*AISB Journal* **1**(6), © SSAISB, 2005
http://www.aisb.org.uk

# A Computational Model of the Trait Impressions of the Face for Agent Perception and Face Synthesis

Sheryl Brahnam

Department of Computer Information Systems, Missouri State University
901 South National, Springfield, MO, 65804, USA
*shb757f@smsu.edu*

**Abstract**

This paper reports a first attempt at developing a computational model of the trait impressions of the face for embodied agents that accommodates the social perception and social construction of faces. Holistic face classifiers, based on principle component analysis (PCA), were trained to match the human classification of faces along the bipolar rating extremes of the following trait dimensions: adjustment, dominance, warmth, sociality, and trustworthiness. Although results were marginally better than chance in matching the perception of dominance (64%), classification rates were significantly better than chance for adjustment (71%), sociality (70%), trustworthiness (81%) and warmth (89%). A second exploratory study demonstrates how PCA models of trait classes could be used by agents to generate faces. Novel faces were synthesized by probing specific PCA trait attribution spaces. Human subjects were then asked to rate the synthesized faces along a number of trait dimensions, and it was found that the synthesized faces succeeded in eliciting predicted trait evaluations.

## 1 Introduction

The semiologist, Magli (1989) has remarked, "upon seeing a face, we immediately produce a symbolic framework that confronts us with a complex and ancient cultural experience" (p. 90). A recurrent theme in the fables, proverbs, and histories, both oral and written, of cultures as diverse as the Egyptian, African, Chinese, and European is that the face is inscribed with signs that reveal the essence of a person's inner soul (Frey, 1993). Although many modern people scoff at such notions and recite such maxims as "Don't judge a book by its cover," there is considerable evidence in the person perception literature that people are predisposed to form impressions of a person's social status, abilities, dispositions, and character traits based on nothing more than that person's facial appearance. Furthermore, there is evidence that these judgments influence and guide people's behavior towards others, especially in situations that are ambiguous or where little information about a person is known (Hochberg and Galper, 1974). As the historian Frey (1993) recently noted, "To this day, the quest to read a person's inner world from her outer appearance has lost nothing of its momentum . . . it seems that the advent of the 'Age of Television' has given additional impetus to the age-old fascination with human appearance" (p. 64).

People are not just caught up in evaluating other people's faces; they are equally preoccupied with managing their own appearances. One would be hard pressed to name one culture that has not required its members to modify their faces in some way. The psychologist Ligget (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). The need to mark a person's social status, to proclaim skill in hunting and in war, and to put one's best face forward at a business meeting are some of the many motives behind facial elaborations. The face, more than any other part of the body, stands for and is identified with the social self, and so important are the social consequences of the appearance of the face that many people are willing to endure enormous pain and expense to manage the messages sent by their faces. It seems that the French poet Henri Michaux may have been right when he wrote, "We lead an excessively facial life" (quoted in Landau, 1989, p. 234).

Once embodied agents enter the social arena, they will be expected to understand the cultural language of the face and not just short-term surface communications and behaviors, such as eye blinking, gazing, head tilting, facial gestures, and emotional expressiveness, which forms the focus of current research into agent faces (Pelachaud and Poggi, 2002). As essential as this research is, research that explores the underlying morphology, or the look of the face, is also important. Sproull et al. (1996) have demonstrated, for instance, that morphological shifts in the facial appearance of embodied agents affect users in ways that mirror findings in the person perception literature, and Donath (2001) and others have cautioned researchers to consider carefully the facial appearance of their agents. Unfortunately, there is no way to predict during design time all the circumstances, tasks, and people the agent will encounter. Thus, there is no way to equip an agent with an embodiment that will function optimally in all situations.

Although people today have recourse to plastic surgery and a host of cosmetic aids, there is a limit to the extent that people can shape their faces for social purposes, but embodied agents do not share this limitation. There is no reason to assume that a particular agent's embodiment must be singular or static or that it must be designed offline by human beings. Facial morphology could be as expressive a channel of communication for embodied agents as are emotional facial displays. Like countless others who each morning prepare their faces to meet the demands of their day, so embodied agents could learn to construct social masks that are appropriate for the situations they encounter, the users they meet, and the tasks they need to accomplish.

To participate in the social world, embodied agents will also need to know how to evaluate the human faces they encounter. Rather than use a predefined set of interaction tactics and practices, for instance, the cultural information found in the user's face could serve the agent as a basis for formulating a more personalized and realistic initial interaction strategy that could then be adjusted as further information about a user is obtained. At the very least, predicting how other human beings would react to a person's facial appearance could produce interaction strategies that mimic the more natural interaction styles of human beings. Understanding the social language of the face will also allow agents to participate in such common activities as commenting on the appearance of others in ways that are realistic and appropriate. This could enhance the agent's believability and acceptability. Being able to perceive faces as people perceive them could also make embodied agents more sensitive in their encounters.

But what is the social language of the face? How can it be modeled for embodied agents? An ongoing area of investigation in social psychology revolves around understanding the facial characteristics that contribute to the formation of impressions about a person's character. As reviewed in section 2, it has been found that large clusters of character traits are strongly associated with attractiveness judgments, emotional displays, age, and gender (Zebrowitz, 1998). As a result, recent research in this area has focused almost exclusively on investigating the facial characteristics of attractiveness, emotion, age, and gender and the role these characteristics play in the attribution process. Research aimed at directly exploring morphological characteristics that trigger very specific attributions has all but been neglected, primarily because this line of investigation has mostly been feature based, has produced contradictory results, and has not lent itself to theory building.

In section 3, it is argued that a psychologically viable model of the trait attribution process is not necessary for embodied agents. Rather, since agents need to perceive faces in terms of the impressions they produce, it would be best to model specific traits directly using holistic face recognition techniques, such as principle component analysis (PCA), or equivalently, linear autoassociative neural networks. As noted in section 4, these techniques have already proven successful at classifying faces according to identity (Turk and Pentland, 1991b), emotion (Padgett and Cottrell, 1998), gender, and age (Valentin et al., 1994a), characteristics that are strongly correlated with impression formation. Thus, it is reasonable to expect that these classifiers will succeed in modeling the human classification of faces into specific trait attribution classes. Another advantage in using holistic face classification techniques is that they lend themselves to face synthesis (Vetter and Poggio, 1997) and, thus, could be used by agents to generate faces with a high probability of making specific impressions on users.

Two studies are reported in this paper that use PCA to model the trait impressions of the face. The objective of the first study was to model the trait impressions of facial morphology. As described in detail in section 5, PCA classifiers were trained to classify faces that were rated either high or low within the five trait dimensions of adjustment (adjusted/unadjusted), dominance (dominant/submissive), warmth (warm/cold), sociality (social/unsocial), and trustworthiness (trustworthy/untrustworthy). A second exploratory study, presented in section 6, demonstrates how PCA classifiers can be used to create novel faces calibrated to produce specific trait impressions. The results and some limitations of the two studies are discussed in section 7, and the paper is concluded in section 8 by noting some of the contributions of these studies and by offering directions for future research.

