion,

agreeableness, and neuroticism⁵ (Matthews and Deary). Table 3.1 lists the Big Five along with some of their many descriptors.

⁵ Note how the first letter of each trait spells the word *ocean*, a mnemonic often given to assist people in remembering the Big Five.

Table 3.1. The Big Five Personality Trait Factors.

<i>Trait</i>	<i>Descriptors</i>
I. Openness to experience (or culture, or intellect)	intelligent – unintelligent sophisticated – unsophisticated creative – uncreative curious – uninquisitive
II. Conscientiousness	organized – disorganized hardworking – lazy reliable – unreliable responsible – undependable
III. Extroversion	extroverted – introverted dominant – submissive adventurous – cautious sociable – reclusive
IV. Agreeableness	kind – unkind generous – stingy warm – cold unselfish – selfish
V. Neuroticism (or Emotional Stability)	stable – unstable relaxed – tense calm – angry unemotional – emotional

Aside from the problem of determining a comprehensive set of traits, there is another problem that involves an uncertainty regarding what is actually being measured when people judge others using trait descriptors. As Hampson (1995) notes, “When raters judge rates on a series of personality traits, is this the study of the impressions of personality of the rateses (the actors), or is it a study of the impressions of personality held by the raters (the observers)?” (p. 27) The social constructivists do not see this

ambiguity in what is being judged as a major problem but rather as further evidence in support of their contention that personality is constructed out of the dual perspectives of an actor and an observer.

3.3.2 THE PERSPECTIVE OF THE OBSERVER

The observer provides a lay perspective that is concerned with the ordinary impressions of personality. Reflections on people's personalities are a major part of everyday discourse. Most languages, including English, are replete with terms that describe the dispositions of others. An early analysis of the English language, for instance, yielded over 18,000 words used as personality descriptors (Allport and Odbert, 1938). Although these terms can be factored down to the Gigantic Three and the Big Five, thus suggesting a convergence between the perspectives of the psychologist and of the lay person (Matthews and Deary, 1998), investigations into the ways in which observers characterize personality focus primarily on how personality is perceived and interpreted (Hampson, 1988a, 1995).

Personality is inferred by an observer from the actor's behaviors, appearance, and possessions. It is believed that these observables are abstracted and classified into various personality categories. These personality categories are then extended to

descriptions of the actor's personality as a whole. Being late for a meeting, for example, is typically classified as an *unpunctual* behavior, and someone who shows a consistent pattern of being late is considered an *unpunctual* person (Hampson, 1988b, 1995). Once the actor has been categorized, the category serves to both explain and predict subsequent behavior.

Research shows that observers use a small subset of the personality descriptors available to them, a subset that typically reflects traits that the observers possess and value (Hampson, 1988b). This has led some to remark that there are as many lay theories of personality as there are individuals (Kelly, 1955). Despite these individual differences, there is considerable evidence in the person perception literature demonstrating that people actually do share a similar set of general beliefs regarding the categorization of actions and people (Schneider, 1973). As is the case with trait theories from the point of view of the actor, it is important to isolate a comprehensive set of personality categories used by people when categorizing others. Rosenberg (1977) has conducted an extensive study of this subject. Employing free-response methods, he has determined seven broad categories that are used to characterize others: intellectual competence and achievement, maturity, attractiveness, integrity, sociability, concern for others, and psychological stability. Table 3.2 lists these categories along with positive and negative examples. It is interesting to note that Rosenberg's personality categories include one that is solely based on the appearance of the actor, namely, attractiveness. More will be said about

these categories in chapter 5, which uses a variation of Rosenberg's categories in classifying the trait impressions of the face.

The key difference to note in the construction of personality from the perspectives of the observer and actor is that from the perspective of the actor personality *influences* the actor, his behaviors and mental events, whereas from the perspective of the observer personality is *perceived* and *interpreted*.

Table 3.2. Major Categories for Traits and Trait-like Feelings.

<i>Category</i>	<i>Positive Descriptors</i>	<i>Negative Descriptors</i>
Intellectual competence and achievement	Intelligent, talented, knows a lot	Unintelligent, ignorant
Maturity	Mature	Immature, childish
Attractiveness (male and female)	Handsome, beautiful, attractive, shapely, good physique	Ugly, fat, unattractive

Table 3.2 Continued. Major Categories for Traits and Trait-like Feelings.

<i>Category</i>	<i>Positive Descriptors</i>	<i>Negative Descriptors</i>
Integrity	Honest, sincere, truthful	Dishonest, insincere, liar, phony
Sociability	Friendly, warm, witty, social sensitivity	Unfriendly, cold, dull, social insensitivity
Concern for others	Kind, generous, sympathetic	Unkind, unsympathetic
Psychological stability	“Together”	Neurotic
Reprinted from the 1976 <i>Nebraska Symposium on Motivation</i> by permission of the University of Nebraska Press. Copyright © 1977 by the University of Nebraska Press.		

3.3.3 THE PERSPECTIVE OF THE SELF-OBSERVER

The dramaturgical model regards personality as a social construction that combines an actor's performance and presentation with an observer's impression of that presentation. In other words, personality is the product of a negotiation process (Hampson, 1988b). First, an actor presents an image of himself that is understood and accepted to a limited degree by an observer. What follows in the course of interaction is a series of continual adjustments in both the presentation of the actor and the interpretation of the observer. For actors to adjust their presentations to accommodate the observer requires an ability for imagining and assessing the impact personality presentations have on others. This results in the formation of a self-reflective social self, or what Cooley (1964) has aptly

labeled the *looking-glass self*. He explains, “As we see our face, figure, and dress in the glass, and are interested in them because they are ours, and pleased or otherwise with them according as they do or do not answer to what we should like them to be; so in imagination we perceive in another’s mind some thought of our appearance, manners, aims, deeds, character, friends, and so forth, and are variously affected by it” (p. 184).

What is the objective of the actor’s presentations? There is evidence suggesting that people aim to produce impressions of personality that are not only consistent with honest self-assessments (self-verification) but which also conform to ideals (self-presentation) (Deaux and Major, 1987). It is not surprising that most people wish to be seen as positively as possible and adjust their behavior to show their best side. What might be surprising to learn is that people also present themselves in ways that conform to the expectations of the observer’s stereotypes, both positive and negative (Goffman, 1967). In this regard there is the famous description William James (1918) presents of the daily permutations of self as it moves through various social settings:

Properly speaking, *a man has as many social selves as there are individuals who recognize him* and carry an image of him in their mind . . . we may practically say that he has as many different social selves as there are distinct *groups* of persons about whose opinion he cares. He generally shows a different side of himself to each of these different groups. Many a youth who is demure enough before his parents and teachers, swears and swaggers like a pirate among his “tough” young friends. We do not show ourselves to our children as to our club-companions, to our customers as to the laborers we employ, to our own masters and employers as to our

intimate friends. From this there results what practically is a division of the man into several selves . . . (p. 293)

As seen from the quote above, people expect and appreciate participants who at least attempt to accommodate the needs and expectations of others. When a person holds fast to a self-image that strains a situation, tensions mount, and people become noticeably uneasy. Moreover, if this behavior is intrinsic to an individual, that person may quickly find himself socially alienated (Goffman, 1967).

3.3.4 INTERRELATION OF THE THREE PERSPECTIVES

The three perspectives of the actor, observer, and self-observer are interwoven. Since the actor adjusts his personality to the observer, the observer may have as much influence on the actor as have traits or drives. Yet the extent to which an actor is willing and able to manage his self-presentations is predetermined by psychological factors (Bem and Allen, 1974). Similarly, the act of adjusting personality is a behavior that is observable and evaluated by others, yet how it is judged depends in part on the traits possessed by the observer. So the three perspectives are inextricably entangled. Nevertheless, by including alongside the perspective of the actor the observational perspectives, the dramaturgical model succeeds in highlighting both the social complexion of personality and its complexity.

3.4 ARTIFICIAL PERSONALITY FROM THE PERSPECTIVE OF THE ACTOR

Given the three perspectives provided by the dramaturgical model, it is possible to organize current research in artificial personality according to the perspective that dominates the work. As is the case with the psychological study of personality, artificial personality mostly revolves around the actor in that an agent's behavior is instantiated based on internal traits, goals, attitudes, current emotional state or a combination of these or similar *personality components* within the agent. In other words, the concern is with the *expression* of personality.

The level of abstraction involved in the realization of personality provides a convenient classification scheme for models of artificial personality that focus on the actor. At the lowest level, personality is handcrafted in systems that integrate scripted and reactive behaviors. As highlighted in section 3.4.1, a major concern at this level is developing animation techniques that enhance the believability of the agents (Loyall, 1997a). At a slightly higher level of abstraction, personality is expressed as a dynamic interplay between the personality traits of a character and the traits associated with social roles. In particular, the Virtual Theater Project at Stanford University (Hayes-Roth, van Gent, and Huber, 1997) has explored the dynamic relation of personality and roles in their role-

constrained improvisational systems, two of which are discussed in section 3.4.2. Finally, at the highest levels of abstraction, personality is represented as goals or emotions. These systems tend to be more complex because they are oftentimes directed towards offering a deeper, more realistic psychological model of personality. Goal-based systems of artificial personality are presented in section 3.4.3 and emotion-based systems in section 3.4.4.

3.4.1 SCRIPTING SYSTEMS

In scripting systems, personality is authored either very indirectly by the writing of weighted action preferences, plans, and other behaviors for nearly every conceivable situation in which the agent might find itself or more directly by the definition of simple trait profiles for both the virtual agent and a wide range of behaviors. In the later case, these trait profiles are used during run-time as filters to constrain the decision process when selecting behaviors to instantiate (Prendinger and Ishizuka, 2001). No matter the approach taken, a character's personality, as in the traditional arts, remains the expression of an author's vision and talents.

A major goal in scripting systems is ensuring that the animation and rendering techniques employed by these systems enhance rather than subvert the intentions of the authors and

the believability of the virtual agents (Blumberg, 1996; Loyall and Bates, 1993). *Believability* is a term often encountered in this context and refers to the human tendency to suspend disbelief within artistic domains so as to enjoy the portrayal of a separate imaginary being (Loyall and Bates, 1997; Reilly, 1996). Believability has little to do with realism, per se. It is doubtful, for instance, that anyone would mistake Mickey Mouse for a real mouse, but he is certainly believable as a character. A major problem with believability is that it is fragile. This is especially the case for virtual agents where technical limitations in animation often disrupt the suspension of disbelief by drawing attention away from the character to the medium conveying it.

Scripted systems intended to enhance believability loosely fall into one of two categories. First is the line of research that explores animation techniques that contribute to the overall believability of the agents. Considerable effort, for instance, has been devoted to investing virtual agents with humanlike navigation and locomotion (Reich, 1997; Thalmann, Noser, and Huang, 1997; Towns, Voerman, Callaway, and Lester, 1998). Effort has also been directed towards making virtual agents less rigid and mechanical in presentation. Perlin (1997), noting that the bodies of living beings are in a constant state of flux, has developed methods for providing virtual agents with realistic body and facial noise. Other important animation techniques that augment believability include lip-synching (Lewis and Parke, 1987; Porter and Susman, 2000), hand gripping (Thalmann et al., 1997), and a variety of deictic behaviors such as gazing (Cassell, Pelachaud, Badler,

Steedman, Achorn, Becket, Douville, Prevost, and Stone, 1994; Cassell, Steedman, Badler, Pelachaud, Stone, Douville, Prevost, and Achorn, 1994; Vilhjálmsón and Cassell, 1998; Vilhjálmsón, 1997), head nodding, and pointing (Lester, Towns, Callaway, Voerman, and Fitzgerald, 2000; Towns et al.). Yet another important component of believability involves the timing, coordination, and production of composite animation (Goldberg, 1997), or multiple coherent actions (Johnson, Rickel, and Lester, 2000; Loyall, 1997b; Maes, Blumberg, and Pentland, 1997; Towns et al.). These involve the performance of simultaneous actions: walking, for example, while chewing gum and waving goodbye.

A second, more recent line of research, investigates animation and language expressiveness as a medium for individuating personality or expressing specific types. Badler, Reich, and Webber (1997) have explored locomotion styles insofar as they convey emotional and attitudinal information. Cassell, Pelachaud, et al. (1994) and Cassell, Steedman, et al. (1994) have suggested ways that facial animation can express specific traits, and a number of researchers have examined the possibility of conveying an agent's personality through syntactic and semantic choices in utterance and pitch expressiveness (André, Klesen, Gebhard, Allen, and Rist, 1999; Ball and Breese, 2000; Bates, Loyall, and Reilly, 1991; Chin, 1996; Loyall and Bates, 1997; Reilly, 1996; Walker, Cahn, and Whittaker, 1997). André et al. (1999), for example, have programmed their extroverted agents to talk loudly and take the initiative in speaking. Yet another

way to singularize personality is to employ a technique animators are particularly fond of, and that is to provide the agent with quirks of speech (Reilly, 1996). Bugs Bunny, for example, reveals his abrasive personality by greeting people with a slightly irritating, “Eh, whazzup Doc?” and Elmer Fudd exposes his lack of intelligence and incompetence as a hunter with his babyish pronunciations: “Be vewy, vewy quiet! I’m hunting wabbits.” Pauses, restarts, and other breakdowns in language (Loyall and Bates, 1997) also have the effect of singularizing personality, as do dialects (Binsted, 1998). Similarly, quirks in behavior accentuate a character’s individuality (Loyall and Bates, 1991; Reilly, 1996, 1997). Due to a programming error, for instance, a character developed by Bates (1994) ended up with a nervous tic. This was found to greatly increase the believability of the character for those who interacted with it.

3.4.2 ROLE-BASED SYSTEMS

The Virtual Theater Project at Stanford University (Doyle and Hayes-Roth, 1998; Hayes-Roth, van Gent, and Huber, 1997; Isbister and Hayes-Roth, 1998; Rousseau, 1996; Rousseau and Hayes-Roth, 1996) is developing virtual agents that play characters with specific personalities that are then directed to improvise within the narrow confines of stereotypical roles. What is interesting in role-constrained improvisation is the dynamic interplay between the personality of the characters and the roles being played. Some personality traits are better suited to specific roles than are others. Good improvisational

actors, including virtual actors, exaggerate those traits within their character that are relevant to a particular role, be those traits complementary or antagonistic.

A pair of roles that Hayes-Roth, van Gent, and Huber (1997) have explored is that provided by a master and his servant. These roles have clearly defined behaviors: the servant must wait at his station until summoned by the master. Depending on the personality of the character, however, different role-appropriate behaviors might be performed in addition to those prescribed by the role. A particularly subservient character, for instance, might volunteer to do more than is actually required, whereas a more dominant servant might perform the bare minimum of duties, and those with obvious reluctance. Similarly, some masters are brutal, while others are easily intimidated. Extreme mismatches in personality, of course, result in role reversals. There is, in fact, a strong theatrical tradition, played out in many variations, of a servant usurping the position of the master (Johnstone, 1987).

In line with improvisational acting techniques, the personality traits employed by Hayes-Roth, van Gent, and Huber (1997) are not psychological but rather theatrical and relevant to the role. In the master-servant scenario, for instance, traits indicative of status are most relevant, and status can best be conveyed theatrically through demeanor, relationship, and space. *Status in Demeanor* refers to a character's style of comporting

himself: his posture and manner. High status is signified through erect posture and a calm manner. Low status is typically signaled by poor posture and a nervous manner. *Status in Relationship* refers to the way characters align themselves. It is not unusual for high status individuals to ignore others. People with low status, in contrast, typically pay close attention to others, especially those of higher status. *Status in Space* is indicated by a character's willingness to abuse objects and occupy space.

Hayes-Roth, van Gent, and Huber, (1997) have experimented with virtual agents playing these roles using two versions of a scenario entitled "While the Master's Away . . ." In the first version, a servant takes advantage of the master's absence to roam freely about the room until the servant is surprised by the master's entrance, at which point the servant returns to his station. The servant's position or role requires that Status in Relationship remain low, but Status in Demeanor can differ depending on the personality of the servant. A submissive servant might crouch or scurry hurriedly away when caught by the master, whereas a servant of higher status might remain unruffled and remove himself with dignity. Status in Space does not depend on the servant's personality but changes in relation to the master's presence: the servant remains at his station located somewhere in the peripheral when the master is present or roams the space freely when the master is absent. In the first version, both Status in Relationship and Status in Space remain the same for the master; Status in Demeanor, however, can vary to indicate a more or less dominant master. In the second version of the scenario, roles are reversed. In this case, a

servant with high Status in Demeanor maintains high Status in Space, even when the master returns. To convincingly accomplish this role reversal, the master must have a Status in Demeanor that diminishes as the action progresses.

Other improvisational settings investigated by the Virtual Theater Project include a virtual bar (Isbister and Hayes-Roth, 1998) and a cyber café (Rousseau and Hayes-Roth, 1997). In the virtual bar a virtual bartender serves drinks to human customers who are represented by their avatars (Isbister and Hayes-Roth). An interesting finding in this project is that users expected the virtual bartender to possess characteristics stereotypical of a bartender: the interaction logs reveal that the virtual bartender had to initiate all conversations and was required to play both peacekeeper and confidant.

3.4.3 GOAL-BASED SYSTEMS

A number of important aspects of personality cannot be defined by writing low-level weighted action preferences or by constructing trait profiles for agents playing specific roles (Rizzo et al., 1999). Both an altruistic agent and a selfish agent might play the role of bartender and both might be characterized as happy, extroverted, and intelligent, but what differentiates a selfish agent from an altruistic one is in the nature of its goals; an altruistic agent chooses those goals that assist others, whereas a selfish agent is more

concerned with benefiting itself, even at the expense of others. An explicit, goal-based model of personality is required to model these aspects of personality.

Carbonell (1980; 1981) was the first to develop a goal-based theory of personality. He developed it as a method for automating the analysis of stories. In the systems he devised, variations in personality are deviations from a culturally normative person who is represented using a goal-tree. This standard tree assigns averaged priorities to a set of goals. Different personality types are recognized by variations from this norm as well as by the intensity of their reactions to goal successes and failures.

A goal-based system developed for believable agents that was inspired by the work of Carbonell is that of Rizzo, et al. (1999). Unlike Carbonell, however, their work does not represent personality as deviations from a prototypical person. Instead, they model personality using the flat, psychological taxonomy of goals developed by Ford (1992). In Ford's system, abstract goals, such as the goal of entertainment, subsume more specific goals pursued by such behaviors as avoiding work or going out. Focusing on help-giving behaviors, Rizzo et al. employ the following relevant goals in Ford's taxonomy: resource provision, resource acquisition, social responsibility, belongingness, image, entertainment, safety, and understanding (Rizzo et al.). The personality of an agent is expressed at a high level by assigning different priorities to each of these seven abstract

goal clusters. During run-time, a plan-based architecture that includes a generative planner as well as a reactive planner automatically produces behaviors that reflect the agent's personality.

In the experiments performed by Rizzo et al. (1999), five different personalities are defined in terms of help-giving: *the altruist* who cares about others and serves their interest even when it is not to its advantage to do so, *the normative* who is willing to help others as long as such behavior is appropriate to some social norm, *the selfish* who helps others only if doing so offers opportunities for self advancement, *the spiteful* who enjoys interfering with the plans of others, and *the suspicious* who is afraid of exploitation and usually refuses help if the motives of other agents are not clear. As an example of how Ford's taxonomy is used in this system to represent and express personality, consider the selfish agent. It prioritizes resource acquisition and image and devalues resource provision and social responsibility. As a result, a selfish agent performs those behaviors that enhance the agent's image and maximize its resource stores.

Not all aspects of personality can be modeled using a goal-based approach. Typically not covered are emotional and attitudinal attributes. It is difficult to define, for example, meekness or moodiness in terms of goal strategies, although reactions to goal

attainments, as alluded to above, can be accounted for by goal-based systems (Carbonell, 1981).

3.4.4 EMOTION-BASED SYSTEMS

Emotion is oftentimes conflated with personality. For developers of virtual agents, providing an agent with an emotional system is considered equivalent to providing an agent with a personality. This is because emotion enlivens an agent and increases believability by adding social interest (Bates, 1994). In systems where the two are distinguished, either personality is defined entirely from within a model of emotion, as is the case with Bates, Loyall, and Reilly (1992a), Moffat (1997), and Reilly (1996), or personality profiles are used to bias emotional expression in the same way that these profiles bias other behaviors, as with André, Rist, and Müller (1998) and Prendinger and Ishizuka (2001).

For Moffat (1997) personality and emotion differ only in duration and focus. Thus, each can be defined in terms of the other: emotion can be viewed as a momentary perturbation in personality that is triggered by specific events, and personality can be characterized as emotion that is unfocused and time invariant. Since several computational models of emotion are already in development (most notably the **Artificial Concern REalization**

System [*ACRES*] [Frijda and Moffat, 1994; Swagerman, 1987] and the *Affective Reasoner* [Elliott, 1992]), Moffat suggests that artificial personality be modeled in terms of emotion.

Extending the theoretical ideas of Frijda (1986), Moffat (1997) explores personality, as it is related to emotion, in a system he has developed called *Will*, so named because it is an autonomous agent with motivations. *Will* is composed of a number of modules: a Perceiver, a Planner, an Executor, a Predictor, a Charge, and an Emotor. The Charge module attaches a concern-relevance value (an emotional charge) to concerns, which are simply items stored in memory. When a concern has enough charge it becomes the center of focus. The Emotor module consists of a set of appraisals and action tendencies. According to Frijda, it is the appraisal categories that produce the various emotions. *Will* includes the following appraisal categories: Valence (the degree of pleasantness of the situation), Un-/Expectedness (the degree a perceived event was expected to happen), Control (the extent the agent can manipulate the situation), Agency (the agent responsible for the event), Morality (the moral ramifications of the event), Probability (the likelihood of another event occurring), and Urgency (the importance of reacting to the current situation). The Emotor module consists of a set of action tendencies. These concern the effects of emotion: some emotions impel agents to approach, flee, or fight, some hurt or help others, and some strengthen or weaken a resolution to accomplish goals.

How is personality defined and expressed in Will? For Moffat (1997) personality is “a consistent reactive bias” within what he calls “the fringe of functionality” (p. 134). In other words, personality consists of a set of biased parameters within a reactive (emotional) system. The extreme points of bias are calculated by determining the values at either end of the parameter that result in system failure. The values in between constitute the “fringe of functionality.” Within the Charge module, the number of concerns, their nature, and relative importance are variable parameters. Within the Emotor module, the number and nature of the appraisal categories and action tendencies are variable, as are the relative strengths of the action tendencies. Moffat contends that as these parameter biases vary, so does the personality of Will.

Moffat (1997) admits, however, that modeling personality affectively, as he suggests, does not account for certain aspects of personality that might be important. Consider for a moment what emotion best reflects intellectual competence. It might be argued that intellectuality is a generalized dampening of emotional expressiveness or that intellectual competence is not in fact a personality trait but something else. Would the same argument hold for integrity? If integrity is considered a personality trait, then what generalized mood would best characterize an honest individual? According to Moffat, Will can account for only half the traits in the Big Five: conscientiousness, neuroticism,

and possibly extroversion. It is interesting to note Moffat's explanation of Will's failure to express the traits of openness, agreeableness, and, to a limited degree, extroversion: Will fails because it lacks an awareness of other agents as social beings, that is, Will is missing the *observational component* of personality.

Another theoretical model of emotion is that developed by Ortony, Clore, and Collins (Ortony, Clore, and Collins, 1988), hereafter referred to as the *OCC model*. It has successfully been used in the development of emotional systems for a number of virtual agents (Bates et al., 1992a; Elliott, 1992; Lee, 1999; Reilly, 1996). It is especially attractive to the agent community because it offers an intuitive explanation of the generation of emotion and an easy method for computing emotional intensities (Lee). Emotions in the OCC model are the result of appraisals of events (as they relate to a person's goals), the self (as it relates to others and according to a set of standards), and objects (as they relate to a set of basic attitudes).

Few have attempted to express personality from within an OCC model. Some who have experimented along this line include Lee (1999) and Martinho and Paiva (1999). Perhaps the most comprehensive model of personality developed from an OCC model, however, is that of Bates, Loyall, and Reilly (1992a) and Reilly (1996). In their emotion system, called *EM*, emotions are constructed, following the OCC model, from goals, the approval

of others, other emotions, and hardwired attitudes. Happiness and sadness, for instance, are generated as a result of goal attainment or failure, while fear and hope are based on a belief in the goal's success or failure. Other emotions, such as pride, shame, reproach, and admiration are based on whether the action is approved or disapproved by other agents. This is determined by an agent's standards. Anger, gratitude, and remorse are composites of other emotions. For example, gratitude is a combination of happiness and admiration. Finally, love and hate are based on attitudes that are initially hardwired but adaptable in the course of interaction.

In EM, personality is expressed, as is the case with Will, by manipulating the emotional parameters within the "fringe of functionality." The importance of specific goal attainments and the nature of the preset attitudes and standards can vary as can the actions selected given the presence of a particular emotion. Fear, for example, can be expressed by fleeing, fighting, or freezing (Bates, Loyall, and Reilly, 1992b). In addition, the propensity to react emotionally and in specific ways can differ (Reilly, 1996). Finally, an agent's estimate of the likelihood of success in obtaining a goal can be skewed optimistically or pessimistically (Bates et al., 1992b).

Of course, emotion can be seen as one behavior among many that can be constrained by personality. This is the approach taken by André et al. (1998), Poggi and Pelachaud

(2000), Stern, Frank, and Resner (1998), and Prendinger and Ishizuki (2001). Most successful, as far as users are concerned, is the Petz product line of *P.F. Magic* (Stern et al., 1998). *P.F. Magic* endows virtual cats and dogs with distinct personality profiles that filter their emotional expressiveness. It is typical in these systems for trait profiles to control the intensity and decay rate of each emotion as well as the emotion activation threshold. They differ from systems like Will in that personality is not expressed exclusively in terms of emotion but rather through a number of behaviors. Poggi and Pelachaud, for example, have developed a system where personality profiles control both goal-selection and emotional expression.

3.5 ARTIFICIAL PERSONALITY FROM THE PERSPECTIVE OF THE OBSERVER

There has been an ongoing shift in focus away from the perspective of the actor towards the observer. Bates (1992a) and Laurel (1993) have led the way by taking inspiration from the arts. Laurel uses the metaphor of a theatre in her discussions of interface agents. This metaphor includes an audience along with an actor. As mentioned earlier, Bates (1992b) has examined how Disney animators breathe life into their characters in an attempt to understand what makes them believable, and Disney animators are concerned less with realism than with the impression of realism. Here is a lesson they write about learning in the making of *Bambi*, “If we had drawn real deer in *Bambi* there would have

been so little acting potential that no one would have believed the deer really existed as characters. But because we drew what people imagined a deer looks like, with a personality to match, the audience accepted our drawings as being completely real” (Thomas and Johnston, 1981, p. 332). A repeated theme in the arts is that what matters in the portrayal of character is the response of the audience.

In a similar vein, Isbister (1994) has advanced the observational perspective in her research on intelligent agents by noting that the *perception* of intelligence is as important as the inner workings of a brain. She writes, “. . . assessing whether a computer agent is *intelligent* is not a matter of deciding whether it has human intelligence or mimics human neurological functioning. Rather, it is a matter of assessing whether the agent possesses the cues that lead a person to label someone/thing as *intelligent*” (paragraph 54). It is a short step to extend this idea to any other trait an agent might possess. But it is Churchill et al. (2000) and de Rosis and Castelfranchi (1999) who have explicitly introduce the observational perspective in artificial personality. Churchill et al. have used Hampson’s metatheory of personality to outline an approach that includes the observer. In this regard, they have proposed that agents pass the “lay personality psychologists test.” Likewise, de Rosis and Castelfranchi have noted that personality is both *generated* and *recognized*, but they have gone beyond merely recognizing the role of the observer to actually developing agents that are capable of assessing the personalities of other agents (Castelfranchi, de Rosis, and Falcone, 1997).

Having an agent observe another agent's personality, however, is only one possibility. There are two others: an agent can observe a user and a user can observe an agent. Reviewed in this section is research related to artificial personality that entails all three observational viewpoints.

3.5.1 AGENT OBSERVING AGENT

In any given situation, a rational agent decides to behave by reasoning about its own mental states, or first order beliefs and goals. A socially intelligent agent considers as well the states of other agents, or second order beliefs and goals. Although some researchers have created social agents that do not represent others, the behaviors of these agents tend to be elementary (Mataric, 1992; Wavish, 1992). To interact in humanlike societies, to form relationships, and to have satisfying social encounters, agents need to reason about other agents.

Castelfranchi, de Rosis, and Falcone (1997) have developed a multiagent system called *GOLEM* where agents, in addition to having a number of social capabilities, have distinct personalities. *GOLEM* formalizes and implements different kinds of social cooperation in an environment where the agents work together to transform a world. Agents in

GOLEM have a goal, knowledge of the present state of the world, know-how about world actions, and a personality that is defined in terms of a number of delegation behavioral stereotypes: *lazy*, *hanger-on*, *delegating-if-needed*, *never-delegating*, *hyper-cooperative*, and so forth. Each agent is entitled to perform world-transforming and world-controlling actions. They can also perform communication acts. At the beginning of play, the agents introduce themselves to the community by describing their know-how and goals. These introductions may also include information regarding the agent's personality. If such information is omitted, other agents will reason about the agent's personality by observing its behaviors. Castelfranchi et al. (1997), provide the following example of the reasoning process an agent might go through to abduct the personality of a hyper-cooperative agent: "If the other agent accepts to perform a delegated action and I know that there is a surface-conflict between that action and what I presume (or know) to be its domain-goal, then I may assume that it is a hyper-cooperative" (p. 20). Knowledge of an agent's personality is useful in abducting an agent's plans and capabilities. For example, an agent could reason that if another agent is presumed to be a hyper-cooperative and if it refuses a delegation, then it can be assumed that that agent is not able to perform the action.

Another approach is that taken recently by Gmytrasiewicz and Lisetti (2001). In their system, personality is characterized in terms of emotions and their transitions; specifically, personality is defined as a finite state machine that consists of a set of

emotions, environmental inputs, and an emotional transformation function. Agents, given an initial state and an environmental input, can use this model to probabilistically predict an agent's emotional state; and, over time, the personality of an agent can be learned.

3.5.2 AGENT OBSERVING USER

Virtual agents, by the very fact that they are embodied, are expected to behave in socially appropriate and intelligent ways. Successful social interaction with a user requires that a virtual agent observe and respond to the user. At the very minimum, the agent must keep track of the user's location in order to direct its gaze appropriately when speaking. Tracking the user's eye gaze also helps choreograph conversational turn taking and attentiveness (Johnson et al., 2000). Truly satisfying social interactions, however, require that virtual agents do more than physically track people.

One area of intense interest is in the development of virtual agents that recognize the user's emotional expressions. Several research groups have experimented with attaching sensors to users that map physiological states, such as blood pressure and galvanic skin responses, to emotional states (Picard, 1997). Ark, Dryer, and Lu (1999) have succeeded in discriminating emotions less obtrusively by attaching the sensors to the user's mouse.

Others have investigated the visual recognition of emotional facial displays (Essa and Pentland, 1997; Pantic and Rothkrantz, 1997; Vyzas and Picard, 1998). Visual systems are probably preferable as they do not rely on physical contact and observe users more naturally.

Only recently, however, have virtual agents been designed to recognize a user's personality. Ball and Breese (2000) are the first to have explored this possibility. Their systems recognize a number of emotions and personality types as these are expressed along two dimensions: dominance, measured by a user's disposition towards controlling or being controlled, and friendliness, measured by a user's tendency to be warm and sympathetic.

By and large, Ball and Breese have concentrated on recognizing personality through language. Within a Bayesian network, the following user behaviors are represented: paraphrase selection, base pitch, pitch variability, speech speed, and energy. The exact input settings that reflect different personality types and emotional expressions are predetermined by the investigators, whose decisions are informed by psychological studies that have examined personality and emotion as it is expressed in language. Positive word choices, for example, are associated with friendliness; loud speech and talkativeness are indicative of extroversion. The researchers also mention that their

systems are capable of recognizing visual indicators of emotion and personality, as revealed, for example, by posture and emotional facial expressions, but thus far no details have been presented regarding this aspect of their systems. Nonetheless, the program they outline takes a good first step towards furnishing an agent with an observational model of a user's personality.

3.5.3 USER OBSERVING AGENT

It is generally recognized that users respond to computers as social entities—that *Computers Are Social Actors (CASA)* (Reeves and Nass, 1996). The CASA paradigm claims that any social science finding which concerns human-to-human attitudes or behaviors holds for human-to-computer interactions (Nass, Isbister, and Lee, 2000). There is now ample evidence in support of CASA. Reeves and Nass, for instance, have taken a number of social science studies, substituted the word "computer" for "human" in the original hypothesis, replaced one or more humans with computers in the method of the study, provided the computer with characteristics associated with humans (for example, language, human roles, a voice, or some other form of embodiment), and then performed the same experiments to determine whether the social principles still hold. In case after case, their subjects, many of whom were sophisticated computer users, applied exactly the same social rules of conduct to computers as they did to people.

