Distance Based Visual Cues to Interpersonal Trust

Ronnie Bryan

Computation and Neural Systems Program

California Institute of Technology

Abstract

This thesis examines the role of interpersonal spacing in determining the visual appearance and emotional response to images of faces. We present new methods for isolating the distance-dependent perspective projection as a visual feature, while controlling for confounding variables such as emotional expression. In behavioral experiments, we demonstrate the relevance of viewing distance to implicit social judgments, notably trust behavior in which real money was at stake. Finally, we provide tools for classifying face images according to viewing distance, and manipulating face images to simulate their appearance at different distances and different levels of trustworthiness.

Table of Contents

Chapter 1: Background and Introduction

The visual world is full of human faces: some you know, but most you don't. How do you decide which faces to pursue further interaction with, and which to avoid? Without detailed information about every person's past behavior, it is nearly impossible to sift through the many faces in a crowd -- yet we do. The growing literature about how we process faces has evolved from asking whether or not faces are “special” (Diamond & Carey, 1986) to where in the brain face processing occurs (Kanwisher, McDermott, Chun, 1997; Haxby, Gobbini, Furey Ishai, Schouten, and Pietrini, 2001; Gauthier, Skudlarski, Gore, & Anderson, 2000), to how face processing influences social inferences and interaction (Todorov, Mandisodza, Goren, and Hall, 2005; Duarte, 2009). This thesis focuses on this most recent evolution, linking the visual processing of face information to the biologically primitive regulation of approach and avoidance behavior.

1.1 Rapid Visual Inference From Faces

The emerging research indicates that rapid and spontaneous visual computations may have a profound influence on the behavioral choices we make regarding other people. It has been found that very brief exposures (39ms) to face images results in consistent evaluations of threat (Bar, Neta, and Linz, 2006). Similarly brief (100ms) evaluations of face images results in consistent evaluations of trustworthiness (Willis & Todorov, 2006). Do rapid evaluations of this kind influence real world decision-making?

Rapid attributions of competence from naïve viewers correlated with the outcome of Senate races in the United States (Todorov et al, 2005). Viewers with no previous knowledge of the candidates in various US Senate races were shown brief (1000ms) exposures to images of the candidates' faces. These viewers then were asked to evaluate their first impressions of these faces along a variety of socially relevant dimensions, such as competence, age, attractiveness, and familiarity. It was found that perceived competence, but not the other traits, was a reliable indicator of the election outcome.

More recent work has replicated the finding that perceived competence is predictive of vote garnering, and has also demonstrated that such effects are sensitive to the conditions of the race, such as whether or not the candidate is the incumbent or the challenger, and whether the voter is partisan or independent (Atkinson, Enos, and Hill, 2009). Another study found that even the competence judgments of children could be used to predict election results, and that such judgments were indistinguishable from those made by adults (Antonakis & Dalgas, 2009). In a brain scanning experiment, it was shown that faces of the losing candidate, but not the winning candidate, caused reliable activations in brain regions such as the insula (Spezio, Rangel, Alvarez, O'Doherty, Mattes, Todorov, Kim, and Adolphs, 2008; however, cf. Rule, Freeman, Moran, Gabrieli, Adams, and Ambady, 2010). This result implies the use of negative information from first impressions may be more influential than positive information, a result that has been replicated in a behavioral task (Mattes, Spezio, Kim, Todorov, Adolphs, and Alvarez, 2010). This thesis supports the idea that negative information can exert a significant influence on interpersonal attributions.

Similar attributions of trustworthiness correlated with interest rates received by small businesspeople in an online network (Duarte 2009). Viewers with no previous experience with the small business owners listed on an online lending site (prosper.com) were asked to evaluate images of these loan applicants on dimensions such as perceived trustworthiness, attractiveness, wealth, and ethnicity. It was found that perceived trustworthiness, but not the other traits, was a reliable indicator of whether the loan request would be fulfilled and how favorable the interest rate would be.

Are there physical cues in the appearance of a face that contain reliable information about a person's intentions and attitudes? This question has been dealt with throughout history in various forms (including the proto-sciences physiognomy, phrenology and craniometry) and has remained controversial throughout. It seems uncontroversial to say, at a minimum, that there is a real human tendency to make quick judgments about other people based simply on their visual appearance, and that it is important to uncover the basis of these judgments.

1.2 Facial Inference Theories

A useful division of the theoretical work so far on face attributions is to focus on the morphable face features independently of the static face features. The changeable face features relate primarily to systems subserving the representation of expression, while the static face features relate primarily to structural properties of the face shape (Oosterhof & Todorov, 2008). Trustworthiness inferences correspond to the perception of harmful intentions and dominance inferences correspond to perception of the capacity to carry them out. In the computerized implementation of this model, morphing a face along the dimension of trustworthiness resulted in the percept of a change in the valence of the expression, and morphing a face along the dimension of dominance resulted in the percept of a change in the babyfacedness and masculinity of the face.

The perception of morphable face features, such as the curvature of the mouth or the openness of the eyes, may interact with the perception of trustworthiness (Oosterhof & Todorov, 2009). Participants who viewed the animation of a face image to wear either a happy or an angry expression were more likely to rate trustworthy looking faces as more happy and untrustworthy looking faces as more angry. The appearance of these features is highly dependent on the muscle tone of the underlying tissue, and therefore has a direct connection a person's intentional state at a particular moment in time. Even if the face muscles are all fully relaxed, a naive observer may perceive a slightly downturned mouth to be an indication of unhappiness or bad intentions. In another study, it was shown that the output of a computer vision system trained to classify facial expression could then be used to predict the personality traits of even neutral faces (Said, Sebe, and Todorov, 2009). This result demonstrates a computational basis for believing that a system designed to evaluate expression may also underlie perceptions of personality traits.

The perception of structural face features, such as babyfacedness and masculinity, also interact with the evaluation of trustworthiness. Adult faces that retain the proportions and shape of youthful faces are seen as more trustworthy (Zebrowitz, 1997). The overgeneralization of traits associated with infants, such as harmlessness, to individuals with facial proportions resembling infants (large eyes, round cheeks, high forehead, and thin eyebrows) holds across a variety of age ranges and is independent of perceived attractiveness (Zebrowitz & Montpare, 1992).

Faces that exhibit a high degree of masculinity, as it influences the face's width-to-height ratio, receive lower investments in an economic trust game and are less likely to return those investments (Stirrat & Perrett, 2010). The width to height ratio of a face relates to prenatal testosterone and has also been linked with a variety of physical and behavioral traits, such as finger digit ratios and sexual orientation (Williams, Pepitone, Christensen, Cooke, Huberman, Breedlove, Breedlove, Jordan, & Breedlove, 2000). Interestingly, it is also reported that CEOs whose faces have a greater width-to-height ratio also achieve superior financial performance (Wong, Ormiston, and Haselhuhn, 2011).

Finally, it should be remembered that average faces are commonly reported as being more attractive (Zebrowitz & Montepare, 2008). Faces that are atypical, by contrast, are more likely to activate the amygdala, a brain region associated with threat perception (Said, Dotsch, and Todorov, 2010).

All of these avenues of study shed light and open up inquiry on factors that may influence people's selection between two different faces; for example, they imply that you are more likely to trust a less masculine familiar face with youthful features and a more than average upturned mouth. These inferences are rooted in physiologically plausible behavioral mechanisms, such as hormones and expressions.

1.3 Interpersonal Distances

In this thesis, I will present research on a third avenue that holds constant both the structural and morphable face features by examining the role of different perspectives on the same face. In contrast to other work, our approach is essentially relational, taking into account information about the spacing between the stimulus and the observer, and not just information about the stimulus itself. The physical spacing between people is a key indicator of the relationship between them, and, as I will explain, importantly influences the appearance of the face as well.

It is easy to see the visual effect of the vertical spacing between people: just adjust the tilt of your head up or down as you look in the mirror. In the absence of head tilt, however, the relative height of an observer to a face will determine the face's appearance in a similar manner. Height is an important variable in assessing threat and physical dominance, and a growing literature points to ways in which even modern society is sensitive to it (Judge, & Cable, 2004; Pawlowski, Dunbart, & Lipowicz, 2000). For an example, see Chapter 3, section 3, figure 1.

What is perhaps less easy to see is that the horizontal spacing between people also influences facial appearance. Obviously the greater the distance that separates an observer from the face, the smaller the face will appear. Even controlling for the change in size, there are subtle, but noticeable, differences in facial appearance. The changes due to viewing distance arise from the fact that the three-dimensional structure of the head will cast a different 2-D perspective projection at different viewing distances. The parts of the face that are closest to the observer, such as the nose, appear relatively larger when viewed at a closer distance as compared to the parts of the face that are farthest from the observer, such as the ears. The sides of the face also become less visible at closer distances, giving faces the appearance of a lesser width. For an example, please see Chapter 3, section 2, figure 1.

Changes along both axes of interpersonal distance, vertical and horizontal, result in differences in the appearance of the face image. It is possible that these changes will be correlated with the existing theoretical mechanisms described above. As I will describe in Chapter 2, the expressions worn by the faces will be held constant in these experiments, and all comparisons will be made between different images of the same individual. This leaves little explanatory room for the aforementioned theoretical mechanism related to facial expression. In the case of horizontal distance, we noted that the face appears narrower as the distance is reduced, giving it a reduced width-to-height ratio. If width-to-height ratio is the predominant theory, this implies that closer faces should be seen as less masculine, and therefore more trustworthy. As we will see in Chapter 3, this is not the case, suggesting that an alternative theoretical mechanism may be more explanatory. In the case of vertical distance, the face appears narrowest at a direct viewing angle, resulting in an increased width-to-height ratio for shorter and taller faces. As we will see in Chapter 3, the direct viewing angle is seen as the most trustworthy, opening the door for a contributory role for width-to-height ratio to play in explaining vertically mediated changes in trust behavior.

I will present a series of experiments demonstrating that people are sensitive to physical spacing parameters when they make quick judgments about other people, although they are almost always unaware of these changes. I will also present some methods for automated computer algorithms to estimate viewing distance, and therefore estimate the behavioral responses associated with different viewing distances and different faces. I will show that it's possible to manipulate facial appearance in order to change the implied viewing distance, and thus the associated behaviors. I will also show that the anatomical features used to estimate camera distance can also be used to estimate behavior in a trust game, and can be adjusted to manipulate related trait inferences.

1.4 Face Vision

It's obvious why cognitive sciences have discussed the "special" status of faces at length (Diamond & Carey, 1986). Our exponentially growing population inundates the senses with social information, most prevalently images of other people. That our visual cortex is so dominant in surface area mirrors the experiential saliency of the presence of other people in our environment. Humans are able to extract an enormous amount of information from an instant of visual exposure to the world, and a significant proportion of the images we are likely to encounter contain faces. Faces are extremely important predictors of future events, and contain information critical to our survival and wellbeing.

The branch of perceptual science that has come to be known as high-level vision concerns itself with processing of important classes of objects that project 2-D images in a complex and context-sensitive way; that is to say, almost all objects. Most of the time, the term high-level vision refers to object centered coordinates, as proposed by David Marr (Marr D., 1982). This is the sense that I use the phrase in the thesis. Object-centered coordinates are considered "high-level" because they rest on more basic information about the location of the object in the image. As a face moves across the visual field, the positions of the facial features may remain constant relative to each other, but they move relative to the surround. In other words, the perceived identity of the face doesn't change as it changes positions, although the neurons in the first layers of processing them are constantly changing.

High-level vision has merged with computational vision most notably through the work of Marr's collaborator Thomas Poggio, who has provided a detailed and biologically inspired model of how successive layers of a neural network can give rise to position invariant representations (Serre, Wolf, Bileschi, Riesenhuber & Poggio, 2007). The model has become a standard because it can be trained to mirror the behavior of neurons in ventral temporal cortex that neurophysiologists have shown to be responsive to high-level features, like those forming faces (Gross & Schonen, 1992). The model is limited in important ways, such as the lack of feedback connections or reciprocating connections with the amygdala (Adolphs, 2004), but it provides a tractable framework for thinking about high-level features.

Face processing has been an important and active area within high-level vision. There are many open controversies on the representational scheme used by the mind to store and process face information, as well as questions about the anatomical substrate in the brain. We will focus on the open questions about mental representations and leave aside, at least for now, how these questions may in the future bear on issues of anatomical specialization.

The dominant cognitive model of face representation proposes that faces occupy points in a high dimensional metric space with the average face located at the origin (Valentine, 1991). Many times assumptions about the viewing conditions of face images determine an interesting part the result. For example, faces must be registered to a common frontal view template under common lighting conditions. To these assumptions we will add the necessity that face images must also be captured at the same photographic distance, since the 2D projection depends critically on this value. It is the suggestion of this thesis that our visual analysis of facial features contains an implicit theory of the three dimensional structure of the face and the viewing distance.

A face can provide information about distant and unseen objects, but it can also provide direct evidence about the future behavior of the most complex stimuli in our sensory environment, other people. The informative content and immediate relevance of a particular facial image become magnified the closer the person stands who reflects it. Perhaps the most intimate and consequential human exchanges have taken place at whisper distances.

1.5 Proxemics

The scientific study of interpersonal distance, called proxemics by the anthropologist Edward Hall, developed in parallel to the development of cognitive science in the 1960s. Hall proposed a quantitative theory of interpersonal attraction roughly analogous to gravitational attraction. He also specified a specific set of threshold distances that delineate qualitatively different kinds of social interaction. The distances chosen to test in this thesis work is motivated by Hall's qualitative threshold separating ‘social space’ from ‘personal space.’ We build on his energetic formulation with notions of high-level visual representations of the three-dimensional shape of the human face.

Verhoff (2008) showed that human perception depends on viewing distance, and suggested that computer face recognition algorithms could benefit from explicitly modeling perspective projection. It is one goal of this thesis to provide support for the notion that human vision is sensitive to this cue, even in the absence of overt knowledge.

Perception of interpersonal distance influences social behaviors (Hayduk, 1983). Notably, interpersonal distance is associated with arousal (Patterson, 1976). Patterson proposed a context dependent model of physiological arousal as a response to interpersonal intimacy, which was defined to include interpersonal distance, gaze direction, posture, expression, and verbal information. Patterson offered the arousal model primarily as a modification of Equilibrium Theory (Argyle & Dean, 1965), which failed to explain participants' occasional increases in affiliative behavior in response to greater interpersonal intimacy displayed by a confederate. Patterson describes a number of changes associated with decreases in interpersonal distance, including galvanic skin response, negative affect, and fidgeting. According to Patterson's model, changes in intimacy that result in negative emotion lead to compensation reactions, while those that lead to positive emotion lead to reciprocity. An example of a change in intimacy that results in positive emotion, and therefore reciprocity, is an increase in eye contact, which reportedly leads people to display more positive facial expressions.

In order to apply Patterson's model to the present experiments, it is important to distinguish the contexts in which a decrease in personal space is likely to result in a negative emotion. The example given by Patterson himself (page 240) is a personal space violation of a stranger, which is predicted to result in distancing behavior. In a negative context such as this, Patterson's model agrees with Argyle and Dean's Equilibrium Theory. A recent corroboration of this prediction is provided by analysis of stress reactions of mass transit passengers during different crowding conditions: passengers who experienced decreases in interpersonal distance secreted more cortisol, reported greater stress, and afterwards exhibited poorer performance on a cognitive task (Evans & Wener, 2007).

One proposed function of such distancing is self-protection, which finds support in the distancing behavior displayed by abused children (Vranic, 2003). Vranic employed a stop distance methodology in which a confederate slowly approached the participant from one of four directions until told to stop. Children who were the victims of physical abuse displayed significantly greater stopping distances than those who were not, especially when the confederate was male. These results may be interpreted within a framework that treats personal space as a self-protection mechanism (Dosey & Meisel, 1969). In this experiment, a similar stop distance paradigm revealed that participants in a high-stress condition were more likely to display a need for greater personal space. It has also been shown that music rated as having a negative affect causes participants to exhibit greater stopping distances in this task (Tajadura-Jiméne, Pantelidou, Rebacz, Västfjäll, Tsakiris, 2011). Another, related, proposed function of personal space is to manage stress and aggression (Evans & Howard, 1973).

1.5.1 Proxemics and Trust

Finally, it should be noted that there is an observed relationship between interpersonal distance and interpersonal trust (Jourard & Friedman, 1970). In this study, the investigators actually study the disclosure of personal information, as measured by the length of recorded vocal responses to personal questions. Self-disclosure is found to be closely related to interpersonal trust (Wheeless & Grotz, 1977). Three groups of distances were analyzed: with the experimenter out of the room, with the experimenter in the room without making eye contact, and with the experimenter in the room making eye contact. The role of eye contact in this study is to increase interpersonal immediacy, which is a more general construct that includes interpersonal distance and other associated behaviors. As the level of immediacy increased, the investigators found that the female participants (N=8), but not the males (N=8), showed a significant reduction in the length of their vocal responses. This finding shows a first hint that the stress and the self-protective responses associated with decreases in interpersonal distance may manifest as a decrease in interpersonal trust.

The results obtained by Jourard & Friedman suggest a positive relationship between distance and trust in the context of self-disclosure to an unknown individual. They further investigate this relationship by attempting to manipulate the familiarity with the investigator, but our interest lies primarily in the relationship between distance and trust in the original context of a lack of familiarity.

This thread of investigation was continued by Johnson & Dabbs (1976), who more systematically controlled the distance between the participant and the experimenter. In their experiment, subjects were divided into three groups, each of which sat at a different distance to the experimenter: 18" (close), 36" (middle), or 54" (far). The paradigm used was very similar to Jourard & Friedman, with the main dependent variables being disclosure time, measured in seconds. The results of this study showed a strong decrease in disclosure time in the group assigned to sit close to the experimenter (248.3s), but no significant difference between the middle and the far group (381.6s and 378.6s, respectively). This basic result has been replicated by numerous other groups, establishing a clear relationship between self-disclosure and experimenter distance (Hansen & Schuldt, 1982), though in some cases it is found to interact with gender (Skotko & Langmeyer, 1977).

1.5.2 Distance Cues

In all of these studies of proxemics, distance is manipulated by actually physically placing the experimenter at different intervals from the subject. Naturally, there are many sensory factors that correlate with interpersonal distance. With the methodologies used in these studies, it is impossible to weigh the relative contribution of any single factor, or to know with certainty if a factor is necessary or sufficient to produce the measured effect.

The most obvious cue relating to interpersonal distance is the field of view of the percept. As the experimenter gets farther away, a greater extent of their body is visible. Similarly, as the experimenter approaches, their face and body become larger and they can be seen at a higher resolution. Unless the experimenter is highly trained, it is also possible that they will wear slightly different facial expressions, maintain different levels of eye contact, or otherwise dynamically interact and respond to feedback from the subject. Other sensory cues may contribute as well. They are more audible at closer distances, which may affect the pitch and tone of their voice as they attempt to communicate. They may be easier to smell, which may affect the valence of the interaction in an unpredictable way. All of these variables point to the necessity for a more controlled paradigm. In addition, many of these variables are easy to consciously perceive. When the size of the face image varies, we are aware of the change, but other features of the face remain constant, such as the identity and expression. This invariance to size is a hallmark of high-level features (similar to invariance to position) and provides us with a robust representation in a highly variable environment. In face-centered coordinates, the size of a face image is not relevant to computing its identity or expression because the image has been registered to a common template. The size of the image is relevant for computing the distance, but if that size cue is removed (as it often is in modern media representations), then the role of the remaining cues may be more significant.

The studies described in this thesis will not investigate any of the manifold cues to interpersonal distance that may be imagined, but instead will focus on the distance dependent warping due to perspective projection. Compared to the multiple cues at play in previous research, this cue is completely isolated, so we can be confident that it is causally responsible for any effects we observe. Moreover, as the cue is subtler and more difficult to distinguish, it also opens the door to the possibility that it may be processed implicitly, without conscious awareness. Although people are often not aware of the differences in visible facial proportions as they vary with distance, it is still possible that they may influence behavior. Such an automatic sensitivity to perspective projection may perhaps be an adaptive response against exploitation by people too eager to enter our personal space.

Given these many emotionally relevant correlates of interpersonal distance, we predict that trust will be influenced by the visual percept of personal space violation, even if that visual percept contains only a subtle, possibly implicit cue. In order to demonstrate this, our study will require more controlled testing conditions, many more participants, and an updated behavioral testing paradigm borrowed from economics.

1.6 Economics

Behavioral economics emerged after key studies from Kahneman and Tversky demonstrated deviations in human behavior from the rational model of classical economics. Since then, psychological and economic sciences have been merging into a compromise discipline that attempts to describe the nuanced ways in which people assign value to objects in their environment. Face perception is relevant to this enterprise since economic activity may be affected by the inferences that people make about each other based solely on information contained in a person's appearance. Often this is as seen as evidence for the influence of a separate emotional system that operates in parallel to and competes with the rational system. I prefer to think of emotional processing as embodying a rational system that operates over a different set of factors than classical economic theory takes into account. This thesis examines the role of one such factor, interpersonal distance. We ask how might interpersonal distance influence downstream psychological processes related to decisions about economic activity.