## 2 Person Perception Literature on the Trait Impressions of the Face

As mentioned in the Introduction, psychological research aimed at directly exploring morphological characteristics that trigger very specific trait attributions has virtually been neglected in large part because this approach has not lent itself to theory building. Although several theories have been advanced to explain why it is that certain facial characteristics consistently elicit specific personality impressions, 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 survival (McArthur and Baron, 1983). Recognizing an angry face, for example, triggers lifesaving 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* [italics mine] 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. Two other overgeneralization effects that have received less attention but are nonetheless significant are based on emotion and gender (Alley, 1988; Symons, 1979).

The trait associations and morphological characterizations of each of these overgeneralization effects are summarized below. Included in the summary are descriptions of some of the more important models of facial attractiveness, maturity, gender, and emotion.

## 2.1 Attractiveness Halo Effect

It is popularly believed that social benefits accrue to those who are most attractive, and current research supports this claim. People respond positively to attractiveness and associate it with positive character traits. Attractive people are considered more socially competent, potent, healthy, intellectually capable, and moral than those less attractive. They are also perceived as being psychologically more adapted (Langlois et al., 2000). Facial abnormalities and unattractiveness, in contrast, elicit negative responses and are associated with negative traits (Langlois et al., 2000). Unattractive people are considered less socially competent and willing to cooperate (Mulford et al., 1998). They are also considered more dishonest, unintelligent, and psychologically unstable and antisocial. Unattractive people are often ignored and, if facially disfigured, avoided (Bull and Rumsey, 1988). The unattractive are also more likely to be objects of aggression (Alcock et al., 1998) and to suffer abuse (Langlois et al., 2000).

What are the morphological characteristics that make a face attractive? To date there is no theory of attractiveness that is generally accepted. Nonetheless, contemporary research into facial attractiveness indicates that straightness of profile (Carello et al., 1989) and closeness to the average (Langlois and Roggman, 1990) are some important factors in attractiveness judgments.

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 determining 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, even among children, for orthognathic or straight profile shapes (Carello et al., 1989; Magro, 1997; Lucker and Graber, 1980). Least attractive is the prognathic (Carello et al., 1989).

Particularly noteworthy, in terms of its potential for adjusting the impressions of agent faces, is the finding that facial attractiveness increases as faces are moved closer to the average (see Figure 1). 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 appeared notably more attractive. Little was done with his observation until 1952, when 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, with 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. Their findings have more recently been confirmed by Langlois, Roggman, and Rieser-Danner (1990), Rubenstein et al. (1999), and Rhodes and Tremewan (1996).


Figure 1: Increased Attractiveness of Averaged Faces. A face (1.0) moved towards the mean (0.0) of 220 randomly generated faces increases in attractiveness. The faces were generated from facial features in the composite program FACES by InterQuest and Micro-Intel, and the average face was computed by averaging the pixel gray scale values of the faces.

## 2.2 Facial Maturity Overgeneralization Effect

Perhaps no face is more capable of eliciting a favorable response than that of a baby. 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 babies (Zebrowitz, 1998). Babyfaced people are universally attributed childlike characteristics. They are perceived to be more submissive, naïve, honest, kindhearted, weaker, and warmer than others. They are also perceived as being more helping, caring, and in need of protection (Berry and McArthur, 1986). Mature-faced individuals, in contrast, are more likely to command respect and to be perceived as experts (Zebrowitz, 1998).

The morphological characteristics 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 (Zebrowitz, 1998).

Figure 2: Negative (Left) to Positive (Right) Cardioidal Strain Transformations. Reproduced from Pittenger and Shaw (1975), p. 376. Copyright ©1975 by the American Psychological Association. Reprinted with permission.

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 2 illustrates the morphological characteristics of facial maturity. The craniofacial profile shapes were produced using a cardiodial strain transformation developed by Todd and Mark (1980). Applied to standard profile shapes, a positive application of the transformation has been shown to approximate real growth (Todd et al., 1981; Todd and Mark, 1980). 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 (Zebrowitz, 1998; Alley, 1983). 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, 1986), and elicit stronger desires to nurture and protect (Alley, 1983).

Examining Figure 2, it can be observed that an extreme negative cardioidal transformation results not only in the most youthful but also the most retrognathic profile shape. Similarly, an extreme application of a positive cardioidal transformation produces the most mature looking and prognathic profile shape. As noted above, profile shape is related to attractiveness judgments, and there is some evidence that the cardioidal transform influences attractiveness judgments as well as judgments regarding facial maturity (Jones, 1995).

## 2.3 Gender Overgeneralization Effect

The gender overgeneralization effect is strongly correlated with facial maturity (Zebrowitz, 1998). Female faces, more than male faces, tend to retain into adulthood the morphological characteristics of youth (Enlow and Hans, 1996) and are more likely to be ascribed characteristics associated with babyfacedness: female faces are thought to be more submissive, caring, and in need of protection. Similarly, male faces, tending to be morphologically more mature, are perceived as having the psychological characteristics typically associated with mature-faced individuals: male faces are thought to be more dominant, intelligent, and capable.

## 2.4 Emotion Overgeneralization Effect

While many social psychologists believe that facial impressions of character are related in part to the morphological configurations that characterize emotional displays, the overgeneralization effect of emotion has not received as much attention as some of the other overgeneralization effects. Nonetheless, there is evidence supporting the idea that morphological configurations suggestive of emotional expressions play a role in the formation of trait impressions. Take smiling for instance. People react positively to smiling faces and find them disarming and thus not very dominant (Keating et al., 1981a). As illustrated in Figure 3, facial dominance significantly declines where even a slight smile is discernible (Mueller and Mazur, 1996). As would be expected, faces where the lips naturally turn upwards are likewise viewed more positively; such faces are considered friendly, kind, easygoing, and nonaggressive (Secord et al., 1954). In a similar vein, faces that have features indicative of anger or hostility, e.g., low-lying eyebrows, thin lips, and withdrawn corners of the mouth, are perceived to be more threatening, aggressive, and dominant (Keating et al., 1981b).

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 describes any facial behavior, including emotion. Recently, a number of emotion recognition systems have been developed that use FACS (Bartlett, 1998; Donato et al., 1999; Essa and Pentland, 1997). There is also a growing body of research concerned with synthesizing emotional displays in artificial faces (Massaro, 1997; Picard, 1997; Waters and Terzopoulos, 1992).

Figure 3: Illustration of the Overgeneralization Effect of Emotion. In the two images, only the lips differ. These faces were generated using FACES by InterQuest and Micro-Intel.

# 3 Problems with Indirectly Modeling the Trait Impressions of the Face

After reviewing the person perception literature on the overgeneralization effects, it might seem that one effective way to model the trait impressions of the face for agent perception and for face synthesis would be to do so indirectly by modeling facial attractiveness, maturity, gender, and emotion. Certainly an agent could alter the social impact of its face by moving it either further or closer to the average or by applying the cardioidal strain transformation or by *freezing* certain emotional displays. Although building agents that learn best how to adapt their faces using these techniques is a research area worth investigating, there are a number of problems with an indirect approach to modeling the trait impressions of the face.