Noting the importance of personality in traditional media and psychology, Nass et al. (2000) have applied the same methodology in determining how users judge and behave with interactive computer characters that exhibit personality. In one study, both verbal and nonverbal cues (posture, body movement, word choice, and sentence structure) indicative of extroversion in human beings were displayed by virtual agents. Their subjects had no trouble successfully labeling the introverted and extroverted cues exhibited by the interactive characters. Nass et al. have also found that users respond to the personalities of computers in predictable ways. Just as people prefer others who have a personality that is similar to their own, so users prefer a computer with a complementary personality (Reeves and Nass, 1996). Users are also attracted to computer interfaces that exhibit the same traits that make people more agreeable. Computer interfaces that flatter, for instance, are particularly liked (Reeves and Nass), as are computer interfaces that exhibit a sense of humor (Morkes, Kernal, and Nass, 2000).

3.6 ARTIFICIAL PERSONALITY FROM THE PERSPECTIVE OF THE SELF-OBSERVER

As noted in the last section, the evidence is overwhelming that the human-computer relationship is fundamentally social, and the more the computer interface is embodied, the more it is expected to behave in socially appropriate ways. (Ball and Breese, 2000; Nass, Steuer, Henriksen, and Reeder, 1993; Reeves and Nass, 1996). Insofar as the

expression of personality is concerned, most people adjust their presentations to accommodate the personalities and needs of others as well as to fulfill the requirements dictated by social roles. Are embodied agents expected to conform to this practice as well? Reeves and Nass have produced strong evidence in support of such a claim. They find that interfaces that attempt to adjust their presentations to suit the personality of individual users are consistently judged more favorably. Even in situations where the agent fails to produce the desired effect, as long as it exhibits an inclination to adapt, users give it credit for trying (Reeves and Nass).

Research that explores the adjustment of a virtual agent's personality in social settings—as with research stemming from the observational perspectives in general—is sparse. But the need for such research is evident. In the human-computer interface community, one area of research activity concerns user interface preferences. It has been found, for instance, that users are affected differently by animated agents (Rickenberg and Reeves, 2000). Some users prefer them, some loathe them, and some need them. There has been found, in fact, a strong correlation between the user's personality type (extroverted versus introverted) and the preference for virtual agents. Similarly, Resnick and Lammers (1985) have shown that users with low self-confidence are more likely to need a humanizing interface.

Few, however, have attempted to incorporate real time observation of the user in the development of virtual agents that then adapt their personalities to suit the user. Ball and Breese (1999), as noted above, are the first to have developed virtual agents capable of diagnosing the user's personality; and, although Ball and Breese have reported informally experimenting with agents that match the user's personality or even oppose it, the effect this has on human-computer interaction has not been investigated. Likewise, De Carolis, de Rosis, and Pizzutilo (2000) have experimented varying their agent's helping style based on the user's personality in several of their documentation systems. In their systems, adaptation to the user is based on personality-related conditions that trigger either a task-oriented approach or an object-oriented approach. Novice users, for example, are provided a dominant, task-oriented agent, whereas expert users are presented a submissive, object-oriented agent. As with Ball and Breese, De Carolis et al. have yet to evaluate the effectiveness of their system.

Finally, it should be mentioned that the virtual agents in the master-servant scenarios developed by Virtual Theater Project at Stanford University (Hayes-Roth, van Gent, and Huber, 1997) and reviewed in section 3.4.2 negotiate personality to a limited degree—but not with the user. Their agents perceive the Status in Demeanor in other agents, and this perception affects subsequent displays of Status in Relationship and Status in Space.

3.7 SOME PROBLEMS WITH CURRENT RESEARCH AND THE NEED FOR THE OBSERVATIONAL PERSPECTIVES

In the development of artificial personality, it is essential that the perspectives of the actor, the observer, and the self-observer be taken into account. Isolating any one perspective at the expense of others causes problems and limits possibilities. That this is so is evident in the current state of affairs where many troublesome issues are a direct result of a one-sided concentration on the actor.

One problem that arises when the observational perspectives are not taken into consideration is the production of agents that are inflexible and insensitive. Like Archie Bunker, these agents may express distinct traits, a clearly recognizable personality—and for that reason they may be considered successful—but they may also upstage other people, making them feel belittled and uncomfortable. People are not asked to interact with an Archie Bunker, and they probably would not enjoy it if they were, but they are asked to interact with virtual agents. Socially adept agents, like socially adept people, know how to step back when needed and blend in with the crowd.

But perhaps the most serious problem today concerns evaluation methodology, or rather a lack of it. Many have noted that system assessment is more often than not overshadowed by the eagerness of developers to describe their systems and to unveil their agents (De

Carolis et al., 2000; Nass et al., 2000). Typically, an evaluation of the effectiveness of these systems is left to some future study or reported anecdotally. In those cases where an evaluation study is conducted, it is oftentimes unclear what precisely is being assessed. This is especially the situation with authoring systems. When users judge the personality expression of a given agent, for example, one may well wonder whether the user is actually evaluating the system as a whole or whether the user is evaluating the talents of an individual author. To a large extent, the problem with system evaluation boils down to the fact that standard experimental research protocols are not being employed (Nass et al.). Proper evaluation requires an assessment of user responses to multiple virtual agents, and in the case of authoring systems, virtual agents created by more than one author. The key idea here is that *proper evaluation requires the inclusion of the observer*. What is essential to realize is that this oversight on the part of researchers to evaluate their systems reflects more than an overriding enthusiasm for agent development: it reveals once again that what is not being fully appreciated is the vital part the observer must play in the *construction* of personality.

Finally, even when the importance of the observer is recognized, artificial personality is rarely designed with the observational perspectives explicitly in mind and, thus, the full potential offered by the observational perspectives is not recognized. This is unfortunate, as the observational perspectives are rich in research possibilities as demonstrated in the

work reviewed in sections 3.5 and 3.6, as well as in the next chapter, where the benefits of modeling the perception of physical personality are outlined.

CHAPTER 4: MODELING THE PHYSICAL PERSONALITY OF THE FACE

Because society is never a *disembodied spectacle*, we engage in social interaction from the very start on the basis of sensory and aesthetic impressions. The look of the other person is the *prima facie* ground of our knowledge of him or her.

O'Neil

4.1 SUMMARY

Although actors convey personality through a multitude of expressive modalities, there is one particular manifestation of personality that immediately impresses itself upon the observer. It is a facet of personality that might best be labeled the *physical personality*, since it comprises those aspects of appearance that produce an initial impression of personality, along with a concomitant set of reactions and expectations in observers. Unlike psychological models of personality that are focused on the actor, models of physical personality require that attention be placed squarely on the observational perspectives, or the *perception* of personality. Discussed in this chapter are some reasons virtual agents should be endowed with the ability of perceiving physical personality, especially the trait impressions of the face. The chapter ends by listing four requirements that a model of the physical personality of the face must satisfy if it is to accommodate the specific needs of virtual agents.

4.2 INTRODUCTION

The review of the reference literature in artificial personality that is provided in chapter 3 reveals that most research thus far has taken the perspective of the actor. An actor expresses personality in a variety of ways. As Goffman (1969) has remarked, “Individuals, like other objects in this world affect the surrounding environment in a manner congruent with their own actions and properties. Their mere presence produces signs and marks. Individuals, in brief, exude expressions” (p. 4). A major research endeavor in artificial personality has been to furnish virtual agents with a means of performing actions within a virtual world that exude signs and marks expressive of a *coherent* individualism. The literature review clearly demonstrates that within the last decade significant advances have been made in this direction.

As argued in chapter 3, however, artificial personality is best characterized as a construction composed equally of all three perspectives: that of the actor, that of the observer, and that of the self-observer. Research that is centered on the observer and self-observer, however, is not focused on the expression of personality as much as it is on the attribution process. Consequently, the questions that need to be addressed from the observational perspectives are concerned with how an observer forms an impression of personality. One method of arriving at a judgment of an actor’s personality, as described in chapter 3, is to reason about the behavior that is being observed (Castelfranchi, de Rosis, and Falcone, 1997). Reasoning is especially useful in understanding actions

within complex social settings. But there is also an approach that is best suited to situations where little information is known about the actor, aside from what is immediately discernable in the actor's voice and physical appearance. It is an approach that is based more on the immediate apperception of personality. As the sociologist O'Neill (1985) explains, "It is through our senses that we first appreciate and evaluate others, immediately shaping our own positive, pleasurable, and trusting responses, or else our negative, fearful, and avoiding reactions. What we see, hear, and feel of other persons is the first basis for our interaction with them" (p. 22). This aspect of personality, stressing as it does the perception rather than the expression or possession of personality, might best be called the *physical personality*, as opposed to the inner *psychological personality* of the actor.

The term *physical personality* is actually one that is borrowed from drama theory (Strickland, 1956) where it refers to the casting limitations imposed by an actor's physical type. The Central Casting Corporation of Hollywood, for example, would hardly label a fat balding middle-aged man a *Romeo*, no matter how romantic he might profess to be *deep inside*. An actor's body type circumscribes the dramatic roles it is possible for him to play. Behind this superficial dramaturgical notion of physical personality, however, lies something much more profound, and that is the general tendency of people to read character into a person's physical form. The idea that character is reflected in the body is universal and goes back to earliest antiquity.

Attempts to decode morphological features and physical blemishes as signs of character are found in the fables, proverbs, and histories of cultures as diverse as the Egyptian, the Indian, the African, and the Chinese. It is an idea that is deeply embedded in Greco-Roman and European culture. Although often scoffed at and refuted, even in the modern world, the old Aristotelian notion that the soul is somehow mirrored in the body consistently reappears as a topic of scientific, philosophic, and aesthetic discourse. As the historian Frey (1993) has recently observed, “To this day, the quest to read a person’s inner world from her outer appearance has lost nothing of its momentum. In fact, judging from the number of publications that have appeared in the field of nonverbal communication since the 1950s it seems that the advent of the ‘Age of Television’ has given additional impetus to the age-old fascination with human appearance” (p. 64).

The purpose of this study is to address the observational stance by modeling the perception of physical personality. Although recently several studies and position papers have noted that the physical appearance of virtual agents plays a significant role in human-computer interaction (Donath, 2001; Sproull, Subramani, Kiesler, Walker, and Waters, 1996), no researcher to date has suggested that virtual agents be provided with a

means of perceiving physical appearances in terms of the impressions they create⁶ or that an observer's impressions of personality be directly modeled. The goal of this study is to begin investigating this perceptual possibility. As physical personality covers a wide range of features—body type, posture, hair styling, eye color, and skin textures—the scope of this study is narrowed to an examination of only those impressions of character that the morphology of the face elicits. For not only does the face reveal evidence regarding the age, sex, physical condition, and current emotional state of an actor, but there is also compelling evidence in the person perception literature that facial morphology provides observers with a plethora of clues, whether accurate or not, regarding an actor's personality, disposition, and attitudes (Zebrowitz, 1998).

The person perception literature has much to say about the trait impressions of the face. It is a topic that has produced within the last couple of decades a considerable body of research. As this literature is of value to the present study, it is reviewed in the next chapter. The remainder of this chapter is divided into two sections. Section 4.3 discusses some of the advantages a model of the trait impressions of face might offer virtual agents.

⁶ As reported in chapter 3, Ball and Breese (2000) have attempted to develop agents that *perceive* personality by analyzing some of the physical characteristics of the user's voice. It should be pointed out, however, that their research effort was directed at arriving at a more or less accurate diagnosis of a speaker's personality. The *impressions of personality* that people form from pitch patterns and other physical characteristics of the voice were not modeled or investigated in their study.

Section 4.4 concludes the chapter by listing four requirements that a model of the trait impressions of the face must satisfy if it is to accommodate the specific needs of virtual agents.

4.3 REASONS FOR MODELING THE PHYSICAL PERSONALITY OF THE FACE

The primary reason for modeling the physical personality of the face is that such a model would furnish agents with a rudimentary social awareness sufficient enough to allow agents to participate in the construction of personality. Human observers are continuously caught up in the physical appearances of others, and people are equally preoccupied with managing their own appearances. If virtual agents could *perceive* physical personality in ways that mimic human observers, then these *virtual observers* could also respond with a degree of convincing social realism to the physical personality of others. Moreover, virtual agents that are able to see themselves as human observers see them could learn to adjust their facial forms to suit the nature of their roles and their dealings with specific types of individuals. Discussed in this section are some of the benefits the ability to read and to make faces would offer virtual agents.

4.3.1 FACE READING

First, a model of the trait impressions of face would provide agents with a more perceptive basis for dealing with users. Human observers orient themselves in their dealings with others by making use of the information that people bodily advertise about themselves. As pointed out by the sociologist O'Neill (1985), the human body is “a symbolic system whereby members communicate to one another their age, gender, marital status, sexual availability, social standing, and the like” (p. 23). The face, perhaps more than any other part of the body, is especially inscribed with cultural meaning. As discussed below, people rely on the information provided in the physical personality of the face when making sense of ambiguous situations or when meeting others for the first time. If virtual agents were endowed with a perception system capable of perceiving the physical personality of faces, this information could be used by the agents in understanding the intentions of people and in formulating plans of interaction. Rather than predefine an initial set of interaction tactics and practices within an agent, for example, the cultural information visible in the face might serve the agent as a basis for formulating an initial interaction strategy that could then be adjusted as further information about a user is obtained. At the very least, an interaction strategy based on physical appearances would mimic the more natural interaction style of human beings.⁷

⁷ An agent might also express its personality by the way it filters and processes facial cues.

Second, providing agents with a means of perceiving faces in terms of their trait impressions would enable agents to participate in the common social occupation of *evaluating others*. People spend a significant amount of time not only gossiping but also worrying about what others are saying behind their backs. If agents could read faces, it would be possible for them to comment about the appearances of others. This leads naturally to the idea of agents functioning as *smart mirrors*. Imagine an agent that could provide honest, albeit delicate, answers when asked how old, attractive, intelligent, and trustworthy the user appears.

4.3.2 MAKING FACES

Research conducted by both academia and industry indicate that interface design is evolving towards more personalized interfaces. Communication with the user is increasingly being mediated by virtual agents (André, Rist, Müller, 1998; Badler, Palmer, and Bindiganavale, 1999; Biocca, 1997). Because faces are such powerful social stimuli (Fridlund, 1994), more attention must be paid to properly portraying faces than to other aspects of embodiment.⁸ As Donath (2001) cautions, “. . . introducing faces into a mediated communication system must be done carefully, for the face is replete with social cues and subtle signals; a poorly designed facial interface sends unintended,

⁸The construction of Woody, the main character in the animated film *Toy Story*, illustrates how important the face is when compared with the rest of the body. In total, Woody's body consisted of slightly over 700 degrees of freedom—200 were dedicated to the face alone.

inaccurate messages, doing more harm than good” (p. 374). For this reason, it is not surprising that artists are typically given the task of designing the physical appearance of virtual agents. If virtual agents could predict the impressions their physical appearance makes on users, however, they could take over the job of the artist and learn to render appearances that best announce their intentions and that elicit specific responses from users. Just as human beings take great pains to enhance their facial appearance for social effect, agents capable of perceiving the physical personality of faces could consider the impressions their facial morphology produces and adjust their faces accordingly.

As mentioned above, people use their faces to broadcast vital social information about themselves to others. Aside from age and status, faces also include information about a person’s intentions. In this regard, Fridlund (1994) has remarked, “We are, it holds, thoroughly social, so that when we gaze at other's faces, we see not revelations of soul, character, or emotion—but declarations of their intentions towards us, and reflections of ours towards them” (p. xii). It is often the case, however, that the intentions behind statements and actions are ambiguous. A number of studies in the person perception literature show that people rely on facial impressions when trying to make sense of these ambiguous situations (Hochberg and Galper, 1974). Given this propensity for human beings to read intentions in faces, agents capable of manipulating their faces along specific trait lines could use their faces to physically advertise their intentions and thus help the user understand the nature of the agent’s roles. An agent’s face will either

clarify its intentions or complicate them.⁹ As anyone needing directions in an unfamiliar city knows, some people look more approachable than others. If an agent is given the task of assisting others, it could publicize its function by generating a face people would automatically be willing to approach. If an agent is expected to offer advice or function in a recommender system, it could alter its features to look especially convincing. If an agent is in a situation where it could be attacked, it could learn to adjust its face to limit user abuses.¹⁰ Just as people learn to control the impressions their faces make, so agents could learn to prepare social masks that are suited to their tasks.

Furthermore, given the fact that people treat media in the same way they treat other human beings (Reeves and Nass, 1996), it is also reasonable to assume that the time-honored adage “first impressions are lasting,” will likewise hold for embodied agents. In this regard, modern research gives credence to folk psychology. Studies show that the characterological impressions of a person’s face not only persist but also deepen over

⁹ The literature on face interfaces provides mixed results when it comes to evaluating the effectiveness of face interfaces (Koda and Maes, 1996; Walker, Sproull, and Subramani 1994). Dehn and Mulken (2000) attribute these conflicting reviews to variations in the kinds of animation used, the type of interfaces the virtual agents are being compared to, differences in measuring effects, and the task domain. Another question that needs to be addressed, however, is to what degree are these variations in reports on the effectiveness of face interfaces due to how the appearance of the agent enhances or conflicts with user expectations.

¹⁰ A particularly interesting response to virtual embodiment is that it appears to lend itself to user abuses (see, for example the report of abuse in [Isbister and Hayes-Roth, 1998]). This is an issue that will eventually have to be addressed by the virtual agent community.

time (Mathes, 1975). What happens when first impressions are proven wrong? Reactions can be retaliatory. As reported in the next chapter, Zebrowitz and McDonald (1991) have found that when defendants that were assumed to be innocent based on their physical appearance were proven at fault, the court treated them significantly more harshly. This strongly suggests that not properly announcing intentions or violating expectations initiated by the physical personality of embodied agents could produce negative repercussions. Agents that are designed to adjust their faces might do a better job in the end than human artists since agents could dynamically learn from their mistakes how best to adapt their faces to particular tasks or to specific social environments.

Admittedly, there are social implications in exploiting negative facial stereotypes. It is not advisable, however, to ignore the impressions the morphology of faces are bound to elicit in others. Rather than perpetuate facial stereotypes, it is possible that virtual agents could actively counter facial profiling, just as some human artists do, by quietly probing and challenging the user's largely unconscious and perhaps unjustified responses to certain facial configurations.

4.4 MODEL REQUIREMENTS

After considering some of the advantages a personality perception system would offer virtual agents, it is possible to list a set of four requirements that a model of the physical personality of the face must satisfy.

First, the model must be capable of handling a comprehensive set of trait impressions—at least those that are typically employed by people when characterizing others. Rosenberg's (1977) personality categories, discussed in section 3.3.2, might provide useful guidelines in this regard.¹¹

Second, the model must be predictive. This would furnish agents with a perceptual system that would enable them to perceive faces much as human beings do in terms of their trait impressions. As suggested above, this would provide an agent with a rudimentary sense of social self-awareness that would allow virtual agents to function as social mirrors for others.

¹¹ It should be pointed out that it is possible that a definition of personality that is useful to developers modeling personality from the perspective of the actor might not be of value to research centered on the observational perspectives. Reilly (1996), for instance, believes that personality defined in terms of traits is inappropriate for authoring systems. Traits, however, may be the only way to go when approaching the observation of personality, since it appears that observers tend to *type* others by classifying them into trait categories (Rosenberg, 1977).

Third, the model must be capable of modeling trait impressions of the face formed by either the general population or possibly specific groups of users. There is evidence that an actor's personality can influence face impressions (Hampson, 1988b; Rabin, 1951). Modeling specific users may be of interest to developers of virtual agents that are intended to serve special groups. Furthermore, perception is yet another behavior which may be exploited to reflect the personality of the virtual agent.

Fourth, a model must be capable of interfacing with mechanisms for managing the physical presentation of the virtual agent. Ideally, the model itself would be capable of synthesizing faces that produce specific impressions. Endowing virtual agents with a perceptual system capable of perceiving faces in terms of their trait impressions could, however, satisfy this requirement. As noted above, just as human beings often look in the mirror when adjusting their facial impressions, a trait impressions system could be used by virtual agents to provide feedback when creating or selecting their own physical personality or when modifying their facial presentations along specific trait lines.

Although not listed as a requirement, because virtual agents function autonomously and interactively, it is assumed that the model will be capable of functioning in real time without human intervention.

CHAPTER 5: TRAIT IMPRESSIONS OF THE FACE

To be plain with you friend, you don't carry in your countenance a letter of recommendation.

Dickens, *Barnaby Rudge*

5.1 SUMMARY

There is considerable consistency in people's impressions of faces. One major theory advanced to explain this consistency is that the perception of facial features has adaptive value and that trait impressions are based on those facial qualities that demand the greatest attention for the survival of the species, namely, physical fitness, age, and emotional state. Faces that possess features that are indicative of these qualities are believed to have pronounced overgeneralization effects (Zebrowitz, 1998). A modification of Rosenberg's trait categories (1977) is used in this chapter to explore clusters of traits associated with five of the most important overgeneralization effects: attractiveness, facial maturity, gender, physical fitness, and emotion. A discussion of the morphology that triggers these overgeneralizations is also presented. Finally, the prospect of indirectly modeling the trait impressions of the face by using existing models of attractiveness, facial maturity, and emotion is evaluated. It is concluded that that indirect methods fail primarily by not offering models that are predictive.

5.2 INTRODUCTION

Fair or not, certain facial characteristics give rise to personality trait impressions in others. Literature and history are full of accounts where people are judged, for good or for ill, according to their facial features. One famous story is that of Charles Darwin (1959), who was nearly rejected passage on the HMS *Beagle* because the captain, as Darwin recounted, “. . . doubted whether anyone with my nose could possess sufficient energy and determination for the voyage.” Although most today would scoff at the captain’s methods of character assessment and recite such maxims as "never judge a book by its cover," evidence abounds that people not only judge others based on their facial features but also believe that the face provides valuable clues regarding a person’s character (Liggett, 1974). There is even a growing body of evidence validating the accuracy and consistency of these judgments (Albright, Kenny, and Malloy, 1988; Berry, 1990, 1991a; Berry and Brownlow, 1989; Berry and Wero, 1993; Cunningham, 1986; Funder, 1987; McArthur and Baron, 1983; McCauley, Jussim, and Lee, 1995; Rind and Gaudet, 1993; Watson, 1989; Zebrowitz, 1998; Zebrowitz and Montepare, 1992; Zebrowitz, Voinescu, and Collins, 1996).

This chapter reviews the current literature on trait impressions of the face. It is concerned, however, less with the accuracy of trait impressions than with consensus or interjudge reliability. Interjudge reliability has been reported for a large number of traits, a fact that has impressed several investigators (Albright et al., 1988; Berry and Wero, 1993; Hochberg and Galper, 1974). Strong consensus across cultures and races has also been observed (Albright, Malloy, Dong, Kenny, and Fang, 1997; Berry and Wero, 1993; Cunningham, Roberts, Barbee, and Druen, 1995; Keating, Mazur, and Segall, 1981a; McArthur and Berry, 1987). Of course, it is possible that not all traits exhibit the same consistency. One of the tasks in this chapter, therefore, is to determine which traits or trait categories exhibit the highest degree of consensus. Furthermore, since the ultimate goal of this dissertation is to model the trait impressions of the face, those facial features and characteristics that contribute most to the formation of trait impressions are of particular interest. Theoretical considerations can best guide an investigation into this area. Although several theories have been advanced to explain trait impressions of the face, one major theory is that the perception of facial features has adaptive value and that those trait impressions that have the most influence are based on those facial qualities that demand the greatest attention for the survival of the species (Alley, 1988; Berry and McArthur, 1985; Guthrie, 1970; Jones, 1995; Mark, Shaw, and Pittenger, 1988; McArthur, 1982; McArthur and Baron, 1983; Zebrowitz, 1998). Recognizing an angry face, for example, may trigger appropriate fight/flight responses or conciliatory behaviors. It is theorized that faces that are similar in structure to angry faces elicit

similar, albeit milder, responses. As Zebrowitz (1998) explains, “We could not function well in this world if we were unable to differentiate men from women, friends from strangers, the angered from the happy, the healthy from the unfit, or children from adults. For this reason, the tendency to respond to the facial qualities that reveal these attributes may be so strong that it is overgeneralized to people whose faces merely resemble those who actually have the attribute” (pp. 14-15).

Two of the most researched overgeneralization effects are the attractiveness halo effect and the facial maturity overgeneralization effect. Other overgeneralization effects have received less attention but are nonetheless significant. They are based on emotions, infirmities, and gender (Alley, 1988; Enlow and Hans, 1996; Symons, 1979; Zebrowitz, 1998). Both the associated traits and the facial qualities that influence or trigger these overgeneralizations are examined in this chapter. The last section evaluates the prospect of indirectly modeling trait impressions of the face by altering the morphological characteristics associated with attractiveness, facial maturity, gender, and emotion. It is concluded that indirect methods, although certainly worthy of future exploration, fail primarily by not offering models that are predictive.

5.3 THE ATTRACTIVENESS HALO EFFECT

One would be hard pressed to name one culture that did not in some way encourage its members to alter the appearance of their faces. As Liggett (1974) has observed, “The desire to alter the face is universal; in every culture and in every age examples of facial elaboration can be found” (p. 46). Liggett goes on to note that, although religious motives and a need to mark social status are factors in facial elaboration, enhancing the aesthetic appeal of the face is paramount. He writes, “All peoples, sophisticated as well as primitive, seem prepared to go through almost unbelievable suffering in pursuit of the purely local ideals of their particular society. Beauty must be pursued at whatever price, because it confers on its possessor profound social influence, power and respect” (p. 46).

Modern research supports Liggett’s claim that social benefits accrue to those who are most attractive. Positive traits are commonly associated with attractiveness and negative traits with facial abnormalities and unattractiveness (Eagly, Ashmore, Makhijan, and Longo, 1991; Feingold, 1992; Langlois et al., 2000). In section 5.3.1 these traits are examined, and in section 5.3.2 some of the morphological characteristics are delineated that are known to contribute to contemporary judgments of attractiveness.

5.3.1 THE TRAIT IMPRESSIONS OF ATTRACTIVENESS

People historically and across many cultures have associated the good with the beautiful. Yet even though cultural artifacts have strongly suggested the existence of an attractiveness stereotype, the subject was barred from psychological research until fairly recently (Berscheid and Walster, 1974; Bull and Rumsey, 1988; Dion and Berscheid, 1972; Jones, 1995). As Aronson (quoted in Dion and Berscheid) in 1968 noted, “It may be that, at some levels, we would hate to find evidence indicating that beautiful women are better liked than homely women—somehow this seems undemocratic” (p. 286). One of the first studies to break the ice and examine the attractiveness stereotype was Dion and Berscheid. They had a group of introductory psychology students rank yearbook photographs of women according to their attractiveness. Another group of students then judged the same set of photographs along 27 personality dimensions that broadly measured social desirability, happiness, and occupational success. The results conclusively supported a beauty-is-good stereotype (Dion and Berscheid). Since then, there has been an explosion of research on attractiveness.¹²

Of particular importance are the quantitative meta-analyses of the literature on attractiveness by Eagly et al. (1991), Feingold (1992), and most recently, Langlois, et al. (2000). Eagly et al. investigated the personality attributions of attractiveness by linking

¹² For book length reviews see Bull and Rumsey (1988), Hatfield and Sprecher (1986), Jackson (1992), and Patzer (1985)

distinctions between social categories to various dimensions of personality perception. These were partitioned, using a slight modification of Rosenberg's system (1977), mentioned in section 3.3.2, into the following categories: social competence, intellectual competence, concern for others, integrity, psychological stability, maturity, and potency.

Since these categories provide a comprehensive and convenient method for organizing a comparison of the various trait impressions associated with specific overgeneralization effects, a brief description of each category is in order. Social competence consists of the skills, traits, and outcomes bound up with sociability. Some example social traits would be extroversion, talkativeness, and a liking for fun. Social outcomes encompass generalized reactions to facial characteristics such as popularity. Intellectual competence focuses on those skills, traits, and results that concern intellectual ability as well as motivation. Concern for others assumes sociability but differs from the social competence category in its focus on the welfare of others. It would include such positive traits as generosity and sensitivity to others and a lack of such traits as egotism and vanity. Integrity is bound up with honesty, that is, truth telling and conformity to laws and norms. Adjustment is associated with normal vs. abnormal psychological functioning, happiness, well being, and self-esteem. Maturity concerns responsibility but was not treated as a category by Eagly et al., as few studies on attractiveness have addressed it. The last category, potency, is related to self-assurance, dominance, and leadership.

Eagly et al. (1991) hypothesized that an analysis of the literature would reveal that attractive people are strongly stereotyped as good in social competence and achievement, moderately stereotyped as strong and well adapted, and weakly stereotyped as good at intellectual activities. As expected, the results showed a strong association of positive traits with attractiveness, with the strongest effect in the category of social competence. Attractive people are also considered more adapted, potent, and intellectually competent than initially predicted. In their meta-analysis, attractiveness had little effect on concern for others and integrity.

Using a different sample of studies, Feingold (1992) independently replicated these results. This analysis, however, employed eight categories of dependent variables: sociability (which included extroversion, need for affiliation, and friendliness), dominance (assertiveness, and ascendancy), sexual warmth (heterosexual responsiveness and need), modesty (including its opposite, vanity), character (genuineness, sincerity, trustworthiness, morality, and kindness), general mental health (emotional stability, adjustment, happiness), intelligence (academic ability, IQ, brightness), and social skills (social adeptness, social competence, and poise). Feingold found attractiveness to be strongly correlated with positive social skills and mental health, and, to a lesser degree, with dominance and intelligence. Ashmore and Longo (1995), noting the remarkable

similarity in the results of the two meta-analyses, took Feingold's categories and mapped them onto Rosenberg's categories as used by Eagly et al. How well the findings in the two studies correspond is illustrated in Table 5.1, which replicates Ashmore and Longo's tabular summary of the meta-analyses.

A recent meta-analysis by Langlois et al. (2000) corroborates Ashmore and Feingold by reporting a strong positive correlation between attractiveness and interpersonal competence, social appeal, and adjustment. However, their analysis indicates that adults and children are judged more positively for occupational competence and intellectual competence than reported by Ashmore and Feingold.

Table 5.1. Mean Weighted Effect Size by Attribute Content in Two Meta-Analyses of the Physical Attractiveness Stereotype.

<i>Eagly, Ashmore, Makhijani, & Longo, 1991</i>		<i>Feingold, 1992</i>	
<i>Attribution</i>	<i>d</i>	<i>Attribution</i>	<i>d</i>
Social Competence	.68	Social Skills	.88
		Sexual Warmth	.78
		Sociability	.46
		Social Competence	.71
Adjustment	.52	General Mental Health	.50
Potency	.49	Dominance	.54
Intellectual Competence	.46	Intelligence	.31
Integrity	.13	Character	-.04
Concern for others	.01	Modesty	.34
Reproduced from Ashmore and Longo (1995), p. 69. Copyright © 1995 by the American Psychological Association. Reprinted with permission.			

In general, the literature strongly supports an positive attractiveness halo effect, but some studies have reported negative associations, especially in terms of concern for others (Ashmore and Longo, 1995; Dermer and Thiel, 1975). Dermer and Thiel, for instance, found that attractive people are sometimes viewed as self-centered, vain, and egotistical. These negative attributions as well as some inconsistencies in the literature may partly be explained by differences in types of attractiveness, a subject that has been explored by Ashmore, Solomon, and Longo (1996), Solomon, Ashmore, Longo (1992), and Berry

(1991b). Two of these studies, (Solomon, et al., 1992) and (Ashmore et al., 1996), had different populations, the first a group of beauty experts in the fashion world and the second a group of introductory psychology students, rank photographs of attractive women. Both studies isolated similar categories of attractiveness. These included the following: cute (non-sexual youthful attractiveness), sexy (overtly sexual attractiveness), and trendy (attractiveness based on more fashionable dress and hairstyles). In general, Ashmore et al. (1996) found that the social traits attributed to attractiveness in the meta-analysis by Eagly et al. (1991) held for the cute subtype but varied from their findings in the other attractiveness subtypes, at least in terms of female attractiveness. The sexy type of attractiveness, for instance, was viewed by most women as slightly not so socially adept. In a similar vein, Berry (1991b) investigated differences in attractiveness trait attributions that varied along the lines of perceived facial maturity. In this study, attractive individuals with babyish faces were generally thought to be more honest, sincere, and warm than attractive individuals with more mature faces. Furthermore, attractive babyfaced men, or “pretty boys,” were perceived to be less powerful than attractive males with more mature faces (see section 5.4 for a more detailed discussion). In addition to these, there are other stereotypes of attractive people that are typically associated with negative traits. Two of the most common are the Nordic “ice queen,” which is viewed as cold and antisocial, and the “dumb blonde,” which is negatively associated with intelligence (Ashmore and Longo, 1995). As Ashmore and Longo (1995)

have remarked, most social stereotypes include a number of subclasses that account for exceptions to the rule, the attractiveness stereotype is no different in this regard.