Trust is one very important mediating variable in the psychology of economics. Trust is the foundation of cooperative activity and what allows us to lower our defenses in the face of the temptation to exploit one another. Trust has been operationalized in economics (Berg, Dickhaut & McCabe, 1995; Camerer, 2003; Houser, 2006) to demonstrate a wide variety of influences upon it, including visual cues from faces (Duarte 2009). A typical economic trust game goes something like this: imagine that you have an opportunity to triple an investment of up to $100. The only catch is that once your investment becomes tripled, it is up to another person to decide how much will be returned to you, and how much will be kept by them. As you might suspect, not all people will be equally trustworthy in terms of returning your fair share of the investment. Maximizing your return in this game requires a leap of trust in another person's willingness act fairly. Minimizing your loss by choosing not to trust your partner comes at the expense of the opportunity to triple your money. This game is a schematized and simplified version of economic exchange in the real world, but it captures critical dynamics about trust and cooperation. Can we predict people's behavior in this game based solely on visual cues from facial appearance? Is sensitivity to three dimensional vision cues about interpersonal distance one such cue?

According to Hall's proxemic thresholds, there is a critical cut off between what is considered social distance (greater than 4') and personal distance (1.5' to 4'). Violations of personal space often cause discomfort, distress, and sympathetic arousal (Patterson, 1976). It stands to reason that in the face of such discomfort, sensitivity to risk and/or loss may be heightened. Such sensitivity to loss would be rational given that violation of personal space may be a prelude to exploitation. In the general theory of proxemics as applied to animals, critical distances are defined in terms of risk of attack, based on the speed of a typical predator (attack zone) and the visual awareness of the animal (flight zone). Most ethologists believe these survival instincts play a diminished role in the human brain, but it may be the case that we have a heightened sensitivity to the social and personal spaces to accompany the increased neural resources dedicated to processing social information. It may be that the greatest risks we take now involve interactions at the handshake distance.

1.7 Face Trait Inferences of Interest

The rationale for studying how camera distance may influence attributions of trust is based in the previously surveyed literatures of personal space and behavioral economics. Personal space violations are thought to result in negative affect and feelings of perceived threat, which we hypothesize negatively influences the socially and economically important attribution of trustworthiness.

Attributions of competence are included as a complementary personality dimension due to its relevance to voting behavior (eg. Todorov 2005). As the second factor in the two-factor model described in Todorov (2008), we felt it was important to include ratings of competence, although prior studies often included only one or the other. Additionally, there seems to be a theoretically connection between competence and distance in the sense that highly competent individuals such as leaders exhibit greater personal space, and people with less social competence may inadvertently violate personal space and make others feel uncomfortable. Therefore, if there is an effect on competence, we hypothesize that faces viewed from a greater distance will be seen as more competent.

Attractiveness is included a trait grounded in physical appearance, so may be more easily accessible for participants to give quick impressions. There is an enormous literature surveying the determinants of attractiveness, but our interest is primarily in observing an additional measure of valenced trait inference.

Heaviness is included based on the observation that the farther photographed faces reveal more of face surface area, and therefore appear wider. This is a trait grounded in physical appearance, so it serves as a test of the subjects' ability to perceptually distinguish the close and far faces and correctly interpret the instructions of the task.

Age is included as a control question to which we expect no influence of camera distance, since the images are equated for resolution. We include this variable in our battery as a comparison to show that not every trait is influenced by perspective projection.

Averageness is included as a control question to determine if the distortions introduced by perspective projection primarily influence judgments by making the faces appear atypical. There is a large literature on the role of typicality and averageness in face processing generally, and specifically as it relates to perception of attractiveness and trust. We expect that faces photographed closer will appear less average, but we also expect that this perception will not entirely explain the results of our experiment.

We also ask about camera distance to test if participants are able to tell which images are close and which are far. The results of this question will address the issue of whether or not these attributions are being made consciously or unconsciously.

In addition to these traits, we will also control the conscious awareness of interpersonal distance with explicit cues, such as verbal information or size. The motivation behind this manipulation is to determine if even the mere suggestion of interpersonal closeness, without the accompanying change in perspective projection, is sufficient to influence attributions of trust.

1.8 Computer Vision

In parallel to the increased attention faces have received within the psychological community, the computer sciences have produced a variety of robust face processing algorithms. A metric space representation of faces similar to Valentine (1991) has been implemented using principal components analysis (Sirovich & Kirby, 1987; Turk & Pentland, 1991). Principal components analysis recasts high dimensional data into a lower dimensional space by projecting that data onto the linear combination of dimensions guaranteed to capture the most variance. This algorithm is called Eigenfaces due to its use of the eigenvectors of the covariance matrix. The technique cleverly saves computational power by analyzing the smaller of the two possible square matrices that can be produced from multiplication of the input rectangular matrix and its transpose. The technique was an important advance, but the calculation of eigenvectors is an unnecessary computational expense that can be sidestepped using the more sophisticated technique of singular value decomposition.

Size-invariant face detection has been implemented using a boosting algorithm that cascades over a novel image representation (Viola & Jones, 2004). This technique opened the door for very fast face detection, a problem of central importance to registering image datasets into the rigid template necessitated by techniques that operate over face-centered coordinates. Any automated method that wishes to be robust to changes in the location of the face within an image will benefit from having a front-end face detector such as that produced by Viola & Jones. The method employs an image basis set similar to wavelets, but restricts its attention to regions most likely to contain a face after a fast initial scan of the image. An iterative multi-scale cascade results in very good performance, though, like Eigenfaces, is somewhat limited in its robustness to viewing angle.

Recognition invariance to changes in viewing angle has been achieved using elastic graph models (Wiskott, Fellous, Kruger, and von der Malsburg, 1997). These models use local image information in the form of Gabor patches to learn the location of important anatomical keypoints on the face. This technique has the appeal of representing the image in terms of a configuration of anatomical locations, which comports nicely to notions of holistic processing popular in the cognitive literature. Additionally, changes in viewing angle are accommodated through flexibility in the expected distances between adjacent anatomical keypoints. As the face rotates, some distances expand while others contract. This technique is similar to the method presented in chapter 4 in the sense that it takes advantage of changes in the location of anatomical locations due to changes in viewing position.

Discrimination between sub-classes of faces has been achieved with modern machine learning algorithms, such as support vectors (Moghaddam & Yang, 2002) or multi-level neural networks (Phung & Bouzerdom, 2007). These techniques offer an advance over systems such as Eigenfaces by considering that the subspace containing the important face information is likely not to be linear. In the case of support vector machines, complex nonlinear functions are better modeled by considering more heavily the margins between classes in image space. Multi-level neural networks also offer the capacity to model arbitrarily complex functions, though in both cases there is a danger of over-fitting noisy data.

As mentioned before, biologically inspired neural network models of object categorization have been implemented as well (Serre et al, 2007). This model, which serves as the standard of its kind, is notable for producing human-like performance, robust to changes in the position of the object within an image, by implementing a neural network modeled on the cascade of feed-forward processing observed along the ventral pathway in the temporal lobe. Such a model would be a good candidate for training on classes of images of varying distances, if what we are interested in is biological plausibility. This task is left for future work.

These algorithms each attempt to model a specific sub-problem within the vast field of face processing, from the detection to the recognition of the face. To these modular units, this thesis adds one more: the classification of a face according to the viewing distance. Given the behavioral importance of viewing distance described in the thesis, it seems only complete to suggest a computational basis by which it may be estimated from images.

1.9 Neural Substrates

Much of the preceding discussion is illuminated by reference to the neural structures thought to be involved. The new data presented in this thesis are behavioral, but it is important to note the implications for our understanding of the brain.

One of the central controversies in how the brain processes face information has to do with whether or not it is "modular" or "distributed" (Kanwisher, McDermott, Chun, 1997; Haxby, Gobbini, Furey Ishai, Schouten, and Pietrini, 2001). This controversy centers on how a patch of cortex on the ventral surface of the temporal lobe responds to images of faces and other objects. According to one view, there is a highly specialized region within the fusiform gyrus that is specialized for face processing in which most of the neurons are selective for faces (Tsao, Freiwald, Tootell, and Livingstone, 2006). According to the other view, the area outside of this patch activates in a distributed pattern that also contains information about the face category (Haxby, Bryan, and Gobbini, 2006). The two views are not incompatible. Alignment of two brains on the basis of the distributed pattern that results from watching a movie gives rise to greater overlap in their functionally defined regions that are responsive to faces (Sabuncu, M., Singer, B., Conroy, B., Bryan, R., Ramadge, P., and Haxby, J., 2010). Perhaps some combination of these approaches will be able to distinguish between the subtly different appearances presented by the close and far faces described in this thesis.

In case the ventral temporal cortex is not the locus of representation of the visual basis of interpersonal distance, there is another sense of the word “distributed” as it used to describe face processing in the brain: rather than considering a pattern in one region, it is important to note that many areas of the brain work together to process faces (Haxby, Hoffman, and Gobbini, 2002). In some sense, this view is also consistent with the more recent “face patch” discoveries (Tsao, Moeller, and Freiwald, 2008), in that several regions are implicated in being face responsive, including: inferior occipital, inferior temporal, superior temporal, amygdala, and anterior temporal. It may be that the visual warping associated with interpersonal distance may be computed within such a distributed network of face regions. The response profiles of these face patches seems to indicate a division of labor processing different configural features associated with changes in appearance due to identity or face direction (Freiwald, Tsao, and Livingstone, 2009).

Social information processing in the brain seems to be part of a system that extends from the distributed face processing system. Multi-modal regions such as superior temporal sulcus (STS) seem to subserve both the analysis of expressive facial movement (Haxby, Hoffman, and Gobbini, 2002) and more general biological motion (Gobbini, Koralek, Bryan, Montgomery, and Haxby, 2007). Since the motion toward or away from a dyadic partner involves both communicative intent expressed through visual changes in the facial appearance as well as biological motion, it stands to reason that this region may be expected to respond to the changes in interpersonal distance described in this thesis. According to one model of the area (Giese & Poggio, 2003), optic flow detectors from area MT may provide STS with the inputs necessary to recognize biological motion. The detection of optic flow and biological would be a critical component of any biologically realistic model of interpersonal distance perception. Moreover, STS's involvement with expression makes it a candidate for trustworthiness judgements, according to the overgeneralization hypothesis (Engell, Todorov, and Haxby, 2010). However, brain scanning studies do not seem to find that the STS is responsive to the trustworthiness of faces (Said, Haxby, and Todorov, 2011).

The amygdala is another area responsive to emotionally salient stimuli (Haxby et al, 2002), and one that has a well-documented response to trustworthiness (Winston, Strange, O'Doherty, and Dolan, 2002; Said et al 2011). Greater activation in the amygdala follows the viewing of faces that are considered to be untrustworthy, though it is not clear yet if the response is linear or nonlinear. Given the role the amygdala plays in fear perception as well (eg. Adolphs & Tranel, 2003), this area is a prime candidate for involvement in the processing of interpersonal distance. Additionally, the amygdala is known to have massive reciprocal connections to the visual areas discussed earlier (Adolphs, 2004).

1.10 Summary

This thesis tests the hypothesis that participants are sensitive to changes in facial appearance that imply a violation of personal space. We go further to suggest that such differences are too subtle to be consciously perceived but nonetheless affect behavior. We also test various follow-up questions, such as how the effect compares to overt manipulations of interpersonal distance and what kinds of attributions are most sensitive to viewing distance. Given the psychological importance of facial closeness established by these behavioral experiments, the rest of the thesis is devoted to image analysis tools designed to estimate and manipulate distance based cues in natural images.

Chapter 2. Stimulus Acquisition and Behavioral Testing Methods

2.1 Stimulus Methods

2.1.1 Horizontal Spacing

The appearance of a three dimensional object depends on the viewing distance. Most people have an intuitive understanding of this by considering that distance causes objects to appear smaller, but when different parts of the same object are at different distances to the viewer, the effects are somewhat harder to visualize. In particular, for faces, the nose is closer to the camera than the ears. For a certain range of distances, this has a strong effect on the 2-D projection made, and thus the appearance of the face (see figures 1 & 2).

In order to study the role of perspective distortion due to viewing distance, it is possible to acquire stimuli in a number of ways. The first round of stimuli were collected by Professor Pietro Perona, who took photographs of volunteers outside at various distances. An example of these stimuli are shown in the figure, but there remained an outstanding question: did people make subtly different facial expressions when they saw that the camera was closer compared to farther? If so, would these differences influence the results of behavioral experiments? While these questions remain interesting to investigate, we felt the first order of business was to eliminate the potential confound due to expression. Therefore, we went about designing a stimulus capture method that would allow us to remove as a factor the facial expressions worn by participants in the close and far photographs (see figure 3).

Camera Set up

Our apparatus uses a half-silvered mirror to reflect and transmit the face image into two cameras that are aligned to the same optical axis. One camera was placed behind the mirror, and the other placed a distance off to the side at a right angle. As can be seen from the figure, the total distance the light traveled to reach the camera behind the mirror is 1.5 feet, and the total distance light traveled to reach the far camera is 4.5 feet. To ensure the cameras were aligned correctly, we used a digital laser measurer to measure the distance to each side of the lens to millimeter precision. To calibrate the alignment, we pointed a level-lined laser into the first camera while monitoring the output the camera through the computer. When the camera is aligned to the horizontal, the entire image is filled with the color of the laser. We performed the same procedure on both cameras simultaneously to ensure they were aligned to the same optical axis.

Since each camera was set to a different focal length, there is the possibility that optical distortion might influence the appearance of the images. Short focal length distortion (sometimes called wide angle lens distortion) mostly affects images in the periphery and does so according to a predictable pattern known as barrel distortion. We photographed a checkerboard image to check the degree of image distortion for the focal length used. We found that the amount of distortion was negligible in the central area in which faces were photographed (see figure 4). Additionally, independent analysis (http://www.imaging-resource.com/PRODS/G10/G10A4.HTM) confirms that the distortion is extremely small (0.9% barrel distortion at 28mm, 0.1% pincushion distortion at 140mm). In any case, all of the lens distortion occurs at the periphery (see figure 4b), so it cannot possibly account for the large physiognomic differences in a face-centered portrait image.

Image Normalization

After acquiring the images, we then had to normalize them to keep any overt cues from distinguishing them. The far image, captured after reflection on the mirror, was left-right flipped to restore the original orientation. The close image was downsampled and resized to match the resolution and dimensions of the far image. Both images were converted to grayscale and set to the same luminance and contrast (for schematic of the process, see figure 5). The resultant images had the same intra-ocular distance, luminance, resolution, contrast, and expression. Each image was also rotated so that the eyes are aligned horizontally and are placed at the vertical center of the image. After processing, the only difference that remained was the geometrical warping due to perspective projection, and subtle differences in focus, resolution and highlighting.

In order to eliminate the subtle lighting differences, we devised a way to warp the far image so that it matched the shape of the close image, but retained its original lighting information. By annotating each face on 115 anatomically important locations, we could calculate a warp field to apply to the image so that these important keypoints would be in the configuration of the close face, but the pixel information would come from the far face. This final control step ensured that the only difference between the far and close image was the geometrical difference due to perspective projection (see figure 6). Such a transformation of the image is only possible when both the near and far faces are photographed simultaneously along the same optical axis; otherwise, anatomical landmarks might not lie in strict registration.

2.1.2 Vertical Spacing

The appearance of a three dimensional object also depends on the vertical viewing angle. Objects viewing from slightly above or below will cast a different 2-D image. Since there exists a good deal of variation in the heights at which people stand, it stands to reason that learning algorithms in the brain have picked up on this statistical regularity as well.

In order to test the likely hypothesis that vertical viewing angle influences social perceptions, it is possible to capture stimuli in a number of ways. The easiest way is to take a series of photographs at different vertical elevations with respect to the portrait subject. This method, however, suffers from the same confound as taking a series of photographs at different horizontal displacements: the portrait subject is liable to alter their emotional expression as a response to their explicit knowledge of the camera position. Therefore, as it was before, it is necessary to design a photographic apparatus to simultaneously take photographs at multiple viewing angles.

The usage of multiple cameras at different heights introduces a new confound: eye gaze direction. When taking photographs at two simultaneous distances, the eyes gaze forward in the same direction for both images, but when taking photographs at multiple vertical angles, the gaze is noticeably averted. To eliminate this confound, we digitally manipulated the gaze direction of the stimuli in order to always be directed at the camera. This process required taking three photographs, instructing the subject to hold their head still and moving only their eyes to each of the three cameras. We then transferred the photographic information from the eyes from these three images onto the three simultaneously recorded images from a single acquisition. This combination of simultaneous photographs and digital manipulation ensured that all the images displayed the same expression and the same gaze direction.

We collected two datasets using this technique: one with 6 Caucasian male faces and one with 9 Caucasian male faces. These images are used for an experiment described in Chapter 3.

For both photographic methods, we recruited volunteers from the pool of subjects that also participated in behavioral testing. These subjects, therefore, were familiar with the rules of the economic trust game and were able to tell us what their own responses would be as the trustee. They were instructed to wear a completely neutral expression, devoid of any emotional information. However, given that we can capture emotional expression from multiple perspectives simultaneously and that expression is very likely to influence social perceptions, we opted to collect positively and negatively valenced stimuli for use in future studies. Therefore, the end result of the stimulus acquisition portion of this work is a large body of facial images that contain simultaneously captured emotional expressions from different viewing positions. For a table of stimuli available for future research, see table 1.

2.2 Behavioral Testing Methods

2.2.1 Participants

Participants were recruited using a combination of approaches, primarily through advertisement on Craigslist, brainscience.caltech.edu, through posted flyers throughout the Pasadena area, and through Amazon Mechanical Turk, which allowed people to participate from their personal computers. Eligible participants were between the ages of 18 and 55, had no history of neurological or psychiatric illness, and had not previously participated in the study. See Table 2 for detailed demographic information for each experiment.

The reason for replicating the experimental results in three separate testing conditions is that it provides broader support for the relevance of perspective projection in everyday life. Subjects who came to the lab participated under very controlled conditions, with their heads resting in a chinrest at a fixed distance to the screen. They were the only subjects in the room and received individualized verbal instruction from the experimenters. These conditions ensured that all subjects had roughly the same experience. At the opposite end of the experimental control / ecological validity trade-off, subjects who participated through Amazon Mechanical Turk did so under a wide variety of testing conditions. We had no experimental control over the size of the display, their proximity to the display from trial to trial, or any personal assurance that they even understood the instructions. To gauge their ability to follow instructions, we had to rely on their ability to successfully complete a series of surveys before receiving compensation.

2.2.2. Experimental Design

To assess the emotional response to these stimuli, we used two paradigms. First, we simply asked participants to rate the stimuli on a 1-7 scale on a variety of characteristics, including Attractiveness, Competence, Trustworthiness, Age and Weight. We also included a block where participants made explicit judgments about the distance to the camera. Second, we asked an independent group of participants to play an economic trust game imagining the people depicted in the stimuli as partners. Thus we sought a direct and an indirect way to ask participants about trust, and received data concerning both explicit judgments and real-world behavior.

All of the paradigms used are "Within-Subject" for the purposes of maximizing statistical power. In other words, each participant views each stimulus face in both the far and the close condition. The possibility of order effects are addressed through counterbalancing: for each participant, a randomly chosen half of the stimulus faces (9 faces) are shown in the close condition for the first half of the study, followed by the far condition for the second half. The other half of the faces (9 faces) is shown in the opposite presentation order (first seen as far, then as close). Thus each half of the experiment contains an equal number of close and far faces, and an equal number of faces are seen first close then far as are seen first far then close. This design allows us to calculate an difference score for each participant, and to perform group level statistics over those scores, while avoiding the possibility that the order of presentation can drive the effect.

The counterbalancing design also helps avoid experimental transparency by ensuring that no face is seen in both conditions consecutively or in the same half of the experiment.

Whether a participant gave direct ratings or played an economic trust game, the design of the experiment was the same. However, the trait rating design included 3 rating scales, so this portion of the study took 3 times as long to complete. The order of trait rating was counterbalanced such that each trait occurred in first, middle, and last an equal proportion of times across participants.

Participants were instructed to rate the faces according to their first impression, immediate gut reaction, and to make their decision quickly. They were told that faces would appear more than once, that we were interested only in how the facial appearance struck them at that particular moment in time, and that it was okay to change their minds over the course of the study. They were also told that all the faces would be Caucasian males because we wished to avoid the complicating effects of race or gender and that all responses should be made relative to their experience with this particular demographic group.

Each participant rated each face in both conditions, separated in time. For each experiment, analysis proceeded by examining the difference between their CLOSE and FAR ratings using a 2x2 mixed ANOVA, with factors viewing distance (within subject factor) and gender (between subjects factor). A post-hoc t-test was then used to examine the difference in the CLOSE and FAR means. In all cases, the data were normalized to a 100-point scale for ease of interpretation as a percentage of the range. However, the analyses were also performed in parallel on the z-transformed data to ensure that the results were the same for data adjusted to meet the Gaussian distribution assumptions of the ANOVA test.

Behavioral methods for testing the role of vertical viewing angle closely parallel those described above for testing the role of camera distance. However, instead of there being two testing conditions, there are three. Therefore, instead of dividing the stimulus set into halves, we divide it into thirds and counterbalance according to the same logic. Each third of the experiment contains an equal proportion of *direct*, *above*, and *below* faces, and each face is seen an equal number of times in each viewing condition.