A major problem concerns the difficulty of using models of attractiveness, facial maturity, gender, and emotion to predict, or to classify, faces in terms of the traits they elicit. In other words, these models would not readily provide embodied agents with *perceptual systems* capable of decoding the impressions faces make on human observers. An exception to this concerns the overgeneralization effect of emotion. As noted above, a number of systems have been developed that recognize and produce emotional facial displays. Most emotion recognition systems, however, utilize FACS, which describes surface facial behaviors more than it describes facial morphology. Recently emotion recognition classifiers have been developed that are based on holistic face recognition techniques (Cottrell and Metcalfe, 1991; Padgett and Cottrell, 1998). Even though these classifiers take into account all the information available in pixel representations of faces, not enough is known about the relationship of the overgeneralization effect of emotion and the person perception of the face to utilize this technology in this task domain. Furthermore, in terms of production and recognition, it is doubtful that morphological characteristics in common with emotional displays can account for a significant range of traits. What emotional display, for example, best reflects honesty or intellectual competence?

This last point highlights a number of other problems with trait associations and the overgeneralization effects. First, the overgeneralization effects are associated less with individual traits than with clusters of traits. Knowing, therefore, which facial characteristics to alter in order to shift facial impressions along specific trait lines would require a much more refined understanding of the overgeneralization effects. Second, it is possible that some trait impressions may be due to facial configurations that are not accounted for by the overgeneralization effects. Third, no single overgeneralization effect accounts for a comprehensive set of traits. What is a comprehensive set of traits? 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, maturity, attractiveness, integrity, sociability, concern for others, and psychological stability. Others have modified his categories to include potency, or dominance (Feingold, 1992; Eagly et al., 1991). To encompass this representative set of traits in developing facial perception systems for embodied agents, all the facial configurations association with the various overgeneralization effects would need to be addressed.

A better approach to take in modeling the trait impressions of the face for embodied agents is to focus directly on the perception of those facial features that give rise to specific trait impressions. It has already been remarked that psychological studies have recently steered away form this line of research because this approach has failed to produce viable psychological theories of the trait attribution process. However, a model of the trait impressions of the face for embodied agents need not be as comprehensive and as capable of explaining the attribution process as psychological models need to be. Focusing on the *perception* of traits in faces using, for example, holistic face classification techniques would allow the classifier to discover the relevant features in trait formations. Other advantages in using holistic face recognition technologies to model the trait impressions of the face are presented in the next section.

# 4 PCA Face Representation and Classification

Isolating the features 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. Much of the visual information contained within a face is highly redundant. What varies is but a small set of relations between features and small differences in textures, complexions, and shapes. 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). 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 (Valentin et al., 1994b; Burton et al., 1993).

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 et al., 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 PCA in the statistical literature. Since a linear autoassociative neural network is equivalent to finding the principal components of the cross-product matrix of a set of inputs, it is sometimes referred to in the literature as a PCA neural network (Oja, 1992; Diamantaras and Kung, 1996). 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.

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, (see, for instance, Kosugi, 1995).

Since each image contributes to each of the eigenvectors, the eigenvectors 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, 1991a). Some examples of eigenfaces are shown in Figure 4.

Figure 4: First 10 Eigenfaces of 220 Randomly Generated Faces. The eigenfaces of 220 randomly generated faces are ordered left to right, top to bottom, by magnitude of the corresponding eigenvalue.

Individual images can be projected onto the face space and represented exactly as weighted combinations of the eigenface components (see Figure 5). The resulting vector of weights that describe each face can be used in data compression and in face classification. 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,b).

To date, no face classification methods have been applied to the task of classifying faces according to the traits they produce. 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 section 2), it is reasonable to assume that PCA can be employed to this end. Both linear autoassociative neural networks and PCA have successfully been used to classify faces according to gender (Valentin et al., 1997; O'Toole and Deffenbacher, 1997), age (Valentin et al., 1994b), and facial expression (Cottrell and Metcalfe, 1991; Padgett and Cottrell, 1998). What is more, they are simple, well understood, and capable of generating novel images from within the eigenface coordinate system (Turk and Pentland, 1991a; Beymer et al., 1993).

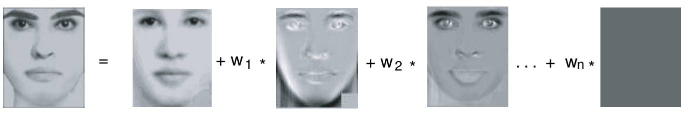


Figure 5: An Illustration of the Linear Combination of Eigenfaces. The face to the left can be represented as a weighted linear combination of eigenfaces.

# 5 Modeling Trait Impressions of the Face Using PCA

This section describes a model of the perception of traits in faces using PCA. The traits modeled were a modification of Rosenberg's (1977) factor analysis of significant trait descriptors, namely, psychological adjustment (adjusted/unadjusted), dominance (dominant/submissive), sociality (social/unsocial), trustworthiness (trustworthy/untrustworthy), and warmth (warm/cold). For definitions of the traits used in this study, the reader is referred to table 8 in the appendix.

## 5.1 Overview

As illustrated in Figure 6, modeling the trait impressions of the face using PCA was a two-step process. The objective of step 1, Data Preparation, was to obtain sets of faces clearly representative of ten bipolar trait descriptors (adjusted/unadjusted, dominant/submissive, social/unsocial, trustworthy/untrustworthy, warm/cold) of the five trait dimensions of adjustment, dominance, sociality, trustworthiness, and warmth. In Step 2, PCA Modeling, these attribution class sets were used to train and to test a separate PCA for each trait dimension.

## 5.2 Data Preparation

The objective of the data preparation process was to prepare faces for PCA classification. As illustrated in Figure 6, this step involved the following: A) generation of the stimulus faces, B) an experiment assessing the trait impressions of the stimulus faces, and C) division of the stimulus faces into trait class sets.

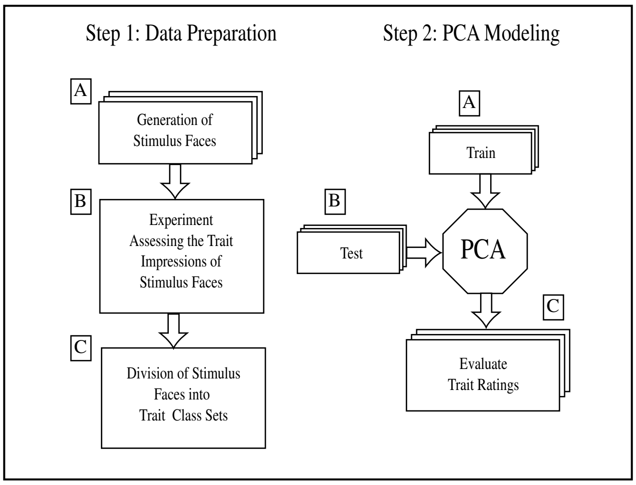


Figure 6: Modeling the Trait impressions of the Face. Note: Although technically PCA is not trained, perhaps because of the equivalency of autoassociative neural networks and PCA, the term *training* is often used in the face recognition literature, (see, for instance, Turk and Pentland, 1991a).

### 5.2.1 Generation of Stimulus Faces

In order to model the trait impressions of the face, it was necessary to acquire a suitable set of stimulus faces. In the person perception literature, stimulus faces are of three types: photographs of faces, drawings of faces, and faces pieced together using facial composite products such as Identi-Kit (Bruce, 1988). No database of faces known to elicit specific trait impressions has been developed for psychological comparison studies. Researchers are required to develop their own datasets of faces.

In contrast, numerous facial databases have been developed to test face classification algorithms. Wegener-Knudsen et al. (2002) provides a comprehensive review of available face databases. However, because these databases have been developed primarily to evaluate face identification techniques, these databases typically 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.