Few studies have focused on the trait attributions associated with normal unattractiveness (Stevenage and McKay, 1999). Nonetheless, the literature on attractiveness consistently demonstrates that negative traits are attributed to the unattractive. Such people are considered less socially competent, honest, intelligent, and psychologically stable. In particular, studies have demonstrated a strong negative association between unattractiveness and integrity [for a review, see Bull and Rumsey (1988)]. Dion (1972), for example, asked undergraduate students to rank a set of facial photographs of children according to honesty. The least attractive were thought least honest. Several studies have reported possible juror and lineup biases against unattractive individuals (Bull and Rumsey; Efran, 1974), and significantly more facial abnormalities have been found among prisoners than in the population at large—sixty percent versus ten percent (Bull and Rumsey). In general, deviant behavior of all types is more likely to be attributed to the unattractive than to others. Dunkle and Francis (1996), for instance, instructed subjects to pick out homosexuals in a set of photographs. The most unattractive men and women were more frequently selected, replicating the findings of an earlier study by Unger, Hilderbrand, and Madar (1982). Miller, Gillen, Schenker, and Radlove (1974) demonstrated that unattractive people are thought capable of more brutality and manipulation. When subjects were shown photographs of people who had administered

shock to other subjects in obedience to an experimenter's directive, those least attractive were thought to have administered the highest shock levels. Unattractive people are also considered less socially competent and desirable (Bull and Rumsey), and they are expected to be less cooperative (Mulford, Orbell, Shatto, and Stockard, 1998). Bull and Rumsey report that the peers of unattractive boys expect them to perform more aggressive and antisocial acts. Social reactions to unattractive people are also more negative (Langlois et al., 2000). Unattractive people are often ignored (Bull and Rumsey) and, if facially disfigured, avoided (Bull and Rumsey; Houston and Bull, 1994). People are also more prone to display more inconsiderate, hostile, and antisocial behaviors towards unattractive people. Unattractive girls, for example, are punished more severely than other children (Alley and Hildebrandt, 1988; Langlois et al., 2000), and unattractive children are more likely to be abused (Bull and Rumsey). In an experiment by Alcock, Solano, and Kayson (1998), subjects were paired, with one of the pair being instructed to spill water on the other. Those who were unattractive received more aggressive reactions. Unattractive people are also thought to be less intelligent and, in general, receive lower evaluations from teachers and other judges of performance (Alley and Hildebrandt, 1988; Bull and Rumsey; Cash and Trimer, 1984).

Aside from the correlation of unattractiveness with negative social skills, the other correlation rankings of traits in Table 5.1 may not apply to unattractiveness. As noted above, both Eagly et al. (1991) and Feingold (1992) found a weak relationship between

integrity and attractiveness. In general, attractive people are not thought significantly more honest than others. However, as also reported above, there is a body of literature showing that unattractive people are considered less honest. Perhaps there is an association of lack of adjustment or mental health with a greater propensity to break the rules.

5.3.2 THE FACIAL QUALITIES OF ATTRACTIVENESS

Some believe that there is no objective standard of beauty, that it is entirely in the eye of the beholder. To some degree this is true. In general, faces that are more familiar or similar to the viewer are considered more attractive (Penton-Voak, Perrett, and Peirce, 1999), and there is evidence that moral judgments influence opinions of beauty. For instance, a face reported to be that of a lifeboat rescuer was found to be significantly more attractive than the same face reported to be that of a murderer (Shepherd, Ellis, McMurrin, and Davies, 1978). But does that mean that there are no objective measures of beauty or that beauty is entirely individual? The strong consensus in judgments of attractiveness would seem to indicate otherwise. As both Feingold (1992) and Patzer (1985) have pointed out, research into attractiveness relies almost exclusively on this consensus, and this reliance is justified. Attractiveness judgments have been found to be consistent not only over time and within groups, but cross-culturally, cross-rationally, and across age groups (Alley and Hildebrandt, 1988; Cunningham et al., 1995; Jones, 1995; Langlois, Riggman, and Rieser-Danner, 1990; Langlois, Riggman, and Vaughn, 1991;

Lucker and Graber, 1980; Perrett, May, and Yoshikawa, 1994; Rubenstein, Kalakanis, and Langlois, 1999). There is even evidence of cross-species similarities in attractiveness judgments (Cunningham, 1986). Although this consensus suggests that some underlying facial qualities are responsible for many attractiveness judgments, pinning down the facial qualities and characteristics that make faces appealing has proven a difficult task. To date, there is no theory of attractiveness that is generally accepted (Carello, Grososky, Shaw, Pittenger, and Mark, 1989). Nonetheless, contemporary research into facial attractiveness has reported evidence indicating that proportion, symmetry, straightness of profile, closeness to population average, and sex specific facial features are important factors in attractiveness judgments.

Attractive faces are proportional, straight, and symmetric

The oldest literature on facial attractiveness, as well as many ancient and modern aphorisms, suggests that facial attractiveness depends on such qualities as symmetry, straightness, and proportion. Medieval artists, for instance, divided the face into sevenths and found those faces most pleasing that had distances between features that measured one seventh. Among the ancient Greeks, proportion (*summetria*), especially that of the golden section, was the defining characteristic of beauty. The golden section is defined such that the ratio of the whole, x , to a larger part, y , equals the ratio of y to the smaller part, z : $x/y = y/z$. As any art student knows, a face that conforms to this ideal has features equally distributed into three horizontal sections. The measure from the hairline to

browbridge, from browbridge to the base of the nose, and from the base of the nose to the tip of the chin are all equal to one third the total length of the face. The width of a well-proportioned face is thought to conform to the golden proportion by measuring approximately two-thirds the length (Alley and Hildebrandt, 1988; Zebrowitz, 1998).

Although there is some evidence that today beautiful faces, especially for Westerners, conform to the golden section (Bruce, 1988; Ricketts, 1982; Zebrowitz, 1998), other investigators have found little correspondence between the golden section and attractiveness (Carello et al., 1989). In fact, some decidedly unattractive faces can have proportions that are golden (Zebrowitz, 1998). Nonetheless, many agree that the golden section or some similar set of ratios probably plays an important role in attractiveness judgments (Alley and Hildebrandt, 1988).

One alternative offered to that of the golden section is the set of geometrical relationships that are formed by normal facial growth. These relationships allow for the most efficient masticulation and have been associated with attractiveness judgments (Carello et al., 1989). Normal growth results in facial profiles that are relatively straight, with *straightness* depending on the individual's bone structure and thus being unique to the individual. Physical anthropologists have identified three types of facial profiles that depend on measures of straightness: the orthognathic, retrognathic, and prognathic

(Enlow and Hans, 1996). The three types can be spotted by noting the position of the chin in terms of a vertical line that drops down along the upper lip and which is perpendicular to a horizontal line that extends outward from the eyeball. A chin that is inside the vertical line produces a retrognathic profile, whereas a chin that extends outside the line, along with the nose, is prognathic. Many studies have demonstrated a preference for orthognathic or straight profile shapes (Carello et al., 1989; Lucker and Graber, 1980; Magro, 1997). Least attractive is the prognathic (Carello et al., 1989). Magro (1997) explains this by noting that derived features in general, that is, features that have developed out of the recent evolutionary changes of bipedalism, enhanced intelligence, reduced sexual dimorphism, manual dexterity, and an omnivorous diet, are considered more beautiful than older, more primitive features. According to Magro, the orthognathic profile is derived, whereas the prognathic is more primitive. Even children exhibit a clear preference for the orthognathic.¹³ In studies on children's preferences, profile straightness is the most significant and reliable measure affecting their attractiveness judgments (Lucker and Graber, 1980; Zebrowitz, 1998).

Related to proportion is symmetry. In general, human beings demonstrate a marked preference for all things symmetrical as is evidenced by the use of symmetry in so much art and architecture. This preference for symmetry extends to the face as well, where it

¹³ Magro (1997) believes Barbie is so popular precisely because her features are entirely derived.

has been found that symmetrical faces are considered more attractive than asymmetrical faces (Grammer and Thornhill, 1994; Mealey, Bridgstock, and Townsend, 1999; Perrett et al., 1999). When speaking of asymmetry in the human body, it is important to note that there are three types: directional asymmetry, which refers to a population wide and directionally constant form of asymmetry (for example, the placement of the human heart), antisymmetry, which refers to a population wide and directional varied asymmetry (for example the spin of a human hair whorl), and fluctuating asymmetry, which consists of random deviations from perfect symmetry. Fluctuating asymmetry forms a statistically normal distribution and is the type of asymmetry that is associated with the perception of attractiveness (Mealey et al.).

One problem with studying fluctuating asymmetry in the human face is the difficulty of isolating and measuring it. Consequently, there are conflicting reports regarding the importance of symmetry in attractiveness judgments. Because fluctuating asymmetry is a deviation from perfect symmetry, attempts have been made to study its effects on attractiveness by artificially creating perfectly symmetric faces, or *chimeras*. Kowner (1997), for instance, tested perfectly symmetric faces by creating their mirror reflections; they were rated less attractive than faces with ordinary fluctuating asymmetry. However, Perrett et al. (1999), suspecting that Kowner's subjects were negatively responding to blemishes introduced in the process of creating the chimeras, reported them to be judged more attractive once the blemishes had been removed.

In yet another study, Mealey et al. (1999) attempted to control for some of the covariates of symmetry by using subjects that were monozygotic twins, that is, genetically, but not developmentally, identical. In this study, the more symmetric twin was thought to be the more attractive.


It should be noted that the preference for symmetrical faces does not extend to expressiveness. Asymmetrical smiling faces have been rated as more attractive than symmetrical smiling faces (Kowner, 1996). This could be explained by the fact that the left hemisphere of the brain is more expressive than the right. Asymmetrical smiles may, therefore, be perceived as more genuine and thus more attractive (Kowner).

Attractive faces are average

One of the first to create and explore average faces was Francis Galton (1878), who did so by ingeniously superimposing photographs of more than one face. His major objective was to obtain a representation of various classes of people: criminals, the healthy, the ill, and the famous. To his surprise, the composites or averages appeared notably more attractive. Little was done with his observation until 1952, when David Katz (1952) maintained that composites are more beautiful than the individual faces

comprising them by virtue of the fact that they are closer to the average. The first systematic study to lend support to his claim, however, had to wait until 1990, when Langlois and Roggman (1990), using digitized photographs of student faces, demonstrated not only that composites are thought more attractive but also that perceived attractiveness increases as more and more faces are averaged (the average being computed arithmetically, using the gray scale pixel values of the constituent images). A year later, Langlois, Roggman, Musselman and Acton (1991) produced additional evidence that this preference for the average is exhibited by infants as well as adults (see also Langlois, Roggman, and Vaughn [1991]). Their findings have more recently been confirmed by Langlois, Roggman, and Rieser-Danner (1990), Rubenstein et al. (1999), and Rhodes and Tremewan (1996).

Figure 5.1. Increased Attractiveness of Averaged Faces.



Note. A face (1.0) morphed along a line towards the mean (0.0) of the 220 randomly generated stimulus faces in Appendix D increases in attractiveness. The mean face was computed by averaging the pixel gray scale values of the 220 faces.

Numerous theories have been advanced to explain the attractiveness of average faces. Langlois et al. (1990) have suggested that average faces are attractive because human beings process sensory data by creating prototypes or averages and a preference for the prototype has the evolutionary advantage of stabilizing selection (see also Rubenstein et al. [1999] and Young and Bruce [1998]). Others have suggested that composites are preferred because they are more familiar (Alley and Cunningham, 1991), youthful (Alley, 1992; Grammer and Thornhill, 1994), and symmetrical (Alley and Cunningham, 1991), and that these qualities contribute more significantly to judgments of attractiveness. Yet others deny that average faces are attractive but rather contend that the process of making photographic composites by averaging pixels introduces artifacts such as blurring, smoothing, poor resolution, and shape bending, which are responsible for the increased attractiveness ratings (Pittenger, 1991). Langlois, Roggman, and Musselman (1994) have responded to these criticisms by demonstrating that a preference for the average is not related to youthfulness and symmetry. On the issue of familiarity, they have argued that the connection between attractiveness and familiarity in terms of average faces supports their theoretical position that the preference for average faces stems from the fact that they are perceived as prototypes. They note that association has been observed in other categories of experience as well (Langlois et al., 1994). As to the argument that artifacts introduced in the image-making process are responsible for the increased attractiveness ratings, Langlois et al. have been careful to point out that a similar degree of blurring was

introduced in the original images and that poor resolution was not a problem (Langlois et al., 1994; Langlois, Roggman, and Musselman, 1991).

Rhodes and Tremewan (1996) entirely bypassed the issues surrounding pixel-based averaging by extracting facial shapes. Using a caricature generating computer program, caricatures, produced by moving facial shapes further from the mean, were rated less attractive than the veridical faces, and face shapes pushed closer to the mean were rated more attractive. Pollard, Shepherd, and Shepherd (1999), however, tested *real* faces whose distance from the average was measured by summing the z-scores of a number of frontal face measurements; those faces closest to the average score were not rated significantly more attractive. Thus, although there is evidence that *averaged* faces are attractive, what exactly is average about these faces is very much in debate (Pittenger, 1991).

Sex specific facial features of attractiveness

Many of the features associated with attractiveness are gender specific. In many species, females prefer the most powerful males, that is, those with the highest status (Cunningham, Barbee, and Pike, 1990). In this regard, human females apparently do not differ from the females of other species, as studies have demonstrated that women prefer dominant looking males (Alley, 1992; Keating, 1985). Keating, for instance, creating

facial composites using Identikit, found men with features associated with dominance (thick eyebrows, thin lips, and square jaws) more attractive than men with less dominant features (thin eyebrows, full lips, and round jaws). However, there is also evidence suggesting that women, as well as men, find *feminized* faces to be more attractive (Rhodes, Hickford, and Jeffery, 2000). Research into male preferences shows that those females who have features associated with approachability, namely, a large smile, high-set eye brows, and dilated pupils, are considered attractive (McArthur and Apatow, 1984). Cunningham (1986) has shown that the mature features of high cheekbones, small chins, and narrow cheeks combined with the neonate features of large eyes and small nose are judged to be most attractive. That youth is a significant factor in male preference was demonstrated by Buss and Barnes (1986). In a large study of the mate preferences of males in twenty-seven diverse countries, Buss and Barnes found a strong universal tendency for males to prefer young partners.

Other factors influencing attractiveness

There are a number of other factors that influence attractiveness judgments. Some of these include facial expression, facial prominence, and context (Alley and Hildebrandt, 1988; Archer, Iritani, Kimes, and Barrios, 1983; Penton-Voak et al., 1999). Just as positive traits are associated with attractiveness and negative traits with unattractiveness, positive and negative emotions also affect attractiveness. For instance, smiling faces are perceived to be more attractive than frowning faces (Alley and Hildebrandt). Another

influence on facial attractiveness is facial prominence in photographs, with prominence being measured by how much the face fills the photographic space. Faces that are more prominent are judged to be more attractive, intelligent, and ambitious (Archer et al.). Yet another factor in attractiveness judgments is context. Studies have been reported where subjects judged a set of female faces to be less attractive when presented with a group of extremely attractive female subjects (Alley and Hildebrandt).¹⁴

5.4 THE OVERGENERALIZATION EFFECT OF FACIAL MATURITY

Perhaps no face is more capable of eliciting a favorable response than that of a baby, as is illustrated in the following account of an encounter between a North Korean guard and a small child:

We had been sternly ordered to keep our hands down and to refrain from speaking to the North Korean guards at the far side of the divided meeting room in Panmunjom, the border town straddling North and South Korea. The guards did look ominous, and the shootout that had recently occurred when a Russian attempted to escape through this room to South Korea made me take these instructions seriously. Suddenly, my son Loren waved his hand, and his high-pitched “hi” chimed across the room. I turned toward the nearest North Korean guard, expecting to see his automatic weapon trained on us. Instead, it was a large grin that was leveled at my son. The transformation in this “enemy” soldier’s stony face brought tears to my eyes. He found my baby totally disarming. (Zebrowitz, 1998, p. 64)

¹⁴ Context influences highlight a concern with the experimenter’s degree of attractiveness as this could influence subjects’ attractiveness ratings [Alley, 1988 #694].

Human beings and animals alike are disarmed and entranced by a youthful face (Lorenz, 1970-71; Zebrowitz, 1998). Even infants show a preference for them (McCall and Kennedy, 1980). In the same way that a variety of neonate markers trigger nurturing behaviors in animals (Lorenz), so may certain facial characteristics of a baby's face provoke care taking behaviors in humans (Frodi et al., 1978). The favorable response to a baby's face is not just reserved for babies, however, but is generalized to adults whose faces resemble those of a baby (Zebrowitz, 1998). Universally, *babyfaced* people are attributed childlike characteristics (Berry and Brownlow, 1989; Berry and McArthur, 1985; McArthur and Berry, 1987). They are perceived to be more submissive, naïve, honest, kindhearted, weaker, and warmer than more mature-faced individuals (Berry and McArthur, 1985, 1986; McArthur and Berry, 1987). Section 5.4.1 examines the trait impressions of babyfacedness. Section 5.4.2 delineates some of the perceptual characteristics that give rise to these impressions.

5.4.1 TRAIT IMPRESSIONS OF FACIAL MATURITY

Since, first and foremost, infants are perceived to be weak and in need of protection, McArthur and Apatow (1984) predicted that people with babyish faces would be perceived as weak. Manipulating schematic drawings to produce high and low degrees of babyfacedness, they found, as expected, strong correlations between babyfacedness, physical weakness, and low social dominance. These correlations held for photographs

of faces as well (Berry and McArthur, 1985). Mueller and Mazur (1996) have more recently corroborated these findings. In fact, babyfaced men are perceived to be weaker than mature-faced women (Friedman and Zebrowitz, 1992). Studies investigating the age effects of babyfacedness show these effects to be independent of perceived age and attractiveness (Berry and Brownlow, 1989; Berry and McArthur, 1985). Extremely babyfaced infants, for example, are seen as physically weaker, more submissive, and more dependent (Zebrowitz and Montepare, 1992). Furthermore, even though babyfaced older adults are perceived to be younger, they are also considered weaker than those with more mature, older looking faces (Zebrowitz, 1998).

Honesty is another trait strongly correlated with babyfacedness (Zebrowitz et al., 1996) (Berry and Zebrowitz-McArthur, 1988; McArthur and Apatow, 1984; Zebrowitz and McDonald, 1991). Berry and Zebrowitz-McArthur took fictitious criminal reports and attached to them photographs of either a mature-faced or a babyfaced male. When college students were asked to render judgments based on the reports, the babyfaced males were consistently considered more innocent than the mature-faced males except in offenses that involved negligence; in these cases, the association of babyishness with irresponsibility had an adverse effect (Zebrowitz, 1998). In addition, it was discovered that when babyfaced individuals violated expectations of innocence and honesty, they were treated significantly more harshly. In another study, an examination of actual small claims cases showed that babyfaceness had an impact on the outcome (Zebrowitz and

McDonald). In cases involving intent, only forty-five percent of the more babyfaced defendants were found at fault compared to ninety-two percent for the more matured faced defendants. Again, these results obtained independently of the degree of attractiveness and age as well as quality of support.

Another strong trait correlated with babyfacedness is lack of credibility. Babyfaced individuals are considered naïve and gullible (Berry, 1991a; Berry and McArthur, 1985). In one study investigating the occupational effects of babyfacedness, babyfaced actors in commercials were less often given roles as experts and more often asked to provide testimonials (Zebrowitz, 1998). When it comes to perceived expertise, babyfaced men are found to rank similarly to women at the lower end of the scale (Zebrowitz).

Utilizing Rosenberg's categories, the literature reviewed thus far would suggest that babyfaced individuals are perceived as being high in integrity but low in intellectual competence and potency. What about the other categories? There is strong evidence that babyfaced individuals are rated high in social competence and concern for others. When children were asked to pick the person most likely to be kind and share things with friends, they showed a high preference for babyfaced children (Montepare and Zebrowitz-McArthur, 1989). In a study investigating effects of sex on facial maturity, babyfaced men were seen to be as warm as mature faced women (Friedman and

Zebrowitz, 1992). People are also more likely to confide in people with babyish faces and show more physical affection towards them (Berry and Landry, 1997; Zebrowitz, 1998). They are also perceived to be more helpful and caring. When a group of college students matched photographs of women and men in their fifties to a number of occupations, babyfaced women were placed more in the helping professions as teachers and nurse's aides (Collins and Zebrowitz, 1995). Although in this study short men, not babyfaced men, were placed in similar professions as babyfaced women, in another study babyfaced men were considered to be as likely as women to care for children (Friedman and Zebrowitz).

Table 5.2 summarizes the trait associations of babyfacedness and compares them with attractiveness. Although a meta-analysis of babyfacedness is not available, the literature suggests that babyfacedness is positively correlated with social competence, concern for others, potency, and to some degree maturity. The literature does not report much in terms of adjustment. As discussed in section 5.3.1, attractiveness is positively correlated with social competence, psychological stability or adjustment, potency, and intellectual competence. Attractiveness is not associated with concern for others, and its relationship to maturity has not been sufficiently investigated.

Table 5.2. Correlation Between Attributes, Attractiveness, and Babyfacedness.

<i>Attribution</i>	<i>Attractiveness</i>	<i>Babyfacedness</i>
Social Competence	Positive	Positive
Intellectual Competence	Positive	Negative
Concern for others	Neutral	Positive
Integrity	Neutral	Positive
Adjustment	Positive	N/A
Maturity	N/A	Negative
Potency	Positive	Negative


5.4.2 PERCEPTUAL CHARACTERISTICS OF FACIAL MATURITY

The facial features that mark a baby's face are large eyes relative to the rest of the face, fine, high eyebrows, light skin and hair color, red lips that are proportionally larger, a small wide nose with a concave bridge, and a small chin. The facial features are also placed lower on the face (Berry and McArthur, 1985; Zebrowitz, 1998).¹⁵ Across cultures, these same facial configurations characterize babyfaced adults (Berry and McArthur; McArthur and Berry, 1987). Even a single babyish facial feature can affect trait impressions (Zebrowitz). For instance, in various studies, when eye size was manipulated, larger eyes were considered warmer, weaker, more honest, incredulous, and submissive (McArthur and Apatow, 1984; Paunonen, Ewan, Earthy, Lefave, and

¹⁵ Lorenz (1970-71) has noted many similarities between the neonate features of humans and animals.

Goldberg, 1999; Zebrowitz et al., 1996). A change in nose size produced similar results (McArthur and Apatow). In contrast, other studies have demonstrated that the more mature facial features of a broad forehead, strong eyebrows, thin lips, lowered eyebrows, and narrow eyes are considered dominant and strong (Keating, 1985; Keating et al., 1981a; Keating, Mazur, and Segall, 1981b; Mueller and Mazur, 1996). As an example, the cover of the October 9, 2000, edition of the *New Yorker* uses these differences for excellent artistic effect (see Figure 5.2).

Figure 5.2. A Media Example of The Overgeneralization Effect of Babyfacedness.



Note. On the cover of *The New Yorker*, October 9, 2000 issue, these vapid babyfaced beauty queens from middle America are contrasted with the sexually savvy, mature faced, Ms. America contestant from New York. Observe how the large eyes and protruding forehead (accentuated by the shadow cast by the bangs) of the babyfaced contestants contrast with the narrow eyes and receding forehead (accentuated by the hair pulled back into a bun) of Ms. New York. Original artwork by Ian Falconer. © 2000 The Condé Nast Publications Inc. Reprinted with permission. All Rights Reserved.

Other significant age-related differences in faces concern developmental changes in craniofacial profile shape. Of particular note are differences in the relative size of the

brain capsule and the slant of the forehead in relation to the chin. The infantile cranium is proportionally much larger than the fully mature cranium, and the infantile forehead protrudes whereas the adult forehead recedes. Another important characteristic is a dramatic increase in jaw size. Figure 5.3 illustrates these differences. The craniofacial profile shapes were produced using a cardioidal strain transformation. Applied to standard profile shapes, this transformation has been shown to approximate real growth (Todd and Mark, 1980; Todd, Mark, Shaw, and Pittenger, 1981). As would be expected, studies on the trait attributions of profiles that vary in the degree of cardioidal strain applied are consistent with findings on facial maturity (Alley, 1983; Zebrowitz, 1998). As craniofacial profile maturity decreases, so do perceived alertness, reliability, intelligence, and strength (Berry and McArthur, 1986). Moreover, infantile profile shapes are more loveable, less threatening (Berry and McArthur), and elicit stronger desires to nurture and protect (Alley).

Figure 5.3. Cardioidal Strain Transformation .



Note. Applied to Standard Profile Shapes, the Cardioidal Strain Transformation Approximates Real Growth. Reproduced from Pittenger and Shaw (1975), p. 376. Copyright © 1975 by the American Psychological Association. Reprinted with permission.

As noted throughout section 5.4.1, the trait impressions of babyfaced adults are similar to traits stereotypically attributed to females (Friedman and Zebrowitz, 1992). This may be related to the fact that the female face tends to retain into adulthood more of the morphological characteristics of youth (Gray, 1985). Female heads and facial features are smaller and the eyes tend to be protrusive. This makes the eyes look disproportionately larger, as are the eyes in a baby's face. Female eyebrows are also thin and arched, sex specific characteristics that are further accentuated in Western culture through the use of make-up and eyebrow plucking (Brownmiller, 1984). Male faces, in contrast, tend to have the mature facial characteristics of a large nose, dominant jaw, angular cheekbones, deep set eyes, and a pronounced brow; the last two characteristics tend to make the eyes look even smaller (Enlow and Hans, 1996; Friedman and Zebrowitz, 1992).

5.5 THE OVERGENERALIZATION EFFECT OF PHYSICAL FITNESS

To some degree the characterological impressions of physical fitness are related to impressions of attractiveness. The facial characteristics of attractive men reported in section 5.3.2, for instance, require high levels of testosterone, which only healthy men can accommodate as such levels lower immunocompetence (Folstad and Karter, 1992). Yet another important component in attractiveness associated with health is symmetry. In many animals, it is indicative of phenotypic fitness and plays an important role in mate attraction (Thornhill and Gangestad, 1993). Humans may be similarly predisposed to prefer symmetrical facial shapes because they advertise the prospect of healthy offspring (Mealey et al., 1999; Thornhill and Gangestad). In fact, low levels of fluctuating asymmetry have been associated with lesser susceptibility to toxins and infections (Polak and Trivers, 1994) and high levels with such disorders as schizophrenia and Down's syndrome (Markow and Wandler, 1986; Peretz et al., 1988). Facial characteristics, such as fluctuating asymmetry, that are indicative of genetic irregularities also produce negative trait impressions (Shaw, 1988; Symons, 1979). In Western culture, for instance, a crooked face is often linked to a crooked character (Zebrowitz et al., 1996); and, in Chinese culture, asymmetrical features indicate not only a person with a personality that is out of balance but quite possibly someone with a karmic debt (Kohn, 1996).

Other chromosomal aberrations such as Down's syndrome, Kabuki syndrome, Smith-Magenis syndrome, and so forth, which are associated with such functional defects as deafness, low intelligence, and congenital heart anomalies, are frequently marked by a distinctive set of facial abnormalities (Allanson and Cole, 1996; Allanson, Greenberg, and Smith, 1999; Allanson and Hennekam, 1997; Allanson, Hennekam, and Ireland, 1997; Allanson, Ledbetter, and Dobyns, 1998; Olney, Schaefer, and Kolodziej, 1998). Faces with features that resemble some of these markers, including above average fluctuating asymmetry, may also create impressions, such as low intelligence, that are associated with these disorders (Symons, 1979; Zebrowitz, 1998).¹⁶

5.6 THE OVERGENERALIZATION EFFECT OF EMOTION

The effect on trait impressions of morphological configurations suggestive of emotional states has not received as much attention as the overgeneralization effects discussed thus far, but there is evidence suggesting such morphological configurations play a significant role in the formation of trait impressions. Lorenz (1970-71), for instance, noted that people often inappropriately respond to features in other animals that suggest emotional states in human beings. A camel is a good case in point as it lifts its nose high in the air and half closes its eyes in a facial gesture most would recognize as contempt. Because

¹⁶ There is evidence, however, suggesting that the craniofacial youthfulness of genetic disorders such as Down's syndrome is partially responsible for low intelligence associations (Fidler and Hodapp, 1999).

people strongly react to this gesture, they have trouble viewing the camel as anything other than an unfriendly and highly disdainful creature. Human faces that structurally resemble features associated with emotional states likewise are regarded in ways appropriate to the emotional display. Take smiling for instance. People react positively to smiling faces and find such faces disarming and thus not very dominant (Keating et al., 1981a; LaFrance and Hecht, 1995; Mueller and Mazur, 1996). In fact, facial dominance significantly declines where even a slight smile is discernable (Mueller and Mazur). As expected, faces where the lips naturally turn upwards are likewise viewed more positively; such faces are considered friendly, kind, easygoing, and nonaggressive (Secord, Dukes, and Bevan, 1954). In a similar vein, faces that have features indicative of anger or hostility, for example, low-lying eyebrows, thin lips, and withdrawn corners of the mouth, are perceived to be more threatening, aggressive, and dominant (Keating, 1985; Keating et al., 1981b).

Figure 5.4. An Illustration of the Overgeneralization Effect of Emotion.



Note. A face with lips that naturally turn upwards (left) is perceived similarly to smiling faces, that is, as low in dominance, whereas a face with lips that turn downwards (right) is perceived as more threatening. In the two images, only the lips differ. These faces were generated using *FACES* by InterQuest and Micro-Intel.

The morphological characteristics of various emotional displays are well understood due in large part to the facial action coding system (*FACS*) developed by Ekman and Friesen (1978). *FACS* is a coding system for accurately describing any facial behavior, including emotion. However, since the person perception literature is rather barren in studies exploring the emotion overgeneralization effect, a complete review of *FACS* and other literature describing the morphology of emotional expression will not be presented. For a recent book on the subject of *FACS*, see Ekman and Rosenberg (1997). The interested reader may also wish to refer to chapters 3 and 4, which discuss some of the pros and cons in modeling artificial personality in terms of emotion, and chapter 6, which describes holistic face classification techniques that have successfully been used in identifying emotional displays.

5.7 INDIRECTLY MODELING TRAIT IMPRESSIONS OF THE FACE

After reviewing the person perception literature on the trait impressions of the face, it might seem that one effective way to model the trait impressions of the face for virtual agents would be to do so indirectly by modeling facial attractiveness, facial maturity, and emotion. There are several problems with this indirect approach. First, although facial maturity and emotion have successfully been modeled (using, in the case of facial maturity, the craniofacial profile shape transformation and, in the case of emotion, FACS), the characteristics of facial attractiveness, as seen in section 5.2, are not so clearly understood. Yet Rosenberg (1977) in his factor analysis of trait descriptors finds attractiveness an important trait category. Second, since facial attractiveness and facial maturity are associated with large clusters of traits, using these models to alter faces along the lines of specific trait impressions would prove difficult and clumsy. Third, no single model accounts for the overgeneralization effects of attractiveness, facial maturity, gender, and emotion. Knowing which overgeneralization model to employ would require a much more refined understanding of the types of traits that can and cannot be manipulated by altering the facial characteristics associated with the various overgeneralization effects. Finally, and most importantly, the models reviewed above, with the exception of emotion, are not suited to the task of classifying or predicting trait impressions of the face; that is, they do not address personality *perception*. It is true that FACS has been used in successful face recognition systems (Bartlett, 1998; Bartlett et al.,

1996; Essa and Pentland, 1997). Unfortunately, not enough is known regarding the overgeneralization effects of emotion to use face recognition systems based on FACS. Furthermore, as was pointed out in chapter 3, modeling personality solely in terms of emotion restricts the range of traits that can be modeled. What emotional displays, for instance, reflect intelligence or honesty? It will be conceded, however, that using existing models of facial attractiveness, facial maturity, gender, and emotion could provide virtual agents with effective adjunctive means of altering their facial presentations and merits further investigation.

For the purpose of this study a better approach to take in modeling the trait impressions of the face is to focus on the perception of those features that give rise to specific trait impressions. Explored hereafter is the possibility of modeling trait impressions of the face using standard holistic face classification techniques. Not only are these techniques predictive, but some are also capable of face synthesis.

CHAPTER 6: PCA FACE CLASSIFICATION TECHNIQUES

As the language of the face is universal, so it is very comprehensive. It is the shorthand of the mind, and crowds a great deal in a little room.

Jeremy Collier

6.1 SUMMARY

It is difficult to isolate those important facial features that hold the keys to an understanding of how faces can be classified. Holistic approaches, such as linear autoassociative neural networks, which allow the classifier system to identify relevant features a posteriori, have been shown to outperform systems that determine important feature sets a priori. It has been demonstrated that classifying faces using a linear autoassociative neural network is equivalent to finding the principal components of the cross-product matrix of a set of faces and reconstructing the faces as a weighted sum of eigenvectors. Since the principal components of an image can be derived statistically or by using an autoassociative neural network, both methods are reviewed in this chapter and their equivalency demonstrated.

6.2 INTRODUCTION

The face, as the English bishop Jeremy Collier (1650-1726) long ago noted, does indeed crowd a great deal of information in a little room. A face reveals the age, sex, and to some degree the physical condition of a person as well as the current emotional state and focus of attention. As noted in the last chapter, it forms impressions, whether grounded in fact or not, regarding a person's character that significantly influence the behavior of others. It is the center of oral communication, is highly plastic and expressive, and plays a vital role in the development of human relations and the self. That the face is able to convey such a wealth of information is all the more astonishing given the remarkable similarity between faces. Much of the visual information contained within a face is highly redundant. What varies is but a small set of relations between features and differences in textures, complexions, and shapes.