Stimulus Set	Number of Faces	Viewpoint	Expression	Gaze Direction	Studies
MDP1	18 males	Frontal	Neutral	Direct	Chapter 3
	18 males	Frontal	Smiling	Direct	-
	18 males	Frontal	Angry	Direct	-
	18 males	Frontal	Neutral	Averted	
	18 males	Frontal	Smiling	Averted	
	18 males	Frontal	Angry	Averted	
	18 males	3/4	Neutral	Direct	
	18 males	3/4	Smiling	Direct	
	18 males	3/4	Angry	Direct	
	18 males	3/4	Neutral	Averted	
	18 males	3/4	Smiling	Averted	
	18 males	3/4	Angry	Averted	
MVAP1	15 males	Frontal	Neutral	Direct	Chapter 3
	15 males	Frontal	Smiling	Direct	
	15 males	Frontal	Angry	Direct	

Table1.

Sample of stimulus sets collected using the simultaneous acquisition method (Caucasian males). In addition to the stimulus sets listed above, we also have collected the following stimulus sets for Multiple Distance Photography (MDP): 18 Caucasian Female, 8 Asian Female, 8 Asian Male, 5 African-American Males, 5 African American Females, and 8 other, all in each of the 12 viewing conditions displayed above.

Study	N	Age (mean ± SEM)	Demographics
MDP EXP 1	23	33.26 ± 2.92	17 female. (7 Caucasian, 6 Asian, 4 Hispanic, 1 African-American, 5 other.)
EXP 1b	45	25.91 ± 1.18	35 female. (34 Caucasian, 6 Asian, 3 Hispanic, 2 African-American.)
EXP 1c	37	26.38 ± 1.45	23 female. (37 Caucasian.)
EXP 2	27	23.93 ± 1.09	17 female. (15 Caucasian, 7 Asian, 1 Hispanic, 1 African-American, 3 other.)
EXP 3a	268	31.5 ± 0.62	148 female. (205 Caucasian, 22 Asian, 13 Hispanic, 14 African-American, 14 other.)
EXP 3b	70	30.32 ± 1.3	27 female. (53 Caucasian, 6 Asian, 4 Hispanic, 3 African-American, 4 other.)
EXP 3c	60	32.15 ± 1.48	27 female. (48 Caucasian, 4 Asian, 3 Hispanic, 3 African-American, 2 other.)
EXP 3d	253	31.83 ± 0.64	143 female. (193 Caucasian, 21 Asian, 12 Hispanic, 13 African-American, 14 Other.)
EXP 3e	134	31.46 ± 0.88	68 female. (100 Caucasian, 12 Asian, 9 Hispanic, 4 African-American, 9 other.)
MVAP EXP 1	80	27.91 ± 1.14	54 female. (34 Caucasian, 22 Asian, 10 Hispanic, 3 African-American, 11 other).
EXP 2	23	20.27 ± 0.61 .	8 female. (8 Caucasian, 12 Asian, 2 Hispanic.)
EXP 3	16	31.94 ± 1.99	12 female. (16 Caucasian).

Table 2.

Demographic characteristics of all study participants.

Figures

Figure 1. Example stimuli collected at FAR (4.5 ft), left, and CLOSE (1.5 ft), right. Both images are acquired at the same instant, ensuring their emotional expressions are the same.

facial representation depends on interpersonal distance in at least 2 ways

visual angle of head, at distance d,
$\theta_h = atan(h_w / (h_p + n_p + d))$

visual angle of nose, at distance d,
$\theta_n = atan(n_w / (n_p + d))$

proportional representation of nose:

θ_n / θ_h

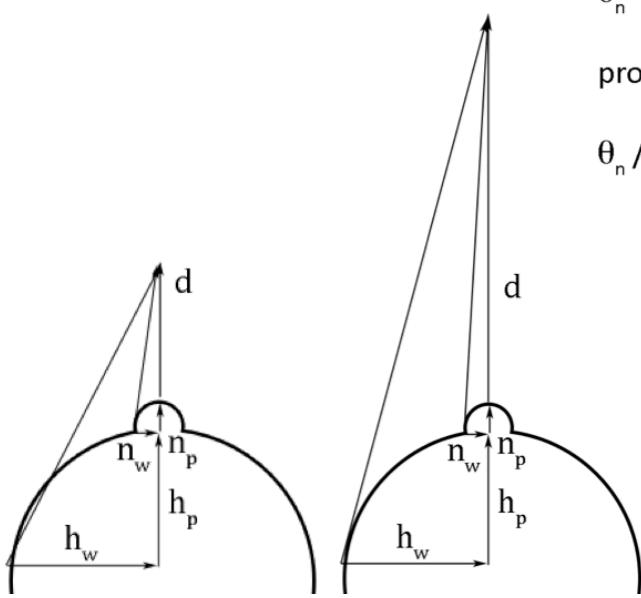

Facial parameters
h_w = head width ~ 147mm
h_p = head protrusion ~ 147mm
n_w = nose width ~ 36mm
n_p = nose protrusion ~ 18mm

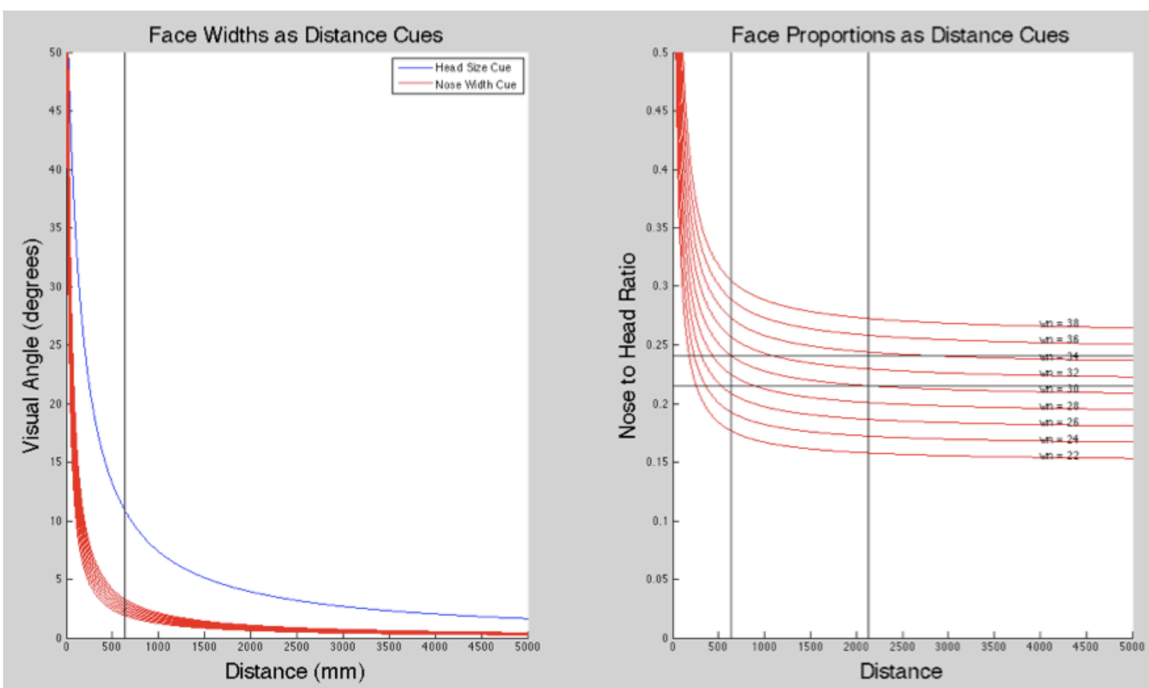

Figure 2. Illustration of geometry behind differences in facial appearance. The proportional representation of the nose, compared to the head, decreases with viewing distances according to the formula derived above, and shown below. Many different possible nose-widths are shown for illustration that the shape of curve is very similar regardless of the parameters of the face's shape.

Figure 3. Photographic set up. Seen to the right is the FAR camera, also visible in the reflection in the mirror. Behind the mirror lies the CLOSE camera. The aluminum panels are present to provide additional illumination.

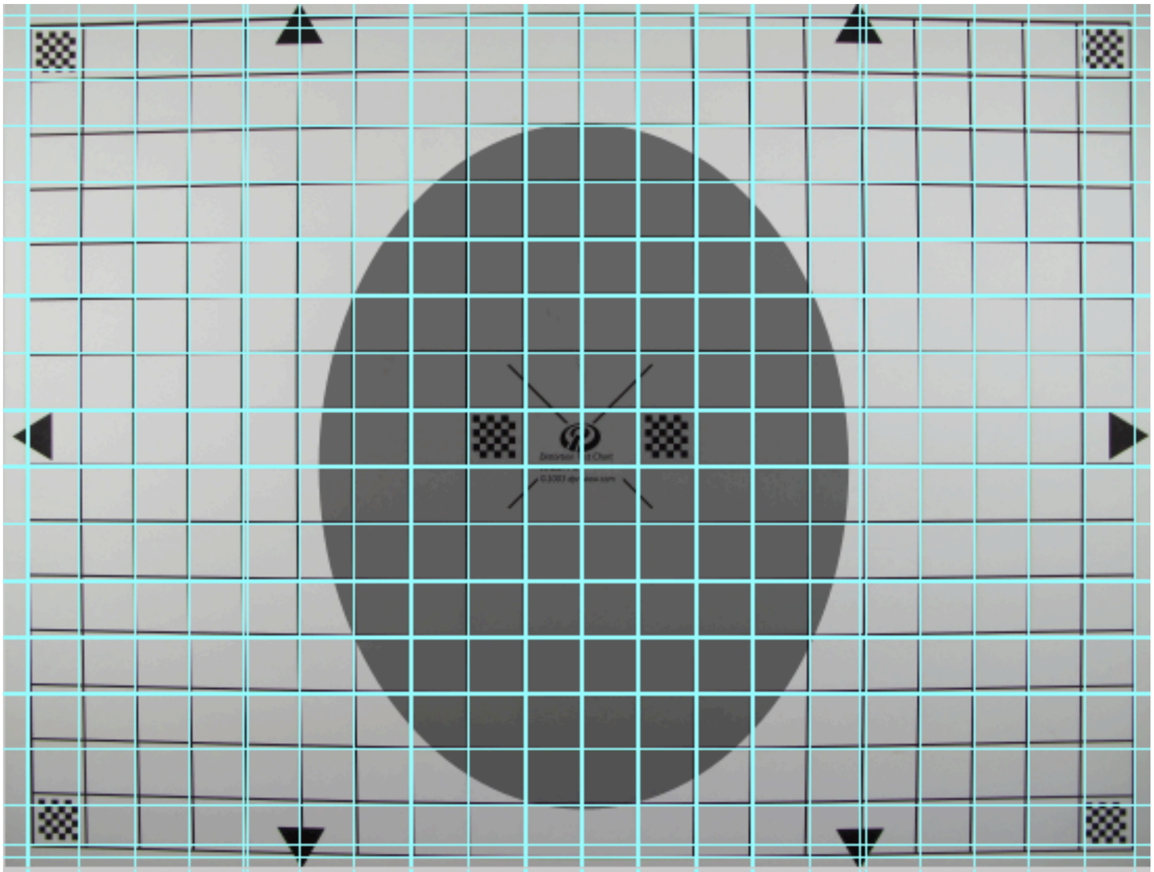

Figure 4. Test patterns. As can be seen, when grid pattern is photographed with our camera at the CLOSE distance (reproduced from www.imaging-resource.com/ with grid overlaid), the magnitude of lens distortion is negligible, especially in the region of interest. This photo was taken with the exact same camera (Canon G10) and focal length (18mm) as was used in the CLOSE condition of the experiment. The FAR condition displays a similarly negligible lens distortion.

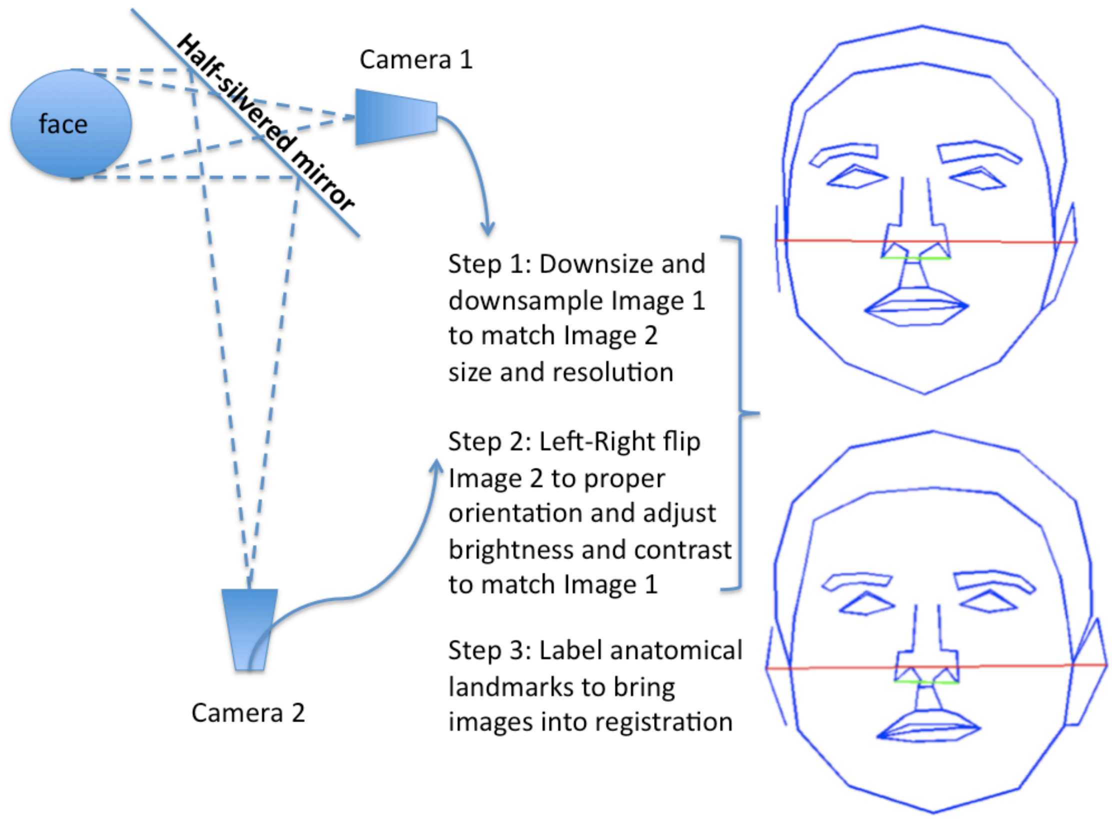

Figure 5. Stimulus creation process. A face is imaged simultaneously from to cameras using a half-silvered mirror. Both images are then normalized to equal size, resolution, brightness, and orientation. Finally, anatomical landmarks are labeled

Figure 6. Digital manipulation of images according to photographic distance. The anatomically labeled images can be digitally adjusted in their proportions to simulate how they would appear at different distances. The face on the left is the original FAR face from figure 1. The face in the middle is the original CLOSE face. The face on the right takes the FAR face (on the left), adjusts the proportions to match those of the CLOSE face (in the middle), and overlays it onto the same background. This manipulation is used in Experiment 2 in Chapter 3 Section 2.

Chapter 3. Behavioral Impact

3.1 Introduction

In this chapter I will introduce several initial results we have found using the stimuli discussed in Chapter 2. We use two basic paradigms: ratings of personality traits and a behavioral economics game. We use both behavioral testing methods in order to find converging evidence that the viewing distance and vertical viewing angle to a face image impacts the psychological processes underlying interpersonal inferences, especially those related to trust decisions.

In addition to the explicit ratings described above, we also sought to determine if the effect we observe generalizes to real-world conditions. As a step in that direction, we implemented an economic trust game in which the participants were incentivized with real money. After playing a round with computer partners, participants were told "Now we will play the same game again, but with a few changes. Instead of seeing differently colored rectangles to represent your investment partner, now you will see images of people's faces. You may have a first impression, immediate gut reaction about whether or not you would like to invest with them. That is what we want you to pay attention to when you make your decision. One trial will be selected at random to determine a real payout. … We've asked the people who appear in the study how much they would actually keep and return for each possible investment amount you can make, and we will use these responses in addition to your investment to determine how much you would make in this game. Then we will give you a percentage of this amount. Treat every trial as if real money were at stake. "

3.2 Influence of Interpersonal Distance

Perspective distortion from interpersonal distance:

an implicit visual cue for social judgments of faces

Ronnie Bryan[1], Pietro Perona[1,2], Ralph Adolphs[1,2]

[1]Computation and Neural Systems Program

[2]Division of Biology

California Institute of Technology

Pasadena, CA 91125

Correspondence to:

Ronnie Bryan (ronnie@caltech.edu)

1200 E California Blvd M/C 136-93, Pasadena, CA, 91125

Acknowledgments: We thank Kathy Lawrence and Pranay Kothari for help with data collection. Funded by grants from NIH and the Moore Foundation (R.A.).

The basis of social judgments, such as trust, derived from visual information in faces remains an open problem in social cognitive science. Many explanations relating to facial appearance have been proposed, suggesting a basis for inferring why some faces are viewed as more trustworthy than others. Here we investigated the contribution of a novel cue: the change of appearance due to the perspective distortion that results from viewing distance. We found that photographs of faces taken from within personal space elicit investment in an economic trust game and lower ratings of trustworthiness, competence, and attractiveness, compared to photographs taken from a greater distance. This effect was replicated across multiple studies controlling for facial image size, facial expression and lighting, and was not explained by face width-to-height ratio, explicit knowledge of the camera distance, or how average the faces are perceived. These results demonstrate a novel facial cue influencing social judgments as a function of interpersonal distance, and moreover one likely to be processed implicitly.

Introduction

We glean a wealth of socially relevant information from faces in the blink of an eye: a person's attractiveness, competence, threat, identity, gender, emotion, and trustworthiness, to mention a few. For example, reliable judgments of trustworthiness can be made from faces viewed for 100ms (1). Multiple factors influence such judgments. The perceived valence of the neutral face, for example, is thought to influence trait attributions by activating brain systems tuned to facial expression (2). The structural width-to-height ratio of a face has been shown to be a reliable indicator of the testosterone and untrustworthy behavior (3). Similarly, features such as the roundness of the cheeks and the apparent size of the eyes, termed babyfacedness, may influence trust by activating representations related to the perception of age (4). These avenues of investigation all attempt to explain why some individuals are perceived as more or less trustworthy than others on first glance. Yet there is one important ecological cue that, to our knowledge, has not received such intensive study: the perspective distortion as a function of physical distance. The change of appearance of an individual with viewing distance is studied here as independent of other factors such as facial expression and is demonstrated to be sufficient in itself for influencing behavior relating to trust and other inferred personality traits. Our approach expands investigation from analysis of the appearance of a face to analysis of the relationship between a viewer and the stimulus.

Three-dimensional objects, such as the human face, produce on the retina a two-dimensional image via perspective projection. The image varies with distance from the center of projection, even when equated for size (see Figure 2a); e.g., the nose looks relatively larger and the ears smaller as the distance decreases (5). Such differences may be modeled as a distance-dependent image warp or distortion (see Figure 1). This effect may have been utilized in portrait paintings not only to induce distance percepts but also to manipulate how viewers feel about the face (6).

Ever since Edward Hall's seminal book on the topic (7), interpersonal distance and personal space have been highlighted as ubiquitous and potent determinants of a wide variety of social behavior (8). Notably, interpersonal distance is associated with arousal (9), self-protective behavior (10), privacy (11), emotional valence (12, 13), management of stress and aggression (14), and interpersonal trust (15). In each of these studies, interpersonal distance is manipulated in an ecologically valid way, that is, participants are observed reacting to a confederate standing at an experimentally determined distance. The result is that the observed changes may result from any or all of the many multi-modal perceptions that accompany a change in interpersonal distance. For example, the size of the face is smaller and the visibility of the body is greater at greater distances. These studies demonstrate the efficacy of interpersonal distance at eliciting a variety of emotional responses relevant to trust.

Reading faces for socially relevant traits such as trust may occur automatically and may elicit reliable ratings after a very brief exposure (16, 1). These findings imply that there may be a system that implicitly evaluates trustworthiness. Interpersonal distance is a potent variable influencing social behavior (7, 8, 17), and is related to activity in the amygdala: even the knowledge of interpersonal closeness causes an increase of activity in this brain structure (18). The amygdala is also a critical structure for the automatic evaluation of threat (19; 20), facial valence information (21), and trustworthiness of faces (22; 23; 24). Therefore, we hypothesized that the distance-dependent perspective projection of a face might be a cue for social judgments, especially those related to trust.

Since interpersonal distance is known to influence a variety of traits, we investigated a broad set of questions in these experiments. Participants not only performed a trust game, but also rated faces on dimensions of apparent trustworthiness, competence, attractiveness, age, weight, averageness, and animal-likeness. These traits were selected because of they have been shown to be important to social decision-making (trust, competence, attractiveness, and age) or might vary with distance in a predictable way (weight, averageness, and animal-likeness).

We investigated the connection between perspective projection and trust in three experiments that obtained social judgments (ratings) as well as measured trust behavior in terms of the amount of money participants were willing to invest in a person whose face they saw (see Table 1 for summary of experiments). The first experiment used photographs taken from different distances, while controlling the size and facial expression of the stimuli; the second used synthetically warped face images to eliminate possible confounds in highlights and focus; the third explored a number of follow-up questions with a larger subject sample tested over the internet. All effects are reported as the difference of the behavioral response to far and close face stimuli. Although participant gender was not a factor of interest in our study, all findings were followed up with exploratory ANOVAs that included participant gender as a possible factor.