To model the trait impressions of the face, it was important to develop a large set of faces that were representative of a broad range of facial types and features. Furthermore, since the objective of this study was to model the trait impressions of *facial morphology*, the faces also needed to be as neutral in facial expressions as possible, and have such incidentals as hairstyle and accessories removed. Developing a proper database of faces for this task is a complex issue and is discussed further in section 8.

For this initial study, 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. With FACES, it was possible to generate randomly a fairly large number of unique faces by manipulating individual facial features. Moreover, by using specific sets of facial features, faces could be reduced to their basic morphological elements without having to block out features, such as hair, with tape or markers as is typically the case with cognitive and machine recognition studies involving photographs of people's faces.

Two hundred and twenty stimulus faces were generated in step 1.A using FACES. The features selected for constructing the stimulus faces included only the full set of 512 eyes, 541 noses, 570 lips, 423 jaws, 480 eyebrows, and 63 foreheads. Excluded were all sets of facial lines, hair, and accessories.

These images were then cropped (see Figure 7) in such a way that missing hair was less noticeable. This did remove forehead information; but, as these were frontal images of faces, the significance of profile head shape noted in section 2.2 was not relevant here. Care was taken to retain eyebrows, however, because they have been found to contribute to both impressions of gender and of facial maturity (Yamaguchi et al., 1995). A final alteration in the images concerned complexion values, which were set to the value of 190 in a gray scale of 256 values to reduce the effects of race.

Figure 7: Examples of Stimulus Face

#### 5.2.2 Experiment Assessing Trait impressions of Stimulus Faces

Once the stimulus faces were generated, they were evaluated in step 1.B by human subjects as detailed below.

*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.

*Dependent Measures*. Each subject judged a set of 20 faces, randomly selected from the 220 stimulus faces, along the five trait dimensions, using a 7-point bipolar scale. Each image was judged by 10 subjects. The order of the bipolar trait descriptors (adjusted/unadjusted, warm/cold, social/unsocial, dominant/submissive, trustworthy/untrustworthy) was randomized as were the association of the bipolar descriptors with the anchor values of 1 and 7. Subjects were also given trait definitions and, in some cases, behavioral potential questions modeled after Berry and Brownlow (1989) and Zebrowitz and Montepare (1992). Refer to Table 8 in the appendix for the term definitions and the behavioral potential questions. A 7-point scale, rather than a 3-point scale, was used because the general consensus is that it is better to provide research subjects with a gradient of opinion when conducting surveys (Converse and Presser, 1986; Friedman and Amoo, 1999).

*Apparatus*. Desktop computers in a lab setting were used both to display the stimulus faces and to administer the questionnaires.

*Results*. Table 1 presents the mean ratings of the 220 faces for each trait dimension and the standard deviations. In general, the impressions elicited by the stimulus faces were slightly skewed towards low facial warmth and high adjustment, dominance, sociality, and trustworthiness.

As subjects were not required to judge the entire set of 220 faces, a complete analysis of human variance is not presented. It is important to stress that the objective of this ex-

Table 1: Rater Means and Standard Deviations of the Stimulus Faces

| Trait Dimension Descriptor | Dimension Means | Standard Deviation |
|---|---|---|
| Adjusted | 4.03 | 0.82 |
| Dominant | 4.16 | 0.85 |
| Social | 4.07 | 0.97 |
| Trustworthy | 4.00 | 0.87 |
| Warm | 3.94 | 1.01 |

periment was not to do yet another psychological study regarding the person perception of faces. This is a topic that has been well researched, and people of different ages, genders, races, and cultures have been shown to be remarkably consistent in their judgments (Albright et al., 1997; Zebrowitz et al., 1993). Rather, the experimental design was geared solely towards obtaining human judgments of the 220 faces in order to extract those few faces that unambiguously elicit the specific traits explored in this study.

#### 5.2.3 Division of Stimulus Faces into Trait Class Sets

In most face classification tasks, such as classifying faces by identity, gender, and race, 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 many faces fail to elicit strong opinions and by the fact that human beings, while consistent in their ratings, are not in total agreement.

In this study, faces were divided in step 1.C, based on their average rating, into three classes: low (with a mean range of 1.0 - 2.9), neutral (with a mean range of 3.0-4.9) and high (with a mean range of 5.0 - 7.0). As a PCA classification of faces with weak attributions is irrelevant for that trait dimension, that is, the classification is not unambiguous, neutral faces were excluded from the PCA training and testing sets. In addition, faces were pruned from the low and high classes that had a standard deviation greater than 1.5 or that had 50% or more ratings marked neutrally or in the opposite class. Thus, only those few faces that elicited strong impressions were used to develop the PCA models.

Table 2 lists the total number of images selected to form the high and low attribution class sets for each of the five trait dimensions. The total number of images is greater than 220 because some images produced significant trait impressions along more than one dimension.

### 5.3 Step 2: PCA Modeling

Once a suitable dataset was developed, five separate PCAs, one for each of the five trait dimensions, were trained and evaluated using MATLAB (MathWorks, 2000). Outlined below are the operations involved in training and testing the PCAs for each trait dimension. The reader should refer to Turk and Pentland (1991a) for additional details.

#### 5.3.1 Training

Training a PCA, in step 2.A, requires three operations: 1) randomly dividing the trait class sets into separate training and testing sets; 2) calculating the eigenvectors from the

Table 2: Number of Images Selected for the Trait Attribution Classes

| Trait Dimension | Attribution Class | Number |
|---|---|---|
| Adjustment | Low | 11 |
| | High | 12 |
| Dominance | Low | 11 |
| | High | 14 |
| Sociality | Low | 12 |
| | High | 14 |
| Trustworthiness | Low | 10 |
| | High | 15 |
| Warmth | Low | 14 |
| | High | 14 |

training set; and 3) calculating the distribution of each class within the face space.

*Operation 1.* The two attribution classes of images (high and low) for each dimension were merged and divided into a training set of images and a testing set of images, with an equal number of images from both classes (high and low) represented in the training and testing sets.

*Operation 2.* The eigenvectors were computed using the following algorithm:

1. Reshape the training images into column vectors, which together form the matrix $\mathbf{\Gamma}$. Let $\mathbf{\Gamma}_k$ represent the column vector of face *k*.

2. Normalize the column vector for each face *k* in the training set of *M* images:

$$\mathbf{\Phi}_{\boldsymbol{k}} = \mathbf{\Gamma}_k - \mathbf{\Psi},\ where\ \mathbf{\Psi} = \frac{1}{M}\sum_{k}^{M}\mathbf{\Gamma}_k \tag{1}$$

3. Compute the eigenfaces using singular value decomposition:

$$\mathbf{\Phi} = \mathbf{USV}^T \tag{2}$$

   where **S** is a diagonal matrix whose diagonal elements are the singular values, or *eigenvalues*, of $\mathbf{\Phi}$, $\mathbf{V}^T$ is the transpose of **V**, and **U** and **V** are unary matrices. The columns of **U** are the eigenvectors of $\mathbf{\Phi\Phi}^T$, and are referred to as *eigenfaces*, as they are face-like in appearance. The columns of **V** are the eigenvectors of $\mathbf{\Phi}^T\mathbf{\Phi}$ and are not used in this analysis.