Isolating the features, therefore, that hold the keys to an understanding of how faces can be processed, whether by human beings or by machines, has proven a difficult task. Historically, the bulk of research has relied on measuring the relative distances between important facial key points: eye corners, mouth corners, nose tip, and chin edge (Brunelli and Poggio, 1993). Many modern researchers have continued to classify faces using facial feature measurements (Burton, Bruce, and Dench, 1993; Samal and Jyengar, 1992). Although this approach has the advantage of drastically reducing the number of variables, a major drawback is the difficulty in determining the best set of key points to

measure (Burton et al.; Valentin, Abdi, O'Toole, and Cottrell, 1994). Burton, Bruce, and Dench, for instance, have demonstrated the difficulty in discovering a useful set of features for discerning sex. An alternative approach is to process faces holistically (Brunelli and Poggio, 1993). Holistic techniques, such as template matching, preserve much of the information contained in the original images and are often preferred because they allow a classifier system to discover the relevant features a posteriori. Furthermore, template approaches have been shown to outperform feature-based systems (Lanitis, Taylor, and Cootes, 1997).

Two related forms of template matching that have achieved considerable success at classifying faces are linear autoassociative neural networks and a technique based on what is known as the Karhuen-Loève expansion in pattern recognition or *Principal Component Analysis (PCA)* in the statistical literature. Autoassociative memories associate input patterns with themselves (Valentin, Abdi, and O'Toole, 1994). Kohonen (1977) was one of the first to use a linear autoassociative neural network to store and recall face images. Sirovich and Kirby (1987) were the first to apply PCA to the data compression of faces and succeeded in economically representing faces in terms of an eigenpicture coordinate system. Turk and Pentland (1991a) adapted their techniques into what has now become a popular method of face classification.


Although other template matching techniques, for instance, Fisherfaces (Belhumeur, Hespanha, and Kriegman, 1997), Gabor Wavelet Representation (Philips, Wechsler, Juang, and Rauss, 1998), and Independent Component Analysis (Bartlett, 1998), are superior at some classification tasks, it is outside the scope of this study to discuss the relative merits of these and other face recognition and classification methods. The objective of this study is to investigate the feasibility of modeling trait impressions of the face. To date, no face classification methods have been applied to this task. However, because of the structural similarities between female faces and baby faces, the relation of attractiveness to average faces, and emotional expression to morphological facial characteristics that resemble the expressions associated with various emotions (see chapter 5), it is reasonable to assume that a linear autoassociative neural network or PCA can be employed to this end. Both have successfully been used to classify faces according to identity (Turk and Pentland, 1991a; Turk and Pentland, 1991b), gender (O'Toole and Deffenbacher, 1997; Valentin, Abdi, Edelman, and O'Toole, 1997), age (Valentin et al., 1994), race (O'Toole, Abdi, Deffenbacher, and Bartlett, 1991), and facial expression (Cottrell and Metcalfe, 1991; Padgett and Cottrell, 1998). What is more, they are simple, well understood, and, in the case of PCA, may be capable of generating novel images from within the eigenpicture coordinate system (Beymer, Shashua, and Poggio, 1993; Hancock, 2000). Linear autoassociative neural networks are described more fully in section 6.3, and PCA as a face classification technique is presented in section 6.4. That both are equivalent is demonstrated in section 6.5. As a result, the term PCA is

generally used in the literature to describe both approaches (Valentin, Abdi, O'Toole, and Cottrell, 1994).

6.3 AUTOASSOCIATIVE NEURAL NETWORKS AND FACE CLASSIFICATION

A linear autoassociative neural network or memory consists of one neural network layer, where each neuron in the layer, as shown in Figure 6.1, is connected with every other neuron but is associated with only one input element.

Figure 6.1. Architecture of a Linear Autoassociative Neural Network.



Note. All neurons are interconnected and the input pattern is associated to itself.

A linear autoassociative neural network can be described by the matrix equation

$$\mathbf{y} = \mathbf{W}\mathbf{x}, \quad (6.1)$$

where $\mathbf{y} = [y_1, y_2, \dots, y_N]^T$ is an output vector, $\mathbf{x} = [x_1, x_2, \dots, x_N]$ is an input vector, and

$\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_N]^T$ and $\mathbf{w}_i = [w_{i1}, w_{i2}, \dots, w_{iN}]^T$ describe the synaptic weights.

Three steps are involved in creating an autoassociative memory for faces:

Step 1: let $\mathbf{I}(x, y)$ be a two-dimensional array of intensity values of size $N \times N$.

Represent $\mathbf{I}(x, y)$ as a single point, that is, a one-dimensional pixel-vector \mathbf{x} of size N^2 .

Step 2: Let each element of the pixel-vector be an input to a neuron of the autoassociative neural network.

Step 3: Train the autoassociative neural network using a learning rule.

Kohonen (1977) uses a simple normalized Hebbian learning rule, where \mathbf{W} is obtained by successively autoassociating each face vector \mathbf{x}_k as follows:

$$\mathbf{W} = \sum_{k=1}^k \mathbf{x}_k \mathbf{x}_k^T, \quad (6.2)$$

As in Equation 6.1, the k^{th} face from the memory is obtained by premultiplying the face vector \mathbf{x}_k by \mathbf{W} . Therefore, the quality of the image reconstruction can be measured by simply computing the cosine of the angle between vectors \mathbf{y}_k and \mathbf{x}_k :

$$\cos(\mathbf{y}_k, \mathbf{x}_k) = \frac{\mathbf{y}_k^T \mathbf{x}_k}{\|\mathbf{y}_k\| \|\mathbf{x}_k\|}, \quad (6.3)$$

where $\|\mathbf{x}_k\|$ is the Euclidean norm of the vector \mathbf{x}_k . Perfect reconstruction would result in a cosine of 1.

The Hebbian learning rule (Equation 6.2) is capable of recalling a set of input vectors as long as they are mutually orthogonal; nonorthogonal vectors produce interference that reduces the correct response of the system (Valentin, Abdi, O'Toole, and Cottrell, 1994). Millward and O'Toole (1986) improved the performance of the linear autoassociative neural network using the Widrow-Hoff (Duda and Hart, 1973) or Delta Rule (McClelland

and Rumelhart, 1986). This algorithm corrects the difference between the actual output and the desired output by iteratively changing the weights \mathbf{W} as follows:

$$\mathbf{W}^{(t+1)} = \mathbf{W}^{(t)} + \eta(\mathbf{x}_k - \mathbf{W}^{(t)}x_k)\mathbf{x}_k^T, \quad (6.4)$$

where η is a learning constant and input vector \mathbf{x}_k is the k^{th} face and is randomly chosen.


Millward and O'Toole (1986) have shown that an autoassociative memory can classify faces by examining the response of the model using the cosine between the input and output vectors. Using pairs of faces, one known by the system and the other unknown, the face classified as *recognized* is that face that has the highest cosine value. Further refinements of this and other neural network models of face classification can be found in Valentin, Abdi, O'Toole, and Cottrell, (1994).

6.4 PCA AND FACE CLASSIFICATION

The central idea behind PCA is to find an orthonormal set of axes pointing in the direction of maximum covariance in the data. In terms of facial images, the idea is to find the orthonormal basis vectors or the eigenvectors of the covariance matrix of a set of images, with each image treated as a single point in a high dimensional space. It is

assumed that the facial images form a connected subregion in the image space. The eigenvectors map the most significant variations between faces and are preferred over other correlation techniques that assume every pixel in an image is of equal importance (Kosugi, 1995). Since each image contributes to each of the eigenvectors, they resemble ghostlike faces when displayed. For this reason, they are oftentimes referred to in the literature as *holons* (Cottrell and Fleming, 1990) or *eigenfaces* (Turk and Pentland, 1991a), and the new coordinate system is referred to as the *face space* (Turk and Pentland). In what follows, *eigenvectors* and *eigenfaces* are used interchangeably. Some examples of eigenfaces are shown in Figure 6.2.

Figure 6.2. The First 10 Eigenfaces of 220 Randomly Generated Faces.




Note. The eigenfaces are ordered left to right, top to bottom, by magnitude of the corresponding eigenvalue, of 220 randomly generated faces (see Appendix D, Figure D.1 for the original set of images and chapter 8 for details on the generation process).

Individual images can be projected onto the face space and represented exactly as weighted combinations of the eigenface components (see Figure 6.3). The resulting vector of weights that describe each face can be used in data compression, face classification, and synthesis. Data compression relies on the fact that the eigenfaces are ordered, with each one accounting for a different amount of variation among the faces. Compression is achieved by reconstructing images using only those few eigenfaces that account for the most variability (Sirovich and Kirby, 1987). This results in dramatic reduction of dimensionality. Classification is performed by projecting a new image onto the face space and comparing the resulting weight vector to the weight vectors of a given class (Turk and Pentland, 1991a; Turk and Pentland, 1991b). Finally, since projection is onto a low-dimensional face space, even images that look nothing like a face produce, as


Figure 6.4 illustrates, face-like reconstructions (Turk and Pentland, 1991a). Thus, it is possible to probe the face space by projecting a novel set of images onto the face space and then reconstructing. Future research may show that it is possible to take advantage of this property to synthesize faces with specific trait associations.

Figure 6.3. An Illustration of the Linear Combination of Eigenfaces.



Note. The face to the left can be represented as a weighted linear combination of eigenfaces.

Figure 6.4. Examples of Face-Like Reconstructions.



Note. The three images (left) are projected onto the eigenvectors extracted from 220 randomly generated composites. The resulting eigenvalues are then used to reconstruct the image (right). All three images result in face-like reconstructions.

The principal components of a set of images can be derived directly as follows. Again, let $\mathbf{I}(x, y)$ be a two-dimensional array of intensity values of size $N \times N$. $\mathbf{I}(x, y)$ may also be represented as a single point, a one-dimensional vector $\mathbf{\Gamma}$ of size N^2 . Let the set of face images be $\mathbf{\Gamma}_1, \mathbf{\Gamma}_2, \mathbf{\Gamma}_3, \dots, \mathbf{\Gamma}_M$. Let

$$\mathbf{\Phi}_k = \mathbf{\Gamma}_k - \mathbf{\Psi} \quad (6.5)$$

represent the mean normalized column vector for a given face $\mathbf{\Gamma}_k$, where

$$\Psi = \frac{1}{M} \sum_{k=1}^M \Gamma_k \quad (6.6)$$

is the average face of the set.

PCA seeks the set of M orthonormal vectors, \mathbf{u}_k , and their associated eigenvalues, λ_k , which best describes the distribution of the image points. The vectors \mathbf{u}_k and scalars λ_k are the eigenvectors and eigenvalues, respectively, of the covariance matrix

$$\mathbf{C} = \frac{1}{M} \sum_{k=1}^M \Phi_k \Phi_k^T = \mathbf{A} \mathbf{A}^T, \quad (6.7)$$

where the matrix $\mathbf{A} = [\Phi_1, \Phi_2, \dots, \Phi_M]$ (Turk and Pentland, 1991a).

The size of \mathbf{C} is N^2 by N^2 which for typical image sizes is an intractable task (Turk and Pentland, 1991a). However, since typically $M < N^2$, that is, the number of images is less than the dimension, there will only be $N-1$ non-zero eigenvectors. Thus, the N^2 eigenvectors can be solved, in this case, by first solving for the eigenvectors of an $M \times M$ matrix, followed by taking the appropriate linear combinations of the data points Φ (Turk and Pentland, 1991a).

PCA is closely associated with the **Singular Value Decomposition (SVD)** of a data matrix and can be decomposed as

$$svd(\Phi) = \mathbf{USV}^T. \quad (6.8)$$

The elements of the diagonal matrix \mathbf{S} are the singular values of Φ and are proportional to the square roots of the eigenvalues of \mathbf{C} , which are the first M columns of the orthogonal matrix \mathbf{U} .

Faces can be classified by projecting a new face Γ onto the face space as follows:

$$\omega_k = \mathbf{u}_k^T (\Gamma_a - \Psi) \quad (6.9)$$


for $k = 1, \dots, M'$ eigenvectors, with $M' \ll M$, if reduced dimensionality is desired. The weights form a vector $\mathbf{\Omega}_k^T = [\omega_1, \omega_2, \dots, \omega_{M'}]$, which contains the projections onto each eigenvector. Classification is performed by calculating the distance of $\mathbf{\Omega}_k$ from $\mathbf{\Omega}$, where $\mathbf{\Omega}$ represents the average weight vector defining some class (Turk and Pentland, 1991a). Turk and Pentland accomplish this for the task of face identification by simply computing the Euclidean distance $\varepsilon_k = \|\mathbf{\Omega} - \mathbf{\Omega}_k\|^2$.

Since the basis vectors of face space are those that optimally represent face images, the reconstruction of a facial image results in a relatively small reconstruction error compared to the reconstruction of nonfacial images. In other words, the reconstruction error, ε , should be within some threshold, θ_δ . When a face is projected onto the face space, it can lie, as Figure 6.5 illustrates, in four different regions:

1. Near face space and near some face class.
2. Near face space but not near some face class.
3. Distant from face space and near a face class.
4. Distant from face space and not near a known class.

Further refinements of PCA techniques and a review of similar techniques can be found in (Bartlett, 1998; Graham and Allison, 1998; Hancock, Burton, and Bruce, 1996; Liu, 1999).

Figure 6.5. A Simplified Illustration of the Four Results of Projecting an Image onto the Face Space.



Note. In this case, there are two eigenfaces (\mathbf{u}_1 and \mathbf{u}_2) and three classes (Ω_1 , Ω_2 , and Ω_3). Reprinted with permission from Turk and Pentland, 1991b, © 1991 IEEE.

6.5 EQUIVALENCY OF PCA AND AUTOASSOCIATIVE NEURAL NETWORKS

It can be demonstrated that a linear autoassociative neural network that classifies faces is equivalent to finding the principal components of the cross-product matrix of a set of faces and reconstructing the faces as a weighted sum of eigenvectors (Abdi, 1988; Cichocki and Unbehauen, 1993; Diamantaras and Kung, 1996; Oja, 1992; Valentin, Abdi, and O'Toole, 1994).

Since the weight matrix \mathbf{W} is positive semidefinite, it can be expressed as a weighted sum of its eigenvectors:

$$\mathbf{W} = \sum_{\ell}^L \lambda_{\ell} \mathbf{u}_{\ell} \mathbf{u}_{\ell}^T = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T \quad \text{with} \quad \mathbf{U}^T \mathbf{U} = \mathbf{I}, \quad (6.10)$$

where \mathbf{I} stands for the identity matrix, $\mathbf{\Lambda}$ represents the diagonal matrix of eigenvalues, \mathbf{U} represents the matrix of eigenvectors \mathbf{u}_{ℓ} of \mathbf{W} , and L is the rank of \mathbf{W} . Thus, the Widrow-Hoff learning rule (Equation 6.4) at a given time t can be rewritten as follows:

$$\mathbf{W}^{(t)} = \mathbf{U} \mathbf{\Phi}^t \mathbf{U}^T \quad \text{with} \quad \mathbf{\Phi}^t = \left[\mathbf{I} - (\mathbf{I} - \eta \mathbf{\Lambda})^t \right]. \quad (6.11)$$

The autoassociative network converges only if

$$\lim_{t \rightarrow \infty} (\mathbf{\Phi}^t) = \mathbf{0}. \quad (6.12)$$

that is, when η is smaller than $2\lambda_{\max}^{-1}$. Thus, at convergence the weight matrix reduces to

$$\mathbf{W}^{\infty} = \mathbf{U} \mathbf{U}^T. \quad (6.13)$$

CHAPTER 7: RESEARCH OBJECTIVES AND OVERVIEW

7.1 SUMMARY

It is predicted that holistic face classification techniques, specifically PCA, will satisfy the four requirements, presented in chapter 4, of a model of the trait impressions of the face that is suitable for virtual agents. In this chapter, a statement of the objectives and the limitations of the current investigation are provided, as well as a brief outline of the steps involved in training and testing a PCA model of facial trait impressions.

7.2 OBJECTIVE

The goal of this study is to investigate the feasibility of modeling physical personality by modeling the trait impressions of the face. As discussed more completely in section 4.3, to be useful to virtual agents a model of the trait impressions of the face must satisfy the following four requirements: 1) it must be capable of modeling a comprehensive set of trait impressions, 2) it must be predictive, as this would provide virtual agents with a perceptual system enabling them to perceive faces, much as human beings do, in terms of their trait impressions, 3) it must easily interface with mechanisms for creating, selecting, and managing the physical personality of the agent, and 4) it must be capable of

modeling the trait impressions of the face as formed by either the general population or specific groups of users.

It is predicted that PCA face classification techniques will eventually prove successful at modeling comprehensive sets of trait impressions and thus satisfy the first requirement. Because of the structural similarities between female faces and babyfacedness, the relation of attractiveness to average faces, and emotional expression to morphological facial characteristics that resemble the expressions associated with various emotions (see chapter 5), it is reasonable to expect that PCA face classification techniques are capable of modeling the clusters of traits associated with facial maturity and emotional displays. As noted in chapter 6, PCA face classification techniques have already proven successful at classifying faces according to facial expression (Cottrell and Metcalfe, 1991; Padgett and Cottrell, 1998), as well as two characteristics that are strongly correlated with facial maturity, namely, age (Valentin, Abdi, O'Toole, and Cottrell, 1994) and gender (O'Toole and Deffenbacher, 1997; Valentin, Abdi, Edelman, and O'Toole, 1997). Furthermore, because PCA face classification techniques are holistic, they should prove successful at modeling specific traits, even though at this time it is not known with certainty which features are involved in the formation of specific trait impressions.

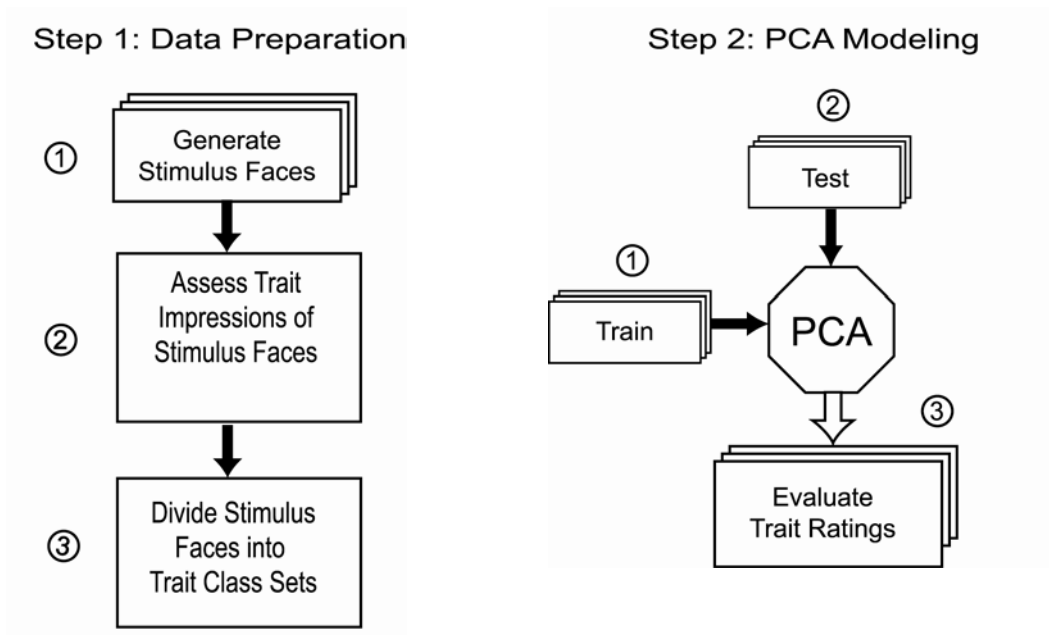
Provided that PCA proves capable of modeling a comprehensive set of trait impressions, the other requirements follow naturally from the properties of PCA. First, PCA is predictive. As a result, it would not only satisfy the second but also the third requirement that the model easily interface with mechanisms for managing physical personality. This is because the PCA would furnish the virtual agent with a perceptual system capable of perceiving the impressions different faces produce. Just as human beings often look in the mirror when adjusting their faces, a trait perception system might provide agents with a similar feedback system when creating their own physical personality or when altering their facial presentations along distinct trait lines. Finally, PCA learns from its training set. Therefore, depending on the composition of the training set, PCA might be able to learn to classify faces according to the impressions that faces evoke in varying groups of users. This would satisfy the fourth requirement. Thus, to reiterate, if it can be shown that PCA is predictive of facial impressions, then the other three requirements will be satisfied. The goal of this study is to show that PCA is predictive of facial impressions.

7.3 LIMITATIONS OF THE CURRENT STUDY

The traits under consideration in this study include seven selected from Rosenberg's trait categories (Rosenberg, 1977). They are listed below in section 7.4 and are primarily based on a modification reported in section 5.3.1 that was used by both Feingold (1992) and Eagly, Ashmore, Makhijan, and Longo (1991) in their meta-analyses of the literature

on the trait impressions of attractiveness. Also included in this study are the overgeneralization effects of attractiveness, facial maturity, and gender. Not explored at this time is the overgeneralization effect of emotion, as it has received little attention in the psychological literature. Finally, although the fourth requirement of modeling other classes of observers may be of interest to developers of virtual agents that are intended to serve specialized groups, this study is limited to modeling the facial impressions of a multicultural population of urban college students.

Figure 7.1. The Two Main Steps Involved in the Experimental Design.



7.4 OVERVIEW OF THE EXPERIMENTAL DESIGN

This section provides an overview of an experiment designed to model the trait impressions of the face using PCA. The experimental design is broken down into two general steps, Data Preparation and PCA Modeling, as represented in Figure 7.1.

7.4.1 STEP 1: DATA PREPARATION

The objective of step 1 is to obtain sets of faces clearly representative of three classes of attributions (high, low, and neutral) along 10 trait dimensions. As described in detail in chapter 8, 220 stimulus faces are randomly generated using a limited dataset of facial features found in the composite software program *FACES* (Freierman, 2000). The trait attributions of each face are then determined by having 10 human subjects classify the faces along the following trait dimensions: adjustment, attractiveness, dominance, facial maturity, intelligence, masculinity, sociality, warmth, trustworthiness, and degree of certainty regarding gender. Once the stimulus faces are rated, the 10 ratings for each face are averaged and ranked from low to high along each dimension. Subsets of faces representative of high, neutral, and low rankings are then selected to form three trait class sets. These trait class sets are then used in step 2 to train and test a separate PCA for each trait dimension.

7.4.2 STEP 2: PCA MODELING

The objective of step 2 is to train PCAs to match human classification of the stimulus faces. For this purpose the class trait sets developed in step 1 are divided into separate training and testing sets. The details of the PCA training and testing procedures are presented in chapter 9. The chapter concludes by presenting two analyses of the PCA models and a study designed to gauge the potential of synthesizing faces from within PCA trait space.

CHAPTER 8: DATA PREPARATION

8.1 SUMMARY

The objective of the data preparation process is to prepare the trait class sets of faces needed to train and test the PCA models. Data preparation is a three step process. In step one, 220 stimulus faces are randomly generated from a database of facial features. In step two, 110 subjects judge the stimulus faces along 10 trait dimensions using a 7-point bipolar scale. In step three, trait class sets are formed for each trait dimension by taking a subset of the 220 faces that are clearly representative of the two bipolar extremes and a neutral class of faces.

8.2 INTRODUCTION

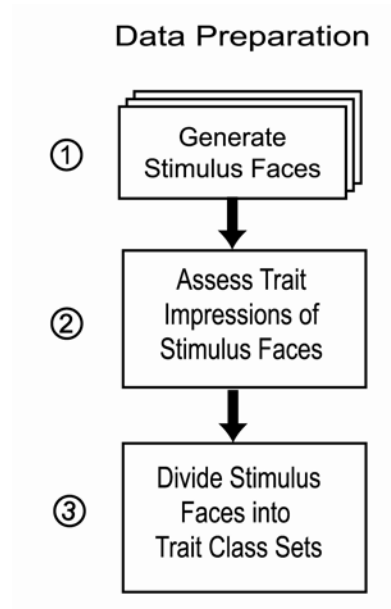
In order to train PCA to classify faces according to their trait impressions, it is necessary to acquire a suitable training set of faces. The FERET database (Philips, Wechsler, Juang, and Rauss, 1998) and others similar to it are available, but they are not appropriate for the purpose at hand. Most of these databases have been developed to evaluate face identification techniques. For this reason, the databases contain numerous photographs of a small set of individuals that vary in pose, lighting conditions, facial expression, and the addition of such occluding accessories as hats and glasses. Varying the noise in multiple images of the same individual is not, however, a requirement for this study. Rather, what is needed is a set of unique images representative of a broad range of facial types that are ideally normalized in terms of rotation, scale, and lighting. Moreover, since the goal is to model the trait impressions of morphology rather than emotion, the faces should be as neutral in expression as possible. Another consideration that needs to be taken into account involves the requirement that a model of facial impressions interface with a means of altering and generating faces.

To satisfy these conditions, permission was obtained to generate faces using the full database of photographs of facial features (eyes, mouths, noses, and so forth) found in the popular composite software program *FACES* (Freierman, 2000), produced by InterQuest and Micro-Intel. Generating faces from this database of features offers the following advantages: 1) a large number of unique faces can be randomly generated and easily

altered by manipulating individual features, 2) the skin tones of the features, in order to facilitate seamless combinations, are somewhat normalized, 3) the facial features in *FACES* are extracted from frontal photographs of people wearing a neutral expression, and 4) the lighting in the photographs is consistent and all backgrounds have been removed. It should be pointed out, however, that, although the facial features are extracted from photographs of individuals from many racial groups and both sexes, the features are not classified according to race or gender. This should not pose a problem as the concern of this project is not with accuracy but with the classification of faces according to the impressions they make.

As illustrated in Figure 8.1, data preparation is a three step process. Step 1, detailed in section 8.3, generates a large set of stimulus faces from the database of facial features found in *FACES*. Step 2, described in section 8.4, has human subjects rate the stimulus faces along 10 trait dimensions. Once the faces have been evaluated, the ratings are averaged and ranked, in step 3, from high to low along one of the bipolar trait descriptors. A subset of faces representative of the bipolar extremes and neutral rankings is then selected to form the trait class sets for that dimension. The details of this selection process are provided in section 8.5.

Figure 8.1. The Three Steps Involved in Data Preparation.



8.3 GENERATION OF THE STIMULUS FACES


A Windows 98 C++ program was written by the principal investigator to randomly generate stimulus faces from the database of facial features found in *FACES*. All complexion values were set to 190 in a gray scale of 256 to reduce the effects of race. One problem encountered in randomly producing faces from the full dataset of features was the production of too many unrealistic faces. Some faces, for instance, ended up with facial hair that covered the eyes while others had nose lines that hung beneath the chin. To reduce these incongruities, the features selected for constructing the stimulus faces were reduced to include only the full set of 512 eyes, 541 noses, 570 lips, 423 jaws, and 480 eyebrows provided in *FACES*. Excluded were all sets of facial lines, foreheads,

hair, and accessories. The elimination of some of the more peripheral features provided the additional benefit of reducing most personality impressions to facial morphology rather than such incidentals as hairstyle and accessories.¹⁷

In addition to reducing the feature set, faces were cropped, as illustrated in Figure 8.2, in such a way that missing hair was less noticeable. Eyebrows were retained because they have been found to contribute to impressions of gender and facial maturity (Brown and Perrett, 1993; Yamaguchi, Hirukawa, and Kanazawa, 1995; Zebrowitz, 1998), and faces that lack eyebrows within cropped faces might form the impression of having very high or light eyebrows. This could bias trait impressions as studies have shown that faces with higher, lighter, and thinner eyebrows are considered more feminine, submissive, and babyfaced (Cunningham, 1986; Keating, Mazur, and Segall, 1981b; Yamaguchi et al.; Zebrowitz and Montepare, 1992).

¹⁷ Modeling the personality impressions of hair styles and accessories is certainly a possibility worth exploring in future research.

Figure 8.2. An Example of an Acceptable Face (Left) and an Unusual Face (Right).



Even with these adjustments in place, numerous faces still appeared highly unrealistic. This was because the stimulus faces were randomly generated from a dataset that included both male and female features as well as a mix of races. In addition, some features belonged to large-bodied and others to small-bodied individuals. With the objective of obtaining an adequate number of acceptable faces with a good mix of features, 1500 faces were randomly generated. This number was arbitrarily chosen. The idea was simply to generate a large supply of faces. Unusual faces,¹⁸ such as the face to the right in Figure 8.2, were then pruned leaving approximately 1000 acceptable faces. From this pool, 220 stimulus faces were randomly selected. This number was based on the availability of human subjects to evaluate the faces.

¹⁸ As judged by the principal investigator.

8.4 AN EXPERIMENT ASSESSING THE TRAIT IMPRESSIONS OF THE STIMULUS FACES

In step 2 of the data preparation process, human subjects judge the 220 stimulus faces generated in step 1. This section describes an experiment that was designed to assess the trait impressions of the stimulus faces.

8.4.1 PARTICIPANTS

One hundred ten (54 male, 56 female) upper level undergraduate students were recruited from a large urban university to judge the stimulus faces. Each student received extra credit in a Computer Information Systems (*CIS*) course for participating in the study. The following demographic information was collected from the subjects: age, gender, socioeconomic status, ethnicity, and major. All were CIS majors, and the average age was 23.79 with a standard deviation of 3.73. The majority (see Figure 8.3) were Asian (57), followed by White (17), African (12), Hispanic (11), and other (13). The median annual income was between \$7,001 and \$15,000 (see Figure 8.4).

Figure 8.3. Ethnicity of (N=110) Participants.





Figure 8.4. Annual Income Distribution in Dollars of (N=110) Participants.



8.4.2 DEPENDENT MEASURES

The judges' perceptions of the stimulus faces were assessed through a 7-point bipolar trait scale along the following seven personality dimensions: dominant/submissive, masculine/feminine, intelligent/unintelligent, social/unsocial, adjusted/unadjusted, warm/cold, and trustworthy/untrustworthy, as well as the following three facial qualities: attractive/unattractive, babyfaced/mature-faced, and male/female, for a total of 10 dimensions.

8.4.3 APPARATUS

In the person perception literature on the face, it is typical for the experimenter to display the stimulus faces using a slide projector. Subjects are then given a paper questionnaire and asked to judge each face along a number of trait dimensions. Since the stimulus faces are computer generated and intended, for the most part, to be viewed on a monitor, desktop computers in a lab setting were used both to display the stimulus faces and to administer the questionnaires.

8.4.4 PROCEDURE

Small groups of subjects were invited to enter a computer lab containing approximately forty computers. The subjects were seated in front of computers, with every second seat

left empty to ensure privacy. A program written in C++ for Windows NT was used to administer the questionnaires. The opening screen, entitled *Welcome to Telling Faces*, requested that the participants in this study wait until instructed to press the *Start* button. They were presented an error message if the button was pressed prematurely. When all were seated, the subjects were asked to type a series of keys that unlocked the opening screen. They were then instructed to press the *Start* button.

The information on the second screen was read aloud by the principal investigator:

Hello, my name is Sheryl Brahnam, and I am the principal investigator of *Telling Faces*.

The ultimate research goal of *Telling Faces* is to enable computers to create faces with personality for such applications as interactive computer games, new user interfaces, and interactive movies on the web.

You are here to help in this endeavor. This program will ask you a number of questions. The first set of questions will be about you—your gender, age, ethnicity, major, and socioeconomic background. All your answers are completely anonymous. The rest of the program will ask you to judge a group of 20 faces according to a number of personality dimensions and facial qualities for a total of 40 questionnaires.

You do not have to participate in this survey and can quit at any time. If you have already participated in this survey this semester, please notify me immediately.

Wait until you are instructed to press the *Next* button.

The third screen presented the subjects with a form for filling out the demographic information (see Figure B.1 in Appendix B for a screen shot of this form). Once the

subjects had answered the demographic questions, they were presented with the fourth screen, which provided the following instructions that were once again read aloud:

Please read and follow all instructions carefully. If you have a question or a problem with the program, raise your hand.

If you want a more detailed explanation of any of the terms used, just click on the term for a definition. Note that we ask for both the gender and degree of masculinity or femininity seen in the face so you may want to review these definitions. Also the term *adjustment* has nothing to do with the image quality but refers to psychological and social adjustment.

Do not be alarmed if you see the same face twice, as they will be presented to you twice. All faces have had the hair removed so they may look a bit unusual. There is no right or wrong answer. We are interested in your “gut” impressions. Make the best ratings you can.

In the upper left-hand corner you will see how many questionnaires you have answered and how many more are left to go. You may get tired. If so, rest a moment. We would rather you exit the program entirely than start randomly selecting answers to finish faster. However, exiting the program prematurely will invalidate all your results. So do not click the "X" button in the caption bar unless it is your intention to exit without finishing.