Experiment 1

Results: In Experiment 1a, faces photographed at the far distance elicited higher investments than those photographed at the close distance: mean investment difference (far faces - close faces) was 3.2±2.1 (95% CI), $t(22)=3.2$, $p<0.01$ (paired t-tests, 2-tailed). Similarly, in Experiment 1b the far faces elicited higher ratings of attractiveness (5.1±1.5, $t(34)=6.8$, $p<0.001$), competence (2.7±1.8, $t(33)=3.0$, $p<0.01$), and trustworthiness (2.8±2.0, $t(35)=2.9$, $p<0.01$) than those photographed at the closer distance (Fig. 1).

We examined the stimulus-by-stimulus correlations between the trait ratings in Experiment 1b among each other, and with the investments made in Experiment 1a.

Among the participants in Experiment 1b, Trust ratings were strongly correlated with competence ratings ($r=0.90$, $p<0.001$) and attractiveness ratings ($r=0.82$, $p<0.001$). Competence and Attractiveness ratings were likewise correlated ($r=0.74$, $p<0.001$). The correlations are so high that the residual trust ratings after regressing out the ratings of attractiveness and competence do not display a significant preference for far faces on their own.

Between the participants of Experiment 1a and Experiment 1b, investments in the trust game were correlated with ratings of trust ($r=0.84$, $p<0.001$), competence ($r=0.86$, $p<0.001$), and attractiveness ($r=0.65$, $p<0.001$). Again, residual investments after regressing out these independent face ratings do not a display statistically significant preference for far faces on their own.

In Experiment 1c, ratings of age and camera distance showed no statistically significant effects of distance (respectively, 0.1 ± 1.2, -3.1 ± 4.3), although ratings of weight revealed that faces photographed farther away appeared heavier (3.9 ± 1.4, $t(35)=5.47$, $p<0.001$). Experiment 1a investment residuals after regressing these ratings out do display a statistically significant preference for faces (regressing out age: mean investment difference = 3.24, $p<0.02$; regressing out distance: mean investment difference = 4.63, $p<0.01$; regressing out weight: mean investment difference = 3.97, $p<0.01$).

Post-experiment debriefing confirmed that none of the participants noticed that face distance was manipulated. Finally, to explore possible gender effects, a 2x2 (participant gender x viewing distance) ANOVA on the trustworthiness ratings confirmed a significant effect of viewing distance ($F(1)=6.68$, $p<0.02$), but failed to find a main effect or interaction of gender ($F(1)<0.3$, n.s.).

Discussion

Faces photographed from within personal space elicited lower monetary investments and lower ratings of trustworthiness, attractiveness and competence than did simultaneously photographed faces from outside of personal space. All three ratings were highly correlated, suggesting that the influence of personal space on social judgments may not be limited to trustworthiness alone.

The finding that the faces appear heavier is consistent with the vertically oblong shape of the human head, which will produce the greatest perspective distortion at the sides. The effect is that the width-to-height ratio is smaller for closer faces, making them appear thinner. The fact that participants rated far faces as heavier confirms they were able to physically distinguish the far faces from the close faces, but when asked explicitly about camera distance, they were not able to answer correctly. When investment amount residuals after regressing out weight ratings continue to display a preference for far faces, indicating that the change in the perception of a face's weight does not fully account for the change in investment that face receives. By contrast, trait ratings of trustworthiness, competence, and attractiveness are candidate mediating variables for investment.

Experiment 2

It is conceivable that subtle differences in highlight and focus between the far and near pictures, independent of distance-induced warping, might contribute to this finding.

More closely photographed faces exhibit a greater sheen on the highlights than do farther faces. Although the global contrast may be equalized by adjusting the dynamic range of the image, the local contrast in face areas that receive more direct illumination may still contain luminance based cues. Similarly, closer facial features such as the nose may be photographed with a slightly different sharpness than the farther features such as the ears due to the varying distance to the lens. We did not notice any differences in focus between these face regions, but Experiment 2 was conducted as safeguard measure anyway.

To completely isolate perspective warp as the factor against these possible confounding variables, we repeated the experiment with synthetically warped faces.

Results: The mean investment difference (far-close faces) was 4.2±2.1 (95%CI), $t(24)=4.2$, $p<0.001$, confirming the effect observed in Experiment 1. Post-experiment debriefing again verified that none of the participants noticed that face distance was manipulated. To explore possible gender effects, a 2x2 (participant gender x viewing distance) ANOVA showed a significant effect of viewing distance: $F(1)=15.76$, $p<0.001$, but no effects of participant gender or interaction with gender ($F<1.3$; n.s.).

Discussion

Experiment 2 confirmed that distance-induced warping alone (perspective projection) influences trust-related investment behavior even when controlling for luminance based cues such as local contrast and focus. This result does not rule out these cues as possible factors, but does show that they are not necessary to obtain the effect we observe. It should be noted as well that the stimuli in Experiment 1 are more ecologically valid than the synthetically produced stimuli in Experiment 2.

The results of Experiment 2 demonstrate that perspective projection warping is sufficient to influence trust behavior, opening the door for the manipulation of images even in the absence of the simultaneous photographic set-up we devised for these experiments.

Experiment 3

Can the effect measured in Experiments 1 and 2 be obtained with explicit distance cues, such as mere verbal information or image size? Might the effect be due to how average (typical) the images appear? Is the effect sufficiently robust to appear outside the laboratory? We explored these questions in Experiment 3.

Experiment 3 addresses underlying issues in the mechanism by which the behavioral effect described in Experiments 1 and 2 might hold. Participants seem not to be aware of any manipulation of camera distance, but do similar results hold if people are consciously aware of distance manipulation? The role of awareness of interpersonal distance is important for suggesting processing stages responsible for the ultimate trustworthiness decision. If explicit knowledge of interpersonal distance is found to be a sufficient factor for explaining trust ratings, the implication is that perspective projection may feed into this system. If, on the other hand, explicit knowledge of interpersonal distance is not sufficient for explaining trustworthiness ratings, the implication is that there may be a subconscious system for social regulation that automatically processes this information. Experiment 3 provides participants with explicit information about interpersonal distance decoupled from the perspective warp: 3b provides verbal information and 3c provides size information.

Another potentially mediating variable that could explain the results of Experiment 1 and 2 is the typicality of the face. Averageness of faces is known to influence a host of cognitive functions (28), including the perception of attractiveness (29), so it is possible that the close faces of Experiment 1 and 2 were seen as less trustworthy simply because they were seen as less average. If participants do in fact view the faces as less average, they should be able to report this perception, as they do in other experiments (30). Experiment 3 obtains averageness ratings to determine if this perception accounts entirely for the effect of viewing distance. Experiment 3 also obtains ratings of animal-likeness as another avenue for participants to indicate that they find the faces to appear unusual.

Results: Experiment 3a replicated the effects observed in Experiment 1 for far-close faces, trustworthiness: 1.6±0.7 (95%CI), $t(238) = 4.3$, $p<0.001$; competence: 1.8± 0.7, $t(244) = 4.8$, $p < 0.001$; attractiveness: 2.6±0.7, $t(238) = 7.5$, $p < 0.001$. As before, 2x2 (participant gender x viewing distance) ANOVAs confirmed a significant effect of viewing distance: trustworthiness: $F(1)=14.4$, $p<0.001$; competence: $F(1)=12.0$, $p<0.001$; attractiveness: $F(1)=44.6$, $p<0.001$), but no effects of participant gender or interaction with gender (all $F<0.6$; n.s.). See Figure 3 for a summary of the results of Experiment 3a-c.

As with Experiment 1, we examined the correlation between the trait ratings and the investments. Investment amounts were highly correlated with these independent ratings of trustworthiness ($r=0.86$, $p<0.001$), competence ($r=0.88$, $r<0.001$), and attractiveness ($r=0.66$, $p<0.001$). Once again, the investment residuals after regressing out each of these ratings did not display a statistically significant preference for face faces.

Experiment 3b showed no effect of explicit verbal information about distance on any rating: trust: -0.2±1.6, $t(65) = -0.22$; competence: 0.2±1.1, ($t(62) = 0.30$; attractiveness: -0.3±1.5, $t(64) = -0.34$, all n.s.

Experiment 3c showed no effect of image size: trust: 0.4± 1.0, $t(53) = 0.78$; competence: -0.1±1.1, $t(56) = -0.12$; attractiveness: -0.2±0.9, $t(53) = -0.38$, all n.s.

Experiment 3d replicated Experiment 1c (heaviness: 2.9±0.6, $t(224) = 9.8$, $p < 0.001$; age: -0.1±0.5, $t(232) = -0.27$, n.s.), although ratings of camera distance now reached statistical significance, likely due to the increased sample size (-3.3±1.3, $t(225) = -5.10$, $p < 0.001$). See Figure 4 for a summary of the results of Experiment 3d-e.

Experiment 3e showed that “far” faces were rated as more Average (1.8±0.5, $t(111) = 3.63$, $p < 0.001$), but not more or less Animal-like (-0.7±1.1, $t(116) = -1.29$, n.s.).

Athough avérageness and trustworthiness ratings across all 36 faces (the 18 close and 18 far versions of each of the 18 individuals) were negatively correlated ($r=-0.36$, $p<0.05$), the residualized trustworthiness ratings, partialling out averageness, still showed a significant effect of distance as before (2.3 +/- 0.72 (SEM), $t(17)=3.3$, $p<0.01$)

A small minority (16.4%) of participants in Experiment 3 indicated in the exit survey that they noticed a change in the face stimuli between trials. Excluding these participants from the analysis did not change any the results significantly.

Discussion

Experiment 3 demonstrated that the influence of perspective distortion is robust even when administered over the internet, where display size and distance to the display are not controlled. Explicit manipulation of perceived distance to the face stimulus through image size or verbal instruction failed to show any effects, indicating that perspective distortion operates through separate processes.

Perceptions of averageness were also influenced by perspective distortion, suggesting the possibility that these might in part mediate the effect on trustworthiness. However, across all of the 36 faces (close and far ones), averageness ratings were in fact anticorrelated with trustworthiness ratings, with the result that partialling out the effect of averageness actually increased the significant of the effect of distance on trustworthiness judgments. We thus conclude that the effect of distance on trustworthiness judgments is not a result of manipulating the averageness of the faces.

Finally, as in Experiment 1, the participants in Experiment 3 were did not accurately judge the camera distance of the stimuli. In both experiments, participants actually were more likely to report the opposite of the correct answer, indicating perhaps some awareness of the manipulation, but not enough knowledge of the domain to make a correct interpretation. Since we included the more straightforward condition of estimating the heaviness of the stimuli, we can check if participants were perhaps just confused about the instructions. As perspective distortion causes the “close” images to have a narrower aspect ratio, participants in Experiment 3 reliably reported that these stimuli were less heavy. This demonstrates that when asked to judge a purely physical trait, participants can demonstrate accurate discrimination on this task.

General Discussion

We report a reliable novel effect, replicated across several different experiments and in separate subject samples: viewers trust faces photographed at 135cm more than those photographed at 45cm. The effect was found in an economic trust game with real money, and in ratings gathered under laboratory conditions as well as over the internet. Geometric warping of the face alone (modeling perspective distortion due to distance) accounted for the effect while controlling for size, expression, resolution, highlights, focus, and explicit knowledge of camera distance.

Faces photographed at the far distance (135cm) were also rated as more average, as well as more competent and attractive. Given that all these ratings are intercorrelated to some extent, it is difficult to determine which of these judgments might possibly be mediating any of the others; for instance, it is plausible that the perceived averageness of the faces in part drives the differences in trustworthiness that we report. However, when controlling for averageness, the effect of distance on trustworthiness judgments in fact increased (Experiment 3), indicating that our distance manipulation does not influence trustworthiness judgments derivatively merely by altering perceived averageness.

It is likely that the cue of perspective distortion from distance usually operates implicitly, as it did in our experiment. Participants were incorrect when asked to judge camera distance, and post-experiment questioning showed that participants were unaware of any manipulation in facial appearance from trial to trial. Given that the far faces, when normalized for inter-ocular distance, are actually a bit greater in area, one possible explanation for the consistently incorrect distance attributions we found is that participants are using a size-based heuristic to guess the size. Since the far faces are wider, they are incorrectly judged to be closer. The implicit nature of our distance cue is intriguing not only because it isolates psychological processes that could otherwise be contaminated by overt reasoning about distance, but also because the two explicit distance cues we examined (image size and verbal information) in fact did not produce effects on trustworthiness judgments.

There is a documented effect of facial masculinity proportions (the face width-to-height ratio) on perceived untrustworthiness (3). However, this is unlikely to account for our finding as the facial width-to-height ratio is actually smaller in our "close" than "far" faces (paired t-test, $t(17)=11.16$, $p<0.001$); if width-to-height ratio were the predominant effect, it would lead to an effect in the direction opposite from what we observed. Face warping from projection distance thus appears to be an independent signal used for social judgments.

The importance of the present findings extend beyond our discovery of a novel social cue from faces. Perspective distortion is perhaps the first implicit cue to interpersonal distance, opening the door for further studies on the underlying psychological processes as well as the brain structures involved in the automatic evaluation of personal space. Attractive aspects of perspective distortion, as a cue to social judgments, are that it has a natural parameterization and that it may be studied in isolation from other cues.

Future applications will be to predict, and to manipulate, viewers' feelings about other people from quantification of the perspective distortion of photographs on the internet, in magazines, and in personal identification documents (5). An important limitation of the findings thus far concerns their generality: the literature documents many variables that interact with personal space. No doubt, there will be effects of gender (31) and familiarity (14), of culture (14, 32), of the expression and of the context in which the face is seen (33, 34), all of which are likely to interact with the perspective factor we isolated here.

Experiment 1 Methods

Subjects: Healthy adult participants were recruited from the local community through posted flyers and Internet ads. Experiment 1a: N = 23, mean age = 33.26 +/- 2.92 (SEM), (17 female, 6 male; 7 Caucasian, 6 Asian, 4 Hispanic, 1 African-American, 5 Other). Experiment 1b: N = 45, mean age = 25.91 +/- 1.18 (SEM), (35 female, 10 male; 34 Caucasian, 6 Asian, 3 Hispanic, 2 African-American), Experiment 1c: N = 37, mean age = 26.38 ± 1.45 (SEM) (23 female, 14 male; 37 Caucasian). Participants in Experiment 1a were non-overlapping with those in Experiments 1b,c whereas all of those in 1c had first participated in 1b.

Stimuli: Participants viewed frontal grayscale photographs of the faces of 18 unfamiliar

Caucasian males, Age=33±12, displaying direct gaze and a neutral expression. For each face, two photographs were taken simultaneously from distances of 45cm and 135cm using a half-silvered mirror (Figure 2a), which ensured that the facial expression would be identical. The distances were chosen to be within and outside of personal space, respectively (7). Camera alignment was confirmed with a digital laser meter; lens distortion of checkerboard test images was negligible. Images were equated for mean contrast, luminance and interocular separation (i.e. face size) and presented for 5s (Experiment 1a) or 2s (Experiments 1b,c) at 11.4 degrees visual angle.

Procedure: Participants were tested individually in the lab and viewed images on a computer monitor using a fixed-distance chin rest. In Experiment 1a, participants played an economic trust game (25), a tool used in behavioral economics (26) that reliably measures trust (27). Participants were given a $100 endowment of which they could invest any portion in a trustee, whose photograph was shown as the stimulus image. The amount invested was tripled and the portion returned to the participant was selected from previously recorded actual choices of the trustees whose faces we had photographed. Participants knew this and were told that one randomly selected trial would be implemented at the very end of the experiment, and would contribute to their actual cash payout. The incentive to participants was thus to genuinely try to estimate the trustworthiness of the trustees whose faces they were shown, in order to maximize their real earnings.

In Experiment 1b, participants rated the faces on Trustworthiness, Competence, and Attractiveness on a 7-point scale (blocked by trait), and in 1c on Age, Weight, and

Distance to the camera (always rated last to avoid the possibility that explicit attention to camera distance might impact other ratings).

Participants viewed all 18 faces twice in each distance condition. Faces were presented in randomized order, but distance pairs were counterbalanced across quarters of the experiment such that half the faces were viewed first in the close condition followed by the far condition. Dollar investment amounts in studies 1 and 2 and raw ratings from all three studies were normalized to a 1-100 scale based on each participant's individual range across all faces.

Experiment 2 Methods

Subjects: N = 27, mean age = 23.93 +/- 1.09 (SEM), (17 female, 10 male), (15 Caucasian, 7 Asian, 1 Hispanic, 1 African-American, 3 Other) recruited from the local community in the same manner as Experiment 1.

Stimuli and Procedure: Photographs of faces from Experiment 1 taken at 135cm were warped to the proportions of those taken at 45cm (Figure 2b). Warping was accomplished by manually labeling 115 anatomical facial locations (including eyes, nose, mouth, ears, and outline) and interpolating using Delaunay triangulation, a standard technique for digital morphing. Thus the location coordinates of major anatomical features are exactly the same for the close faces in Experiment 1 and Experiment 2, but the luminance values are slightly different. The average 2D correlation between the pixel values of a close face in Experiment 1 and its corresponding synthetic warp in Experiment 2 is quite high (r=0.95 +/- 0.004 (SEM)), indicating that the role of these subtle luminance differences may in fact be negligible.

Participants performed the same economic trust game as in Experiment 1a.

Experiment 3 Methods

Subjects: Participants were recruited only from the United States and tested over the internet via Amazon's Mechanical Turk, permitting larger sample sizes (Experiment 3a, N = 268, 148 female; Experiment 3b, N = 70, 27 female; Experiment 3c, N = 60, 27 female; Experiment 3d, N = 253, 143 female; Experiment 3e, N = 134, 68 female).

Stimuli: Experiment 3a, 3d, 3e, and 3f all used identical stimuli as Experiment 1. Experiment 3b used only the "far" stimuli from Experiment 1, but accompanied by a verbal cue to distance before presentation indicating that the person was "standing 1.5 feet in front of you" or "standing 4.5 feet in front of you." Experiment 3c used only the "far" stimuli from Experiment 1, but adjusted the size of the image to take up the entire screen or just half of it.

Procedure: Experiments 3a,b,c obtained the same ratings as in Experiment 1b: trustworthiness, competence, attractiveness. Whereas Experiment 3a showed the identical stimuli as in Experiment 1b (strictly replicating that lab-based experiment), Experiment 3b used only the "far" faces accompanied by a verbal cue to indicate that the person was standing either near or far, and Experiment 3c showed the "far" faces at 2 different screen sizes. Experiments were administered in fixed order, 3a,b,c.

Experiment 3d obtained the same ratings as in Experiment 1c: age, weight, and distance to the camera; Experiment 3e obtained ratings of how average, and how animal-like the faces appeared. These Experiments were also administered in fixed order, 3a,d,e. See Table 1 for more information about all the experiments.

References

1. Willis, J. & Todorov, A. (2006). First Impressions: making up your mind after a 100-ms exposure to a face. Psychological Science, 17, 592-598.
2. Oosterhof, N. N., & Todorov, A. (2009). Shared perceptual basis of emotional expressions and trustworthiness impressions from faces. Emotion, 9, 128-133.
3. Stirrat, M., & Perrett, D. (2010). Valid facial cues to cooperation and trust: male facial width and trustworthiness. Psychological Science, 21, 349-354.
4. Zebrowitz, L. (1997). Reading faces: window to the soul? Boulder, CO: Westview Press.
5. Verhoff, M., Witzel, C., Kreutz, K., & Ramsthaler, F. (2008). The ideal subject distance for passport pictures. Forensic Science International, 178 (2-3): 153-156.
6. Perona, P. (2007). A new Perspective on Portraiture. Journal of Vision 7(9): article 992.
7. Hall, E. (1966). The Hidden Dimension. Garden City, NY: Doubleday.
8. Hayduk, L. (1983). Personal Space: where we now stand. Psychological Bulletin, 94, 293-335.
9. Patterson, Miles L. (1976) An arousal model of interpersonal intimacy. Psychological Review, 83(3): 235-245. doi: 10.1037/0033-295X.83.3.235
10. Vranic, A (2003) Personal Space in Physically Abused Children. Environment and Behavior, 35(4): 550-565. doi: 10.1177/0013916503035004006
11. Li, S. (2007) How far is far enough?: A measure of information privacy in terms

of interpersonal distance. *Environment and Behavior,* 39(3): 317-331.

12. Evans, G., Wener, R. (2007) Crowding and personal space invasion on the train: Please don't make me sit in the middle. Journal of Environmental Psychology, 27(1): 90-94.
13. Tajadura-Jiménez, A., Pantelidou, G., Rebacz, P., Västfjäll, D., Tsakiris, M. (2011) I-space: the effects of emotional valence and source of music on interpersonal distance. PLoS One, 6(10):e26083.
14. Evans, G., & Howard, R. (1973) Personal Space. Psychological Bulletin, 80(4): 334-344.
15. Skotko, V. & Langmeyer, D. (1977). The Effects of Interation Distance and Gender on Self-Disclosure in the Dyad. *Sociometry*, 40 (2): 178 – 182.
16. Hassin R, & Trope Y. (2000) Facing faces: studies on the cognitive aspects of physiognomy. Journal of Personality and Social Psychology, 78(5):837-52.
17. Dosey, M. and Meisels, M. (1969). Personal space and self-Protection. Journal of Personality and Social Psychology, 11(2), 93-97.
18. Kennedy, D. P., Gläscher, J., Tyszka, J. M., & Adolphs, R. (2009). Personal space regulation by the human amygdala. Nature Neuroscience, 12, 1226-1227.
19. Davis, M., Whalen, P.J. (2001). The amygdala: vigilance and emotion. *Molecular Psychiatry*, 6, 13–34.
20. Adolphs, R., Tranel, D., Damasio, H., & Damasio, A. R. (1995). Fear and the human amygdala. *Journal of Neuroscience*, 15, 5879–5891.