*Operation 3.* The distribution within the face space for each of the classes was computed by projecting each training image $\mathbf{\Gamma}_k$ onto the eigenfaces as follows:

$$\omega_k = \mathbf{U}_k^T(\mathbf{\Gamma}_k - \mathbf{\Psi}) \qquad \text{(for } k = 1, ..., M) \tag{3}$$

Let $\mathbf{\Omega}^T = [\omega_1, \omega_1, ..., \omega_M]$, 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, $\mathbf{\Omega}$, for each training class (Turk and Pentland, 1991b).

#### 5.3.2 Testing

Evaluating the system using the testing set of images in step 2.B required two operations: 1) projecting each test image $\mathbf{\Gamma}_k$ onto the face space to obtain $\mathbf{\Omega}_k$ as in Operation 3 above, and 2) determining the best-fit class membership. Best-fit membership was 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 was then averaged and used as an index to evaluate the performance of the system.

#### 5.3.3 Model Evaluation

Because of the small number of images in the trait sets, a cross-validation technique was employed in step 2.C such that only two images from each set were selected to form the testing set, and training and testing were performed as outlined in sections 5.3.1 and 5.3.2. This process was repeated twenty times for each trait dimension. The ratio of right to wrong classifications was used as the classification index, and the twenty classification indexes of each trait were averaged to form the final classification score for that trait.

Table 3: Averaged PCA Classification Scores

| Trait Dimension | Classification Rate |
|---|---|
| Adjustment | .71 |
| Dominance | .64 |
| Sociality | .70 |
| Trustworthiness | .81 |
| Warmth | .89 |

Table 3 displays the classification scores for the five PCAs. All five PCA classification rates were above chance, with trustworthiness and warmth scoring well above chance. Results were not as good for dominance. See section 7 for a more complete discussion of the results of this study.

## 6 Synthesizing Faces with Predicted Trait Evaluations

Reported in this section is a preliminary study conducted to demonstrate the possibility of using PCA to construct novel faces with a high probability of eliciting specific trait impressions. As described in section 6.1, certain stimulus faces were projected onto PCAs trained with stimulus faces that ranked in the first study as either high or low in a trait dimension; they were then reconstructed. This process generated novel faces. As described in section 6.2, predictions were made regarding the impressions the synthesized faces would make on human observers. These predictions were then compared with the evaluations of human subjects.

### 6.1 Face Synthesis

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.

Two notable exceptions are Vetter and Poggio (1997) and a pilot composite system developed by Hancock (2000). In the latter system, shape-free facial information and shape information are extracted and subjected to PCA. Using a genetic algorithm, novel faces are evolved by recombining the eigenfaces of shape-free facial images and then morphing them to any number of shape components.

One of the benefits in using the database of facial features in FACES to produce the stimulus faces used in these studies is that the features were normalized and aligned to facilitate seamless combinations. Since the entire set of stimulus faces were therefore automatically normalized and aligned (the degree of alignment can be seen in the clarity of the eigenfaces in Figure 1), this exploration into face synthesis used a simpler approach to recombine the eigenfaces: face synthesis was performed by probing the appropriate trait attribution space. With PCA, image projection is onto a low-dimensional space (Turk and Pentland, 1991b). For this reason, even images that look nothing like a face, when projected onto a face space, produce face-like reconstructions. In other words, as illustrated in Figure 8, these non-face images serve as a means of probing the face space since the reconstructions combine characteristics of the faces used to define the PCA face space.

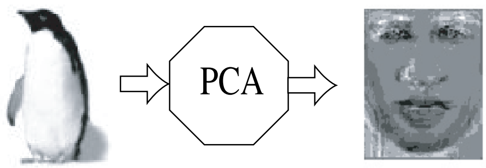


Figure 8: Illustration of Face Space Probing. An Image of a Penguin Projected onto a PCA Face Space Results in a Face-like Reconstruction.

In this study, a subsection of the face space, namely the PCA trait attribution space, was similarly probed. Attribution space probing was accomplished as follows: two PCAs, one for each of the two attribution classes of high and low for each trait dimension, were trained using all images in the appropriate attribution class set. In order to generate novel faces, the PCA attribution spaces needed to be seeded with as many faces as possible. For this reason, all stimulus images in the first experiment with an average rating $\leq 3.0$ within each trait dimension were used to train the low PCA attribution spaces, and all stimulus images with an average rating of $\geq 5.0$ were used to train the high PCA attribution spaces. See Table 4 for the total number of images used to train the PCAs for each of the eight attribution classes.

Face synthesis was performed by probing the two PCA attribution spaces for each of the five trait dimensions. This was accomplished by taking an image in one attribution class set and projecting and reconstructing it using a PCA trained with the images of the opposite attribution class set.

Figure 9 shows two examples of face synthesis using the cold (low) and the warm (high) PCA attribution spaces. On the right of Figure 9, an image classified as cold (top) and an image classified as warm (bottom) were projected onto the opposite PCA attribution space and reconstructed. This resulted in the new images shown on the left. A total of 340 images (every image in one attribution class was projected onto the opposite

Table 4: Number of Stimulus Faces Used For Face Synthesis.

| Trait Dimension Descriptor | Number Rated $\leq 3.0$ (Low Attribution Class) | Number Rated $\geq 5.0$ (High Attribution Class) |
|---|---|---|
| Adjusted | 25 | 26 |
| Dominant | 17 | 37 |
| Social | 45 | 34 |
| Trustworthy | 29 | 34 |
| Warm | 51 | 42 |

PCA attribution space) were synthesized from the eight PCA attribution spaces using this procedure. As the stimulus faces were closely aligned, few artifacts were introduced in the reconstruction process. Compare, for example, in Figure 9, the artifacts introduced in the synthesized faces (right) to the stimulus faces (left). Although some faces produced more artifacts than others, no attempt was made to enhance the synthesized faces.

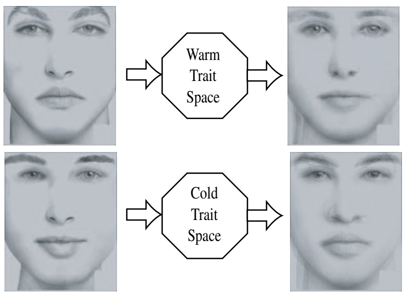


Figure 9: Examples of Faces Synthesized by Probing the PCA Attribution Spaces of Facial Warmth

## 6.2 Human Assessment of Synthesized Faces

One hundred ten synthesized images were randomly selected and rated by ten human subjects from the same pool of subjects used in the first experiment to assess the stimulus faces. The same procedures were also followed as in the first experiment. It was predicted that faces synthesized by probing the low PCA attribution space of a particular dimension would be ranked by the human subjects at the lower end of that trait dimension's rating scale (i.e., $< 3.5$) and that faces synthesized by probing the high PCA attribution space of the same dimension would be ranked at the higher end of the rating scale (i.e., $> 3.5$).

## 6.3 Results

Table 5 shows the average assessment of the faces synthesized from the two PCA attribution spaces for each of the five trait dimensions. In general, faces synthesized by probing the low PCA attribution spaces were rated at the lower end of the scale (average low score is 3.11), whereas faces synthesized by probing the high PCA attribution spaces were rated at the higher end of the scale (average high score is 4.82). Only faces synthesized by probing the low PCA attribution space of adjustment failed to be rated as predicted.