The fifth screen informed the subjects that they could mark the 40 questionnaires that followed at their own pace, but that first they would be presented with a brief slide show of the faces they would be judging. The slide show allowed the subjects to familiarize themselves with the general format of the faces. Each subject then judged a set of 20 faces, randomly selected from the 220 stimulus faces, along the 10 dimensions listed above. Only five dimensions for each face were presented at a time (see Appendix B,

Figure B.2, for a screen shot of an example face questionnaire). This was a restriction imposed by screen size and the need to keep the facial image always visible while it was being rated. The ordering of the faces and the trait dimensions and facial qualities were completely randomized as were the bipolar trait descriptors associated with the end values of 1 and 7. This was done for two reasons: 1) to reduce the likelihood of ratings along one dimension influencing ratings along other dimensions and 2) to encourage the participants to pay closer attention to the terms when marking the forms.

While marking the questionnaires, a subject could review the instructions or any of the term definitions by clicking on the *Instructions* button. The appropriate term definition was also presented whenever a subject clicked on any of the terms. Appendix A, Table A.1 presents the 10 term definitions along with some helpful behavioral potential questions that were fashioned after Zebrowitz and Montepare (1992) and Berry and Brownlow (1989). Appendix B, Figure B.2, shows a sample questionnaire with the *Term Definition/Instruction* screen open.

After the subjects completed the questionnaires, they were thanked for their participation and handed an information sheet that asked them to refrain from discussing the experiment (see Appendix C for a copy of the information sheet).

8.4.5 RESULTS

Table 8.1 presents the mean ratings for each trait dimension as well as the standard deviations. The 10 ratings of the 220 faces (for a total of 2200 ratings per trait dimension) were averaged in terms of one bipolar descriptor to obtain the trait dimension mean. As displayed in Table 8.1, the trait means fell close to 4.00, except for the trait dimension of attractiveness. The low attractiveness ratings were most likely the result of the face generation process as described in section 8.3.

Table 8.1. The 10 Trait Dimension Means, Dimension Standard Deviations, and the Lowest and the Highest Face Means within Each Dimension.

<i>Trait Dimension Descriptor</i>	<i>Dimension Means</i>	<i>Standard Deviation</i>	<i>Lowest Face Mean</i>	<i>Highest Face Mean</i>
Adjusted	4.03	0.82	1.8	6.3
Attractive	3.23	0.90	1.4	5.7
Warm	3.94	1.01	0.5	5.2
Dominant	4.16	0.85	1.7	6.7
Intelligent	4.32	0.65	2.5	5.9

Table 8.1 Continued. The 10 Trait Dimension Means, Dimension Standard Deviations, and the Lowest and the Highest Face Means within Each Dimension.

<i>Trait Dimension Descriptor</i>	<i>Dimension Means</i>	<i>Standard Deviation</i>	<i>Lowest Face Mean</i>	<i>Highest Face Mean</i>
Male	4.56	1.51	1.2	7.0
Masculine	4.29	1.24	1.3	6.5
Mature-faced	4.65	0.89	2.4	6.7
Social	4.07	0.97	1.8	6.3
Trustworthy	4.00	0.87	1.6	5.9

Note. The 10 ratings of the 220 faces were averaged along one of the bipolar descriptors. The highest and lowest face mean is reported for each trait dimension. The dimension mean and standard deviation were computed using the 220 face means for that dimension.

Figure 8.5 shows a frequency histogram for each of the 10 trait dimensions. As the frequency histograms demonstrate, the ratings of most trait dimensions fall along a normal curve. The notable exceptions are attractiveness, as described above, and gender, where the ratings fall more at the two extremes. In general, the faces are also perceived to be slightly more mature-faced, intelligent, and male.

Figure 8.5. Frequency Histograms (A-J) for Each of the 10 Trait Dimensions.

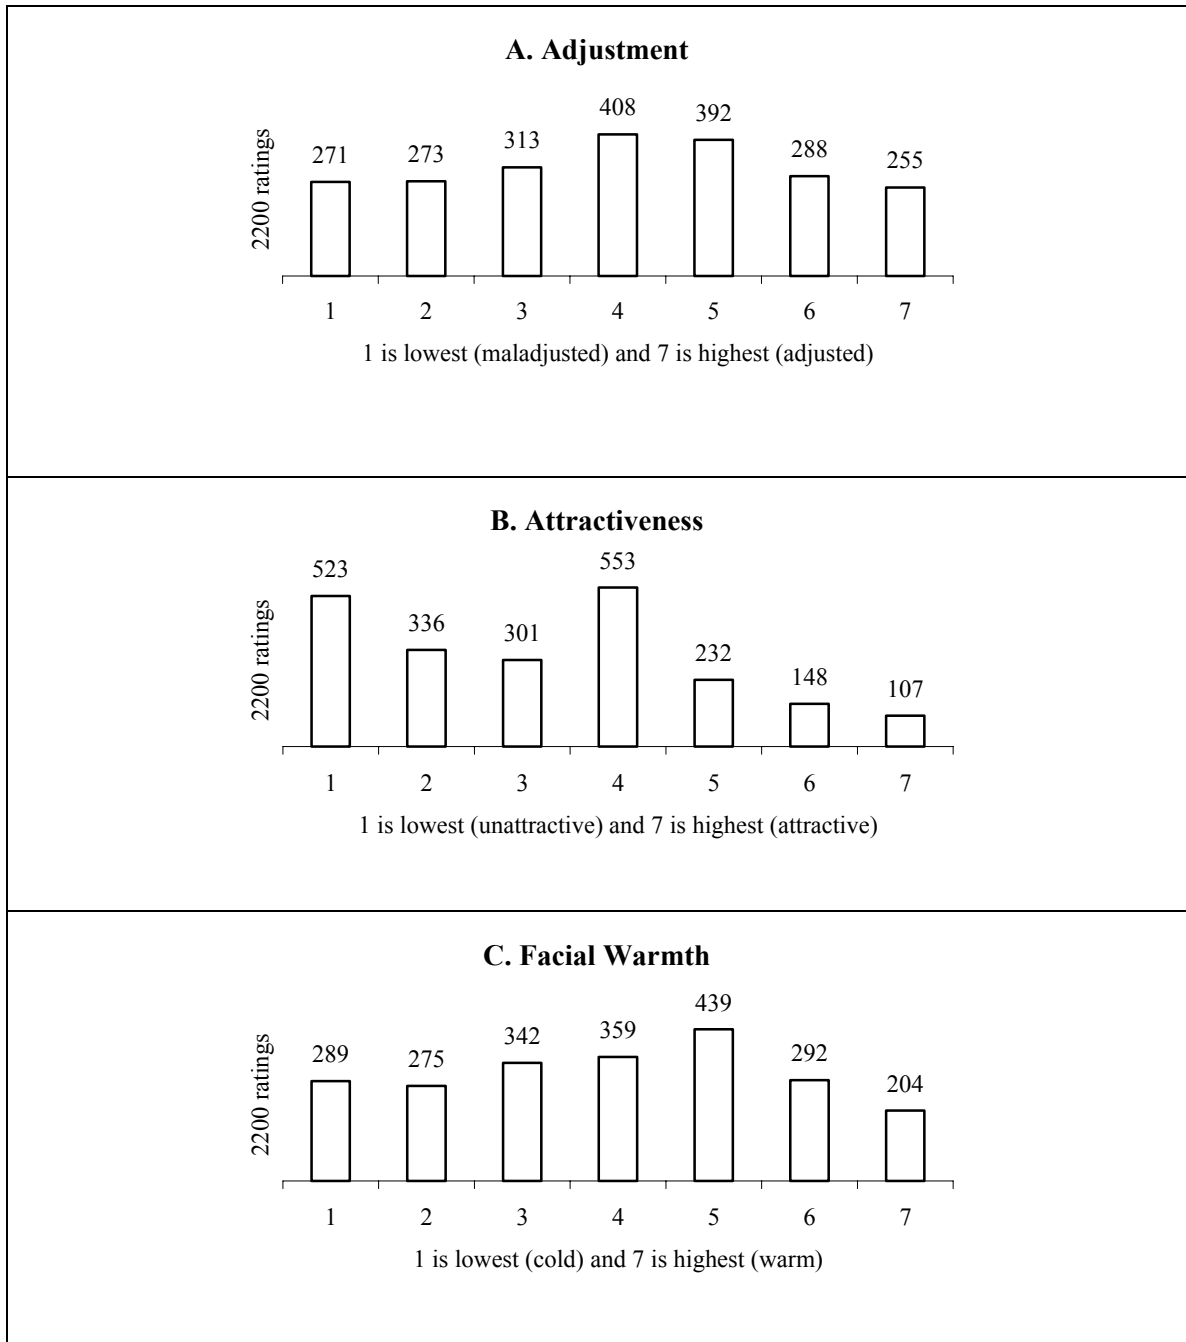


Figure 8.5 Continued. Frequency Histograms (A-J).

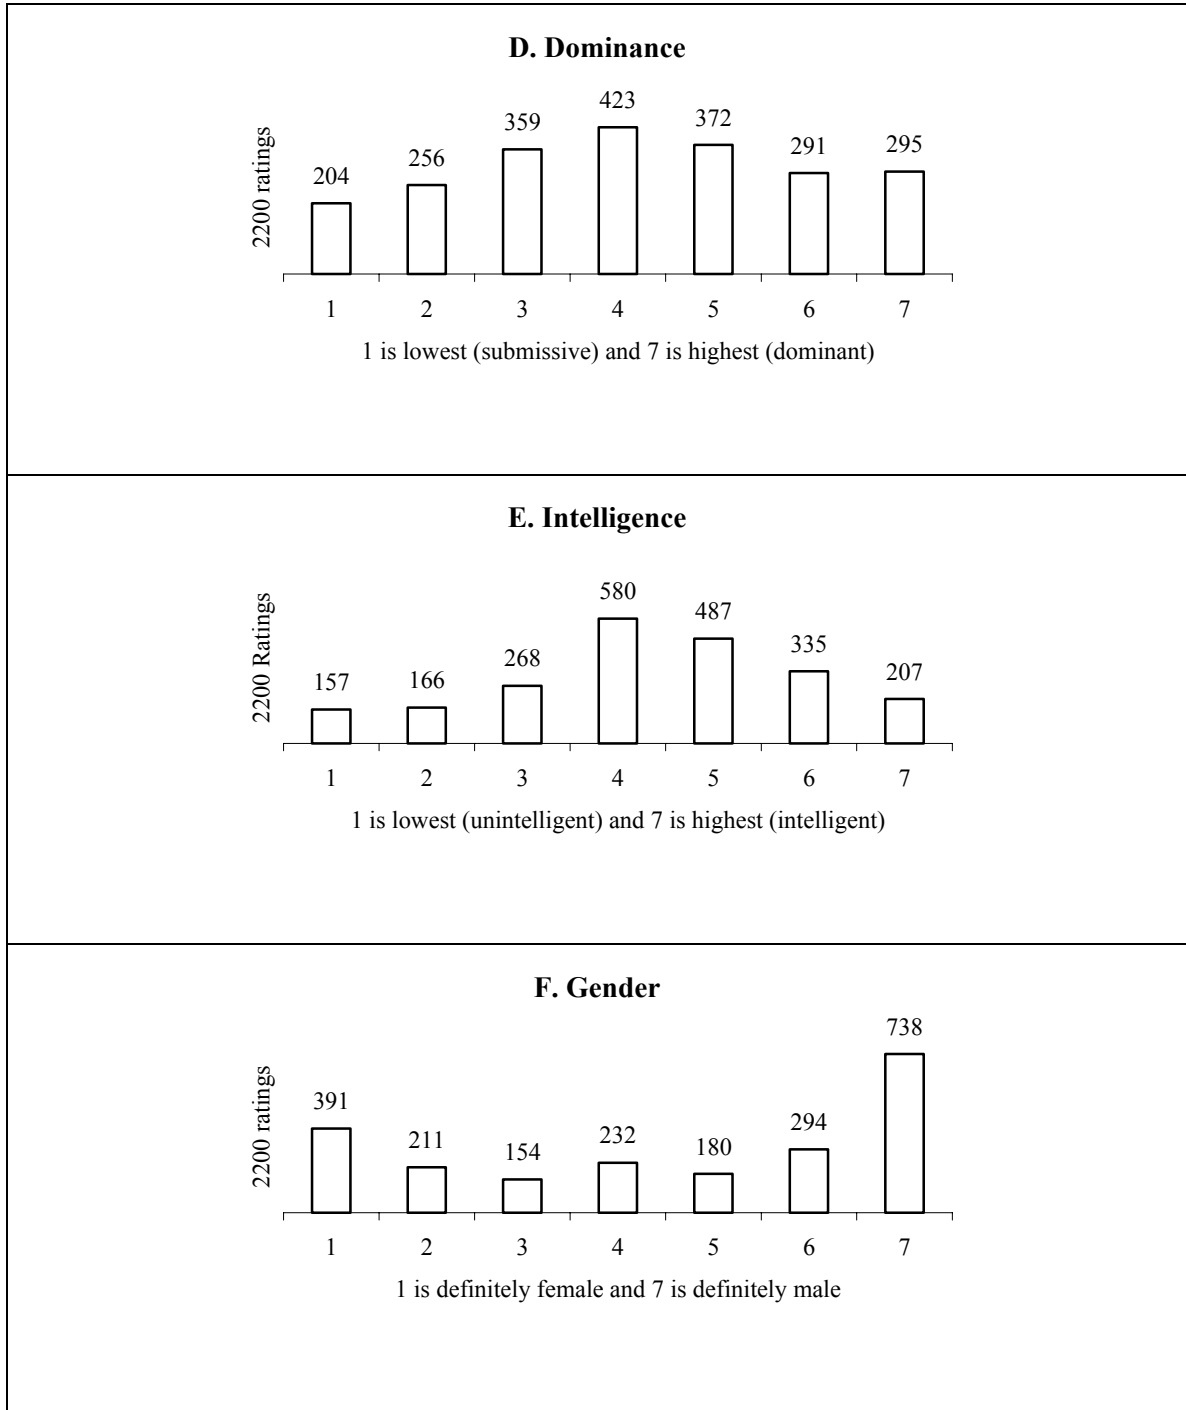


Figure 8.5 Continued. Frequency Histograms (A-J).

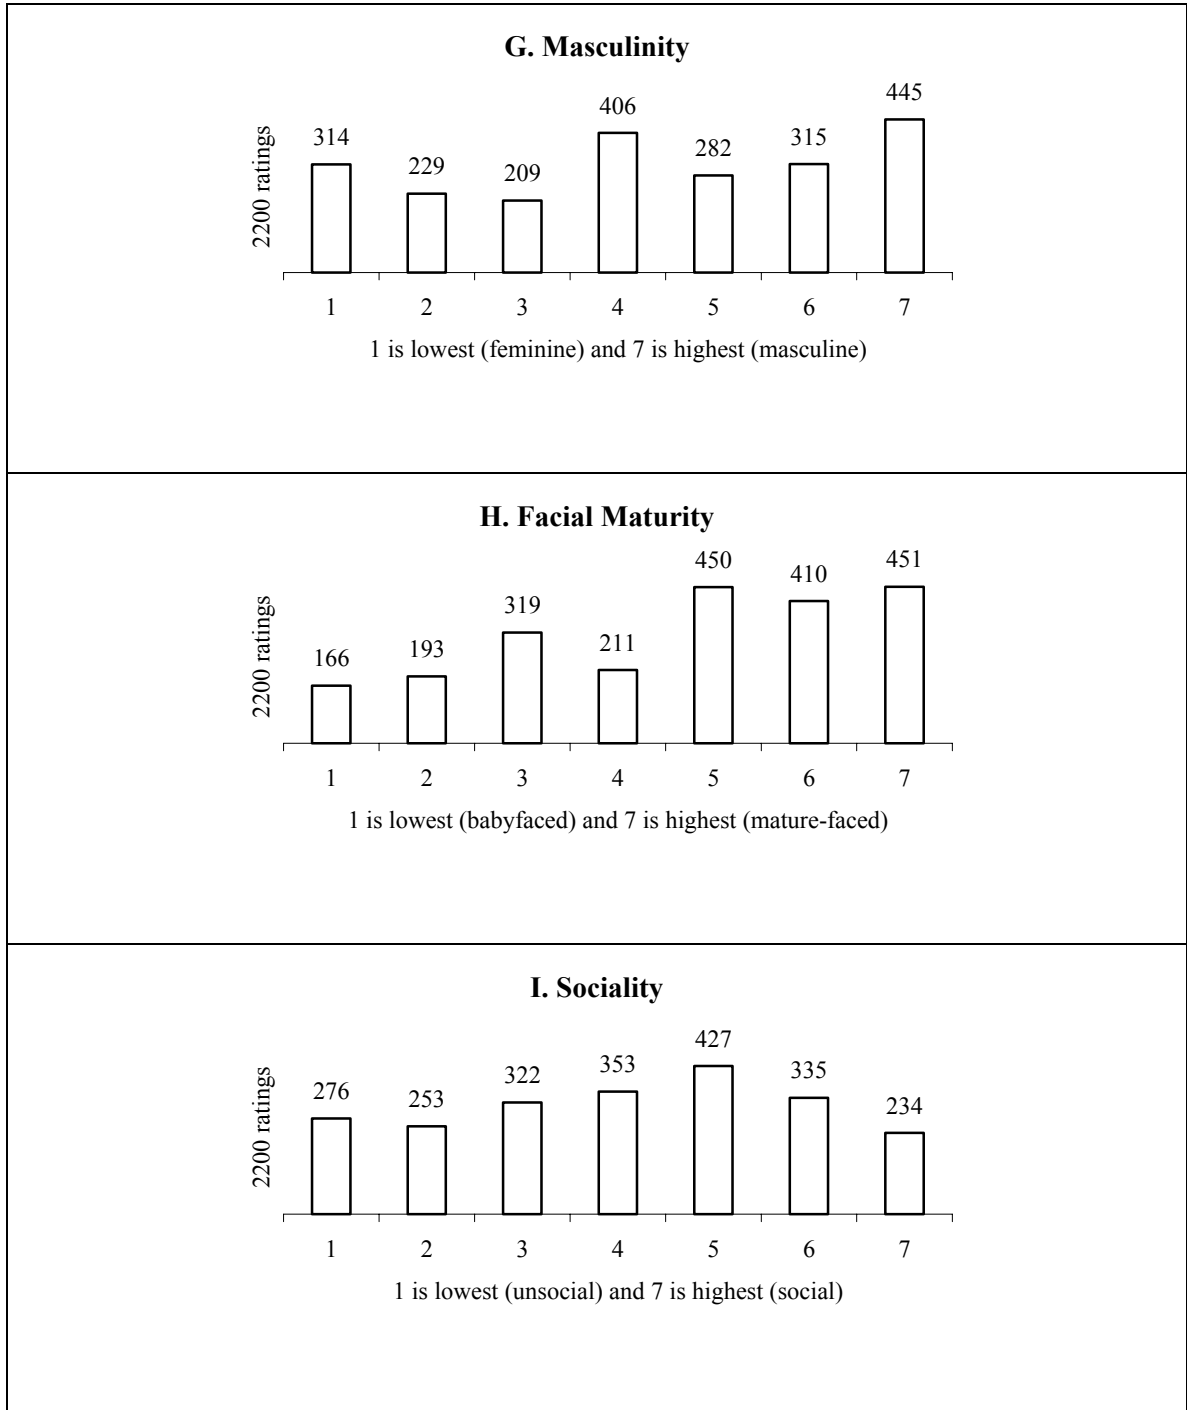
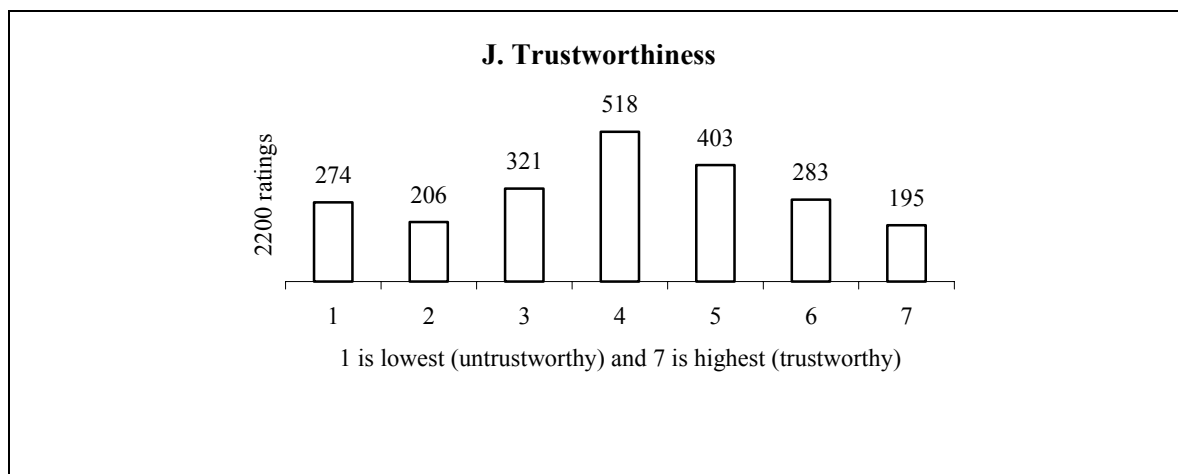


Figure 8.5 Continued. Frequency Histograms (A-J).



Note. Ten judges evaluated each of the 220 faces for a total of 2200 ratings per trait dimension.

8.5 DIVISION OF STIMULUS FACES INTO TRAIT CLASS SETS

In most face classification tasks, such as classifying faces by identity, the division of faces into relevant classes poses few problems as the classes are clearly definable. In the classification task of matching human impressions of faces, however, the division of faces into relevant trait classes is not a straightforward process. It is complicated by the fact that human beings are not in total agreement in their attributions and by the fact that many faces fail to elicit strong opinions. As a PCA classification of faces with weak

attributions is irrelevant for that trait dimension, that is, the classification is not unambiguous, faces with weak attributions should be excluded from the PCA training and testing sets. A PCA classification is relevant only in those cases where the consensus is pronounced and the opinions are strong. Described below is the method that was employed in forming the PCA trait classes. The objective was to obtain class sets composed of stimulus faces that are unambiguously relevant within each of the 10 trait dimensions of dominance, masculinity, intelligence, sociality, adjustment, warmth, trustworthiness, attractiveness, facial maturity, and degree of certainty regarding gender.

The trait class sets for the trait dimensions were obtained by first averaging the 10 ratings for each of the 220 faces and then ranking them from high to low along one of the trait descriptors. Faces were divided, based on their average rating, into three classes: low (with a mean range of 1.0 – 3.0), neutral (with a mean range of 3.1 – 4.9), and high (with a mean range of 5.0 – 7.0). In addition, the faces were pruned from all three classes that had a standard deviation greater than 1.5. The high and low classes were further restricted to include only those faces that tallied five ratings or more when the 1 and 2 point ratings were combined for the low class and the 6 and 7 point ratings were combined for the high class. Faces were also removed from both the low and high classes if the number of 4 point ratings totaled four ratings or more. Removed from the neutral class were all faces that tallied less than eight ratings when the 3, 4, and 5 point ratings were combined.

Table 8.2 lists the total number of faces that were selected for each class set along with the class mean range. The number of representative faces remaining in each class was remarkably low. In part this may have been due to the fact that the faces were constructed by randomly combining facial features. As noted in Table 8.2, intelligence and attractiveness had classes containing fewer than 10 faces. Because these dimensions failed to produce a sufficient number of faces to train and test a PCA, these two trait dimensions were not modeled in this study.

Table 8.2. Trait Class Sets: Mean Ranges and Number of Images Selected.

<i>Trait Dimension Descriptor</i>		<i>Mean Range</i>	<i>Number</i>
Adjusted	Low	1.8 – 3.0	11
	Neutral	3.5 – 4.6	15
	High	5.0 – 6.3	12
Attractive	Low	1.4 – 2.3	40
	Neutral	3.1 – 4.4	14
	High	5.1 – 5.7	8
Dominant	Low	1.7 – 3.0	11
	Neutral	3.5 – 4.5	10
	High	5.6 – 6.7	14

Table 8.2 Continued. Trait Class Sets: Mean Ranges and Number of Images Selected.

<i>Trait Dimension Descriptor</i>		<i>Mean Range</i>	<i>Number</i>
Intelligent	Low	2.5 – 3.0	4
	Neutral	4.0 – 4.3	13
	High	5.2 – 5.9	20
Male	Low	1.2 – 2.8	13
	Neutral	3.2 – 4.9	12
	High	6.3 – 7.0	16
Masculine	Low	1.3 – 2.5	14
	Neutral	3.3 – 4.9	12
	High	6.1 – 6.5	20
Mature-Faced	Low	2.4 – 3.0	11
	Neutral	3.1 – 4.9	27
	High	5.7 – 6.7	17
Social	Low	1.8 – 2.7	12
	Neutral	3.4 – 4.3	19
	High	5.6 – 6.3	14
Trustworthy	Low	1.6 – 2.6	10
	Neutral	3.4 – 4.3	11
	High	5.1 – 5.9	15
Warm	Low	0.5 – 1.7	14
	Neutral	3.3 – 4.8	23
	High	4.1 – 5.2	14

CHAPTER 9: PCA MODELING

9.1 SUMMARY

Detailed in this chapter are the steps taken to train and evaluate the PCA trait models. The performance of PCA classification for the eight trait dimensions of dominance, masculinity, sociality, adjustment, warmth, trustworthiness, facial maturity, and degree of certainty regarding gender is analyzed for three-class, and then for two-class PCAs. It is found that PCAs trained to classify two trait classes significantly outperform PCAs trained to classify three trait classes. The chapter concludes by reporting on a promising exploration in face synthesis.

9.2 INTRODUCTION

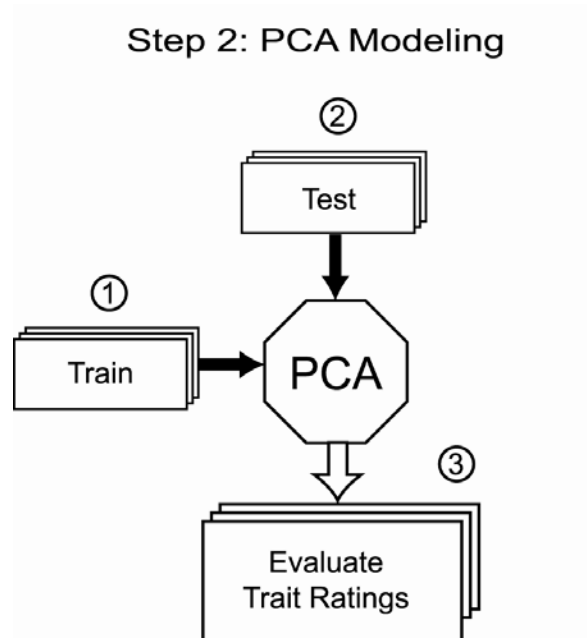
Using MATLAB (The MathWorks, 2000), separate PCAs are trained for each of the following trait dimensions: dominance, masculinity, sociality, adjustment, warmth, trustworthiness, facial maturity, and degree of certainty regarding gender.¹⁹ The three steps required to train and test a PCA are illustrated in Figure 9.1. In step 1, the trait

¹⁹ There are not enough faces in the trait classes of intelligence and attractiveness to train and test a PCA in those two dimensions. Modeling the impressions of intelligence and attractiveness will be explored in a future study.

class sets²⁰ are divided into separate training and testing sets of images, and a PCA is trained using the training set. The operations involved in step 1 are detailed in section 9.3. In step 2, the PCAs are tested using the testing set of images. The operations involved in testing the PCAs are described in section 9.4. In step 3, PCA performance is analyzed. Two analyses of PCA trait models are provided in section 9.5. The first analysis discusses the results of training PCAs with three classes of faces representative of low, neutral, and high ratings along one of the bipolar trait descriptors. The second analysis reports on the results of training PCAs using only images rated extremely low and high in a given dimension. This chapter concludes by reporting on an exploratory study that was designed to gauge the possibility of synthesizing faces with a high probability of producing specific trait impressions from within the PCA trait space.

²⁰ The trait class sets are described in chapter 8, "Data Preparation." Each set is composed of facial images representative of important points (high, neutral, and low trait ratings) along one of the bipolar trait descriptors.

Figure 9.1. The Three Steps Involved in PCA Modeling.



9.3 TRAINING OPERATIONS

A PCA is trained for each of the eight trait dimensions. Training a PCA requires three operations: 1) randomly dividing the trait class sets into separate training and testing sets, 2) calculating the eigenvectors from the training set, and 3) calculating the distribution of each class within the face space. These training operations are discussed in detail below.

Operation 1: Randomly dividing the trait class sets into separate training and testing sets.

Dividing the trait class sets into training and testing sets is a three step process. Let n be the number of classes within a given trait dimension. In step 1, the images within each trait class (obtained as described in section 8.5) are randomly divided into two sets, a class training set and a class testing set. Because training and testing employs repeated sampling due to the small number of images in the class trait sets, the size of the n class testing sets, x , is randomly chosen each iteration to be either two or three, and x is the same for all n testing sets. Testing images are randomly selected and removed from each of the n trait class sets. The remaining images form the training set for that class. In step 2, images are removed from each of the n class training sets in such a way that the size of each class training set, y , is maximized and y is the same for all n training sets. In step three, the n class training sets are merged to form a single training set, and the n class testing sets are merged to form a single testing set. The size of the training set is $y * n$, and the size of the testing set is $x * n$.

Operation 2: Calculating the eigenvectors from the training set.

The eigenvectors are computed using the following algorithm:

1. Reshape test images into column vectors, which together form matrix Φ . Let Γ_k represent the column vector of face k .

2. Normalize the column vector for each face k : $\boldsymbol{\varphi}_k = \boldsymbol{\Gamma}_k - \boldsymbol{\psi}$, where $\boldsymbol{\psi} = \frac{1}{M} \sum_k^M \boldsymbol{\Gamma}_i$.
3. Compute the eigenfaces using singular value decomposition, that is, $[\mathbf{U} \mathbf{S} \mathbf{V}] = \text{svd}(\boldsymbol{\varphi})$. \mathbf{U} are the eigenvectors or eigenfaces of the covariance matrix and \mathbf{S} are the eigenvalues.

Operation 3: Calculating the distribution of each class within the face space.

The distribution within the face space for each of the classes is computed by projecting each training image $\boldsymbol{\Gamma}_\alpha$ onto the eigenfaces as follows:

$$\omega_k = \mathbf{u}_k^T (\boldsymbol{\Gamma}_\alpha - \boldsymbol{\psi}) \quad (9.1)$$

Let $\boldsymbol{\Omega}^T = [\omega_1, \omega_2, \dots, \omega_M]$ be the weight vector that describes the contribution of each eigenvector in representing a face. A representative class vector is obtained by averaging the projected vectors, $\boldsymbol{\Omega}$, for each training class (Turk and Pentland, 1991a; Valentin, Abdi, Edelman, and O'Toole, 1997).

9.4 TESTING OPERATIONS

Evaluating the system using the testing set of images requires two operations: 1) projecting each test image $\boldsymbol{\Gamma}_k$ onto the face space to obtain $\boldsymbol{\Omega}_k$ as in Equation 9.1, and 2)

determining the best-fit class membership. Best-fit membership is determined by calculating the smallest Euclidian distance, d , of $\mathbf{\Omega}_k$ from $\mathbf{\Omega}_j$, where $\mathbf{\Omega}_j$ represents the average weight vector of the training images in some class j . The number of correct classifications is then averaged and used as an index to evaluate the performance of the system.

9.5 PCA EVALUATION

This section evaluates three-class and two-class PCAs trained to classify faces in terms of the following eight trait dimensions: adjustment, dominance, masculinity, sociality, warmth, trustworthiness, facial maturity, and degree of certainty regarding gender. The PCAs were trained and tested using the operations outlined in sections 9.3 and 9.4.

9.5.1 ANALYSIS OF THREE-CLASS PCAS

The three-class PCAs were trained to classify faces into three classes representing low, neutral, and high trait ratings. A jackknife technique (Valentin, Abdi, and O'Toole, 1994) was employed where 20 separate PCAs were trained for each trait dimension using varying combinations of training and testing sets as described in operation 1, section 9.3.

The ratio of right to wrong classifications was used as the classification index,²¹ and the classification indexes were averaged for a final classification score. Table 9.1 displays the averaged three-class PCA classification scores for the eight trait dimensions. In general, the results were poor although consistently better than chance.

Table 9.1. Averaged Classification Scores for the Three-Class PCAs.

<i>Trait Dimension</i>	<i>Classification Score</i>	<i>Trait Dimension</i>	<i>Classification Score</i>
Adjustment	.49	Masculinity	.67
Dominance	.46	Facial Maturity	.42
Warmth	.47	Sociality	.42
Gender	.58	Trustworthiness	.48

Note. The classification score was computed by averaging the performance ratios of 20 PCAs. All the three-class PCAs performed better than chance (.33).

The poor performance of the three-class PCAs is probably due to problems with the neutral class of faces. When the distances between the three averaged class projections within each trait dimension were measured, the class average of the neutral projections was found to be skewed significantly towards the class average of one of the other class

²¹ For a definition of the classification metric used in these analyses refer again to section 9.5.

projections (see Table 9.2 for an example). This may have had a negative impact on PCA performance. Another problem with the neutral class is that it may not represent a clear category. Subjects may have marked faces neutrally either because the faces produced trait impressions that were too faint to judge or because the faces produced conflicting impressions. Randomly generating faces from sets of facial features, as was done in this study, may have affected PCA performance by generating an unusually high number of faces that elicited conflicting trait impressions. It is not known at this time how three-class PCAs would perform if trained with faces not generated randomly from feature sets.