21. Adolphs, R., & Tranel, D. (2003). Amygdala damage impairs emotion recognition from scenes only when they contain facial expressions. *Neuropsychologia*, 41, 1281–1289.

22. Winston

23. Engell, A.D., Haxby, J.V., Todorov, A. (2007). Implicit trustworthiness decisions: automatic coding of face properties in human amygdala. *Journal of Cognitive Neuroscience,* 19, 1508–19.

24. Adolphs, R., Tranel, D., Damasio, A.R. (1998). The human amygdala in social judgment. *Nature*, 393, 470–74.

25. Berg, J., Dickhaut, J. & McCabe, K (1995). Trust, reciprocity and social history. Games and Economic Behavior, 10, 122-142.

26. Camerer, C.F. (2003). Behavioral Game Theory: Experiments in Strategic Interaction. Princeton University Press, Princeton, NJ.

27. Houser, D., Schunk, D., & Winter, J. (2006): Trust Games Measure Trust. Discussion Papers in Economics 38; http://epub.ub.uni-muenchen.de/1350.

28. Tsao, D. & Freiwald,W. (2006) What's so special about the average face? *Trends in Cognitive Science*, 10(9): 391-393.

29. Winkielman, P., Halberstadt, J., Fazendeiro, T., and Catty, F. (2006). Prototypes Are Attractive Because They Are Easy on the Mind. *Psychological Science*, 17(9): 799-806.

30. Rhodes, G. & Tremewan, T. (1996) Averageness, Exaggeration, and Facial Attractiveness. *Psychological Science*, 7 (2): 105-110.

31. Horne, N. (2006) "The influence of biological sex, sexuality and gender role on interpersonal distance." *The British journal of social psychology* 45.3: 579-597.

32. Beaulieu, C. (2004) "Intercultural Study of Personal Space: A Case Study." *Journal of Applied Social Psychology*, 34.4: 794-805. DOI: 10.1111/j.1559-1816.2004.tb02571.x

33. Adams, L. & Zuckerman D. (1991) The Effect of Lighting Conditions on Personal Space Requirements. *The Journal of General Psychology,* 118(4): 335-340.

34. Kramer RM. (2006) The great intimidators. Harvard Business Review, 84(2):88-96, 164.

Figures

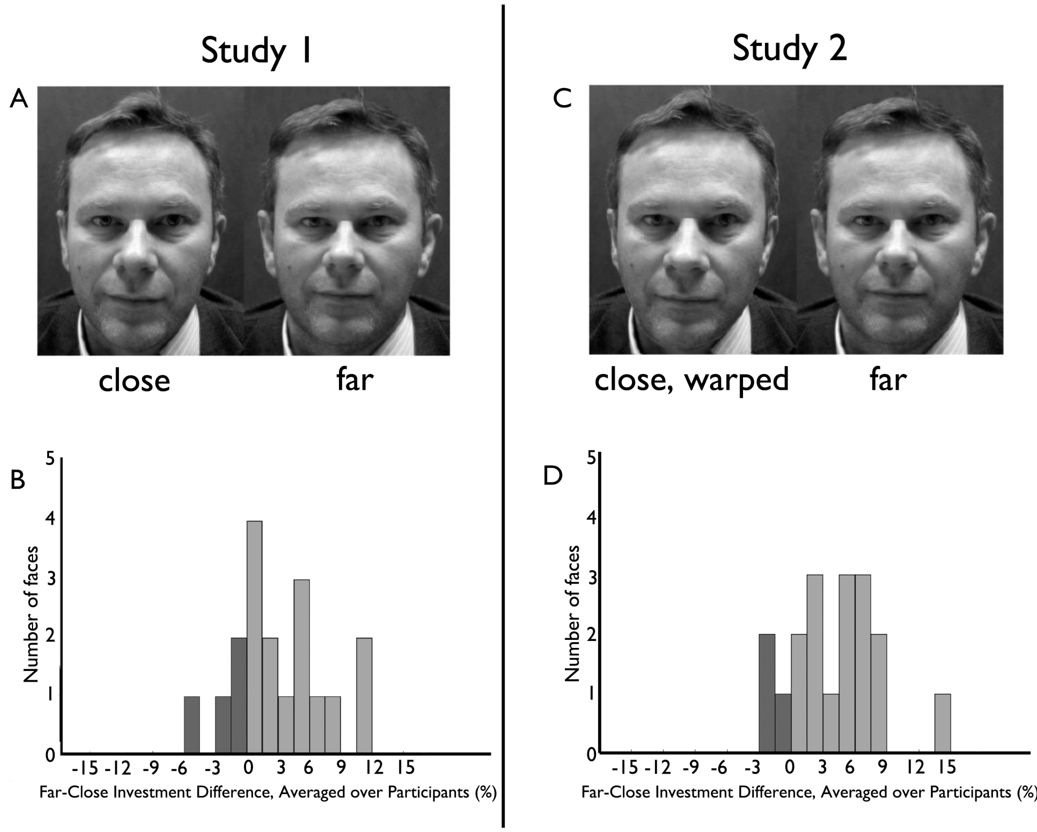

Figure 1 – Perspective distortion from distance influences trust (Experiment 1a and Experiment 2). Histograms show investment difference (far-close) for each face, averaged over all participants. A disproportionately larger number of faces received a positive investment difference (light bars) compared to those receiving a negative investment difference (dark bars).

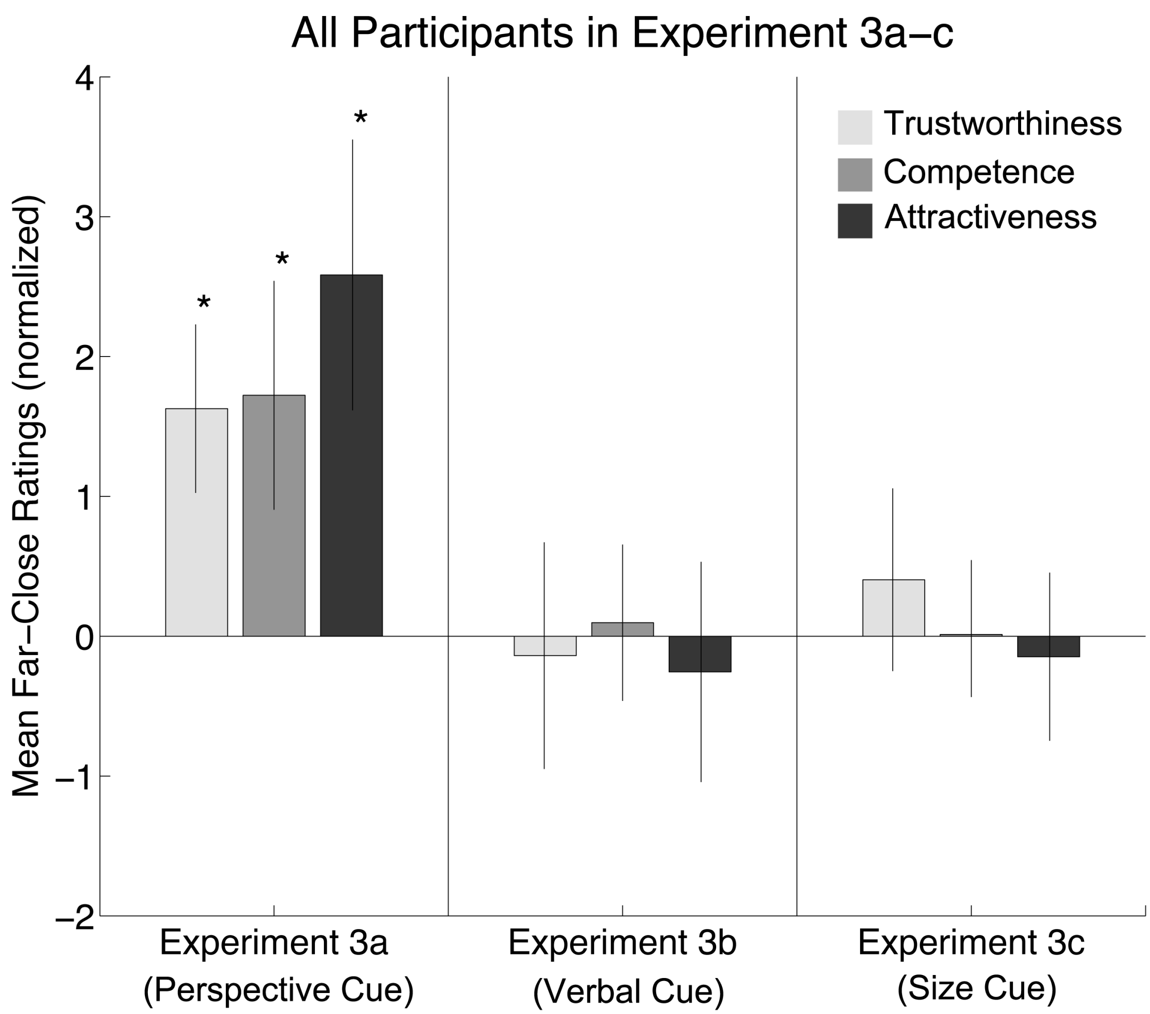

Figure 2 – Social judgments as a function of perspective distortion (Experiment 3a), verbal information (Experiment 3b), and image size (Experiment 3c). In each Experiment, ratings were obtained for Trust (solid black bars), Competence (gray bars), and Attractiveness (white bars). The mean Far-Close score over all participants and stimulus faces is shown on the y-axis (±S.E.M.)

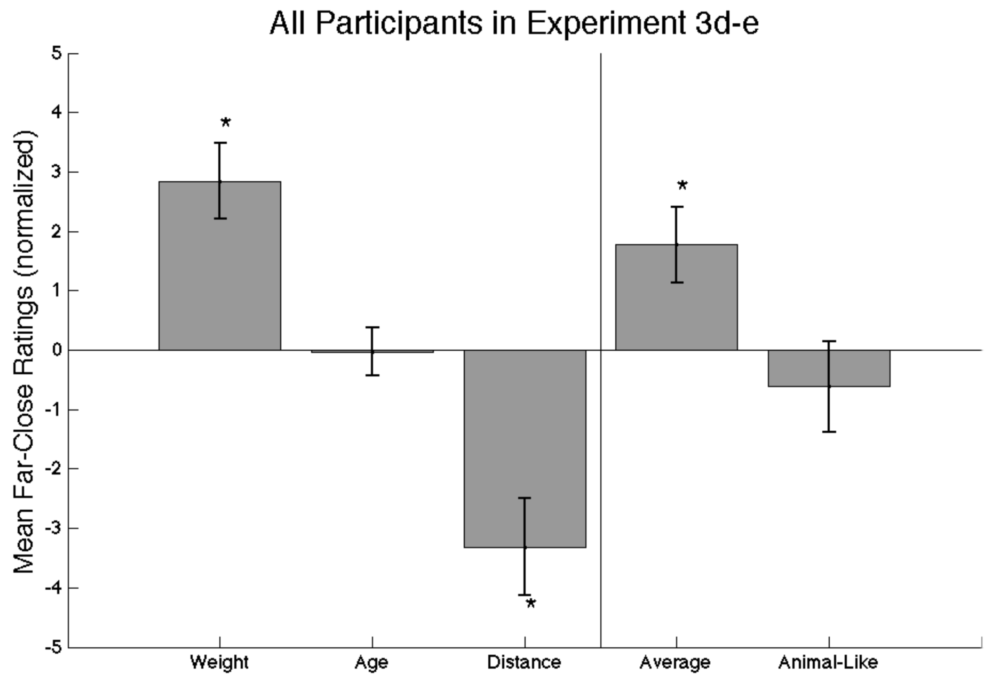

Figure 3 – Additional social judgments from perspective distortion (Experiment 3d-e). Shown are means and S.E.M. for ratings of Heaviness, Age, Distance to Camera (Experiment 3d), Averageness (Experiment 3e).

3.2.2 “Deservingness” vs “Betrayal” Extension

Although we found a fairly reliable effect of camera distance in both the ratings of "trust" and the economic trust game, it's not entirely clear what is driving the results. Of the possible mediating variables to trust, we focus on two in this next study: deservingness and betrayal risk (Bohnet, & Zeckhauser, 2004). We define deservingness to the participants as the following, after playing the economic trust game with the computer: "For this round, please assume that all of the investment partners will keep more than their fair of the money. That is, you are guaranteed to lose money by investing with them. You may, however, for whatever reason, feel like investing anyway, perhaps because you feel the person 'deserves' the money. Please invest according to how much you would not mind losing to the partner." We define betrayal risk as the following: "For this round, please try to judge how likely it is that the person will return your fair share of the investment. 50% means you are unsure, numbers higher than 50% mean you are very sure they will return your fair share, and numbers lower than 50% mean you are very sure they will keep more than their fair share." Aside from these differences in the instructions, everything else about the experiment remained the same.

Results

Both deservingness and betrayal elicited statistically higher responses for far faces than close faces, replicating the initial finding when asked simply about trust, generically. The effect sizes for both, however, are smaller and display some sensitivity to the normalization scheme chosen. Using simply the raw values from the 100 point scale the participants used to respond, the deservingness effect size is 2.0 +/- 0.81 (SEM), 95%CI

$= 3.29$, $t(20) = 2.55$, $p < 0.02$, and betrayal effect size = 1.94 +/- 0.92 (SEM), 95%CI = 3.73, $t(20) = 2.17$, $p < 0.05$. In other words, both measures found an effect of similar size, but the deservingness effect is slightly larger and more consistent across participants.

Discussion: Both trust sub-scales of Deservingness and Betrayal risk elicited slightly weaker responses from participants, indicating that there may be many mediating variables that work together to convert the visual facial input into a trust decision.

Vertical Viewing Angle Experiments: In the section that follows, we describe the results of experiments using stimuli taken at multiple different viewing angles.

Faces viewed at eye-level are perceived as more trustworthy

Ronald Bryan[1], Pranay Kothari[2], and Ralph Adolphs[1,2]

[1]Computation and Neural Systems Program

[2]Division of Biology

California Institute of Technology

Pasadena, CA 91125

Correspondence to:

Ronnie Bryan (ronnie@caltech.edu)

1200 E California Blvd M/C 136-93, Pasadena, CA, 91125

Acknowledgments: We thank Laura Loesch, Dennis Lazar, and Kathy Lawrence for help with data collection. Funded by grants from NIH and the Moore Foundation (R.A.).

Height has been associated with greater perceived physical and social dominance. The visual appearance of a face changes according to the height of the individual relative to the observer. It is likely then that humans subconsciously perceive relative height when viewing faces tilted at different vertical angles, which may influence social judgments about those faces. To test this hypothesis, we generated face stimuli by simultaneously photographing male subjects from three vertical viewing angles: *direct*, *above*, and *below*. These stimuli were tested in an economic trust game in which real money was at stake, as well as in a ratings experiment. The results provide evidence that vertical viewing angle affects how individuals make social judgments from face information.

Introduction

Height is arguably one of the most easily identifiable traits and is associated with greater financial success (Judge et al. 2004), improved health (Hebert et al. 1993), and higher reproductive success (Pawlowski et al. 2000). Given that even brief exposures to images of faces can result in reliable social judgments (Willis & Todorov, 2006), and that interpersonal distance influences these judgments (Bryan, Perona, & Adolphs, submitted) we hypothesize that such judgments should also be sensitive to the relative height of the observer.

The visual appearance of a face varies according to the height of an individual relative to that of the observer in a similar manner as changes in vertical head tilt. Perceptions of computer generated head models (with eyes closed) depend on the head tilt angle, such that faces viewed from below are seen as more dominant (Mignault & Chaudhuri, 2003). We extend this finding by testing perceptions of photographs of people with eyes open and from three viewing angles. By taking three simultaneous photographs, we were able to control for changes in facial expression, and by digitally manipulating the appearance of the eyes, we were able to control for gaze direction.

In Experiment 1, we show that economic behavior in a trust game involving real money depends on the viewing angle of the trustee, such that faces that appear to be the same height as the observer receive the greatest investment amounts. In Experiment 2, we replicate and follow up this effect by collecting ratings along multiple social dimensions, including the socially important traits trustworthiness and competence.

Experiment 1

Methods

Subjects: Healthy adult participants were recruited from the local community. N = 80 (54 female), age = 27.91 +/- 1.14.

Stimuli: To capture three simultaneous photographs from different heights, we constructed an adjustable mount for three Rocketfish 2MP AutoFocus USB Webcams. The cameras were placed 3 inches apart (equal to the standard deviation of height in the American population) and 2 feet from the subject. The middle camera was placed at eye level with the subject. We then captured images with the subject keeping their head straight from three viewpoints: *above*, *direct*, and *below* viewing angles. All facial stimuli were converted to grayscale and aligned such that the inter-ocular line was perfectly horizontal. Additionally, all stimuli were normalized for size by inscribing them within a rectangle of fixed width.

To control for gaze direction, each face model was photographed three times, once looking directly into each camera. The eye information from these photographs was then transferred to copies of the simultaneous photographs, ensuring that all the stimuli would contain an identical neutral facial expression and a direct gaze. Finally, the background information was removed from the photographs, leaving only facial information as a cue to relative height. Six different individuals had their photographs taken with this method to serve as stimuli for the experiment. All stimulus models were Caucasian males (mean age= 20.17+/-0.34 (SEM)).

Procedure: Subjects performed an economic trust game (Berg, Dickhaut, McCabe, 1995), a tool used in behavioral economics (Camerer, 2003) that reliably measures trust (Houser, 2006). They were given a $100 endowment and then rated the trustworthiness of face images by indicating an amount of that money between 5 and 100 dollars they wished to invest with the person depicted on the screen. Preliminary to this investment game, subjects were familiarized to the rules by playing 24 rounds with randomized computer partners. During this practice phase, four partners were available, each represented by a rectangle of a different color. Subjects were instructed that their investment amount would be tripled, and the partner returned a fraction of this new amount to them. Two randomly chosen rectangular partners were given a level of higher 'trustworthiness,' meaning they would usually return more money to the participant than was originally invested. The other two colored rectangles would return less than the original investment. Participants were given five seconds to make their investment decision, and then received feedback on how much their investment turned into. Using this feedback, participants were then told their goal was to determine which colored partners they would prefer to invest in over the course of 24 trials. All participants performed normally on this task, maximizing their investments with the trustworthy partners.

This task was then repeated but with the face stimuli acquired earlier replacing the computerized colored rectangles as partners in the investment game. Additionally, in the face rating round, no feedback on the investment was given, so that subjects could only rely on facial information when making their investment decisions. Participants were explicitly told beforehand that all faces would be of Caucasian males, so race and gender should not be used as factors when figuring out how much to invest. Subjects were instructed that one randomly chosen investment made during this round would be used to determine a real payout based on prerecorded responses from the face stimulus model.

Subjects participated by placing their head in a chinrest at a fixed distance from the screen, resulting in stimuli that appeared at 9.46 degrees of visual angle. Each face stimulus appeared on the screen for 5 seconds, during which time the subjects could increase or decrease their investment amount using the keyboard. At the end of 5 seconds, the investment amount displayed on the screen was recorded as their response, and the face stimulus disappeared. Each participant's data was then normalized to a 100-point scale.

The experiment order was counterbalanced such that in each block, three faces were shown from each viewing angle, and over the course of the experiment, an individual's face image would appear at all three viewing angles in different blocks. The experimental order was randomized for each subject so that no face was consistently seen from one particular viewing angle first.

Results

A 3x2 mixed ANOVA with repeated measures factor *Viewing Angle* (above, direct, and below) and between-subjects factor *Participant Gender* (male and female) reveal a significant main effect of *Viewing Angle* ($F(2)=7.40$, $p<0.001$), but neither a *Gender* ($F(1)=0.51$, n.s.). nor a *Gender x Viewing Angle* interaction was observed ($F(2)=1.96$, n.s.).

Post-hoc testing with a paired (2-tailed) t-test confirmed that the *direct* viewing angle was preferred to the *above* viewing angle; the mean investment difference was 3.84±1.65 (95% CI), ($t(79)=4.65$, $p<0.001$). A slightly weaker, but still significant, preference was observed for the *direct* viewing angle over the *below* viewing angle; the mean investment difference was 1.86±1.79 (95% CI), ($t(79)=2.07$, $p<0.05$).