Table 5: Averaged Trait Ratings of Synthesized Faces and Standard Deviations.

| Trait Dimension | Attribution Space | Average Rating of Synthesized Faces | Standard Deviation |
|---|---|---|---|
| Adjustment | Low | 3.86 | 1.44 |
| | High | 5.12 | 1.38 |
| Dominance | Low | 3.34 | 1.13 |
| | High | 5.00 | 1.30 |
| Sociality | Low | 3.17 | 1.38 |
| | High | 5.40 | 1.53 |
| Trustworthiness | Low | 3.21 | 1.58 |
| | High | 4.69 | 1.42 |
| Warmth | Low | 1.98 | 1.44 |
| | High | 3.90 | 1.35 |

# 7 Discussion

**Study 1: Modeling the Trait Impressions of Faces Using PCA**

The first study modeled the perception of traits in faces using PCA face recognition techniques. Although the results were marginally better than chance in the classification of faces according to the trait of dominance (.64), the PCA classifiers did a fair job matching average high and low human ratings of faces in the traits dimensions of adjustment (.71) and sociality (.70), and a good job matching user ratings of trustworthiness (.81), and warmth (.89).

At present, no hypothesis can be offered to account for the lower dominance classification score. A shortcoming in the first study was the design of the experiment assessing the stimulus faces. Had it been designed to provide a complete statistical analysis of user ratings, such an analysis might have provided some insight into the poorer performance of PCA recognition of high and low facial dominance.

As mentioned in section 5.2.2, extremely high scores were not expected. Unlike the task of classifying faces according to gender, age, and identity, matching human ratings of faces into trait categories is fuzzy. Although there is considerable evidence that people across cultures and age groups are consistent in their ratings of faces, people are not in total agreement. An attempt was made to produce a dataset of faces within each trait attribution class which demonstrated strong consensus ratings, but even so, consensus was not one hundred percent.

**Study 2: Face Synthesis Using PCA**

To the extent that agents increasingly have simulated faces, it is desirable to have the agents design those faces themselves rather than rely on human designers to do so. The second study explored the possibility of generating novel faces from within the attribution spaces of adjustment, dominance, sociality, trustworthiness, and warmth. It was predicted that faces synthesized by probing the low PCA attribution spaces would be rated at the lower end of the scale and that faces synthesized by probing the high PCA attribution spaces would be rated at the higher end of the scale.

Table 6: Average Ratings of Synthesized Faces and Stimulus Faces.

| Trait Dimension | Attribution Space | Average Rating of Stimulus Faces | Average Rating Synthesized Faces |
|---|---|---|---|
| Adjustment | Low | 2.65 | 3.86 |
| | High | 5.35 | 5.12 |
| Dominance | Low | 2.67 | 3.34 |
| | High | 5.52 | 5.00 |
| Sociality | Low | 2.64 | 3.17 |
| | High | 5.44 | 5.40 |
| Trustworthiness | Low | 2.60 | 3.21 |
| | High | 5.30 | 4.69 |
| Warmth | Low | 2.60 | 1.98 |
| | High | 5.34 | 3.90 |

The average trait ratings of the synthesized faces used to develop the various PCA trait attribution spaces are presented in Tables 5 and 6. Except for the trait dimension of adjustment, human subjects rated the synthesized faces as predicted. Table 6 also presents the average ratings of the stimulus faces used to train the PCAs for the ten attribution classes. From this table, the differences between the average ratings of the synthesized faces and the average ratings of the stimulus faces for each trait dimension can be calculated as 2.06 for warmth, 1.44 for adjustment, 1.19 for dominance, 1.22 for trustworthiness, and 0.57 for sociality. In particular, the faces synthesized from within the high and low sociality attribution spaces closely matched the average ratings of the stimulus faces used to train the PCAs. The largest difference was for warmth and adjustment. Table 7 shows the total average of high and low ratings for both the synthesized faces and the stimulus faces used in training the PCAs. The total difference between the average ratings of the stimulus faces and the average ratings of the synthesized faces in the ten attribution classes is 0.52, nearly half a point in the seven point scale. Clearly the synthesized faces elicited trait impressions that closely matched the trait ratings of stimulus faces used to train the PCAs.

Although the results of the second study indicate that it may be possible to generate faces with a high probability of eliciting specific impressions in users, much more work needs to be done in this area. This was an exploratory study into face synthesis within refined face spaces, and because the stimulus faces were highly processed and aligned, PCA synthesis was limited to recombining *shape-free* facial information within the PCA attribution spaces.

Table 7: Total Average of the Synthesized Faces and the Stimulus Faces

| Attribution Class | Total Average of Stimulus Faces | Total Average of Synthesized Faces |
|---|---|---|
| Low | 2.63 | 3.11 |
| High | 5.39 | 4.82 |

# 8 Conclusion

This paper reports a first attempt at developing a computational model of the trait impressions of the face for embodied agents that accommodates the social perception and social construction of faces. Two studies were presented. In the first study, a standard holistic face recognition technique based on PCA was used to match the human classification of faces at the bipolar rating extremes of the following trait dimensions: adjustment, dominance, warmth, sociality, and trustworthiness. Although results were marginally better than chance in the classification of faces according to the trait of dominance, PCA did a good job matching the average high and low human ratings of faces in the trait dimensions of adjustment, sociality, trustworthiness, and warmth. A second study explored the possibility of synthesizing faces intended to elicit particular trait impressions in observers. Using PCA models, 110 faces were synthesized and assessed by human subjects. The results were promising: the difference between the average ratings of the synthesized faces and the average ratings of the stimulus faces used to train the PCAs was found to be slightly over half a point in a rating scale of seven.

The research reported in this paper makes a number of contributions. It is the first research endeavor that not only suggests that embodied agents learn to design their own *socially intelligent embodiment*, or *smart embodiment*, but also indicates how this might be accomplished. This paper also presents the first computational model of the trait impressions of the face, and is further unique in using face recognition technology to classify social, or cultural, perceptions of faces rather than attributes of faces that are factual, such as identity and gender.

There are a number of directions that offer promising avenues for further exploration, some of which take into consideration limitations in the two studies presented in this paper. Particularly important, for both modeling the trait impressions of the face and for smart face synthesis, is the need to develop a database of faces that exhibit strong human consensus in a comprehensive set of trait categories. The creation of this database could be approached in several ways. Large collections of two-dimensional photographs and three-dimensional scans of actual faces could be evaluated, and those that produce marked attribution effects could be assembled into appropriate trait categories. Datasets of faces could also be generated artificially using either a variety of geometrical transformations, such as the cardioidal transform mentioned in section 2.2, or by simply piecing together facial features, either randomly, as was the case in this project, or with an eye towards eliciting specific trait attributions. These artificially generated faces would also need to be evaluated by human subjects.

Each of these approaches offers some attractive benefits. Two-dimensional photographs and facial composites have been widely studied in the person perception literature and present a simpler approach to modeling faces in terms of the traits they elicit than would be offered by three-dimensional scans. An advantage using facial composite programs, whether two-dimensional or three-dimensional, is that the contributions of individual facial features in the attribution process could more easily be investigated. Future studies might even investigate the possibility of designing embodied agents that learn to compose faces that are calculated to produce specific impressions in users by manipulating a relatively small set of facial features.