Table 9.2. Example of Distances between Averaged Class Projections.

<i>Trustworthiness</i>	<i>Class 1 (high)</i>	<i>Class 2 (low)</i>	<i>Class 3 (neutral)</i>
<i>Class 1 (high)</i>	0.0	1243.4	135.94
<i>Class 2 (low)</i>	1243.4	0.0	1107.4
<i>Class 3 (neutral)</i>	135.94	1107.4	0.0

Note. The Euclidian metric was used to measure the distances between the averaged projections of class 1, class 2, and class 3 in the trait dimension of trustworthiness. The averaged class 1 (high trustworthiness) projections and the averaged class 3 (neutral) projections lie close together in the trait space.

9.5.2 ANALYSIS OF TWO-CLASS PCAS

Because the neutral class of faces is problematical, PCAs were trained with two classes of faces. Except for the fact that all neutral images were removed from the training and testing sets of the two-class PCAs, the process of training and testing the two-class PCAs was identical to the process of training and testing the three-class PCAs. As was the case with three-class PCAs, the ratio of right to wrong classifications was used as the classification index for the two-class PCAs, and the same classification metric was employed (see section 9.4 for details). Similarly, the PCA classification indexes were averaged for a final PCA classification score. Table 9.3 displays the averaged two-class PCA classification scores for the eight trait dimensions. The two-class PCAs outperformed the three-class PCAs with classification scores that were significantly better than chance (.50) for all eight trait dimensions.

Table 9.3. Averaged Classification Scores for the Two-Class PCAs.

<i>Trait Dimension</i>	<i>Classification Score</i>	<i>Trait Dimension</i>	<i>Classification Score</i>
Adjustment	.71	Masculinity	.94
Dominance	.64	Facial Maturity	.70
Warmth	.89	Sociality	.70
Gender	.79	Trustworthiness	.78


Note. The classification score was computed by averaging the performance ratios of 20 PCAs. All the two-class PCAs performed better than chance (.50).

9.6 AN EXPLORATION INTO FACE SYNTHESIS FROM WITHIN THE PCA TRAIT SPACES

As noted in chapters 4 and 6, it may be possible to create novel faces with a high probability of producing specific trait impressions by synthesizing faces from within the PCA trait space. Although composite facial systems such as *SpotIt!* (Brunelli and Mich, 1996) have utilized the PCA face space to organize facial features in terms of their similarity, little work has been done in synthesizing faces directly from within the PCA face space. A notable exception is the pilot composite system developed by Hancock (2000) that evolves novel faces by recombining the eigenfaces using a genetic algorithm.

This exploration into face synthesis takes a different approach. In chapter 6, it was noted that it may be possible to probe PCA face space by projecting a novel set of images onto the face space and then reconstructing. Figure 9.2 (adapted from Figure 6.4) demonstrates the effect this has on an image that looks nothing like a face. When projected onto a PCA face space, the image of a penguin produces a face-like reconstruction (Turk and Pentland, 1991a).

Figure 9.2. An Example of Face Space Probing.



Note. The image of a penguin projected onto a PCA face space results in a face-like reconstruction.

This section describes an experiment that was designed to explore the possibility of synthesizing faces with a high probability of eliciting specific trait impressions by probing the appropriate PCA trait space. The experimental design can be broken down into two steps. In step 1, faces are synthesized, as described in section 9.6.1, by probing PCA trait spaces developed from images that were rated high and low along the following five trait dimensions: adjustment, dominance, warmth, sociality, and trustworthiness. In step 2, the synthesized faces are assessed by human subjects using the same experimental design described in section 8.4. The results of the assessments of the synthesized faces are reported in section 9.6.2.

9.6.1 FACE SYNTHESIS

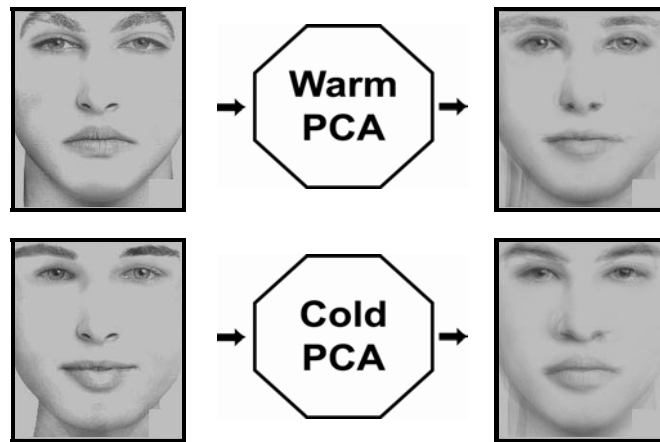
In this experiment, face synthesis is performed by probing two PCA trait spaces per trait dimension, that is, 10 PCAs are trained, one for each of the following bipolar extremes: unadjusted/adjusted, cold/warm, submissive/dominant, untrustworthy/trustworthy, and unsocial/social. In order to generate novel faces, the PCA face space needs to be seeded with as many faces as possible. For this reason, all images within each trait dimension with an average rating ≤ 3.0 along one of the bipolar descriptors are used to train what is hereafter referred to as the *low PCA trait space* for that dimension, and all images with an average score ≥ 5.0 are used to train what is referred to as the *high PCA trait space* for that dimension. Table 9.4 presents the total number of images that were used to train the 10 PCAs.

Table 9.4. Expanded Trait Class Sets for PCAs Used in Face Synthesis.

<i>Trait Dimension Descriptor</i>	<i>Number rated ≤ 3.0</i>	<i>Number rated ≥ 5.0</i>
Adjusted	25	26
Dominant	17	37
Warm	51	42
Social	45	34
Trustworthy	29	34

Novel images were generated by projecting the images rated ≤ 3.0 along one of the bipolar descriptors onto the high PCA trait space for that dimension and reconstructing the images. In the same way, images rated ≥ 5.0 along one of the bipolar descriptors were projected onto the low PCA trait space for that dimension and reconstructed. No attempt was made to remove artifacts introduced in the face synthesis process. Figure 9.2 shows two examples of face synthesis using the cold (low) and the warm (high) PCA trait spaces for facial warmth. Although a total of 340 faces were synthesized from the 10 PCA trait spaces, only 100 images were selected for human subject evaluation (see Table 9.5).

Figure 9.3. Examples of Face Synthesis by Probing the PCA Trait Space.



Note. An image rated cold (top left) is projected onto the warm (high) PCA trait space and reconstructed (top right). An image rated warm (bottom left) is projected onto the cold (low) PCA trait space and reconstructed (bottom right). It is predicted that the top reconstruction will be rated warm by human subjects and that the bottom reconstruction will be rated cold.

Table 9.5. Number of Novel Images Selected for Human Subject Assessment

<i>Trait Dimension</i>	<i>Number Synthesized from the low PCA trait space</i>	<i>Number synthesized from the high PCA trait space</i>
Unadjusted/Adjusted	9	9
Submissive/Dominant	8	8
Cold/Warm	22	22
Unsocial/Social	6	6
Untrustworthy/Trustworthy	10	10

Note. A total of 100 synthesized images were selected.

Before the synthesized faces were rated, the expanded set of faces (that is, all images with an average rating ≤ 3.0 or an average rating ≥ 5.0) was used to train and test a single two-class PCA for each of the five trait dimensions. The training and testing steps outlined in section 9.3 and 9.4 were followed with the exception that the number of testing images was enlarged to five. Table 9.6 displays the averaged classification scores for the expanded two-class PCAs. Since the performance of the expanded two-class PCAs was found to be comparable in performance to the two-class PCAs evaluated in section 9.5.2, it was predicted that faces synthesized by probing the low PCA trait spaces would be rated by human subjects at the lower end of the rating scale (≤ 3.5) and that faces synthesized by probing the high PCA trait spaces would be rated by human subjects at the upper end of the rating scale (≥ 3.5).

Table 9.6. Averaged Classification Rates for the Expanded Two-Class PCAs.

<i>Trait Dimension</i>	<i>Classification Rate</i>
Adjustment	.67
Dominance	.65
Warmth	.76
Sociality	.66
Trustworthiness	.81

Note. The classification rate was computed by averaging 20 PCA classification scores. Unlike the two-class PCAs evaluated in section 9.5.2, the two-class PCAs evaluated here used an expanded set of images. For each dimension, all images with an average rating of ≤ 3.0 or an average rating of ≥ 5.0 were used in training and testing the PCAs. All the PCAs performed greater than chance (.50).

9.6.2 THE TRAIT IMPRESSIONS OF THE SYNTHESIZED FACES

Fifty upper level undergraduate students were recruited from an urban university to judge the 100 synthesized faces in a lab setting using the same C++ program described in section 8.4.3 to administer the questionnaires and the same experimental procedures outlined in section 8.4.4. The students received extra credit in a Computer Information Systems course for their participation. As was the case in the evaluations of the original 220 stimulus faces, each subject judged a set of 20 faces randomly selected from the 100 synthesized images using a 7-point bipolar scale along the following 10 trait dimensions: dominant/submissive, masculine/feminine, intelligent/unintelligent, social/unsocial, adjusted/unadjusted, male/female, trustworthy/untrustworthy, attractive/unattractive, warm/cold, and babyfaced/mature-faced.

The 10 assessments of each face were then averaged for the trait dimension used to synthesize the face. Recall that it was predicted that faces synthesized by probing the low PCA trait space of a particular dimension would have averaged ≤ 3.5 in that dimension, and that faces synthesized by probing the high PCA trait space of a particular dimension would have averaged ≥ 3.5 for that dimension. Table 9.7 shows the average trait ratings for faces synthesized by probing the high and low PCA trait spaces of the five dimensions of adjustment, dominance, warm, sociality, and trustworthiness. Figure 9.4 presents the 10 frequency histograms. Except for faces synthesized by probing the low PCA trait space of adjustment, the human subjects rated the synthesized faces as predicted.

Table 9.7. Average Trait Ratings of the Synthesized Faces.

<i>Synthesized Trait Dimension</i>	<i>PCA Probe Space</i>	<i>Average Rating</i>
Adjustment	Low	3.86
	High	5.12
Dominance	Low	3.34
	High	5.00
Warmth	Low	1.98
	High	3.90
Sociality	Low	3.17
	High	5.40
Trustworthiness	Low	3.21
	High	4.69

Note. The assessments of all faces synthesized by probing the low and high PCA trait spaces were averaged for the trait dimension used in face synthesis.

Figure 9.4. Frequency Histograms of the Ratings of the Synthesized Faces.

Ratings of Faces Synthesized by Probing Low PCA Trait Spaces

Ratings of Faces Synthesized by Probing High PCA Trait Spaces



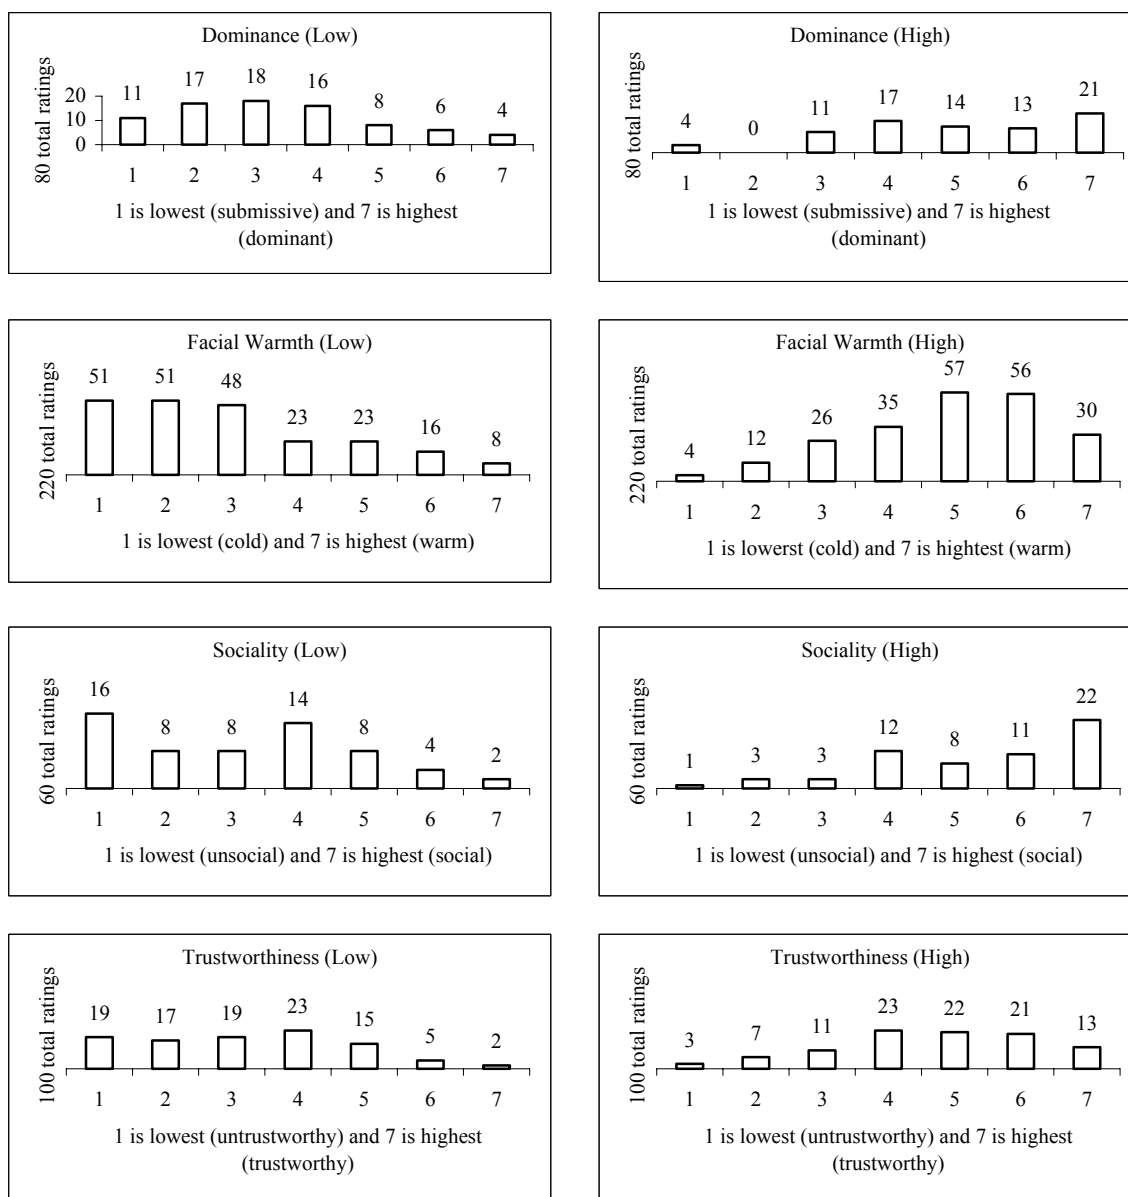


Figure 9.4 Continued. Frequency Histograms of the Ratings of the Synthesized Faces.

*Ratings of Faces Synthesized by Probing Low PCA Trait Spaces**Ratings of Faces Synthesized by Probing High PCA Trait Spaces*

Note. Low PCAs were trained with the stimulus faces that rated ≤ 3.0 , and high PCAs were trained with the stimulus faces with a rating ≥ 5.0 .

CHAPTER 10: CONCLUSION AND DIRECTIONS FOR FUTURE RESEARCH

Let us not hesitate to say it: there should be, deep in every actor, an element of the robot. The function of art is to lead this robot toward the natural; to proceed by artificial means toward an imitation of nature. It is because the violin is a hollow box, like a dead body, that it is so satisfying to furnish it with a soul.

Barrault

The increasing popularity of computer-generated characters in the movies has given rise to a host of desktop characters, or *virtual agents*. Although computer users have generally welcomed the prospect of interacting with virtual agents, more often than not these desktop characters have disappointed by behaving too mechanically and inconsistently. In a word, they have lacked the sparkling personalities that have made computer-generated characters in the movies so endearing and successful. Researchers, recognizing the importance of personality in sustaining user interest and in creating socially engaging virtual agents, have sought ways of endowing virtual agents with a convincing *artificial personality* (Trappl and Petta, 1997).

The intention of this study has been to expand this exciting new area of research by exploring the following: 1) Hampson's (1995) dramaturgical model of personality as a means of understanding and organizing the field of artificial personality, 2) the possibilities offered virtual agents by modeling the *perception of personality*, especially

the perception of the *physical personality* of the actor, and 3) the feasibility of modeling physical personality (both in terms of perception and synthesis) by modeling the trait impressions of the face using autoassociative neural networks or, equivalently, Principal Component Analysis (PCA). Section 10.1 summarizes these explorations and presents a few conclusions. Section 10.2 offers directions for future research.

10.1 SUMMARY AND CONCLUSIONS

As noted in chapter 3, there are many aspects to personality. One pressing concern in this area of research has been to define those aspects of personality that are of central importance for virtual agents. A distinguishing characteristic of virtual agents is that they are expected to interact with users. Unlike the computer-generated characters seen in the movies, virtual agents have the daunting task of sustaining their characterization while functioning autonomously within a highly volatile social environment. Thus, it is essential that artificial personality consider the social aspects of personality.

This study has explored the dramaturgical metamodel of personality developed by the social constructivists (Hampson, 1995) in an attempt to understand how artificial personality might be constructed in the course of social interaction. According to the social constructivists, personality is the product of three perspectives: that of the actor,

which involves the expression of an internal, psychological personality, that of the observer, which entails the perception and interpretation of observable manifestations of personality, and that of the self-observer, which concerns the management of self-presentations. The review of the reference literature presented in chapter 3 reveals that research into artificial personality has primarily been centered on the actor, that is, the expression of personality. The observational components of personality have received far less attention. Recent research has examined ways in which agents can observe users, but most work in this area has been concerned with the agent observing the user's emotional state rather than the user's personality. In addition, agents have been developed which abduct the personalities of other agents (Castelfranchi, de Rosis, and Falcone, 1997), and researchers have begun investigating how users attribute personality to computers (Reeves and Nass, 1996). No work to date, however, has explicitly addressed modeling the observational perspectives. As reported in chapter 3, there are serious drawbacks with the current one-sided focus on the actor. Offsetting this imbalance by emphasizing the observational perspectives would not only resolve some of these issues but also open up promising new areas of exploration.

Modeling the observational perspectives has been addressed in this study by exploring the possibility of modeling a manifestation of personality that immediately impresses itself upon the observer. It is an aspect of personality that is best described as the *physical personality*, since it is based on what is instantaneously projected by the actor's

physical appearance. Unlike psychological models of personality that are focused on the actor, models of physical personality require that attention be placed squarely on the observational perspectives or the *perception of personality*.

Although recently several studies and position papers have noted that the physical appearance of virtual agents plays a significant role in human-computer interaction (Donath, 2001; Sproull, Subramani, Kiesler, Walker, and Waters, 1996), no researcher to date has suggested that virtual agents be provided with a means of perceiving physical appearances in terms of the impressions they create or that an observer's trait impressions be modeled. A major benefit in providing virtual agents with such a model is that it would furnish the agents with a rudimentary social awareness sufficient enough to allow them to participate in the social construction of personality. Human observers are continuously caught up in the physical appearances of others, and people are equally preoccupied with managing their own appearances. If virtual agents could *perceive* physical personality in ways that mimic human observers, then these *virtual observers* could also respond with a degree of convincing social realism to the physical personality of others. Moreover, if virtual agents could see themselves as human observers see them, virtual agents could learn to take over the job of the artist and create for themselves socially intelligent embodiments.

Since modeling every facet of physical personality (clothing, hairstyle, posture, body type, walking style, and so on) is a major enterprise, the focus of this study has been limited to investigating the physical personality of the face. Not only does the face reveal evidence regarding the age, sex, physical condition, and current emotional state of an actor, but there is also compelling evidence in the person perception literature that facial morphology provides observers with a plethora of clues, whether accurate or not, regarding an actor's personality, disposition, and attitudes (Zebrowitz, 1998). Demonstrating the feasibility of modeling the trait impressions of the face would suggest that it is possible to model other aspects of physical personality as well.

As was shown in chapter 4, to be useful to virtual agents a model of the trait impressions of the face must satisfy the following requirements: 1) it must be predictive, as this would provide virtual agents with a perceptual system enabling them to perceive faces, much as human beings do, in terms of their trait impressions; 2) it must be capable of modeling a comprehensive set of trait impressions; 3) it must easily interface with mechanisms for creating, selecting, and managing the physical personality of the agent; and 4) it must be capable of modeling the trait impressions of the face as formed by either the general population or possibly specific groups of users. Although not listed as a requirement, because virtual agents function autonomously and interactively, it is assumed that the model will be capable of functioning in real time without human intervention.

In chapter 5, the possibility of modeling the trait impressions of the face using existing models of facial maturity, attractiveness, emotional expression, and other facial characteristics that are strongly associated with clusters of traits was considered. Particularly promising in this regard is the cardiodal strain transformation (Todd and Mark, 1980; Todd, Mark, Shaw, and Pittenger, 1981), which approximates real growth when applied to standard profile shapes. There is strong evidence, for instance, that as craniofacial profile maturity decreases, so do perceived alertness, reliability, intelligence, and strength (Berry and McArthur, 1986). It was determined, however, that although these indirect models might provide adjunctive means of manipulating the trait impressions of faces, these models are inadequate for virtual agents primarily because they are not predictive.

A more direct approach to take is to model the *perception* of those facial features that give rise to specific trait impressions. Since it is not known at this time which facial features produce specific trait impressions, it is best to approach this problem using holistic face classification techniques. Holistic approaches allow the classifier to identify relevant features a posteriori. It was noted in chapter 6 that one holistic technique, PCA, is particularly promising because it has proven capable of classifying faces according to gender and age, facial characteristics that were shown in chapter 5 to be strongly

correlated with large clusters of traits. Furthermore, as was discussed in chapter 7, PCA models are not only predictive but also capable of face synthesis (Hancock, 2000).

Chapters 7-9 describe an investigation that was undertaken to explore the performance of PCA in this task domain. The main focus was on PCA trait classification, but face synthesis within the PCA trait spaces were also explored. The experimental design for face classification was broken down into two steps. First, as described in chapter 8, 220 stimulus faces were randomly generated using a limited dataset of facial features found in the composite software program *FACES* (Freierman, 2000). The trait impressions of each face were then determined by having 10 human subjects classify the faces along the following trait dimensions: adjustment, attractiveness, dominance, facial maturity, intelligence, masculinity, sociality, warmth, trustworthiness, and degree of certainty regarding gender. Once the stimulus faces were rated for each face, they were averaged and ranked from low to high along one bipolar trait descriptor within each trait dimension. Subsets of faces clearly representative of high, neutral, and low rankings were then used to form three class sets for each trait dimension. These trait class sets were used in step 2 to train and test the PCAs. Perhaps because of the methods used to generate the faces, the number of representative faces in each class was small, between 4 and 40. Because intelligence and attractiveness had classes containing fewer than 10 faces, no attempt was made to model these two trait dimensions.

In chapter 9, the steps taken to train and test PCAs along the eight remaining dimensions were described. Two analyses of PCA performance were presented. The first analysis discussed the results of training PCAs along each dimension using three classes of faces representative of low, neutral, and high ratings along one of the bipolar trait descriptors. It was found that although three-class PCAs performed better than chance, the performance was poor. The poor performance of the three-class PCAs probably stemmed from problems with the neutral class of faces. When the distances between the three averaged class projections within each trait dimension were measured, for instance, the class average of the neutral projections was found to be skewed significantly towards the class average of one of the other class projections. This may have had a negative impact on PCA performance. Another problem with the neutral class is that it may not represent a clear category. Subjects may have marked faces neutrally either because the faces produced trait impressions that were too faint to judge or because the faces produced conflicting impressions. Randomly generating faces from sets of facial features, as was done in this study, may have affected PCA performance by generating an unusually high number of faces that elicited conflicting trait impressions.

The second analysis reported on the results of training PCAs using only images rated extremely low and high within a given dimension. It was found that two-class PCAs

significantly outperformed three-class PCAs. The performance of the two-class PCAs is promising. It suggests that the trait impressions of the face can be modeled by holistic face classification techniques.

Finally, as mentioned above, an exploratory study was undertaken to gauge the possibility of synthesizing faces with a high probability of producing specific trait impressions from within the PCA trait space. This investigation involved two steps. In step 1, faces were synthesized by probing PCA trait spaces developed from images that were rated high and low along the following five trait dimensions: adjustment, dominance, warmth, sociality, and trustworthiness. As described in detail in section 9.6.1, PCA trait spaces were probed by projecting faces not used in the training process onto the trait space and reconstructing. It was predicted that faces synthesized by probing the low PCA trait space of a particular dimension would be ranked by human subjects at the low end of the trait dimension rating scale and that faces synthesized by probing the high PCA trait space of the same dimension would be ranked at the high end of the scale. In step 2, the synthesized faces were assessed by human subjects. Except for faces synthesized by probing the low PCA trait space of adjustment, the human subjects rated the synthesized faces as predicted. This exploratory study suggests that it will be possible to synthesize faces for virtual agents that will succeed in eliciting specific trait impressions.

10.2 DIRECTIONS FOR FUTURE RESEARCH

There are a number of directions that offer promising avenues for future exploration. Chapter 4 has listed a number of ways a virtual agent could make use of a model of physical personality, both in terms of perceiving physical personality and in terms of creating socially intelligent embodiment. Before these application possibilities can be explored, however, more work needs to be done modeling physical personality. Provided in this section are a few directions for future research that are centered specifically on modeling the trait impressions of the face.

10.2.1 BUILDING A SPECIALIZED DATABASE OF FACES

In order to build better models of the trait impressions of the face, it is essential to acquire a large database of faces clearly representative of various traits. It would probably be best to use actual photographs of faces rather than random composites as was done in this study. Faces in the database will need to be judged along a number of representative trait dimensions by human subjects, and demographic data of the subjects will need to be collected.

10.2.2 EXPLORING TRAIT CLASSIFICATION

The results of this study suggest that it may be possible to match the human classification of faces into trait categories. As it was not the intention of this study to perform a complete analysis of PCA in this task domain, there is much work that needs to be done in this area. First, it may be possible to use the classification metric as a means of measuring confidence in the PCA classification of faces into high and low trait categories. Future studies might investigate whether low confidence measures are predictive of faces rated neutrally by human subjects. Second, classification performance could probably be enhanced by exploring other holistic face classification techniques. Especially promising is an approach based on Fisher's linear discriminants (Belhumeur, Hespanha, and Kriegman, 1997), a supervised learning procedure that projects the images onto a subspace that maximizes the between-class scatter and minimizes the within-class scatter of the projected data. Another promising approach is Independent Component Analysis (Bartlett, 1998). It is a generalization of PCA that separates the higher-order moments of the input in addition to the second-order moments. Third, it would be interesting to investigate the impact different eigenvectors have on trait classification.

10.2.3 SYNTHESIZING FACES

Work has just begun in synthesizing faces from within PCA face space. The results reported in section 9.6, suggest that it may be possible to synthesize faces with a high

probability of eliciting specific trait impressions in users. This opens up the exciting prospect of the virtual agent taking over the job of the artist and generating its own socially intelligent embodiment. The possibility of synthesizing faces within face space has only recently been explored (Hancock, 2000). Further explorations in synthesizing faces within trait spaces might begin, for instance, by having human subjects rate faces that were synthesized by applying Hancock's techniques.

**APPENDIX A: TERM DEFINITIONS AND BEHAVIORAL POTENTIAL
QUESTIONS**

Table A.1 below provides the definitions and behavioral potential questions (Berry and Brownlow, 1989; Zebrowitz and Montepare, 1992) that were available to subjects filling out the computerized questionnaires (see chapter 8).

Table A.1. Definitions and Some Behavioral Potential Questions.

<i>Term Definition</i>	<i>Behavioral Potential Questions</i>
<p><i>Attractive, Unattractive, and Average</i> Are used here as they are generally used in every day discourse.</p>	A helpful question might be: "Does s/he look like someone you would arrange a blind date with for a friend who likes good looking men or women?"
<p><i>Babyfaced, Mature-Faced, Uncertain</i> Asks for the degree of certainty in your perception of either babyfacedness or maturity of the face.</p> <p><i>*Babyfaced</i> As the phrase suggests, "babyfaced" indicates there is something that reminds you of a baby or child in the facial features. Both men and women can be babyfaced.</p> <p><i>*Mature-Faced</i> The resemblance to a child or baby is missing in a mature-faced individual.</p>	(None offered)
<p><i>Male, Female, Uncertain</i> We are simply asking you to guess the sex of the individual and how confident you are in your judgment.</p>	(None offered)

Table A.1 Continued. Definitions and Some Behavioral Potential Questions.

<p><i>Dominant, Submissive, Uncertain</i> Here we are looking at how dominating the person is.</p> <p><i>*Dominant</i> Is a person who is most likely to tell other people what to do.</p> <p><i>*Submissive</i> Is a person who usually follows orders and is not very assertive.</p>	<p>A helpful question might be: "Does s/he look like someone who would be the kind of roommate who would comply with most of your wishes about the furniture arrangement, quiet hours, and house rules?"</p>
<p><i>Trustworthy, Untrustworthy, uncertain</i></p> <p><i>*Trustworthy</i> Is a person who is mostly honest and who is not likely to steal, lie, or cheat.</p> <p><i>*Untrustworthy</i> Is a person who is often not honest and who possibly steals, lies, or cheats.</p>	<p>A helpful question might be: "Does s/he look like someone you would ask to watch your backpack while you made a quick visit to the restroom?"</p>
<p><i>Intelligent, Unintelligent, Uncertain</i></p> <p><i>*Intelligent</i> Is a person who is possibly very educated, capable, and interested in intellectual work.</p> <p><i>*Unintelligent</i> Is a person who probably does not value school because s/he was not good at school subjects.</p>	<p>A helpful question might be: "Does s/he look like someone you would learn from when discussing such topics as art, politics, philosophy, or science?"</p>
<p><i>Social, Unsocial, Uncertain</i> Here we are looking at how social the person is.</p> <p><i>*Social</i> Is a person who is very outgoing, extroverted, and who enjoys parties and other social activities.</p> <p><i>*Unsocial</i> Is a person who is introverted, a loner, and who would prefer to stay home rather than go out.</p>	<p>A helpful question might be: "Does s/he look like someone who would attend a school dance or party?"</p>

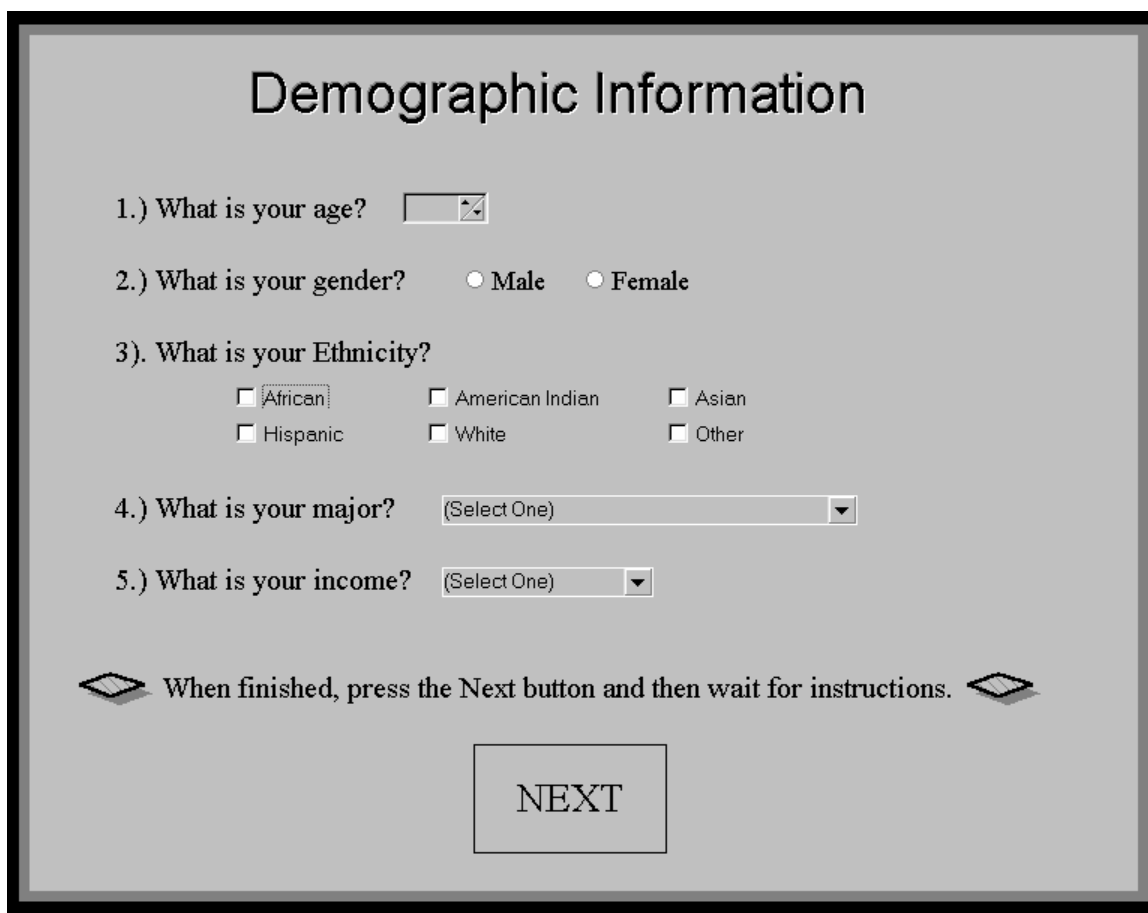
Table A.1 Continued. Definitions and Some Behavioral Potential Questions.