The preference for *direct* over *above* viewing angles was stronger for female participants compared to male participants (mean investment difference of 4.39 vs 2.70, respectively), but post-hoc testing with a 2-sample (2-tailed) t-test found no statistical difference (mean difference 1.69±3.52 [95% CI], $t(78)=0.96$, n.s.). Males exhibited a slightly larger preference for *direct* compared to *below* viewing angle as compared to females (3.21 vs.1.21, respectively), but this effect was not statistically significant (mean difference 2.0±3.83 [95% CI], $t(78)=1.03$, n.s.). Additionally, we observed no significant correlation with the height of the participant, but a negative trend, such that taller participants exhibit a slightly weaker effect ($r = -0.19$, $p = 0.09$).

Discussion

We found that when participants made economic trust decisions in which real money was at stake, they exhibited a statistically significant preference for face images taken from a *direct* viewing angle, as opposed to a viewing angle from *above* or *below*. The effect was strongest when comparing the *direct* viewing angle to the *above* viewing angle, indicating a preference for investment partners of the same height over those who appear shorter. We also observed that *direct* compared to the *below* viewing angle received higher investments, indicating a preference for partners of the same height over those who appear taller. We also observed an initial indication that participant gender may play a role, but we do not have the statistical power in this study to be certain.

In order to further investigate the attributions that may underlie this effect, we conducted Experiment 2, which surveyed a wider set of socially relevant judgments in a more ecologically natural setting.

Experiment 2

Methods

Subjects: Healthy adult participants were recruited from the Caltech campus. N = 22 (8 female), age = 21.27 ± 0.61 (SEM).

Stimuli: The stimuli in Experiment 2 were identical to those in Experiment 1.

Procedure: Experiment 2 was conducted in a computer lab, which did not control the exact distance to the screen subjects sat on each trial. Additionally, subjects were in an environment that included many other people also participating in the experiment. The timing of Experiment 2 was slightly quicker than Experiment 1. Each face appeared for 2 seconds with the relevant trait judgment displayed above it. During the display time no response was recorded. After 2 seconds, the face would disappear, and a rating scale would appear indicating it was time to make a response between 1 and 7.

Participants rated the stimuli on several measures: Trustworthiness, Competence, Attractiveness, Dominance, Happiness, and Anger. Additionally, subjects completed a final block where they rated the Vertical Viewing Angle of the faces. The experiment was divided into three sections: first ratings of Trustworthiness, Competence, and Attractiveness, then ratings of Dominance, Happiness, and Anger, followed at the end by ratings of Vertical Viewing Angle. Each section was counterbalanced to avoid order effects in the same manner as Experiment 1, but now interleaving three trials per viewing angle condition.

Results

We analyzed each personality trait ratings with a 3x2 mixed ANOVA with repeated measures factor *Viewing Angle* (above, direct, and below) and between-subjects factor *Participant Gender* (male and female). If a significant effect was found, we follow up with a paired (2-tailed) t-test between *direct* and both *above* and *below* ratings.

For Trustworthiness ratings, we found a significant main effect of *Viewing Angle* ($F(2)=3.82$ $p<0.05$), but not *Gender* ($F(1)=1.48$, n.s.) or *Viewing Angle x Gender* interaction ($F(2)=1.16$, n.s.). Post-hoc testing confirmed that the *direct* viewing angle was preferred to the *above* viewing angle; the mean rating difference was 6.59±6.23 (95% CI), ($t(21)= 2.20$, $p< 0.05$). We did not find a statistically significant preference for *direct* as compared to *below* (mean investment difference was -0.35±5.53 (95% CI), $t(21)= -0.13$, n.s.). Finally, we observed no correlation with the height of the participant ($r = 0.05$, n.s.).

For Competence ratings, we found a weakly significant main effect of *Viewing Angle* ($F(2)=2.77$, $p<0.08$), but no effect of *Gender* ($F(1)=1.45$, n.s.) or *Viewing Angle x Gender* interaction ($F(2)=2.49$, n.s.). Post-hoc testing showed that the *above* viewing angle was rated as more competent than the *direct* viewing angle (mean difference 5.34±4.48 (95% CI), $t(21)=2.48$, $p<0.05$).

For Attractiveness ratings, we found no significant main effect of *Viewing Angle* ($F(2)=0.16$, n.s.), *Gender* ($F(1)=1.24$, n.s.) or *Viewing Angle x Gender* interaction ($F(2)=0.26$, n.s.).

For Dominance ratings, we found a weakly significant main effect of *Viewing Angle* ($F(2)=2.50$ $p<0.10$), but not *Gender* ($F(1)=0.30$, n.s.) or *Viewing Angle x Gender* interaction ($F(2)=0.22$, n.s.). Post-hoc testing revealed that the *direct* viewing angle was rated as less dominant compared to *above* (mean rating difference was 6.58±5.70 (95% CI), ($t(21)= 2.40$, $p<0.05$).

For Happiness ratings, we found no significant main effect of *Viewing Angle* ($F(2)=0.47$, n.s.), *Gender* ($F(1)=0.72$, n.s.) or *Viewing Angle x Gender* interaction ($F(2)=1.37$, n.s.).

For Anger ratings, we found a significant main effect of *Viewing Angle* ($F(2)=6.95$

$p<0.02$), but not *Gender* ($F(1)=1.77$, n.s.) or *Viewing Angle x Gender* interaction ($F(2)=0.57$, n.s.). Post-hoc testing revealed that the *direct* viewing angle was rated as less angry than the *above* viewing angle; the mean rating difference 6.60 ± 5.85 (95% CI), ($t(21)= 2.34$, $p<0.05$). We found no difference between the *direct* viewing angle and *below* (mean rating difference of 1.84 ± 4.0 (95% CI), ($t(21)= 0.95$, n.s.).

For Vertical Viewing Angle ratings, we found no significant main effect of *Viewing Angle* ($F(2)=0.089$, n.s.), *Gender* ($F(1)=0.41$, n.s.) or *Viewing Angle x Gender* interaction ($F(2)=0.020$, n.s.).

Discussion

We found that vertical viewing angle had a statistically significant effect on ratings of Trustworthiness and Anger, but not Competence, Attractiveness, Dominance, or Happiness. For Trustworthiness ratings, the results replicate those of Experiment 1, indicating that the *direct* viewing angle was rated as more trustworthy than the *above* viewing angle. This effect could be related to perceptions of anger since participants rated the *above* viewing angle as angrier than the *direct* viewing angle, despite the face that the face images were captured at the same instant, and thus had the same exact expression.

We did not find that the *below* viewing angle led to higher ratings of dominance, or that the *above* led to lower ratings of dominance, despite what might be expected from previous literature (Mignault & Chaudhuri, 2003). This is likely due to the fact that we used a much smaller angular deviation (7.12 degrees, whereas Mignault & Chaudhuri tested 10, 20, and 30 degree deviations).

Since we found similar effects for Trustworthiness and Anger, it is possible that these results are correlated. However, we did not find any significant correlation between these ratings, when averaged across face stimuli (Trustworthiness and Anger, r= -0.28, n.s). The fact that Anger ratings and Trustworthiness ratings were influenced in the same way by head tilt suggests that there may be a shared perceptual mechanism underlying both (Oosterhof, N. N., & Todorov, A., 2009).

We also found that when asked to rate the head tilt, subjects performed at chance, indicating that the deviations from a direct viewing angle used in the experiment were not so large as to be noticeable under these experimental conditions.

Experiment 3

Methods

Subjects: Healthy adult participants were recruited from Amazon Mechanical Turk, N = 16 (12 female), age = 31.94 ± 1.99 (SEM).

Stimuli: The stimuli in Experiment 3 were 8 Caucasian males who did not appear in Experiments 1 & 2; age = 31.88 ± 3.48 (SEM). The *above* and *below* cameras were placed at 3 inches from eye-level, and the subjects sat at 3 ft from the camera. The result was a smaller image, as well as a smaller angular deviation (4.76 degrees).

Procedure: Experiment 3 proceeded identically as Experiment 2, collecting ratings of various social judgments. The main difference is that subjects participated under more naturalistic testing conditions, under variable display sizes and environments.

Results

Analysis was conducted in the same manner as Experiment 2. For Trustworthiness ratings, the 3x2 mixed ANOVA not reveal a significant main effect of *Viewing Angle* ($F(2)=2.50$ $p=0.10$), *Gender* ($F(1)=0.20$, n.s.) or *Viewing Angle x Gender* interaction ($F(2)=0.047$, n.s.). However, post-hoc testing (two-tailed paired t-test) confirmed that the *direct* viewing angle was preferred to the *above* viewing angle, as in Experiment 1 & 2: the mean rating difference was 5.50± 5.44 (95% CI), ($t(15)= 2.15$, $p< 0.05$). As in Experiment 2, we did not find a statistically significant preference for *direct* as compared to *below* (mean rating difference was 2.24± 4.66 (95% CI), $t(15)= 1.03$, n.s.). Finally, as in Experiment 1, we found a trend that the preference for *direct* was weaker for the taller participants (r=--0.19, n.s.).

For Competence ratings, we found a significant main effect of *Viewing Angle* ($F(2)=5.15$, $p<0.05$), but no effect of *Gender* ($F(1)=0.029$, n.s.). We did find a significant *Viewing Angle x Gender* interaction ($F(2)= 4.00$, $p<0.05$). Post-hoc testing did not confirm the preference for *direct* to *above* or *below* viewing angles at the group level: respectively, the mean difference was 3.73±5.78 (95% CI), ($t(15)=1.37$, n.s.), and 3.91±5.14 (95% CI), ($t(15)= 1.62$, n.s.). Due to the *Viewing Angle x Gender* interaction, we also performed a post-hoc test for each gender subset to identify the direction of the trend. The small number of males in the study (N=4) exhibited a weakly significant higher preference for the *direct* compared to *above*: the mean rating difference was 13.14±17.23(95%CI), ($t(3)=2.43$, $p<0.10$). The females exhibited no trend rating *direct* as more competent than *above*: the mean rating difference was 0.58±4.87 (95% CI), (t(11)=0.22, n.s.). The males also exhibited a trend toward rating *direct* higher than *below*, with a mean rating difference of 12.38±18.42 (95% CI), (t(3)=2.14, n.s.) while the females did not, with a mean rating difference of 1.08±4.73 (95% CI), (t(11)=0.50, n.s.).

For Attractiveness ratings, we found no significant main effect of *Viewing Angle* (F(2)=1.61, n.s.), *Gender* (F(1)=2.11, n.s.) or *Viewing Angle x Gender* interaction (F(2)=2.30, n.s.).

For Dominance ratings, we found no significant main effect of *Viewing Angle* (F(2)=0.60, n.s.), *Gender* (F(1)=1.24, n.s.) or *Viewing Angle x Gender* interaction (F(2)=0.11, n.s.).

For Happiness ratings, we found no significant main effect of *Viewing Angle* (F(2)=1.20, n.s.), *Gender* (F(1)=3.08, n.s.) or *Viewing Angle x Gender* interaction (F(2)=2.52, n.s.).

For Anger ratings, we found no significant main effect of *Viewing Angle* (F(2)=0.24, n.s.), *Gender* (F(1)=2.24, n.s.) or *Viewing Angle x Gender* interaction (F(2)=0.49, n.s.).

Finally, for Vertical Viewing Angle ratings, we did find a significant main effect of *Viewing Angle* (F(2)=8.28, p<0.002), but not *Gender* (F(1)=2.34, n.s.) or *Viewing Angle x Gender* interaction (F(2)=0.93, n.s.). Post-hoc testing revealed that subjects were accurate at identifying *direct* compared to *below*, with a mean difference in the ratings of 8.66±4.52 (95% CI), (t(14)=4.11, p<0.002). Subjects also exhibited a weakly significant trend toward accurately rating *direct* compared to *above*, with a mean rating difference of -8.61±9.67 (95% CI), (t(14)=-1.91, p<0.10).

Discussion

With an independent set of stimuli and more naturalistic testing conditions, we replicated the finding of Experiment 1 and 2 that the *direct* viewing angle is rated as more trustworthy than the *above* viewing angle. Unlike Experiment 2, subjects in Experiment 3 did display accuracy in identifying the vertical viewing angle of the stimuli, but the effect was stronger for *direct* compared to *below*, the comparison that did not result in differences in trust ratings. For *direct* compared to *above*, the accuracy of the ratings was not as consistent at the group level. This dichotomy suggests the possibility that the participants who were more accurate at identifying the vertical viewing angle were more or less likely to exhibit a difference in their trustworthiness ratings. However, we found no such relationship (Pearson's $r=0.23$, n.s.). Finally, we searched for a role of participant height, but found no significant correlation between height and preference for *direct* viewing angle compared to *above*: pooling across all three experiments yields a negative trend ($r = -0.12$, n.s.).

We also found an early indication that ratings of competence may be sensitive to height in a way that interacts with participant gender. Males are more likely than females to rate faces as less competent when they are viewed from above (two sample t-test, $t(14)=2.27$, $p<0.05$) or below ($t(14)=2.30$, $p<0.05$). Since the stimuli were male faces, this opens up the possibility that a different result may hold for female facial stimuli.

As in Experiment 2, we did not find an influence of viewing angle on ratings of dominance, again most likely to due to the subtle deviations from eye-level. Unlike Experiment 2, we did not find an influence of viewing angle on ratings of anger. We are not sure why these participants would be more accurate at identifying the angle, but would not display the same perceptions of emotional expression.

General Discussion

In three experiments, under varying degrees of control over the testing conditions, and in two different stimulus sets, we found a reliable influence of vertical viewing angle on the social judgments made based on face images. In all three experiments, the *direct* viewing angle was seen as more trustworthy than the *below* viewing angle. In Experiment 1, subjects who played an economic trust game with real money invested higher amounts when shown an image of an individual taken from eye level.

In both Experiments 2 and 3, the role of vertical viewing angle was also found in a separate population who gave explicit ratings of Trustworthiness, as well as several other social judgments. While participants of Experiment 2 also showed an effect for Dominance and Anger judgments, and participants of Experiment 3 also showed an effect for Competence and Viewing Angle judgments, both groups agreed about ratings of Trustworthiness, the construct of interest in this study. The fact that we found differences between the groups who participated under different testing conditions and different stimuli indicates that there may be many avenues for further investigation.

The effect we found indicates that faces that appear to belong to shorter individuals, or whose heads are tilted slightly forward, are seen and treated as less trustworthy. Although we did not find that competence ratings universally followed the same pattern, as would be suggested by the fact that taller individuals are seen as more competent and intelligent (Judge et al, 2004), our results may provide a mechanism by which shorter individuals are penalized in economic transactions, which could explain the greater success of taller individuals in business (Judge et al. 2004) and politics (McCann, 2001). This connection is made more plausible considering that trustworthiness ratings from faces have been shown to correlate with economic parameters such as interest rates (Duarte, J., Seigel, S., & Young, L., 2009).

Future experiments may extend the generality of these findings by testing a greater diversity of stimulus classes (races, ages, and genders), and a greater range of viewing angles. The role of individual differences, such as height and gender, were inconclusive in this study, but larger studies with greater sample sizes may uncover such effects. Additionally, our results provide evidence that not only trustworthiness, but perceptions of competence and emotional expression may be influenced by viewing angle.

References

Berg, J., Dickhaut, J., McCabe, K (1995). "Trust, reciprocity and social history. *Games and Economic Behavior* 10: 122-142.

Bryan, Perona, & Adolphs (submitted). A novel cue to social judgments: perspective distortion of faces.

Camerer, C.F. Behavioral Game Theory: Experiments in Strategic Interaction (Princeton University Press, Princeton, NJ, 2003).

Duarte, J., Seigel, S., & Young, L. (2009). Trust and Credit. AFA 2010 Atlanta Meetings Paper. Available at SSRN: http://ssrn.com/abstract=1343275

Hebert, PR, Rich-Edwards, JW, Manson, JE, Ridker, PM, Cook, NR, O'Connor, GT, Buring, JE, & Hennekens, CH (1993). Height and incidence of cardiovascular disease in male physicians. *Circulation* 88:1437-1443.

Houser, D., Schunk, D., & Winter, J. (2006): Trust Games Measure Trust. Discussion Papers in Economics 38.

Judge, TA, & Cable, DM. (2004) The Effect of Physical Height on Workplace Success and Income: Preliminary Test of a Theoretical Model. *Journal of Applied Psychology* 89(3): 428-441.

McCann, SJ. (2001) Height, societal threat, and the victory margin in presidential elections (1824-1992). *Psychol Rep* 88(3): 741-2

Mignault, A & Chaudhuri, A. (2003) The many faces of a neutral face: head tilt and perception of dominance and emotion. *Journal of Nonverbal Behavior* 27(2):111-132.

Oosterhof, N. N., & Todorov, A. (2009). Shared perceptual basis of emotional expressions and trustworthiness impressions from faces. Emotion, 9, 128-133.

Pawlowski, B, Dunbart, RM, & Lipowicz, A. (2000). Tall men have more reproductive success. *Nature* 403: 156.

Willis, J. & Todorov, A. (2006). First Impressions: making up your mind after a 100-ms exposure to a face. Psychological Science 17: 592-598

Figures

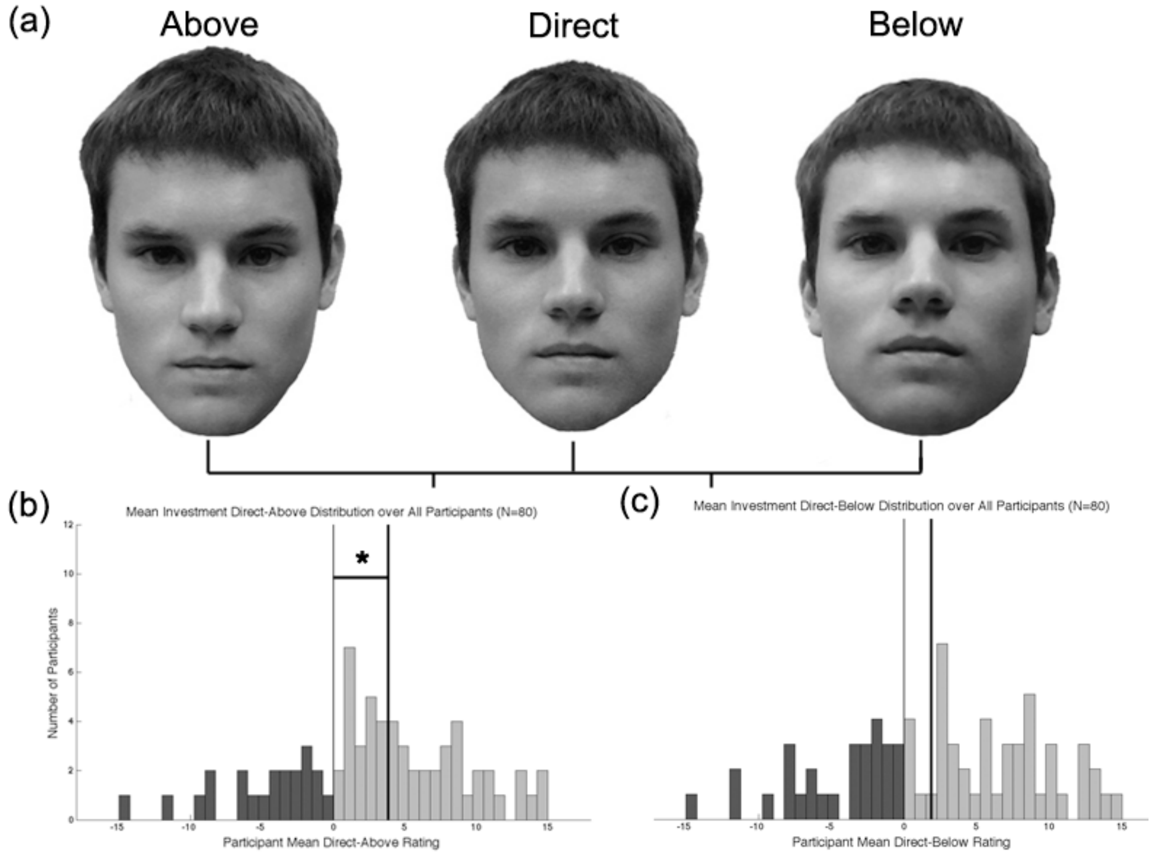

Figure 1 – Example of stimuli and results of trust game (Experiment 1)

Panel (a) shows examples of simultaneous photos with eyes adjusted to direct gaze for three viewing conditions: *above, direct, and below*. Panels (b,c) show histograms of investment differences (direct – above, in panel b; direct – below, in panel c), averaged over all face stimuli. A disproportionately larger number of participants invested more heavily in the *direct* viewing condition, as compared to *above* (light bars in panel b).