Each of these approaches is also problematical. A danger in using a dataset of faces that have been *artificially* produced is that models developed from these faces might oversimplify the problem too much and not model actual faces. These are criticisms that could also be leveled against many psychological studies that use artificially generated faces. It might be thought that using photographs of actual faces would solve these problems. However, photographs are two-dimensional representations, and it could be argued that people form impressions of faces based on multidimensional views of faces. Three-dimensional scans of actual faces also present representational dilemmas. How faces are seen in space for instance could affect viewer ratings. Will the viewer control how the scans are viewed or will the scan move on their own? Even judging films of faces is problematical as the perspective of the camera is typically artificial and stationary.

As stressed in section 3, a psychologically viable model of the trait attribution process is not essential for embodied agents; rather, the focus should be on selecting a dataset of faces that accommodates the particular tasks and the perception capabilities of the agents. Given the fact that faces, no matter how they are represented, are similar in appearance and, unless highly schematized, produce trait impressions in observers (Brunswik, 1947), it is likely the case that any fairly realistic representation of faces will model the *real faces* the agent will encounter as long as those faces are represented to the agent in the same fashion, e.g., as a set of pixels or geometrical shapes.

In addition to developing appropriate datasets to use in modeling the trait impressions of the face for embodied agents, future research will also need to explore additional face classification techniques. This study used PCA because it is capable of face synthesis as well as face classification. However, other face classification techniques have proven superior to PCA. Two face recognition techniques that should be explored in future studies are independent component analysis (Bartlett, 1998), a generalization of PCA that separates the higher-order moments of the input in addition to the second-order moments, and support vector machines (Vapnik, 1995), learning systems that separate a set of input patterns into two classes with an optimal separating hyperplane. Future studies might also explore classifying faces along a given trait dimension into three classes, i.e., a neutral category as well as the bipolar extremes.

As mentioned in the introduction, one of the most interesting possibilities a model of the trait impressions of the face offers embodied agents is the prospect of designing agents capable of creating an embodiment for themselves that is calculated to produce specific effects on users. Towards this aim, new face synthesis techniques from within these models need to be developed. In this study, trait spaces were probed using images of faces that were perceived to be at the opposite extreme of the trait dimension. Future studies might explore simply perturbing the average weight vector for each trait class. In addition, the synthesized faces in this study were reconstructions of eigenfaces, or *shape-free* image components, a process that introduced artifacts that may have made an impact on impression formation. Future studies in face synthesis will need to consider what Hancock (2000) calls *eigenshapes*, or vectors subjected to PCA that describe the outline of the face and its features. Studies also need to be conducted that appraise the degree of novelty that is exhibited by faces synthesized from within the face class spaces.

Finally, the value and practicality of embedding these models in embodied agents need to be evaluated.

## Acknowledgements

Special thanks go to Pierre Cote for granting permission to use the database of facial features found in FACES by Interquest and Micro-Intel.


## Appendix: Term Definitions and Behavioral Potential Questions

Table 8 below provides the definitions and behavioral potential questions (some of which were adapted from (Zebrowitz and Montepare, 1992; Berry and Brownlow, 1989)) that were available to subjects filling out the computerized questionnaires (see section 5).

Table 8: Definitions and Behavioral Potential Questions.

| Term Definition | Behavioral Potential Questions |
|---|---|
| **Adjusted, Unadjusted, Uncertain**<br>Here we are looking at how mentally healthy and adjusted the person is.<br><br>**Adjusted**<br>Is a person who is fairly happy, mentally healthy, and who feels s/he belongs to society.<br><br>**Unadjusted**<br>Is a person who is unhappy or discontent, possibly even mentally ill or troubled, and who feels like an outsider. | (None offered) |
| **Dominant, Submissive, Uncertain**<br>Here we are looking at how dominating the person is.<br><br>**Dominant**<br>Is person who is most likely to tell other people what to do.<br><br>**Submissive**<br>Is a person who usually follows orders, and is not very assertive | 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 arrangements, quiet hours, and house rules?” |

Table 9: Definitions and Behavioral Potential Questions (Continuted)

| Term Definition | Behavioral Potential Questions |
|---|---|
| **Trustworthy, Untrustworthy, Uncertain**<br><br>**Trustworthy**<br>Is a person who is mostly honest and who is not likely to steal, lie, or cheat.<br><br>**Untrustworthy**<br>Is a person who is often not honest and who possibly steals, lies, and cheats | 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?" |
| **Social, Unsocial, Uncertain**<br>Here we are looking for how social the person is.<br><br>**Social**<br>Is person who is very outgoing, extroverted, and who enjoys parties and other social activities.<br><br>**Unsocial**<br>Is a person who introverted, a loner, and who would prefer to stay home rather than go out. | A helpful question might be: "Does s/he look like someone who would attend a school dance or party?" |
| **Warm, Cold, Uncertain**<br>Here we are looking for how approachable the person is. | A helpful question might be: "Does s/he look like someone who would turn a cold shoulder to your attempts at friendly conversation?" |

# References

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), 558–69.

Alcock, D., Solano, J., and Kayson, W. A. (1998). How individuals' responses and attractiveness influence aggression. *Psychological Reports*, **82** (June 3/2), 1435–38.

Alley, T. R. (1983). Infantile head shape as an elicitor of adult protection. *Merrill-Palmer Quarterly*, **29**, 411–27.

Alley, T. R. (1988). Social and applied aspects of face perception. In Alley, T. R., editor, *Social and applied aspects of perceiving faces*, pages 1–10. Lawrence Erlbaum Associates, Publishers, Hillsdale, NJ

Bartlett, M. S. (1998). *Face image analysis by unsupervised learning and redundancy reduction*. Ph.d. dissertation, University of California.

Berry, D. S. and Brownlow, S. (1989). Were the physiognomists right? *Personality and Social Psychology Bulletin*, **15**(2), 266–79.

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.

Beymer, D., Shashua, A., and Poggio, T. (1993). Example based image analysis and synthesis. Technical Report A. I. Memo 1431, MIT.

Bruce, V. (1988). *Recognising faces*. Lawrence Erlbaum Associates, Publishers, Hove and London.

Brunelli, R. and Mich, O. (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), 1042–52.

Brunswik, E. (1947) *Perception and the representative design of psychological experiments* University of California Press, Berkeley and Los Angeles.

Bull, R. and Rumsey, N. (1988). *The social psychology of facial appearance*. Springer-Verlag, New York.

Burton, A. M., Bruce, V., and Dench, N. (1993). What's the difference between men and women? evidence from facial measurement. *Perception*, **22**(2), 153–76.

Converse, J. M., Presser, S. (1986). *Survey Questions: Handcrafting the standardized questionnaire* Sage Publications, London.

Carello, C., Grosofsky, 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), 227–51.

Cottrell, G. W. and Fleming, M. K. (1990). Face recognition using unsupervised feature extraction. In *International Conference on Neural Networks*, pages 322–25. Kluwer Academic Publishers, Dordrecht.

Cottrell, G. W. and Metcalfe, J. (1991). Empath: Face, emotion, and gender recognition using holons. In Touretzky, D., editor, *Advances in neural information processing systems*, pages 564–71. Morgan and Kaufman, San Mateo, CA.

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

Donath, J. (2001). Mediated faces. In Beynon, M., Nehaniv, C., and Dautenhahn, K., editors, *Cognitive Technology: Instruments of Mind: Proceedings of the 4th International Conference*, *Lecture notes in artificial intelligence*, page 373, Springer, Warwick, United Kingdom.