<p><i>Adjusted, Unadjusted, Uncertain</i> Here we are looking at how mentally healthy and adjusted the person is.</p> <p><i>*Adjusted</i> Is person who is fairly happy, mentally healthy, and who feels s/he belongs to society.</p> <p><i>*Unadjusted</i> Is a person who is unhappy or discontent, possibly even mentally ill or troubled, and who feels like an outsider.</p>	(None offered)
<p><i>Warm, Cold, Uncertain</i> Here we are looking at how approachable the person is.</p>	A helpful question might be: "Does s/he look like someone who would turn a cold shoulder to your attempts at friendly conversation?"
<p><i>Feminine, Androgynous, Masculine</i> Here we are not asking the sex of the individual but rather how masculine or feminine the person looks. Androgynous is equally feminine and masculine although maybe fairly clearly male or female.</p>	(None offered)
<p>Note. Some of these questions were adapted from (Berry and Brownlow, 1989; Zebrowitz and Montepare, 1992)</p>	

APPENDIX B: SCREEN SHOTS OF THE COMPUTERIZED QUESTIONNAIRE

Below are two screen shots (minus black background and caption bar) of the computerized questionnaire (see chapter 8). The first screen shot (Figure B.1) is of the Demographic Information form. The second (Figure B.2) presents an example face questionnaire with the Term Definitions/Instructions screen open. In this case, the subject has clicked on one of the terms in Q1 in order to obtain a definition.

Figure B.1. Demographic Information Form.



The screenshot shows a form titled "Demographic Information" with five numbered questions. Question 1 is "What is your age?" with a dropdown menu. Question 2 is "What is your gender?" with radio buttons for "Male" and "Female". Question 3 is "What is your Ethnicity?" with checkboxes for "African", "Hispanic", "American Indian", "White", "Asian", and "Other". Question 4 is "What is your major?" with a dropdown menu showing "(Select One)". Question 5 is "What is your income?" with a dropdown menu showing "(Select One)". At the bottom, there is a instruction: "When finished, press the Next button and then wait for instructions." and a "NEXT" button.

Demographic Information

1.) What is your age?

2.) What is your gender? Male Female

3.) What is your Ethnicity?

African American Indian Asian
 Hispanic White Other

4.) What is your major?

5.) What is your income?

When finished, press the Next button and then wait for instructions.

NEXT

Figure B.2. An Example of a Face Questionnaire with the Term Definitions/ Instructions Screen Open.

The screenshot shows a questionnaire interface with a central image of a woman's face. At the top left is an 'Instructions' button with a downward arrow, and at the top right is a 'Next' button with a rightward arrow. Below the image are five question items (Q1-Q5), each with a 7-point Likert scale and two opposite adjectives. A 'Term Definitions' pop-up window is open over Q2, displaying definitions for 'Baby Faced' and 'Mature Faced'.

Q1. 1 2 3 4 5 6 7
Baby Faced **Uncertain** **Mature Faced**

Q2. 1 7
Adjusted **Maladjusted**

Q3. 1 7
Intelligent **Unintelligent**

Q4. 1 7
Cold **Warm**

Q5. 1 7
Social **Unsocial**

Term Definitions . . .

Baby Faced, Mature Faced, Uncertain
 Asks for the degree of certainty in your perception of either babyfacedness or maturity of the face.

*Baby Faced
 As the phrase suggests, "baby faced" indicates there is something that reminds you of a baby or child in the facial features. Both men and women can be baby faced.

*Mature Faced

APPENDIX C: INFORMATION SHEET HANDED TO HUMAN SUBJECTS

Below is a copy of the information sheet that will be given to each subject as they leave the lab after completing the computerized questionnaire (see chapter 8).

Information Sheet

Thank you again for participating in this research project.

We ask that you not discuss this experiment with anyone.

If you have any questions or concerns, please feel free to contact Sheryl Brahnam, the principal investigator.

You may also want to learn about the results of the questionnaire you filled out today and how the results are being used to produce computer generated faces with personality. All research results and papers will eventually be posted on the Internet (see contact information below).


Sheryl Brahnam's contact information:
Baruch Office: Room 333 46 E. 26th Street
Department: Computer Information Systems
Phone Number: (212) 802-6233
Email: brahnam@tellingfaces.com
Web: www.tellingfaces.com


If you have any questions about your rights as a participant in this study, you can contact Hilry Fisher, Sponsored Research, Graduate School/City University of New York, (212) 817-7523, hfisher@gc.cuny.edu

APPENDIX D: STIMULUS FACES

Figure D.1 below shows the 220 stimulus faces randomly selected from the pruned population of 1500 randomly generated faces (see chapter 8 for more details on the generation process). The actual images used in the experiment designed to assess the trait impressions of these faces (see section 8.3) were jpeg files (204 x 173 pixels) and were not numbered.

Figure D.1. The 220 Stimulus Faces.







REFERENCES

- Abdi, H. (1988). A generalized approach for connectionist auto-associative memories: Interpretation, implications and illustration for face processing. In J. Demongeot (Ed.), *Artificial intelligence and cognitive sciences* (pp. 151-164). Manchester, United Kingdom: Manchester University Press.
- Abelson, R. P. (1963). Computer simulation of "hot" cognition. In S. S. Tomkins and S. Messick (Eds.), *Computer simulation of personality: Frontier of psychological theory* (pp. 277-298). New York: Wiley.
- Abelson, R. P., and Carroll, J. D. (1965). Computer simulation of individual belief systems. *American Behavioral Scientist*, 8, 24-30.
- Albright, L., Kenny, D. A., and Malloy, T. E. (1988). Consensus in personality judgements at zero acquaintance. *Journal of Personality and Social Psychology*, 55(3), 387-395.
- Albright, L., Malloy, T. E., Dong, Q., Kenny, D. A., and Fang, X. (1997). Cross-cultural consensus in personality judgments. *Journal of Personality and Social Psychology*, 72(3), 558-569.
- Alcock, D., Solano, J., and Kayson, W. A. (1998). How individuals' responses and attractiveness influence aggression. *Psychological Reports*, 82(June 3/2), 1435-1438.
- Allanson, J. E., and Cole, T. R. (1996). Sotos syndrome: Evolution of facial phenotype subjective and objective assessment. *American Journal of Medical Genetics*, 79(3), 229-230.
- Allanson, J. E., Greenberg, F., and Smith, A. C. (1999). The face of Smith-Magenis syndrome: A subjective and objective study. *Journal of Medical Genetics*, 36(5), 394-397.
- Allanson, J. E., and Hennekam, R. C. (1997). Rubinstein-Taybi syndrome: Objective evaluation of craniofacial structure. *American Journal of Medical Genetics*, 71(4), 414-419.

- Allanson, J. E., Hennekam, R. C., and Ireland, M. (1997). De Lange syndrome: Subjective and objective comparison of the classical and mild phenotypes. *Journal of Medical Genetics*, 34(8), 645-650.
- Allanson, J. E., Ledbetter, D. H., and Dobyns, W. B. (1998). Classical lissencephaly syndromes: Does the face reflect the brain? *Journal of Medical Genetics*, 35(11), 920-923.
- Alley, T. R. (1983). Infantile head shape as a elicitor of adult protection. *Merrill-Palmer Quarterly*, 29, 411-427.
- Alley, T. R. (1988). Social and applied aspects of face perception. In T. R. Alley (Ed.), *Social and applied aspects of perceiving faces* (pp. 1-10). Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers.
- Alley, T. R. (1992). Perceived age, physical attractiveness and sex differences in preferred mates' ages. *Behavioral and Brain Sciences*, 15(1), 92.
- Alley, T. R., and Cunningham, M. R. (1991). Average faces are attractive, but very attractive faces are not average. *Psychological Science*, 2, 123-125.
- Alley, T. R., and Hildebrandt, K. A. (1988). Determinants and consequences of facial aesthetics. In T. R. Alley (Ed.), *Social and applied aspects of perceiving faces* (pp. 101-140). Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers.
- Allport, G. W. (1961). *Pattern and growth in personality*. New York: Holt, Rinehart, and Winston.
- Allport, G. W., and Odbert, H. S. (1938). Trait names: A psycho-lexical study. *Psychological Monographs*, 47(1).
- André, E., Klesen, M., Gebhard, P., Allen, S., and Rist, T. (1999). *Integrating models of personality and emotions into lifelike characters*. Proceedings of the Workshop on Affect in Interactions: Towards a New Generation of Interfaces in Conjunction with the 3rd I3 Annual Conference, Siena, Italy.
- André, E., Muller, J., and Rist, T. (1996). *The PPP persona: A multipurpose animated presentation agent*. Proceedings of Advanced Visual Interfaces.
- André, E., Rist, T., and Müller, J. (1998). *Integrating reactive and scripted behaviors in a life-like presentation agent*. Proceedings of the Second International Conference on Autonomous Agents.

- Antonio, A. D., Aylett, R., and Ballin, D. (Eds.). (2001). *Intelligent virtual agents: Third international workshop (IVA 2001)*. Berlin and New York: Springer.
- Arafa, Y., Charlton, P., and Mamdani, A. (1998). *Modeling personal service assistants with personality: From metaphor to implementation*. 5th International Conference of the Society for Adaptive behavior (SAB '98): Grounding Emotions in Adaptive Systems, University of Zurich, Switzerland.
- Archer, D., Iritani, B., Kimes, D. D., and Barios, M. (1983). Face-ism: Five studies of sex differences in facial prominence. *Journal of Personality and Social Psychology*, 45, 725-735.
- Ark, W., Dryer, C., and Lu, D. J. (1999). *The emotion mouse*. Proceedings of the Eighth International Conference on Human-Computer Interaction. Stuttgart: University of Stuttgart, Germany.
- Ashmore, R. D., and Longo, L. C. (1995). Accuracy of stereotypes: What research on physical attractiveness can teach us. In L. Yueh-Ting, L. J. Jussim and C. R. McCauley (Eds.), *Stereotype accuracy: Toward appreciating group differences* (pp. 63-86). Washington, DC: American Psychological Association.
- Ashmore, R. D., Solomon, M. R., and Longo, L. C. (1996). Thinking about fashion models' looks: A multidimensional approach to the structure of perceived physical attractiveness. *Personality and Social Psychology Bulletin*, 22(11), 1083-1104.
- Badler, N. I. (2001). Virtual beings. *Communications of the ACM*, 44(3), 33-35.
- Badler, N. I., Bindiganavale, R., Bourne, J., Allbeck, J., Shi, J., and Palmer, M. (1999). *Real time virtual humans*, accessed 2001, from <http://www.cis.upenn.edu/~badler/bcs/Paper.htm>
- Badler, N. I., Palmer, M. S., and Bindiganavale, R. (1999). Animation control for real-time virtual humans. *Communications of the ACM*, 42(8), 64-73.
- Badler, N. I., Reich, B. D., and Webber, B. L. (1997). Towards personalities for animated agents with reactive and planning behaviors. In R. Trappl and P. Petta (Eds.), *Creating personalities for synthetic actors: Towards autonomous personality agents* (Vol. 1195). New York: Springer.

- Ball, G., and Breese, J. (1999). *Emotion and personality in a conversational character*. Proceedings of the Workshop on Embodied Conversational Characters, Tahoe City, CA.
- Ball, G., and Breese, J. (2000). Emotion and personality in a conversational agent. In J. Cassell, J. Sullivan, S. Prevost and E. Churchill (Eds.), *Embodied conversational agents* (pp. 189-219). Cambridge, MA: The MIT Press.
- Ball, G., Ling, D., Kurlander, D., Miller, J., Pugh, D., Skelly, T., Stankosky, A., Thiel, D., van Dantzich, M., and Wax, T. (1997). Lifelike computer characters: The persona project at Microsoft. In J. M. Bradshaw (Ed.), *Software agents* (pp. 191-222). Menlo Park, CA: AAAI Press and The MIT Press.
- Bartlett, M. S. (1998). *Face image analysis by unsupervised learning and redundancy reduction*. Unpublished Ph.D. dissertation, University of California, Berkeley.
- Bartlett, M. S., Viola, P. A., Sejnowski, T. J., Golomb, B. A., Larsen, J., Hager, J. C., and Ekman, P. (1996). *Classifying facial action*. Cambridge, MA: The MIT Press.
- Bates, J. (1992a). *The nature of character in interactive worlds and the oz project* (No. Technical Report CMU-CS-92-200). Pittsburgh, PA: School of Computer Science, Carnegie Mellon University.
- Bates, J. (1992b). Virtual reality, art, and entertainment. *Presence: The Journal of Teleoperators and Virtual Environments*, 1(1)(Winter), 133-138.
- Bates, J. (1994). The role of emotion in believable agents. *Communications of the ACM*, 37(7), 122-125.
- Bates, J., Loyall, A. B., and Reilly, W. S. N. (1991). Broad agents. *SIGART Bulletin*, 2(4), 38-40.
- Bates, J., Loyall, A. B., and Reilly, W. S. N. (1992a). *An architecture for action, emotion, and social behavior*. Proceedings of the Fourth European Workshop on Modeling Autonomous Agents in a Multi-agent World, S. Martino al Camino, Italy.
- Bates, J., Loyall, A. B., and Reilly, W. S. N. (1992b). *Integrating reactivity, goals, and emotion in a broad agent*. Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society, Bloomington, IN.

- Belhumeur, P., Hespanha, J., and Kriegman, D. (1997). Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 19(7), 711-720.
- Bem, D. J., and Allen, A. (1974). On predicting some people some of the time. *Psychological Review*, 81, 506-520.
- Benford, S., Bowers, J., Fahlén, L. E., Greenhalgh, C., and Snowdon, D. (1997). Embodiments, avatars, clones and agents for multiuser, multi-sensory virtual worlds. *Multimedia Systems*, 5, 93-104.
- Berry, D. S. (1990). The perceiver as naive scientist or the scientist as naive perceiver?: An ecological view of social knowledge acquisition. *Contemporary Social Psychology*, 14(3), 145-153.
- Berry, D. S. (1991a). Accuracy in social perception: Contributions of facial and vocal information. *Journal of Personality and Social Psychology*, 61(2), 298-307.
- Berry, D. S. (1991b). Attractive faces are not all created equal: Joint effects of facial babyishness and attractiveness on social perception. *Personality and Social Psychology Bulletin*, 17(5), 523-531.
- Berry, D. S., and Brownlow, S. (1989). Were the physiognomists right? *Personality and Social Psychology Bulletin*, 15(2), 266-279.
- Berry, D. S., and Landry, J. C. (1997). Facial maturity and daily social interaction. *Journal of Personality and Social Psychology*, 72(3), 570-580.
- Berry, D. S., and McArthur, L. Z. (1985). Some components and consequences of a babyface. *Journal of Personality and Social Psychology*, 48(2), 312-323.
- Berry, D. S., and McArthur, L. Z. (1986). Perceiving character in faces: The impact of age-related craniofacial changes on social perception. *Psychological Bulletin*, 100(1), 3-18.
- Berry, D. S., and Wero, J. L. F. (1993). Accuracy in face perception: A view from ecological psychology. *Journal of Personality*, 61(4), 497-520.
- Berry, D. S., and Zebrowitz-McArthur, L. A. (1988). What's in a face? Facial maturity and the attribution of legal responsibility. *Personality and Social Psychology Bulletin*, 14(1), 23-33.

- Berscheid, E., and Walster, E. (1974). Physical attractiveness. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 7, pp. 158-216). New York: Academic Press.
- Beskow, J., and McGlashan, S. (1997). *Olga: A conversational agent with gestures*. Proceedings of International Joint Conferences on Artificial Intelligence Workshop on Animated Interface Agents: Making Them Intelligent, Nagoya, Japan.
- Beymer, D., Shashua, A., and Poggio, T. (1993). *Example based image analysis and synthesis* (No. A. I. Memo 1431). Cambridge, MA: MIT.
- Bickmore, T., and Cassell, J. (2001). *Relational agents: A model and implementation of building user trust*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Seattle, WA.
- Binsted, K., and Luke, S. (1998). *Character design for soccer commentary*. Robot Soccer World Cup II: Proceedings of the Second RoboCup Workshop (pp. 23-35), Paris: Springer-Verlag.
- Biocca, F. (1997). The cyborg's dilemma: Progressive embodiment in virtual environments. *Journal of Computer-Mediated Communication*, 3(2).
- Blumberg, B. M. (1996). *Old tricks, new dogs: Ethology and interactive creatures*. Unpublished Ph.D. Dissertation, MIT, Cambridge, MA.
- Bradshaw, J. M. (Ed.). (1997). *Software agents*. Menlo Park, CA.: AAI Press and The MIT Press.
- Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: The MIT Press.
- Brown, E., and Perrett, D. I. (1993). What gives a face its gender? *Perception*, 22(7), 829-840.
- Brownmiller, S. (1984). *Femininity*. New York: Linden Press and Simon & Schuster.
- Bruce, V. (1988). *Recognising faces*. Hove and London: Lawrence Erlbaum Associates, Publishers.
- Brunelli, R., and O. Mich. (1996). SpotIt! An interactive identikit system. *Graphical Models and Image Processing*, 58(5), 399-404.

- Brunelli, R., and Poggio, T. (1993). Face recognition: Features versus templates. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 15(10), 1042-1052.
- Bull, R., and Rumsey, N. (1988). *The social psychology of facial appearance*. New York: Springer-Verlag.
- Burton, A. M., Bruce, V., and Dench, N. (1993). What's the difference between men and women? Evidence from facial measurement. *Perception*, 22(2), 153-176.
- Buss, D. M., and Barnes, M. (1986). Preferences in human mate selection. *Journal of Personality and Social Psychology*, 50, 559-570.
- Carbonell, J. G. (1980). Towards a process model of human personality traits. *Artificial Intelligence*, 15(1-2), 49-74.
- Carbonell, J. G. (1981). *Subjective understanding: Computer models of belief systems*. Ann Arbor, MI: UMI Research Press.
- Carello, C., Groszofsky, A., Shaw, R. E., Pittenger, J. B., and Mark, L. S. (1989). Attractiveness of facial profiles is a function of distance from archetype. *Ecological Psychology*, 1(3), 227-251.
- Cash, T. F., and Trimer, C. A. (1984). Sexism and beautyism in women's evaluations of peer performance. *Sex Roles*, 10, 87-98.
- Cassell, J., Pelachaud, C., Badler, N. I., 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*. SIGGRAPH '94, Orlando, FL.
- Cassell, J., Steedman, M., Badler, N. I., Pelachaud, C., Stone, M., Douville, B., Prevost, S., and Achorn, B. (1994). *Modeling the interaction between speech and gesture*. Proceedings of the 16th Annual Conference of the Cognitive Science Society, Georgia Institute of Technology, Atlanta, GA.
- Castelfranchi, C., de Rosis, F., and Falcone, R. (1997). *Social attitudes and personality in agents*. AAAI Fall Symposium on Socially Intelligent Agents.
- Cattell, R. B. (1965). *The scientific analysis of personality*. Harmondsworth, United Kingdom: Penguin.

- Cattell, R. B., Eber, H., and Tatsuoka, M. M. (1970). *Handbook for the sixteen personality factor questionnaire (16PF)*. Champaign, IL.: Institute for Personality and Ability Testing.
- Chin, J. P. (1996). *Personality trait attributions to voice mail user interfaces*. CHI '96. Proceedings of the CHI '96 Conference Companion on Human Factors in Computing Systems: Common Ground.
- Churchill, E. F., Cook, L., Hodgson, P., Prevost, S., and Sullivan, J. W. (2000). "May i help you?": Designing embodied conversational agent allies. In J. Cassell, J. Sullivan, S. Prevost and E. Churchill (Eds.), *Embodied conversational agents* (pp. 65-94). Cambridge, MA.: The MIT Press.
- Cichocki, A., and Unbehauen, R. (1993). *Neural networks for optimization and signal processing*. New York: John Wiley & Sons.
- Collins, M. A., and Zebrowitz, L. A. (1995). The contributions of appearance to occupational outcomes in civilian and military settings. *Journal of Applied Social Psychology*, 25(1), 129-163.
- Cooley, C. H. (1964). *Human nature and the social other*. New York: Schocken Books.
- Cottrell, G. W., and Fleming, M. K. (1990). *Face recognition using unsupervised feature extraction*. Proceeding of the International Conference on Neural Networks (pp. 322-325), Paris.
- Cottrell, G. W., and Metcalfe, J. (1991). EMPATH: Face, emotion, and gender recognition using holons. In D. Touretzky (Ed.), *Advances in neural information processing systems* (pp. 564-571). San Mateo, CA: Morgan & Kaufman.
- Cunningham, M. R. (1986). Measuring the physical in physical attractiveness: Quasi-experiments on the sociobiology of female facial beauty. *Journal of Personality and Social Psychology*, 50(5), 925-935.
- Cunningham, M. R., Barbee, A. P., and Pike, C. L. (1990). What do women want?: Facialmetric assessment of multiple motives in the perception of male facial physical attractiveness. *Journal of Personality and Social Psychology*, 59(1), 61-72.
- Cunningham, M. R., Roberts, A. R., Barbee, A. P., and Druen, P. B. (1995). "Their ideas of beauty are, on the whole, the same as ours": Consistency and variability in the

cross-cultural perception of female physical attractiveness. *Journal of Personality and Social Psychology*, 68(2), 261-279.

Darwin, C. A. (1959). *Autobiography*. New York: Harcourt.

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*, 1(7), 573-617.

Deaux, K., and Major, B. (1987). Putting gender into context: An interactive model of gender-related behavior. *Psychological Review*, 94, 369-389.

De Carolis, B. D., de Rosis, F., and Pizzutilo, S. (2000, May). *Automated generation of agent behavior from formal models of interaction*. Proceedings of Advanced Visual Interfaces 2000, Palermo, Italy.

de Rosis, F., and Castelfranchi, C. (1999). *A contribution to cognitive modeling of affective phenomena*, accessed 2002, from <http://www.ai.mit.edu/people/jvelas/ebaa99/derosis-ebaa99.pdf>

de Rosis, F., Covino, E., Falcone, R., and Castelfranchi, C. (2001). Bayesian cognitive diagnosis in believable multiagent systems. In M. A. Williams and H. Rott (Eds.), *Frontiers of belief revision*. Dordrecht, The Netherlands: Kluwer Academic Publishers.

Dehn, D. M., and van Mulken, S. (2000). The impact of animated interface agents: A review of empirical research. *International Journal of Human Computer Studies*, 52(1), 1-22.

Dermer, M., and Thiel, D. L. (1975). When beauty may fail. *Journal of Personality and Social Psychology*, 31(6), 1168-1176.

Diamantaras, K. I., and Kung, S. Y. (1996). *Principal component neural networks: Theory and applications*. New York: John Wiley & Sons, Inc.

Dion, K. K. (1972). Physical attractiveness and evaluation of childrens' transgressions. *Journal of Personality and Social Psychology*, 24, 207-213.

Dion, K. K., and Berscheid, E. (1972). What is beautiful is good. *Journal of Personality and Social Psychology*, 24(3), 285-290.

- Donath, J. (2001). *Mediated faces*. Cognitive Technology: Instruments of Mind: Proceedings of the 4th International Conference, Warwick, United Kingdom.
- Doyle, P., and Hayes-Roth, B. (1998). *Agents in annotated worlds*. Proceedings of the Second International Conference on Autonomous Agents.
- Duda, R. O., and Hart, P. E. (1973). *Pattern classification and scene analysis*. London: Academic Press.
- Dunkle, J. H., and Francis, P. L. (1996). "Physical attractiveness stereotype" and the attribution of homosexuality revisited. *Journal of Homosexuality*, 30(3), 13-29.
- Eagly, A. H., Ashmore, R. D., Makhijan, M. G., and Longo, L. C. (1991). What is beautiful is good, but . . . : A meta-analytic review of research on the physical attractiveness stereotype. *Psychological Bulletin*, 110(1), 109-128.
- Effran, M. G. (1974). The effect of physical appearance on the judgement of guilt, interpersonal attraction, and severity of recommended punishment in a simulated jury task. *Journal of Research in Personality*, 8(1), 45-54.
- Ekman, P., and Friesen, W. V. (1978). *Facial action coding system*. Palo Alto, CA: Consulting Psychologists Press Inc.
- Ekman, P., and Rosenberg, E. L. (Eds.). (1997). *What the face reveals: Basic and applied studies of spontaneous expression using the facial action coding system (FACS)*. New York: Oxford University Press.
- Elliott, C., and Melchoir, E. (1995). *Getting to the point: Emotion as a necessary and sufficient element of story construction*. AAAI Technical Report for the Spring Symposium on Interactive Story Systems, Stanford University.
- Elliott, C. D. (1992). *The affective reasoner: A process model of emotion in a multi-agent system*. Unpublished Ph.D. dissertation, Northwestern University, Evanston, IL.
- Elliott, C. D., and Brzezinski, J. (1999). *DRAFT: Research paradigms for agents as synthetic characters*, accessed 1999, from <http://www.depaul.edu/~elliott/katia/web/>
- Elliott, C. D., Brzezinski, J., Sheth, S., and Salvatoriello, R. (1998). *Story-morphing in the affective reasoning paradigm: Generating stories semi-automatically for use with emotionally "intelligent" multimedia agents*. AGENTS '98. Proceedings of the Second International Conference on Autonomous Agents, Minneapolis, MN.

- Enlow, D. H., and Hans, M. G. (1996). *Essentials of facial growth*. Philadelphia: W. B. Saunders Company.
- Essa, I. A., and Pentland, A. P. (1997). Coding, analysis, interpretation, and recognition of facial expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7), 757-763.
- Eysenck, H. J. (1991). Dimensions of personality: 16, 5, or 3?—Criteria for a taxonomic paradigm. *Journal of Personality & Social Psychology*, 51, 1058-1068.
- Feingold, A. (1992). Good-looking people are not what we think. *Psychological Bulletin*, 111(2), 304-341.
- Fidler, D. J., and Hodapp, R. M. (1999). Craniofacial maturity and perceived personality in children with Down syndrome. *American Association on Mental Retardation*, 104(5), 410-421.
- Folstad, I., and Karter, A. J. (1992). Parasites, bright males, and the immunocompetence handicap. *American Naturalist*, 139(3), 603-622.
- Foner, L. (1993). *What's an agent, anyway? A sociological case study*, accessed 2001, from <http://foner.www.media.mit.edu/people/foner/Julia/Julia-cite.html>
- Ford, M. E. (1992). *Motivating humans: Goals, emotions, and personal agency beliefs*. Newbury Park, CA: Sage.
- Franklin, S., and Graesser, A. (Eds.). (1996). *Is it an agent, or just a program?: A taxonomy for autonomous agents* (Vol. 1193). New York: Springer.
- Freierman, S. (2000, 17 February). Constructing a real-life Mr. Potato Head. Faces: The ultimate composite picture. *The New York Times*, p. 6.
- Frey, S. (1993). Lavater, Lichtenberg, and the suggestive power of the human face. In E. Shookman (Ed.), *The faces of physiognomy: Interdisciplinary approaches to johann caspar lavater* (pp. 64-103). Columbia, SC: Camden House.
- Fridlund, A. J. (1994). *Human facial expression: An evolutionary view*. San Diego: Academic Press.
- Friedman, H., and Zebrowitz, L. A. (1992). The contribution of facial maturity to sex-role stereotypes. *Personality and Social Psychology Bulletin*, 18(4), 430-438.

- Frijda, N. (1986). *The emotions*. Cambridge: Cambridge University Press.
- Frijda, N. H., and Moffat, D. (1994). Modeling emotion. *Cognitive Studies*, 1(2), 5–15.
- Frodi, A. M., Lamb, M. E., Leavitt, L. A., Donovan, W. L., Neff, C., and Sherry, D. (1978). Fathers' and mothers' responses to the faces and cries of normal and premature infants. *Developmental Psychology*, 14(5), 490-498.
- Funder, D. C. (1987). Errors and mistakes: Evaluating the accuracy of social judgement. *Psychological Bulletin*, 101(1), 75-90.
- Galton, F. (1878). Composite portraits. *Nature*, 18, 97-100.
- Gass, W. H. (1979). *Fiction and the figures of life*. Boston: Nonpareil Books.
- Gmytrasiewicz, P. J., and Lisetti, C. L. (2001). *Emotions and personality in agent design and modeling*. Second Workshop on Attitude, Personality and Emotions in User-Adapted Interactions, Sonthofen, Germany.
- Goffman, E. (1967). *Interaction ritual*. New York: Pantheon Books.
- Goffman, E. (1969). *Strategic Interaction*. Philadelphia: University of Pennsylvania Press.
- Goldberg, A. (1997). IMPROV: A system for real-time animation of behavior-based interactive synthetic actors. In R. Trappl and P. Petta (Eds.), *Creating personalities for synthetic actors: Towards autonomous personality agents* (Vol. 1195). New York: Springer.
- Gove, P. B. (Ed.). (1986). *Webster's new international dictionary of the english language unabridged*. Springfield, MA: Merriam-Webster Inc., Publishers.
- Graham, D. B., and Allison, N. M. (1998). *Automatic face representation and classification*. 9th British Machine Vision Conference, University of Southampton, United Kingdom.
- Grammer, K., and Thornhill, R. (1994). Human (homo sapiens) facial attractiveness and sexual selection: The role of symmetry and averageness. *Journal of Comparative Psychology*, 108(3), 233-242.
- Gray, H. (1985). *Anatomy of the human body* (30th ed.). Philadelphia: Lea & Febiger.

- Guthrie, R. D. (1970). Evolution of human threat display organs. In T. Dobzhansky, M. K. Hecht and W. C. Steere (Eds.), *Evolutionary biology* (Vol. 4, pp. 257-302). New York: Appleton-Century Crofts.
- Hampson, S. E. (1988a). *The construction of personality* (Second ed.). London: Routledge.
- Hampson, S. E. (1988b). The emergence of personality: A broader context for biological perspectives. In A. Gale and M. W. Eysenck (Eds.), *Handbook of individual differences: Biological perspectives*. Chichester, United Kingdom: Wiley.
- Hampson, S. E. (1995). The construction of personality. In S. E. Hampson and A. M. Colman (Eds.), *Individual differences and personality*. London and New York: Longman.
- Hancock, P. (2000). Evolving faces from principal components. *Behavior Research Methods, Instruments and Computers*, 32(2), 327-333.
- Hancock, P. J. B., Burton, A. M., and Bruce, V. (1996). Face processing: Human perception and principal components analysis. *Memory and Cognition*, 24(1), 26-40.
- Hatfield, E., and Sprecher, S. (1986). *Mirror, mirror: The importance of looks in everyday life*. Albany: State University of New York Press.
- Hayem, R., Fourmaintraux, T., Petit, K., Rauber, N., and Kisseleva, O. (1998). *Avatars: New fields of implication*. Virtual Worlds: First International Conference, VW '98, Paris.
- Hayes-Roth, B. (1995). An architecture for adaptive intelligent systems. *Artificial Intelligence*, 72(1-2), 329-365.
- Hayes-Roth, B., Johnson, V., van Gent, R., and Wescourt, K. (1999). Staffing the web with interactive characters. *Communications of the ACM*, 42(3), 103-105.
- Hayes-Roth, B., Sincoff, E., Brownston, L., Huard, R. D., and Lent, B. (1995). *Directed improvisation with animated puppets*. CHI '95. Conference Companion on Human Factors in Computing Systems, Denver, CO.
- Hayes-Roth, B., and van Gent, R. (1996). *Improvisational puppets, actors, and avatars*. Proceedings of the Computer Game Developers' Conference, Santa Clara, CA.

- Hayes-Roth, B., and van Gent, R. (1997). *Story-making with improvisational puppets*. AGENTS '97. Proceedings of the First International Conference on Autonomous Agents, Marina Del Rey, CA.
- Hayes-Roth, B., van Gent, R., and Huber, D. (1997). Acting in character. In R. Trapp and P. Petta (Eds.), *Creating personalities for synthetic actors: Towards autonomous personality agents*. Berlin and New York: Springer.
- Hepburn, K. (1999). *Quotes_hepburn_katharine.html*, accessed 2001, from http://www.cyber-nation.com/victory/quotations/authors/quotes_hepburn_katharine.html
- Hietala, P., and Niemirepo, T. (1998). The competence of learning companion agents. *International Journal of Artificial Intelligence in Education*, 9, 178-192.
- Hirsh, H., Hayes-Roth, B., Stern, A., and Murray, J. H. (1998). Trends and controversies: Interactive fiction. *IEEE Intelligent Systems*, 13(6), 12-21.
- Hochberg, J., and Galper, R. E. (1974). Attribution of intention as a function of physiognomy. *Memory and Cognition*, 2(1A), 39-42.
- Hoffner, C., and Cantor, J. (1991). Perceiving and responding to mass media characters. In J. Bryant and D. Zillman (Eds.), *Responding to the screen: Reception and reaction processes* (pp. 63-101). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Houle, M., and Simon. (1997, July). *Social and ethical education in computing using virtual environments*. ACSE '97. Proceedings of the Second Australasian Conference on Computer Science Education.
- Houston, V., and Bull, R. (1994). Do people avoid sitting next to someone who is facially disfigured? *European Journal of Social Psychology*, 24(2), 279-284.
- Huard, R. D., and Hayes-Roth, B. (1996). *Children's play with improvisational puppets* (No. KSL-96-27). Santa Clara, CA: Knowledge Systems Laboratory Department of Computer Science, Stanford University.
- Isbister, K. (1994). *Perceived intelligence and the design of computer character*, accessed 2002, from <http://www.cyborganic.com/People/kath/lccpaper.html>
- Isbister, K., and Hayes-Roth, B. (1998). *Social implications of using synthetic characters: An examination of a role-specific intelligent agent* (No. KSL-98-01).