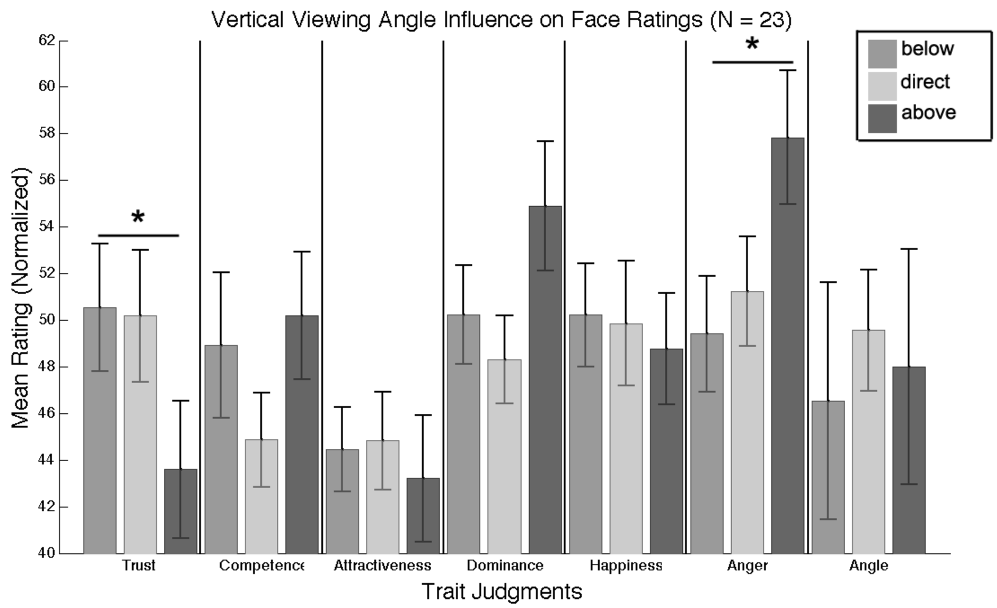

Figure 2 – Trait ratings on other social dimensions (Experiment 2). To follow up the result of Experiment 1, we collected ratings on other social judgments, and found that only Trustworthiness and Anger showed sensitivity to viewing angle.

Chapter 4: Applications of Distance Based Cues

Abstract

Previously, we have shown that the viewing distance to a face influences the image's emotional impact. In this work, we show how measurements taken on the image of a face can be used to estimate the distance from which the face was photographed, and, by extension, how the same measurements can be used to predict and manipulate perceptions of trustworthiness.

In two datasets, taken at different distances and under different conditions, we model how the geometry of the face changes with viewing distance. By training classifiers on these datasets, we achieve accurate distance and trust estimates in out-of-sample images, setting the stage for automated psychological scoring of face images. By manipulating the appearance of the faces according to the learned weights, we are able to change how they are rated by human observers, validating the methods described in this chapter.

1. Introduction

Even when controlling for size, the geometry of the image of a face changes in a predictable way according to viewing distance (Perona, 2007; Verhoff, Witzel, Kreutz, & Ramsthaler, 2008). This systematic change influences decision-making related to approach and avoidance behavior, such as economic trust (Bryan, Perona, and Adolphs, submitted). We hypothesize that by analyzing the geometry of faces in images, it should be possible to estimate distance and predict how they may influence people's perceptions of interpersonal space and trust.

A 3-D head is mapped to a 2-D image by perspective projection in a systematic way. Features closer to the camera, such as the nose, appear larger than those farther away, such as the ears. As the camera moves closer to the subject, the difference in relative magnification between these features increases. Therefore, we reason that configural information about the location of various face features should be an informative cue for predicting photographic distance.

In order to estimate the viewing distance to the stimuli used in these experiments, we will train a regressor based on a wide set of distances. The training stimulus set, which we will call Stimulus Set 1, was collected by Pietro Perona and includes 24 faces viewed at 7 distances each. The test stimulus set, which was collected by Ronnie Bryan, includes 18 faces viewed at 2 distances each, one inside and one outside the theoretical boundary of personal space proposed by Edwin Hall (1966). This chapter will discuss how the configural geometry of landmarks on face imges changes with camera distance, internal cross-validation of machine learning methods, and finally an external generalization estimating the viewing distance for the test stimuli used in the behavioral experiments. We also discuss how these landmark features are relevant to decision making about trust, and conduct a pilot experiment demonstrating the efficacy of these transformations.

2. Methods

2.1 Stimulus Collection

Stimulus Set 1 was collected by taking photographs of volunteers outdoors in natural lighting at seven camera distances (in ft): 2, 3, 4, 6, 8, 12, and 16. The distance was measured between the plane of the face and the sensor in the camera. The images were taken with a zoom lens so that the faces would approximately fill the frame. After calibrating the lens and correcting for the fact that the center of projection moves when the zoom setting is changed, the distances are corrected to (in ft): 1.76, 2.86, 3.93, 6.09, 8.23, 12.46, and 17.00. See Table 1 for a summary of these distances.

These images were then scaled to the same approximate size by cropping the face region and resampled to a common 1080x960 pixel grid. Further normalization steps are taken after facial annotation, described below.

2.2 Anatomical Annotation

Both the stimulus sets described in this chapter were annotated according to the same scheme. The location of 15 landmark features was clicked for each face at each distance, creating a 30 element descriptor for each face containing the horizontal and vertical coordinates. These anatomical features are as follows: the outer point of the left eye, the outer point of the right eye, the outer point of the left ear, the left temple, the top of the head, the top of the forehead, the outer point of the left nostril, the left extent of the mouth, the left jaw, the chin, the outer right ear, the right temple, the right nostril, the right mouth extent, and the right jaw. See Table 2 for a summary of these landmarks and Figure 1 for an example of the annotations.

The position of the eyes was the first annotation collected for each image. This first step serves two functions: to align the face to the horizontal, and to provide a face-centered coordinate system for subsequent annotation. After obtaining eye positions, each image was rotated so that the axis joining the eyes would be horizontal. Subsequent annotations are made on to this aligned image, and the midpoint between the eyes serves as the origin for the face-centered coordinate system (where x1 refers to the horizontal displacement, and x2 refers to the vertical).

We chose landmarks so as to probe points on the head that have different distances from the camera, so that we would obtain a good differential signal for depth. The farthest points on the head that are visible to the camera are the ears. Ears have different shapes; our landmark was chosen to be shape-independent: the outermost point on the ear.

The left temple is defined as the point in vertical alignment with the eyes that marks the position of the side of the head. It is between the ear and the eye.

The top of the head and the top of the forehead landmarks are chosen to be in horizontal alignment with the origin. These two points lie in horizontal alignment, but are at different distances to the camera. The top of the head is the farthest vertical extent of the head, usually the top of any hair. The top of the forehead marks the hairline; two potentially cases of ambiguity can arise. If a face has bangs that occlude the view of the hairline, it is approximated based on the vertical position of an adjacent part of the forehead. If a head is balding, the receding hairline is inspected more closely for a location that serves to approximate it.

The left nostril is the closest point to the camera among the landmarks we chose.

The annotation is defined as the midpoint of the outer nostril. The left corner of the mouth is the next landmark.

The jawline is the point on the jaw lying at a 45 degree angle from the origin. A line is drawn from the origin at a 45 degree angle below the horizontal to the end of the image. Where this line intersects the jaw, the annotation is made.

The bottom of the chin is the next annotation, which together with the top of the head and the outer extent of the ear, denotes a bounding box around the face image. If a face wears a beard, the bottom extent of the beard is chosen for this annotation.

Finally, the right side counterparts of the left side features described above are annotated as well.

For an example of how these anatomical positions vary with distance, please see Figure 1, which shows the locations of these markers on an example stimulus face for seven distances.

2.3 Preparation and Alignment

The anatomical annotations described in the previous section correspond to horizontal and vertical coordinates on the 2-D image. At the first stage of collection, these coordinates are positive values representing distances from the upper left corner (matrix convention). To remove any source of variation due to the location of the face within the image, the coordinates of each face are re-centered according to the centroid of the features. The new coordinates are now 'face-centered'. For each face, let $x1_{centroid}$ be the mean of the horizontal coordinates and $x2_{centroid}$ be the mean of the vertical coordinates. To switch from matrix convention to metric convention, we define the new coordinates as

$$x1c = x1 - x1_{centroid}$$

$$x2c = x2_{cenroid\ -} - x2$$

For the horizontal coordinates, positive values represent a greater distance to the right of the origin. For the vertical coordinates, positive values represent a greater distance up from the origin. Figure 2a shows the values for three example distances from Figure 1. The differences between these coordinates are the systematic variation we exploit for distance estimation. The differences appear quite subtle in this scale, so to better illustrate how a feature changes, we will consider only the deviation from the corresponding feature on the mean face.

To compute the mean face, all the face descriptors must be brought into alignment for scale, position, and rotation using Ordinary Procrustes Alignment (Goodall, 1991). The new face descriptors are defined as

$$x1_{norm1} = s*R_1*x1c - t$$

$$x2_{norm1} = s*R_1*x2c - t$$

where s is the scaling factor, R is the rotation matrix, and t is the translation scalar chosen to minimize the distances between each face and a reference face. Once this is done, the mean face is defined as the average of all the $x1_{norm1}$ and $x2_{norm1}$ face coordinates, $X1_{mean}$ and $X2_{mean}$. The mean face is then rotated by R_2 so that the axis between the eyes is parallel to the horizontal. Each face is now rotated one more time by R_3 to fall into alignment with the mean face, rather than the initial reference face, and the deviations from the mean now constitute the descriptors. Since the faces were initially rotated to the horizontal for annotation, the matrices R_2 and R_3 perform very small rotations.

$$x1_{norm} = R_3*x1_{norm1} - R_2*X1_{mean}$$

$$x2_{norm} = R_3*x2_{norm1} - R_2*X2_{mean}$$

The values of these coordinates for the example stimulus are shown in Figure 2b. For references, Figure 2c shows the entire stimulus set plotted on the x1-x2 axes, with the mean face highlighted in red and the closest and farthest distances of Figure 1 shown in green.

2.4 Geometry Estimation

To demonstrate the estimation problem, consider how well we can estimate the distances of a series of faces such as the ones shown in Figure 1 using the rest of the stimulus set for training. The vectors $x1_{norm}$ and $x2_{norm}$ are each 1x15 and together constitute a statistical descriptor of a single face. The test faces from Figure 1 are arranged into a matrix called $A_{testing}$, which is 7x31, where the first 15 columns are the horizontal displacements, the second 15 columns are the vertical displacements, and the final column contains all ones, for modeling a constant displacement. The seven indices specifying the identities of the test data are a 7x1 vector called testing.

$$A_{testing}=[X1(testing,:)\ X2(testing,:)\ ones(size(testing))]$$

The remaining faces are index by a 161x1 vector called training, used to organize the training data into a 161x31prediction matrix $A_{training}$.

$$A_{training}=[X1(training,:)\ X2(training,:)\ ones(size(training))]$$

If the camera distances for these training data are arranged into a 7x1 vector called ytraining, they may be approximated from the image data by the linear relation

$$y_{training} \sim A_{training} * w$$

where w is a 31x1 vector defining the linear combination of columns in $A_{training}$ that reconstructs $y_{training}$. Regression via least square error minimization is defined by setting

$$w = (A_{training}^{T} * A_{training})^{-1} * A_{training}^{T} * y_{training}$$

The vector w may be computed from the facial landmarks using the previous equation.

In order to use this training data to generalize to other faces, we define a test set of face coordinates in the same format as described above, but including all the faces. For purposes of illustration and ground-truth validation, this example will keep the seven training examples in the test set. Let $A_{testing}$ now be the 7x31 matrix containing the horizontal and vertical coordinates of the seven faces of the testing example and a column of ones. The linear estimator for these camera distances is defined as

$$y_{predict} = A_{testing} * w$$

The actual camera distances are contained in the 7x1 vector $y_{testing}$, and we define the prediction error as the difference of the prediction and the known value

$$error = y_{predict} - y_{testing}$$

3. Results

3.1 Cross Validation

Iterating the preceding analysis over the 24 identities in Stimulus Set 1 results in statistically significant correlation between the $y_{predict}$ and the $y_{testing}$ for all but one of the test identities (average $r = 0.82$, average $p=0.009$). When the predictions across all faces are combined into the single prediction for each test distance, the correlation to the actual distance is somewhat higher, due to the cancellation of errors ($r = 0.88$, $p<0.01$). The combined results of this analysis are shown in Figure 3, which also shows how the estimation is becomes saturated at greater distance. This saturation is due to the nonlinear nature of perspective projection, which requires a transformation of data to correct.

To make the problem linear, we perform a reciprocal transform on the distances in Y. The geometry of perspective projection is such that the angle subtended by a feature is tan-1(width /distance), which gives an inverse relationship to distance. Now instead of predicting the camera distance, we predict 1/(Camera Distance). The performance for each test face improves (average $r = 0.95$, average $p = 0.002$), and the combined results across all identities gives an overall correlation between the estimated distance and the true distance is nearly perfect ($r=0.99$, $p<0.0001$), meaning that the average error is equal to zero (see Figure 4).

However, as can be seen from Figure 5, the error per testing identity is likely to be quite high. Some identities are consistently over-estimated and others are consistently under-estimated. The red lines in Figure 3 correspond to identities that are consistently under-estimated, and the green lines correspond to those that are over-estimated. These systematic errors are important because they pose the question as to whether some physiognomies look more ‘distant’ and ‘trustworthy’ than others. See the section entitled **Predicting Trust Perceptions from Facial Features** for further discussion of this point.

3.2 Generalization to Stimulus Set 2

Once the regressor estimated distance from face landmarks has been competed, we can attempt generalize to another set of images. Stimulus Set 2 consists of 18 individuals photographed at two different distances: 1.5 ft and 4.5 ft. The same 15 anatomical locations are annotated as in Stimulus Set 1, generating a new matrix $A_{testing}$ that is 36x31, where the first 15 columns are the horizontal keypoint coordinates, the second 15 columns are the vertical coordinates, and the final column consists entirely of ones. This matrix $A_{testing}$ is then multiplied by the weight vector w computed through least squares error minimization on Stimulus Set 1. The resulting predictions are, as before, and are compared to the known $y_{testing}$ values stored in a 36x1 vector. Once again, the prediction error is the difference between $y_{testing}$ and $y_{predict}$.

Figure 6 shows the results of this external generalization. The predictions and the true values correlate well ($r=0.69$, $p<0.001$) for these 36 test cases. This result indicates that this method may have general applicability to frontal images of faces. As in the cross-validation experiment, some of the face distances from this generalization test are consistently over-estimated and others are consistently under-estimated. Within each test distance, the regressor gives a distribution of estimates, implying that some faces have a physiognomy that cause them to appear closer than others. For example, a face with a wider nose, flatter ears, and narrower aspect ratio will be estimated as closer to the camera than a face with a narrow nose, wider ears, and a wider aspect ratio.

3.4 Interpreting Stimulus Set 2

Since Stimulus Set 2 was used in a behavioral experiment in which camera distance was found to influence trust, we can potentially better model those trust inferences by using the apparent distance of the face (rather than simply the actual distance). Since we have seen that some face distances are systematically overestimated or underestimated, it may be the case that the facial appearance difference that give rise to these estimation errors may also influence psychological perceptions. In other words, if a face appears close even when it is far, does this impact the investment it may receive?

As we saw in Chapter 3, faces that are photographed close receive lower investments than the equivalent face photographed far. This result can be obtained through a comparison of the mean investment for far faces minus the mean investment from close faces and results in a difference of 3.2 percentage points. If instead, we consider the mean investment to faces that appear far (according to the median estimate of the classifier) minus the mean investment of those that appear close, the difference increases to 4.2 percentage points. Figure 7 shows this comparison. This result indicates that the face distance classifier may be used in place of knowledge of actual distances, and that the estimated distance may be a more reliable indicator of the perception of closeness than the actual distance. The reason this is true is because some faces appear closer than others even at the same distance, causing them to elicit lower investment.

3.4 Predicting Trust Perceptions from Facial Features

The facial landmarks used in the preceding analyses were chosen because they sampled a range of different depths on the facial surface and thus were likely to be displaced on the 2-D image as the camera distance changed, due to perspective projection. Based on prior experiments (see Chapter 3), we have reason to believe that such variation in the appearance of the face results in different psychological responses to the image. In general, we found that the farther faces (those photographed at 4.5 ft) elicited greater trust responses than the close faces (those photographed at 1.5 ft). Given that the predictor variables analyzed above showed good correlation with camera distance, and camera distance demonstrates reliable correlation with perceptions such as trustworthiness, we reason that the basis described here should also capture variation in the psychological responses related to trust.

In order to test if these predictor variables generalize beyond camera distance, we set up the learning problem in the same manner as before. The desired variable to predict now is investment in an economic trust game (described in Chapter 3), which will be defined as the 36x1 vector Y. The predictors are arranged into $A_{training}$ as the identical 36x31 matrix described above. The regression is set up to find the weighting of the facial features w , which are then applied to $A_{testing}$ to get internal cross-validation on the predictive value of these feature combinations.

To get a better sampling of Stimulus Set 2, which is smaller than Stimulus Set 1 (for which we don't have equivalent behavioral data for training), we use iterative resampling. In 100 iterations, Stimulus Set 1 is divided into 32 training faces and 4 test faces. Thus the size of $A_{training}$ is 32x31 and the size of $A_{testing}$ is 4x31. The performance on each iteration, as well as the feature weights themselves, are stored for further analysis.

After 100 iterations, the average correlation between the predicted investment and the actual investment was found to be r = 0.60. This gives hope that this basis set can be used for predicting psychological responses. Figure 8 shows the data plotted behind the 100 regression lines from the iterations. Figure 9 shows a comparison of the weights learned from training on distance and the weights learned from training on investment. The correlation between the two sets of weights is essentially zero, indicating that though the basis is useful for predicting investment, it is a different combination of features that is most useful.

3.5 Manipulating Trust Perceptions from Facial Features

There are two strategies we can use in this basis set to manipulate the perception of trust: manipulate the distance (using the weights learned from Stimulus Set 1) or manipulate the facial proportions associated with trust (using the weights learned from Stimulus Set 2).

For an example of how to manipulate the distance of a photograph, see Chapter 3, Experiment 2. In that procedure, the configural geometry of the close face is applied to the pixel values of the far face to elicit the psychological response of the close face.

In order to manipulate trust directly, we must use the learned relationship between the feature positions and the investments, which is encoded by the weights, w. These weights reflect the magnitude and direction of how each facial feature influences investment in the trust game, so to change how a face will be invested in, it is necessary to add or subtract these weights from the face vector

$$face_{adjusted} = face + factor_{shift} * w$$

Where $factor_{shift}$ is defined as a positive or negative value that controls how much warping should occur in the learned direction. See Figure 10 for examples of how facial appearance can be altered to produce different psychological reactions. For simplicity, we have used a reduced basis set to construct these new faces, symmetrically warping them from the learned associations of the first 10 keypoints (instead of warping right and left independently).

To test if the manipulations to facial appearance resulted in the predicted psychological perceptions, we conducted a small experiment. Each of the 18 faces in Stimulus Set 2 was adjusted to elicit greater investment or lesser investment. Participants (N=23, age=38.6 +/- 3.05 SEM, 10 female) from Amazon Mechanical Turk rated these faces along three dimensions of interpersonal inference: Trustworthiness, Competence, and Attractiveness. When these three intercorrelated measures of valence are combined, the result is that faces adjusted to receive greater investments receive 2.15 percentage points higher valence ratings (paired t-test, $t(22)=2.5$, $p<0.02$). This result confirms that the feature set used to classify distance and investment of face images can also be used to manipulate the perception of those images.

4. Conclusion

We have shown that the statistical regularity in the configural geometry of face images can be used to predict the viewing distance on a stimulus set consisting of images taken at seven different distances and generalized to a stimulus set consisting of images taken at two different distances. This first step demonstrates the tractability of the problem in a low dimensional space of 15 facial locations, and leads the way to automated classification of distance in media images. Since camera distance is found to influence psychological variables such as trustworthiness, such information could be applied in a variety of fields to predict human reactions to facial images. We have shown that the estimated distance to a face is a more sensitive measure of psychological response than the actual distance. We have also shown that this basis set can be used to predict trustworthiness perceptions from facial measurements, regardless of any distance information. As a demonstration of the method, we have used this basis to manipulate the appearance of faces, and thus the psychological reactions to them. The methods described in this chapter could have wide applicability, from using the distance or trust score as a covariate for psychological studies, to estimating and manipulating the appearance of facial images in media and advertisement.

Tables

Image #	Monopod Distance (ft)	Focal Length (mm)	Effective Distance (ft)
1	2	28	1.76
2	3	50	2.86
3	4	65	3.93
4	6	100	6.09
5	8	130	8.23
6	12	180	12.46
7	16	300	17

Table 1. Stimulus Set 1 Photographic Distances.

Annotation #	Anatomical Feature
1	left outer eye
2	right outer eye
3	outer left ear
4	left temple
5	top of head
6	top of forehead
7	left nostril
8	left mouth extent
9	left jaw
10	chin
11	outer right ear
12	right temple
13	right nostril
14	right mouth extent
15	right jaw

Table 2. Anatomical Annotations. For both Stimulus Set 1 and Stimulus Set 2, each face's anatomy was recorded as 15 keypoint features.