Donato, G., Bartlett, M. S., Hager, J. C., Ekman, P., and Seinowski, T. J. (1999). Classifying facial action. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **21**(10), 974–89.

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), 109–28.

Ekman, P. and Friesen, W. V. (1978). *Facial action coding system*. Consulting Psychologists Press Inc., Palo Alto, CA.

Enlow, D. H. and Hans, M. G. (1996). *Essentials of facial growth*. W. B. Saunders Company, Philadelphia.

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), 757–63.

Feingold, A. (1992). Good-looking people are not what we think. *Psychological Bulletin*, **111**(2), 304–41.

Freierman, S. (2000, Feburary 17). Constructing a real-life mr. potato head. Faces: The ultimate composite picture, *The New York Times*, page 6.

Friedman, H. H. and Amoo, T. (1999). Rating the rating scales. *Journal of Marketing Management*, **9**, 114-23.

Frey, S. (1993). Lavater, lichtenberg, and the suggestive power of the human face. In Shookman, E., editor, *The faces of physiognomy: Interdisciplinary approaches to johann caspar lavater*, *Studies in German literature, linguistics, and Culture*, pages 64–103. Camden House, Columbia, SC.

Galton, F. (1878). Composite portraits. *Nature*, **18**, 97–100.

Hancock, P. (2000). Evolving faces from principal components. *Behavior Research Methods, Instruments and Computers*, **32**(2), 327–33.

Hochberg, J. and Galper, R. E. (1974). Attribution of intention as a function of physiognomy. *Memory and Cognition*, **2**(1A), 39–42.

Jones, D. (1995). Sexual selection, physical attractiveness, and facial neoteny. *Current Anthropology*, **36**(5), 723–48.

Katz, D. (1952). Le portrait composite et la typologie. *Ikon (Revue Internationale de Filmology)*, **2**, 207–14.

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), 615–26.

Kohonen, T. (1977). *Associative memory: A system theoretic approach*. Springler-Verlag, Berlin.

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.

Landau, T. (1989). *About faces: The evolution of the human face*. Anchor Books, New York.

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**, 115–21.

Langlois, J. H., Roggman, L. A., Musselman, L., and Acton, S. (1991). A picture is worth a thousand words: Reply to 'on the difficulty of averaging faces'. *Psychological Science*, **2**(5), 354–57.

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), 153–59.

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), 743–56.

Liggett, J. (1974). *The human face*. Stein and Day, New York.

Lucker, G. W. and Graber, L. W. (1980). Physiognomic features and facial appearance judgements in children. *The Journal of Psychology*, **104**(Second Half), 261–68.

Magli, P. (1989). The face and the soul. In Feher, M., Naddaf, R., and Tazi, N., editors, *Fragments for a history of the human body*, volume 2, pages 86–127. Zone, New York.

Magro, A. M. (1997). Why barbie is perceived as beautiful. *Perceptual and Motor Skills*, **85**(1), 363–74.

Massaro, D. M. (1997). *Perceiving talking faces: From speech perception to a behavioral principle*. MIT Press, Cambridge, MA.

McArthur, L. Z. and Baron, R. M. (1983). Toward an ecological theory of social perception. *Psychology Review*, **90**(3), 215–38.

Mueller, U. and Mazur, A. (1996). Facial dominance of west point cadets as a predictor of later military rank. *Social Forces*, **74**(3), 823–28.

Mulford, M., Orbell, J., Shatto, C., and Stockard, J. (1998). Physical attractiveness, opportunity, and success in everyday exchange. *American Journal of Sociology*, **103**(6), 1565–92.

Oja, E. (1992). Principal components, minor components and linear neural networks. *Neural Networks*, **5**, 927–35.

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**, 146–60.

Padgett, C. and Cottrell, G. W. (1998). A simple neural network models categorical perception of facial expressions. In *Proceedings of the 20th Annual Cognitive Science Conference*, Lawrence Erlbaum Associates, Publishers, Madison, WI.

Pelachaud, C. and Poggi, I. (2002). Subtleties of facial expressions in embodied agents. *Journal of Visualization and Computer Animation*, **13**(5), 301–12.

Picard, R. W. (1997). *Affective computing*,1st edition. The MIT Press, Cambridge, MA.

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.

Rhodes, G. and Tremewan, T. (1996). Averageness exaggeration and facial attractiveness. *Psychological Science*, **7**(2), 105–10.

Rosenberg, S. (1977). New approaches to the analysis of personal constructs in person perception. In Land, A. L. and Cole, J. K., editors, *Nebraska symposium on motivation*, volume 24, pages 179–242. University of Nebraska Press, Lincoln.

Rubenstein, A. J., Kalakanis, L., and Langlois, J. H. (1999). Infant preferences for attractive faces: A cognitive explanation. *Developmental Psychology*, **35**(3), 848–55.

Secord, P. F., Dukes, W. F., and Bevan, W. W. (1954). An experiment in social perceiving. *Genetic Psychology Monographs*, **49**(Second Half), 231–79.

Sirovich, L. and Kirby, M. (1987). Low dimensional procedure for the characterization of human faces. *Journal of Optical Society of America*, **4**(3), 519–24.

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.

Symons, D. (1979). *The evolution of human sexuality*. Oxford University Press, New York.

The MathWorks. (2000). *Using MATLAB: The language of technical computing*. The Mathworks, Inc., Natick, MA.

Todd, J. T. and Mark, L. S. (1980). The perception of human growth. *Scientific American*, **242**(2), 132–44.

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.

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. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 586–91, IEEE Computer Society Press, Silver Spring, MD.

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. (1994a). 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), 413–29.

Valentin, D., Abdi, H., O'Toole, A. J., and Cottrell, G. W. (1994b). Connectionist models of face processing: A survey. *Pattern Recognition*, **27**(9), 1209–30.

Vapnik, V. (1995). *The Nature of Statistical Learning Theory*. Springer-Verlag, New York.

Vetter, T. and Poggio, T. (1997). Linear object classes and image synthesis from a single example image. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, **19**, 733–42.

Waters, K. and Terzopoulos, D. (1992). The computer synthesis of expressive faces. *Philosophical Transactions of the Royal Society of London*, **B335**, 87–93.

Wegener-Knudsen, M., Martin, J.-C., Dybkjær, L., Berman, S., Bernsen, N. O., Choukri, K., Heid, U., Kita, S., Mapelli, V., Pelachaud, C., Poggi, I., van Elswijk, G., and Wittenburg, P. (2002). Survey of nimm data resources, current and future user profiles, markets and user needs for nimm resources. sle natural interactivity and multimodality. Technical Report Deliverable D8.1, http://isle.nis.sdu.dk/reports/wp8/.

Yamaguchi, M. K., Hirukawa, T., and Kanazawa, S. (1995). Judgment of gender through facial parts. *Perception*, **24**(5), 563–75.

Zebrowitz, L. A. (1998). *Reading faces: Window to the soul?* Westview Press, Boulder, CO.

Zebrowitz, L. A. and Montepare, J. M. (1992). Impressions of babyfaced individuals across the life span. *Developmental Psychology*, **28**(6), 1143–52.

Zebrowitz, L. A., Montepare, J. M., and Lee, H. K. (1993). They don't all look alike: Individuated impressions of other racial groups. *Journal of Personality and Social Psychology*, **65**(1), 85–101.