Santa Clara, CA: Knowledge Systems Laboratory, Department of Computer Science, Stanford University.

- Jackson, L. (1992). *Physical appearance and gender: Sociobiological and sociocultural perspectives*. Albany: State University of New York Press.
- James, W. (1918). *The principles of psychology* (Vol. 1). New York: Holt.
- Johnson, M. P., Wilson, A., Kline, C., Blumberg, B. M., and Bobick, A. F. (1999). *Sympathetic interfaces: Using a plush toy to direct synthetic characters*. Proceedings CHI '99, Pittsburgh, PA.
- Johnson, W. L., Rickel, J. W., and Lester, J. C. (2000). Animated pedagogical agents: Face-to-face interaction in interactive learning environments. *International Journal of Artificial Intelligence in Education*, 11, 47-78.
- Johnstone, K. (1987). *IMPROV: Improvisation in the theater*. New York: Penquin Books.
- Jones, D. (1995). Sexual selection, physical attractiveness, and facial neoteny. *Current Anthropology*, 36(5), 723-748.
- Katz, D. (1952). Le portrait composite et la typologie. *Ikon (Revue Internationale de Filmology)*, 2, 207-214.
- Keating, C. F. (1985). Gender and the physiognomy of dominance and attractiveness. *Psychological Quarterly*, 48, 61-70.
- Keating, C. F., Mazur, A., and Segall, M. H. (1981a). A cross-cultural exploration of physiognomic traits of dominance and happiness. *Ethology and Sociobiology*, 2, 41-48.
- Keating, C. F., Mazur, A., and Segall, M. H. (1981b). Culture and the perception of social dominance. *Journal of Personality and Social Psychology*, 40(4), 615-626.
- Kelly, G. A. (1955). *The psychology of personal constructs*. New York: Norton.
- Khan, A. F. U. (1995). The ethics of autonomous learning systems. In K. M. Ford, C. Glymour and P. J. Hayes (Eds.), *Android epistemology* (pp. 253-266). Cambridge, MA: AAAI Press and The MIT Press.

- Koda, T., and Maes, P. (1996). *Agents with faces: The effects of personification of agents*. Proceedings of Human-Computer Interaction HCI '96 (pp. 98-103), London.
- Kohn, L. (1996). The looks of Laozi. *Asian Folklore Studies*, 55(2).
- Kohonen, T. (1977). *Associative memory: A system theoretic approach*. Berlin: Springer.
- Kosugi, M. (1995). Human-face search and location in a scene by multi-pyramid architecture for personal identification. *Systems and Computers in Japan*, 26(6), 27-38.
- Kowner, R. (1996). Facial asymmetry and attractiveness judgment in developmental perspective. *Journal of Experimental Psychology: Human Perception and Performance*, 22(3), 662-675.
- Kowner, R. (1997). The perception and attribution of facial asymmetry in normal adults. *The Psychological Record*, 47(3), 371-385.
- LaFrance, M., and Hecht, M. A. (1995). Why smiles generate leniency. *Personality and Social Psychology Bulletin*, 21(3), 207-214.
- Langlois, J. H., Kalakanis, L., Rubenstein, A. J., Larson, A., Hallam, M., and Smoot, M. (2000). Maxims or myths of beauty? A meta-analytic and theoretical review. *Psychological Bulletin*, 126(3), 390-423.
- Langlois, J. H., and Roggman, L. A. (1990). Attractive faces are only average. *Psychological Science*, 1, 115-121.
- Langlois, J. H., Roggman, L. A., and Musselman, L. (1994). What is average and what is not average about attractive faces? *Psychological Science*, 5(4), 214-220.
- Langlois, J. H., Roggman, L. A., Musselman, L., and Acton, S. (1991a). A picture is worth a thousand words: Reply to 'on the difficulty of averaging faces'. *Psychological Science*, 2(5), 354-357.
- Langlois, J. H., Roggman, L. A., and Rieser-Danner, L. A. (1990). Infants' differential social responses to attractive and unattractive faces. *Developmental Psychology*, 26(1), 153-159.

- Langlois, J. H., Roggman, L. A., and Vaughn, L. S. (1991b). Facial diversity and infant preferences for attractive faces. *Developmental Psychology*, 27(1), 79-84.
- Lanier, J. (1995). Agents of alienation. *Interactions*, 2(3), 66-72.
- Lanitis, A., Taylor, C. J., and Cootes, T. F. (1997). Automatic interpretation and coding of face images using flexible models. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 19(7), 743-756.
- Laurel, B. (1990a). Interface agents: Metaphors with character. In B. Laurel (Ed.), *The art of human-computer interface design* (pp. 355-366). Reading, MA: Addison-Wesley Publishing Company.
- Laurel, B. (Ed.). (1990b). *The art of human-computer interface design*. Reading, MA: Addison-Wesley Publishing Company.
- Laurel, B. (1993). *Computers as theatre*. Reading, MA: Addison-Wesley Publishing Company.
- Lee, K. (1999). *Integration of various emotion eliciting factors for life-like agents*. Proceedings of the Seventh ACM International Conference on Multimedia Systems, Orlando, FL.
- Leonard, A. (1997). *Bots: The origin of a new species*. New York: Penguin.
- Lester, J. C., Converse, S. A., Kahler, S. E., Barlow, S. T., Stone, B. A., and Bhogal, R. S. (1997). *The persona effect: Affective impact of animated pedagogical agents*. Conference Proceedings on Human Factors in Computing Systems.
- Lester, J. C., and Stone, B. A. (1997). *Increasing believability in animated pedagogical agents*. Proceedings of the First International Conference on Autonomous Agents (pp. 16-21), Marina del Rey, CA.
- Lester, J. C., Towns, S. G., Callaway, C. B., Voerman, J. L., and FitzGerald, P. J. (2000). Deictic and emotive communication in animated pedagogical agents. In J. Cassell, J. Sullivan, S. Prevost and E. Churchill (Eds.), *Embodied conversational agents* (pp. 123-154). Cambridge, MA: The MIT Press.
- Lester, J. C., Voerman, J. L., Towns, S. G., and Callaway, C. B. (1997). *Cosmo: A life-like animated pedagogical agent with deictic believability*, accessed 2001, from dfki.unisb.de/pub/anina/lester.ps

- Lewis, J. P., and Parke, F. I. (1987). *Automated lip-synch and speech synthesis for character animation*. CHI/GI 1987 Conference Proceedings on Human Factors in Computing Systems and Graphics Interface (pp. 143-147), New York: ACM Press.
- Liggett, J. (1974). *The human face*. New York: Stein and Day.
- Liu, C. (1999). *Statistical and evolutionary approaches for face recognition*. Unpublished Ph.D. dissertation, George Mason University, Fairfax, VA.
- Loehlin, J. C. (1962). The personality of Aldous. *Discovery*, 23, 23-26.
- Loehlin, J. C. (1968). *Computer models of personality*. New York: Random House.
- Lorenz, K. (1970-71). *Studies in animal and human behaviour*. Cambridge, MA: Harvard University Press.
- Loyall, A. B. (1997a). *Believable agents: Building interactive personalities*. Unpublished Dissertation, Carnegie Mellon University, Pittsburgh, PA.
- Loyall, A. B. (1997b). Some requirements and approaches for natural language in a believable agent. In R. Trappl and P. Petta (Eds.), *Creating personalities for synthetic actors: Towards autonomous personality agents* (Vol. 1195). New York: Springer.
- Loyall, A. B., and Bates, J. (1991). *Hap: A reactive, adaptive architecture for agents* (No. Technical Report CMU-CS-91-147). Pittsburgh, PA: School of Computer Science, Carnegie Mellon University.
- Loyall, A. B., and Bates, J. (1993). *Real-time control of animated broad agents*. Proceedings of the Fifteenth Annual Conference of the Cognitive Science Society, Boulder, CO.
- Loyall, A. B., and Bates, J. (1997). *Personality-rich believable agents that use language*. Proceedings of the First International Conference on Autonomous Agents, Marina del Rey, CA.
- Lucker, G. W., and Graber, L. W. (1980). Physiognomic features and facial appearance judgements in children. *The Journal of Psychology*, 104(Second Half), 261-268.

- Maes, P., Darrell, T., Blumberg, B. M., and Pentland, A. (1997). The ALIVE system: Wireless, full-body interaction with autonomous agents. *Multimedia Systems*, 2(2), 105-112.
- Maes, P., and Wexelblat, A. (1996, April). *Interface agents*. Proceedings of the CHI '96 Conference Companion on Human Factors in Computing Systems: Common Ground, Vancouver, Canada.
- Magro, A. M. (1997). Why barbie is perceived as beautiful. *Perceptual and Motor Skills*, 85(1), 363-374.
- Mark, L. S., Shaw, R. E., and Pittenger, J. B. (1988). Natural constraints, scales of analysis, and information for the perception of growing faces. In T. R. Alley (Ed.), *Social and applied aspects of perceiving faces* (pp. 11-50). Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers.
- Markow, T. A., and Wandler, K. (1986). Fluctuating dermatoglyphic asymmetry and the genetics of liability to schizophrenia. *Psychiatry Research*, 19, 323-328.
- Martinho, C., and Paiva, A. (1999). *Pathematic agents: Rapid development of believable emotional agents in intelligent virtual environments*. AGENTS '99. Proceedings of the Third Annual Conference on Autonomous Agents, Seattle, WA.
- Mataric, M. J. (1992). *Designing emergent behaviors: From local interactions to collective intelligence*. Second International Conference on Simulation of Adaptive Behavior, Cambridge, MA.
- Mathes, E. (1975). The effects of physical attractiveness and anxiety on heterosexual attraction. *Journal of Marriage and the Family*, 37, 769-773.
- Matthews, G., and Deary, I. J. (1998). *Personality traits* (First ed.): Cambridge: Cambridge University Press.
- Mayer, F. S., and Sutton, K. (1996). *Personality: An integrative approach*. Upper Saddle River, NJ: Prentice-Hall, Inc.
- McArthur, L. Z. (1982). Judging a book by its cover: A cognitive analysis of the relationship between physical appearance and stereotyping. In A. H. Hastorf and A. M. Isen (Eds.), *Cognitive social psychology* (pp. 149-211). New York: Elsevier North-Holland.

- McArthur, L. Z., and Apatow, K. (1984). Impressions of baby-faced adults. *Social Cognition*, 2(4), 315-342.
- McArthur, L. Z., and Baron, R. M. (1983). Toward an ecological theory of social perception. *Psychology Review*, 90(3), 215-238.
- McArthur, L. Z., and Berry, D. S. (1987). Cross-cultural agreement in perceptions of babyfaced adults. *Journal of Cross-cultural Psychology*, 18(2), 165-192.
- McCall, R. B., and Kennedy, C. B. (1980). Attention of 4-month infants to discrepancy and babyishness. *Journal of Experimental Child Psychology*, 29(2), 189-201.
- McCauley, C. R., Jussim, L. J., and Lee, Y.-T. (1995). Stereotype accuracy: Toward appreciating group differences. In Y.-T. Lee, L. J. Jussim and C. R. McCauley (Eds.), *Stereotype accuracy: Toward appreciating group differences* (pp. 293-312). Washington, DC: American Psychological Association.
- McClelland, J. L., and Rumelhart, D. E. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: The MIT Press and Bradford Books.
- Mealey, L., Bridgstock, R., and Townsend, G. C. (1999). Symmetry and perceived facial attractiveness: A monozygotic co-twin comparison. *Journal of Personality and Social Psychology*, 76(1), 151-158.
- Meuleman, P., Heister, A., Kohar, H., and Tedd, D. (1998). *Double agent: Presentation and filtering agents for a digital television recording system*. CHI '98. Proceedings of the CHI 98 Summary Conference on CHI 98 Summary: Human Factors in Computing Systems, Los Angeles, CA.
- Miksch, S., Cheng, K., and Hayes-Roth, B. (1997). *An intelligent assistant for patient health care*. Proceedings of the First International Conference on Autonomous Agents.
- Miller, A., Gillen, B., Schenker, C., and Radlove, S. (1974). The prediction and perception of obedience to authority. *Journal of Personality*, 42(1), 23-42.
- Millward, R., and O'Toole, A. J. (1986). Recognition memory transfer between spatial-frequency analyzed faces. In H. D. Ellis, M. A. Jeeves, F. Newcombe and A. Young (Eds.), *Aspects of face processing* (pp. 34-44). Dordrecht, The Netherlands: Martinus Nijhoff.

- Moffat, D. (1997). Personality parameters and programs. In R. Trappl and P. Petta (Eds.), *Creating personalities for synthetic actors: Towards autonomous personality agents* (Vol. 1195). New York: Springer.
- Montepare, J. M., and Zebrowitz-McArthur, L. A. (1989). Children's perceptions of babyfaced adults. *Perceptual and Motor Skills*, 69(2), 467-472.
- Morkes, J., Kernal, H., and Nass, C. (2000). Effects of humor in task-oriented human-computer interaction and computer-mediated communication: A direct test of SRCT theory. *Human-Computer Interaction*, 14(4), 395-435.
- Mueller, U., and Mazur, A. (1996). Facial dominance of west point cadets as a predictor of later military rank. *Social Forces*, 74(3), 823-828.
- Mulford, M., Orbell, J., Shatto, C., and Stockard, J. (1998). Physical attractiveness, opportunity, and success in everyday exchange. *American Journal of Sociology*, 103(6), 1565-1592.
- Müller, J. P., Wooldridge, M. J., and Jennings, N. R. (Eds.). (1996). *Intelligent agents III: Agent theories, architectures, and languages*. Berlin: Springer.
- Murray, J. H. (1997). *Hamlet on the holodeck*. Cambridge, MA: The MIT Press.
- Nakatsu, R., Tosa, N., and Ochi, T. (1998, September). *Interactive movie system with multi-person participation and anytime interaction capabilities*. Proceedings of the Sixth ACM International Multimedia Conference on Technologies for Interactive Movies, Bristol, United Kingdom.
- Nass, C., Isbister, K., and Lee, E.-J. (2000). Truth is beauty: Researching embodied conversational agents. In J. Cassell, J. Sullivan, S. Prevost and E. Churchill (Eds.), *Embodied conversational agents* (pp. 374-402). Cambridge, MA: The MIT Press.
- Nass, C., Steuer, J., Henriksen, L., and Reeder, H. (1993). Anthropomorphism, agency and ethopoeia: Computers as social actors. *Human Communication Research*, 19(4), 504-527.
- Negroponte, N. (1995). *Being digital*. New York: Vintage Books.
- Neilson, W. A. (Ed.). (1955). *Webster's new international dictionary of the english language*. Springfield, MA: G. & C. Merriam Company, Publishers.

- Oja, E. (1992). Principal components, minor components and linear neural networks. *Neural Networks*, 5, 927-935.
- Olney, A. H., Schaefer, G. B., and Kolodziej, P. (1998). Kabuki syndrome. *Ear Nose Throat Journal*, 77(9), 734.
- O'Neill, J. (1985). *Five bodies: The human shape of modern society*. Ithaca, NY: Cornell University Press.
- Oren, T., Salomon, G., Kreitman, K., and Don, A. (1990). Guides: Characterizing the interface. In B. Laurel (Ed.), *The art of human-computer interface design* (pp. 367-382). Reading, MA: Addison-Wesley Publishing Company.
- Ortony, A., Clore, G., and Collins, A. (1988). *The cognitive structure of emotions*. New York: Cambridge University Press.
- O'Toole, A. J., Abdi, H., Deffenbacher, K. A., and Bartlett, J. C. (1991). *Classifying faces by race and sex using an autoassociative memory trained for recognition*. Proceedings of the 13th Annual Conference on Cognitive Science, Hillsdale, NJ.
- O'Toole, A. J., and Deffenbacher, K. A. (1997). The perception of face gender: The role of stimulus structure in recognition and classification. *Memory and Cognition*, 26, 146-160.
- Padgett, C., and Cottrell, G. W. (1998). *A simple neural network models categorical perception of facial expressions*. Proceedings of the Twentieth Annual Cognitive Science Conference, Madison, WI.
- Pantic, M., and Rothkrantz, L. J. M. (1997). *Automatic recognition of facial expressions and human emotions*. Proceedings of ASCI 97, Heijen, The Netherlands.
- Patzer, G. (1985). *The physical attractiveness phenomena*. New York: Plenum Press.
- Paunonen, S. V., Ewan, K., Earthy, J., Lefave, S., and Goldberg, H. (1999). Facial features as personality cues. *Journal of Personality*, 67(3), 555-556.
- Penton-Voak, I. S., Perrett, D. I., and Peirce, J. W. (1999). Computer graphic studies of the role of facial similarity in judgements of attractiveness. *Current Psychology: Developmental, Learning, Personality, Social*, 8(1), 104-117.

- Peretz, B., Ever-Hadani, P., Casamassimo, P., Eidelman, E., Shellhart, C., and Hagerman, R. (1988). Crown size asymmetry in males with FRA (X) or Martin-Bell syndrome. *American Journal of Medical Genetics*, 30(1/2), 185-190.
- Perlin, K. (1995). Real time responsive animation with personality. *IEEE Transactions on Visualization and Computer Graphics*, 1(1), 5-15.
- Perlin, K. (1997). *Layered compositing of facial expression*. ACM SIGGRAPH 97 Technical Sketch.
- Perrett, D. I., Burt, D. M., Penton-Voak, I. S., Lee, K. J., Rowland, D. A., and Edwards, R. (1999). Symmetry and human facial attractiveness. *Evolution and Human Behavior*, 20(5), 295-307.
- Perrett, D. I., May, K. A., and Yoshikawa, S. (1994). Facial shape and judgements of female attractiveness. *Nature*, 368(6468), 239-242.
- Pervin, L. A., and John, O. P. (1997). *Personality: Theory and research* (7th ed.). New York: John Wiley & Sons, Inc.
- Philips, P., Wechsler, H., Juang, J., and Rauss, P. (1998). The FERET database and evaluation procedure for face-recognition algorithms. *Image and Vision Computing Journal*, 16(5), 295-306.
- Picard, R. W. (1997). *Affective computing* (1st ed.). Cambridge, MA: The MIT Press.
- Pinhanez, C. S. (1996). *Computer theater*, accessed 2000, from <http://pinhanez.www.media.mit.edu/people/pinhanez/tr378.html>
- Pittenger, J. B. (1991). On the difficulty of averaging faces: Comments on Langois and Roggman. *Psychological Science*, 2(5), 351-353.
- Pittenger, J. B., and Shaw, R. G. (1975). Aging faces as viscal-elastic events: Implications for a theory of nonrigid shape perceptions. *Journal of Experimental Psychology: Human Perception and Performance*, 1, 376.
- Poggi, I., and Pelachaud, C. (2000). Performative facial expressions in animated faces. In J. Cassell, J. Sullivan, S. Prevost and E. Churchill (Eds.), *Embodied conversational agents* (pp. 155-188). Cambridge, MA: The MIT Press.
- Polak, M., and Trivers, R. (1994). The science of symmetry in biology. *Trends in Ecology and Evolution*, 9, 122-124.

- Pollard, J., Shepherd, J., and Shepherd, J. (1999). Average faces are average faces. *Current Psychology*, 18(1).
- Porter, T., and Susman, G. (2000). On site: Creating lifelike characters in pixar movies. *Communications of the ACM*, 43(1), 25-29.
- Prendinger, H., and Ishizuka, M. (2001). *Agents that talk back (sometimes): Filter programs for affective communication*. Second Workshop on Attitude, Personality and Emotions in User-Adapted Interactions, Sonthofen, Germany.
- Rabin, A. I. (1951). The Szondi test. In H. Anderson and G. Anderson (Eds.), *An introduction to projective techniques*. New York: Prentice Hall, Inc.
- Reeves, B., and Nass, C. I. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. Stanford, CA: CSLI Publications and Cambridge University Press.
- Reich, B. D. (1997). *An architecture for behavioral locomotion*. Unpublished Ph.D. dissertation, University of Pennsylvania, Philadelphia.
- Reilly, W. S. N. (1996). *Believable social and emotional agents*. Unpublished Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA.
- Reilly, W. S. N. (1997, February). *A methodology for building believable social agents*. Proceedings of the First International Conference on Autonomous Agents, Marina Del Rey, CA.
- Resnick, P. V., and Lammers, H. B. (1985). The influence of self-esteem on cognitive responses to machine-like versus human-like computer feedback. *The Journal of Social Psychology*, 125(6), 761-769.
- Rhodes, G., Hickford, C., and Jeffery, L. (2000). Sex-typicality and attractiveness: Are supermale and superfemale faces super-attractive. *British Journal of Psychology*, 91(1), 125-140.
- Rhodes, G., and Tremewan, T. (1996). Averageness exaggeration and facial attractiveness. *Psychological Science*, 7(2), 105-110.
- Rickel, J., and Johnson, W. L. (1998). *STEVE: A pedagogical agent for virtual reality*. Proceedings of the Second International Conference on Autonomous Agents, Minneapolis/St. Paul, MN.

- Rickenberg, R., and Reeves, B. (2000). *The effects of animated characters on anxiety, task performance, and evaluations of user interfaces*. Proceedings of the CHI 2000 Conference on Human Factors in Computing Systems, The Hague, Netherlands.
- Ricketts, M. (1982). The biologic significance of the divine proportion and Fibonacci series. *American Journal of Orthodontics*, 81, 351-370.
- Rind, B., and Gaudet, S. (1993). Judging personality traits of adolescents from photographs. *The Journal of Social Psychology*, 133(6), 815-824.
- Rist, T., André, E., and Müller, J. (1997). *Adding animated presentation agents to the interface*. IUI '97. Proceedings of the 1997 International Conference on Intelligent User Interfaces, Orlando, FL.
- Rizzo, P., Veloso, M. M., Miceli, M., and Cesta, A. (1999). Goal-based personalities and social behaviors in believable agents. *Applied Artificial Intelligence*, 13, 239-271.
- Rogers, C. R. (1955). *Client-centered therapy: Its current practice, implications, and theory*. Boston: Houghton Mifflin.
- Rosenberg, S. (1977). New approaches to the analysis of personal constructs in person perception. In A. L. Land and J. K. Cole (Eds.), *Nebraska symposium on motivation* (Vol. 24, pp. 179-242). Lincoln: University of Nebraska Press.
- Rousseau, D. (1996). *Personality in computer characters*. Working Notes of the AAAI-96 Workshop on AI/ALife.
- Rousseau, D., and Hayes-Roth, B. (1996). *Personality in synthetic agents* (No. Report No. KSL 96-21): Santa Clara, CA: Knowledge Systems Laboratory, Department of Computer Science, Stanford University.
- Rousseau, D., and Hayes-Roth, B. (1997). *Interacting with personality-rich characters* (No. KSL-97-06). Santa Clara, CA: Knowledge Systems Laboratory, Department of Computer Science, Stanford University.
- Rubenstein, A. J., Kalakanis, L., and Langlois, J. H. (1999). Infant preferences for attractive faces: A cognitive explanation. *Developmental Psychology*, 35(3), 848-855.

- Samal, A., and Jyengar, P. A. (1992). Automatic recognition and analysis of human faces and facial expressions: A survey. *Pattern Recognition*, 25(1), 65-77.
- Schneider, D. J. (1973). Implicit personality theory: A review. *Psychological Bulletin*, 79, 294-309.
- Secord, P. F., Dukes, W. F., and Bevan, W. W. (1954). An experiment in social perceiving. *Genetic Psychology Monographs*, 49(Second Half), 231-279.
- Shaw, E., Johnson, W. L., and Ganeshan, R. (1999). *Pedagogical agents on the web*. Proceedings of the third international conference on autonomous agents, Seattle, WA.
- Shaw, W. C. (1988). Social aspects of dentofacial anomalies. In T. R. Alley (Ed.), *Social and applied aspects of perceiving faces* (pp. 191-216). Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers.
- Shepherd, J. W., Ellis, H. D., McMurrin, M., and Davies, G. M. (1978). Effect of character attribution on photofit construction of the face. *European Journal of Social Psychology*, 8, 263-268.
- Shneiderman, B., and Maes, P. (1997). Direct manipulation vs. interface agents. *Interactions*, 4(6), 42-61.
- Sirovich, L., and Kirby, M. (1987). Low dimensional procedure for the characterization of human faces. *Journal of Optical Society of America*, 4(3), 519-524.
- Smith, D. C. (1994). KidSim: Programming agents without a programming agent. *Communication ACM*, 37(7), 55-67.
- Solomon, M. R., Ashmore, R. D., and Longo, L. C. (1992). The beauty match-up hypothesis: Congruence between types of beauty and product images in advertisements. *Journal of Advertising*, 21, 23-34.
- Sproull, L., Subramani, R., Kiesler, S., Walker, J. H., and Waters, K. (1996). When the interface is a face. *Human Computer Interaction*, 11, 97-124.
- Stern, A., Frank, A., and Resner, B. (1998). *Virtual petz (video session): A hybrid approach to creating autonomous, lifelike dogz and catz*. Proceedings of the Second International Conference on Autonomous Agents, Minneapolis, MN.

- Stevenage, V., and McKay, Y. (1999). Model applicants: The effect of facial appearance on recruitment decisions. *British Journal of Psychology*, 90(2), 221.
- Stork, D. G. (Ed.). (1997). *HAL's legacy: 2001's computer as dream and reality*. London: The MIT Press.
- Strickland, F. C. (1956). *The technique of acting*. New York: McGraw-Hill Book Company, Inc.
- Swagerman, J. (1987). *The Artificial Concern REalization System ACRES: A computer model of emotion*. Unpublished doctoral dissertation, University of Amsterdam.
- Symons, D. (1979). *The evolution of human sexuality*. New York: Oxford University Press.
- Takeuchi, A., and Naito, T. (1995). *Situated facial displays: Towards social interaction*. Conference Proceedings on Human Factors in Computing Systems, Denver, CO.
- Thalmann, D., Noser, H., and Huang, Z. (1997). Autonomous virtual actors based on virtual sensors. In R. Trappl and P. Petta (Eds.), *Creating personalities for synthetic actors: Towards autonomous personality agents* (Vol. 1195). New York: Springer.
- The MathWorks. (2000). *Using MATLAB: The language of technical computing*. Natick, MA: The Mathworks, Inc.
- Thomas, F., and Johnston, O. (1981). *Disney animation: The illusion of life*. New York: Abbeville Press.
- Thórisson, K. R. (1997). *Gandalf: An embodied humanoid capable of real-time multimodal dialogue with people*. Proceedings of the First International Conference on Autonomous Agents, Marina del Rey, CA: ACM Press.
- Thornhill, R., and Gangestad, S. W. (1993). Human facial beauty: Averageness, symmetry, and parasite resistance. *Human Nature*, 4(3), 237-269.
- Todd, J. T., and Mark, L. S. (1980). The perception of human growth. *Scientific American*, 242(2), 132-144.
- Todd, J. T., Mark, L. S., Shaw, R. E., and Pittenger, J. B. (1981). Issues related to the prediction of craniofacial growth. *American Journal of Orthodontics*, 79, 63-80.

- Tomkins, S. S., and Messick, S. (Eds.). (1963). *Computer simulation of personality: Frontier of psychological theory*. New York: Wiley.
- Towns, S. G., Voerman, J. L., Callaway, C. B., and Lester, J. C. (1998, January). *Coherent gestures, locomotion, and speech in life-like pedagogical agents*. IUI '98. Proceedings of the 1998 International Conference on Intelligent User Interfaces, San Francisco, CA.
- Trappl, R., and Petta, P. (1997). *Creating personalities for synthetic actors: Towards autonomous personality agents*. Berlin and New York: Springer.
- Triandis, H. C. (1989). The self and social behavior in differing cultural contexts. *Psychological Review*, 96, 506-520.
- Turk, M. A., and Pentland, A. P. (1991a). Eigenfaces for recognition. *Journal of Cognitive Neuroscience*, 3(1), 71-86.
- Turk, M. A., and Pentland, A. P. (1991b). *Face recognition using eigenfaces*. IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Silver Spring, MD.
- Unger, R., Hilderbrand, M., and Madar, T. (1982). Physical attractiveness and assumptions about social deviance: Some sex-by-sex comparisons. *Personality and Social Psychology Bulletin*, 8(2), 293-299.
- Valentin, D., Abdi, H., Edelman, B. E., and O'Toole, A. J. (1997). Principal component and neural network analyses of face images: What can be generalized in gender classification? *Journal of Mathematical Psychology*, 41(4), 398-413.
- Valentin, D., Abdi, H., and O'Toole, A. J. (1994). Categorization and identification of human face images by neural networks: A review of the linear autoassociative and principal component approaches. *Journal of Biological Systems*, 2(3), 413-429.
- Valentin, D., Abdi, H., O'Toole, A. J., and Cottrell, G. W. (1994). Connectionist models of face processing: A survey. *Pattern Recognition*, 27(9), 1209-1230.
- Vardi, G. (1999, February). *Navigation scheme for interactive movies with linear narrative*. HYPERTEXT '99. Proceedings of the 10th ACM Conference on Hypertext and Hypermedia: Returning to our Diverse Roots, Edinburgh, Scotland.

- Vilhjálmsón, H. (1997). *Autonomous communicative behaviors in avatars*. Unpublished M.S. thesis, MIT, Cambridge, MA.
- Vilhjálmsón, H., and Cassell, J. (1998). *BodyChat: Autonomous communicative behaviors in avatars*. Proceedings of ACM Autonomous Agents '98, Minneapolis, MN.
- Vyzas, E., and Picard, R. W. (1998). *Affective pattern classification*. AAAI Fall Symposium Series: Emotional and Intelligent: the Tangled Knot of Cognition, Orlando, FL.
- Walker, J. H., Sproull, L., and Subramani, R. (1994). *Using a human face in an interface*. Proceedings of the CHI '94 Conference Companion on Human Factors in Computing Systems, Boston, MA.
- Walker, M., Cahn, J., and Whittaker, S. (1997). *Improvising linguistic style: Social and affective bases of agent personality*. Proceedings of the First International Conference on Autonomous Agents, ACM SIGART, Marina del Rey, CA.
- Watson, D. (1989). Strangers' ratings of the five robust personality factors: Evidence of a surprising convergence with self-report. *Journal of Personality and Social Psychology*, 57(1), 120-128.
- Watt, S. N. K. (1997). Artificial societies and psychological agents. In H. S. Nwana and N. Azarmi (Eds.), *Software agents and soft computing: Towards enhancing machine intelligence* (p. 298). New York: Springer.
- Wavish, P. (1992). Exploiting emergent behavior in multi-agent systems. In Y. Demazeau and E. Werner (Eds.), *Decentralized artificial intelligence* (Vol. 3, pp. 297--310). Germany: Elsevier Science Publishers.
- Wavish, P., and Connah, D. (1997). *Virtual actors that can perform scripts and improvise roles*. Proceedings of the First International Conference on Autonomous Agents, Marina del Rey, CA.
- Wooldridge, M. J. (1996). Agents as rorschach test: A response to franklin and graesser. In J. P. Müller, M. J. Wooldridge and N. R. Jennings (Eds.), *Intelligent agents III: Agent theories, architectures, and languages* (pp. 47-48). Berlin: Springer.
- Wright, P., Milroy, R., and Lickorish, A. (1999). Static and animated graphics in learning from interactive texts. *European Journal of Psychology of Education*, XIV, 203-224.

- Yamaguchi, M. K., Hirukawa, T., and Kanazawa, S. (1995). Judgment of gender through facial parts. *Perception, 24*(5), 563-575.
- Young, A. W., and Bruce, V. (1998). Pictures at an exhibition: The science of the face. *The Psychologist, 11*, 120-125.
- Zebrowitz, L. A. (1998). *Reading faces: Window to the soul?* Boulder, CO: Westview Press.
- Zebrowitz, L. A., and McDonald, S. M. (1991). The impact of litigants' babyfacedness and attractiveness on adjudications in small claims courts. *Law and Human Behavior, 15*(6), 603-623.
- Zebrowitz, L. A., and Montepare, J. M. (1992). Impressions of babyfaced individuals across the life span. *Developmental Psychology, 28*(6), 1143-1152.
- Zebrowitz, L. A., Voinescu, L., and Collins, M. A. (1996). "Wide-eyed" and "crooked-faced": Determinants of perceived and real honesty across the life span. *Personality and Social Psychology Bulletin, 22*(12), 1258-1269.