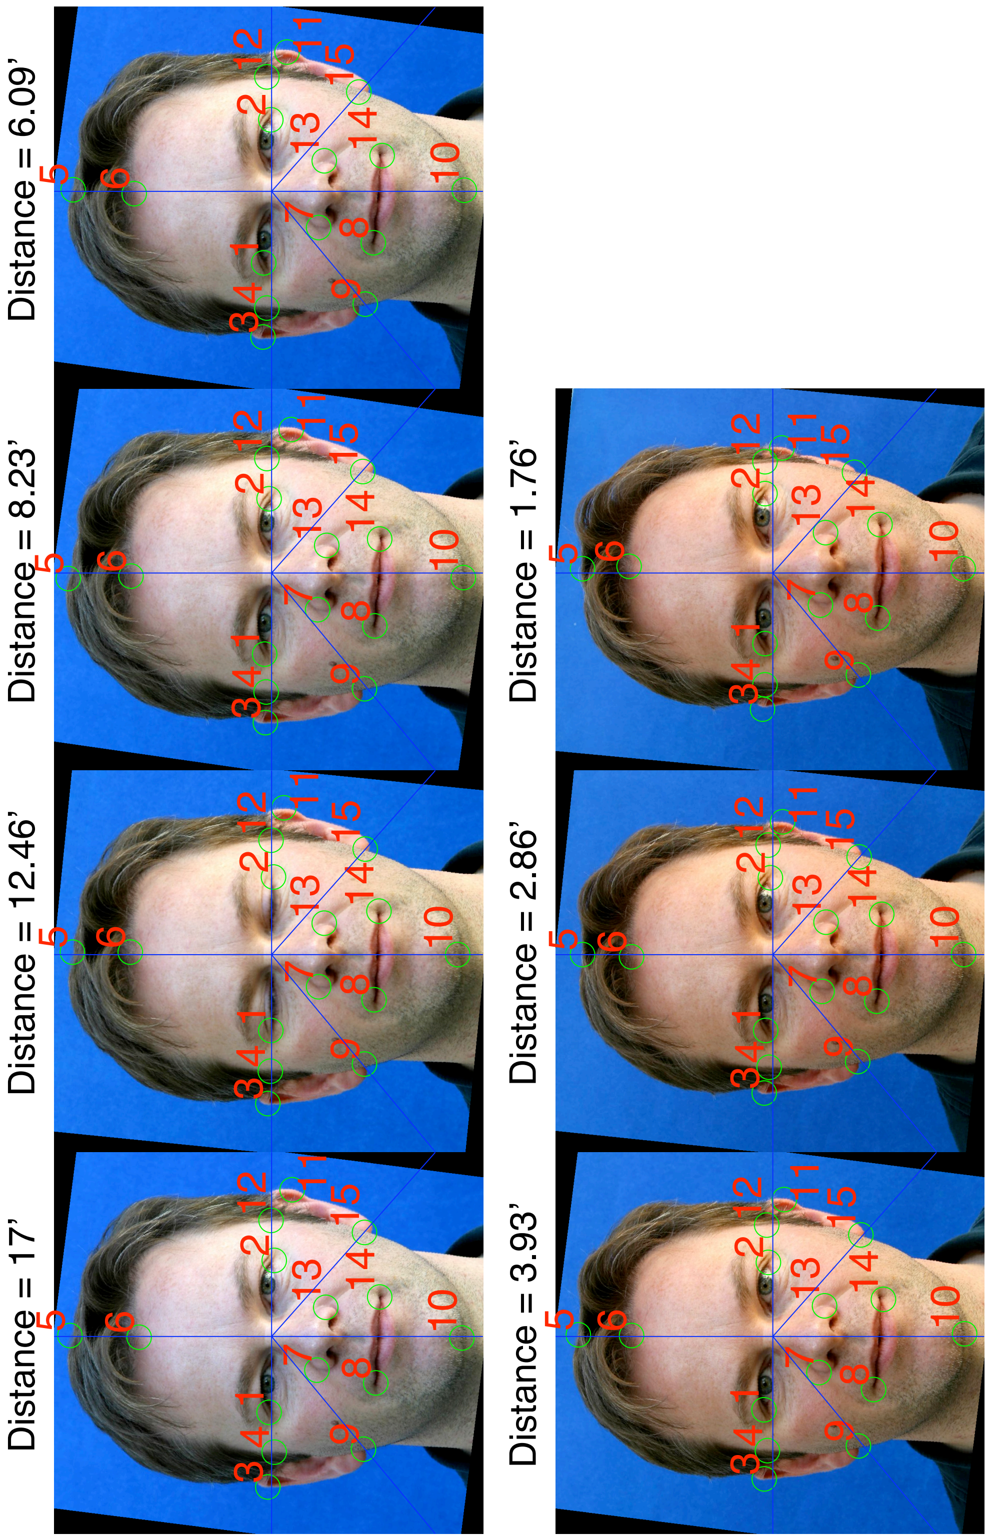
Distance = 17’
Distance = 12.46’
Distance = 8.23’
Distance = 6.09’
Distance = 3.93’
Distance = 2.86’
Distance = 1.76’

1. Example Stimuli and anatomical keypoints

Photographs of the same individual taken at seven different distances (labeled above the image). The faces are rotated so the eyes lie on the horizontal and are scaled to the same size. Fifteen keypoints are labeled to show the features used for classification.

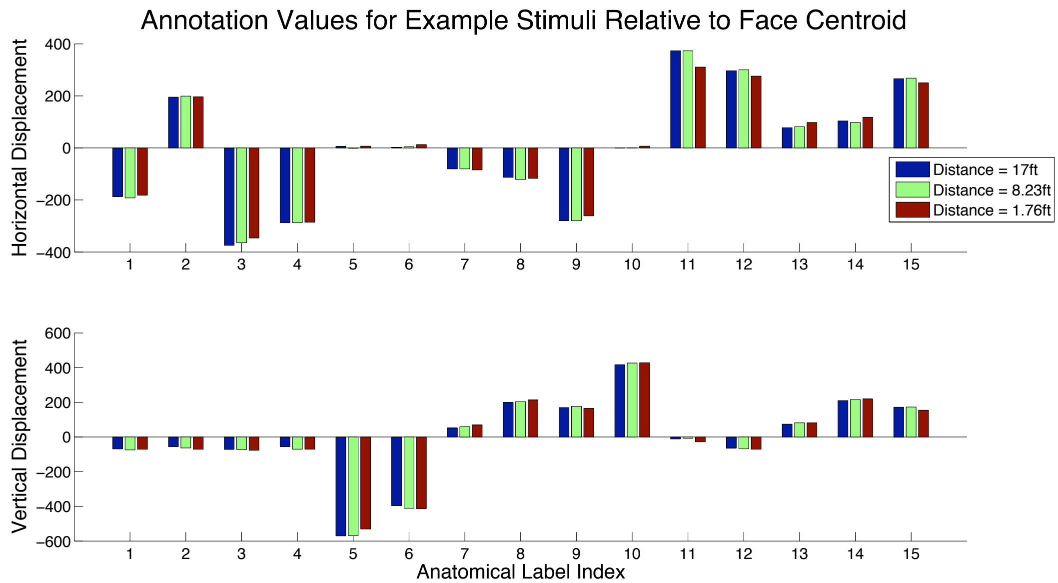

2a. Example Stimuli keypoint coordinate values relative to centroid of features

Raw magnitudes of the keypoint labels. The differences between the blue (1.76 ft), green (3.93 ft), and red (17 ft) are the signal we exploit for classification.

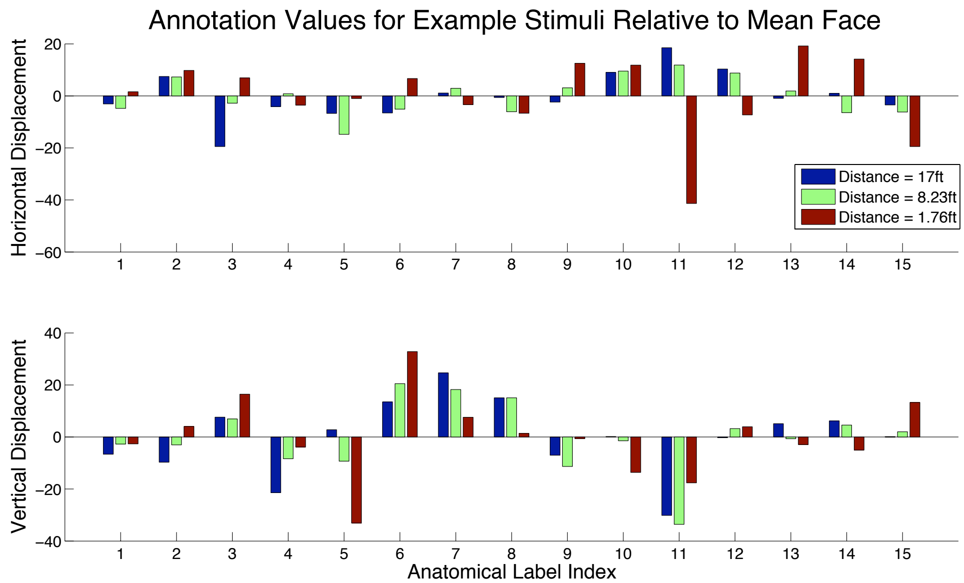

2b. Example Stimuli keypoint coordinate values relative to mean face

Shown are the same coordinate values subtracted by the mean of all the other faces in the stimulus set. These values better illustrate the signal we are searching for.

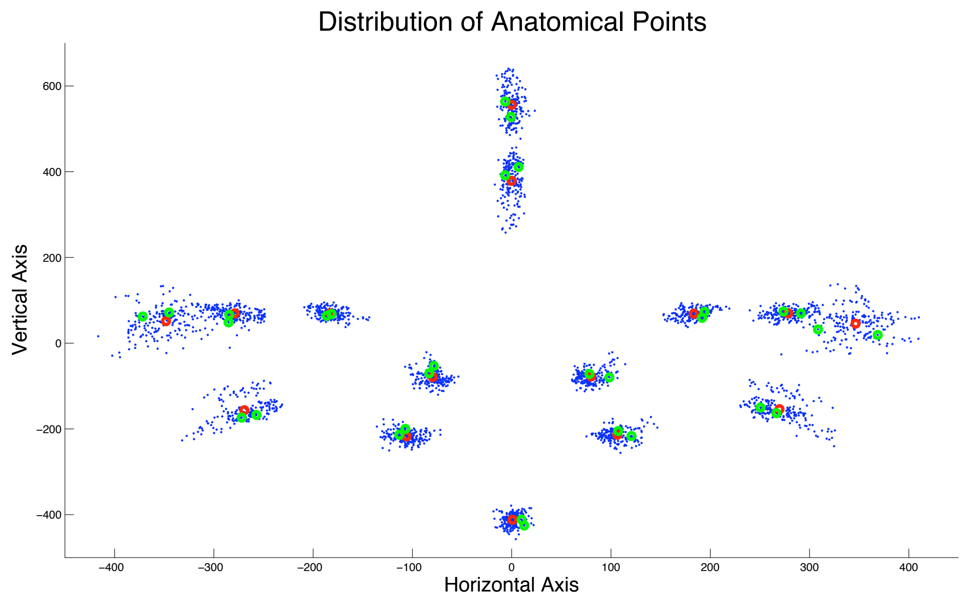

2c. Distribution of Anatomical Keypoints

All faces in the stimulus set are shown above, with the mean face indicated by the red points. The green points show two distances of the example face from Figure 1.

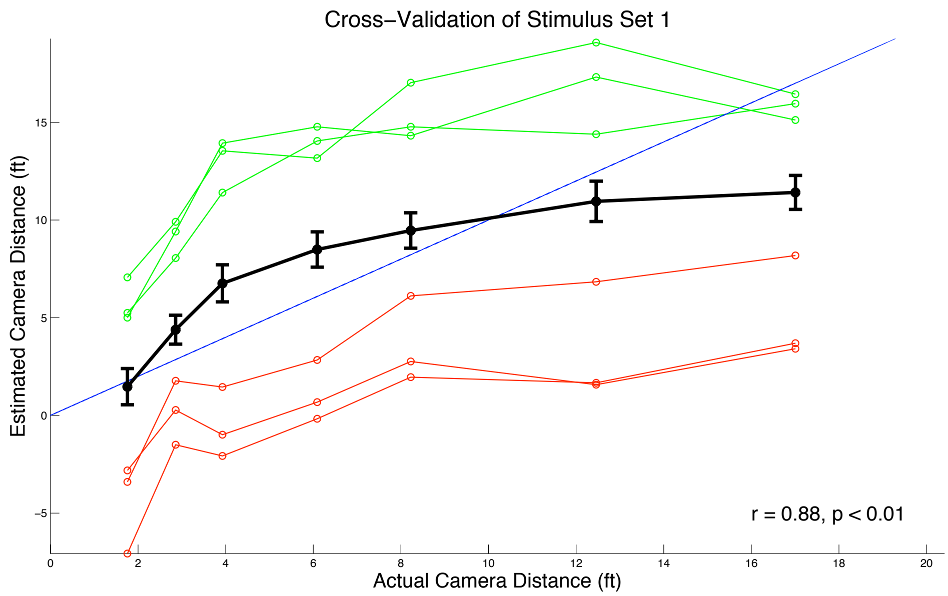

3. All Faces Classification Results

The average classification results when all faces (excluding the test face) are used to train the classifier (error bars represent Standard Deviation). . Red lines are examples of faces that are under-estimated and green lines are examples that are over-estimated.

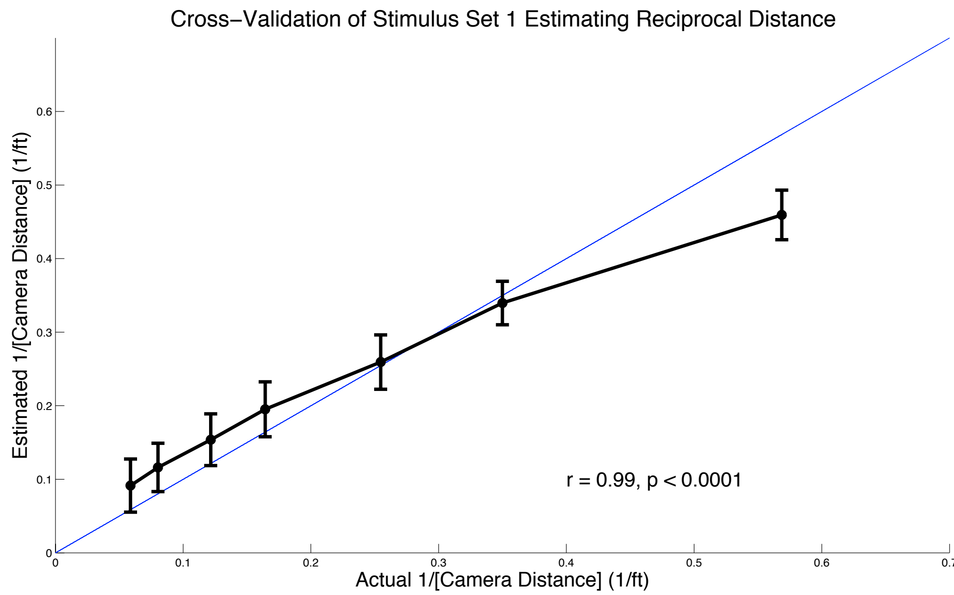

4. Estimating Reciprocal of Distance

Classification is improved when the problem is made linear by taking the reciprocal of distance instead of distance. The reason this helps is because perspective projection influences appearance nonlinearly, so the data must be transformed for linear analysis. (Error bars represent Standard Error)

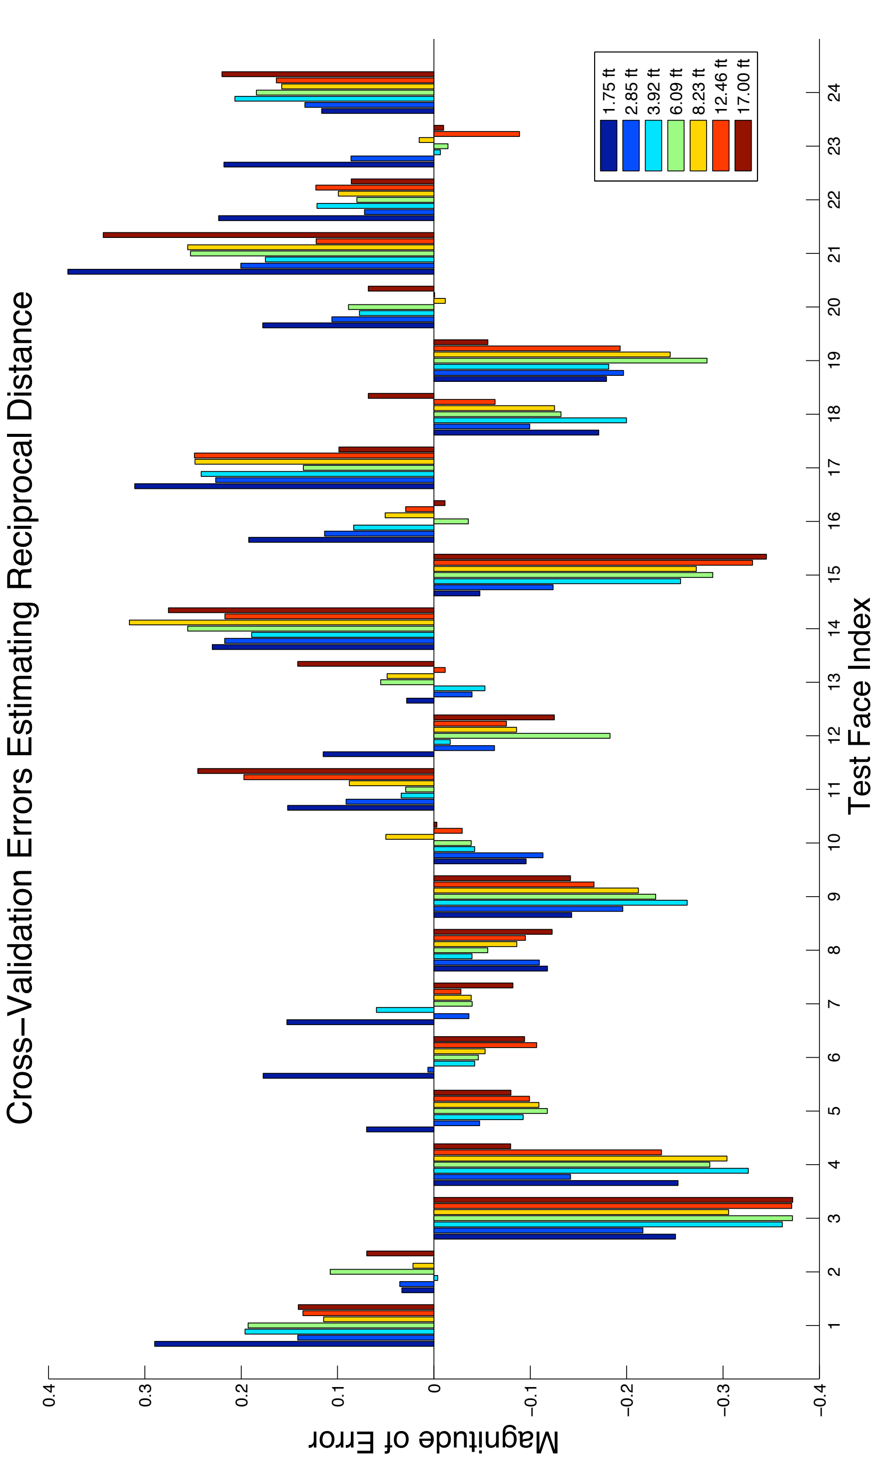

Cross–Validation Errors Estimating Reciprocal Distance
Magnitude of Error
Test Face Index
1.75 ft
2.85 ft
3.92 ft
6.09 ft
8.23 ft
12.46 ft
17.00 ft

5. Distribution of Errors

Classification errors for each face for all distances (when that face is excluded from training and serves as a test data point). Some faces are systematically higher and some are systematically lower.

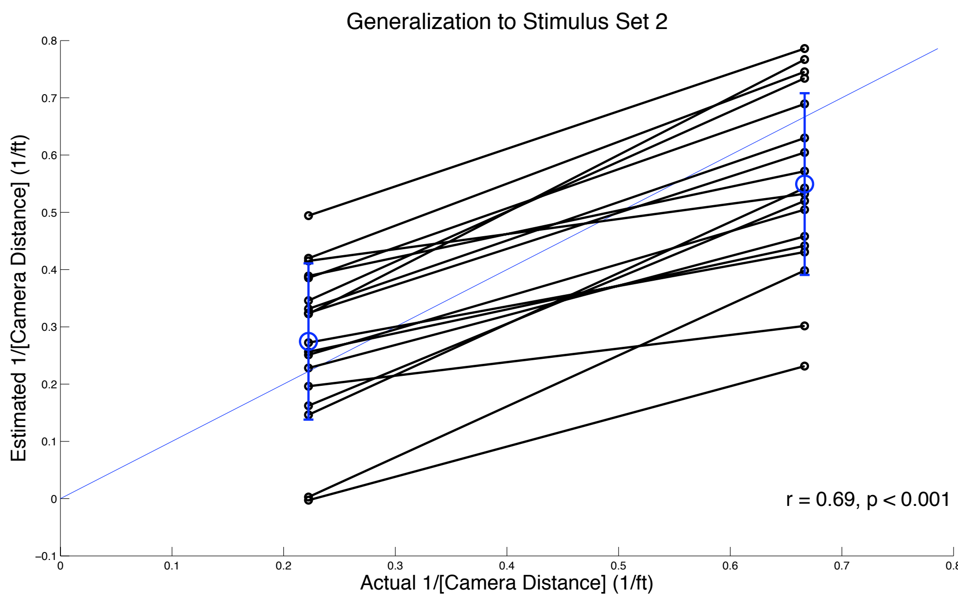

6. Generalization to Stimulus Set 2

The classifier is tested with a generalization stimulus set collected under different conditions and at different distances. As with the training set, some faces are consistently over-estimated and others are consistently under-estimated.

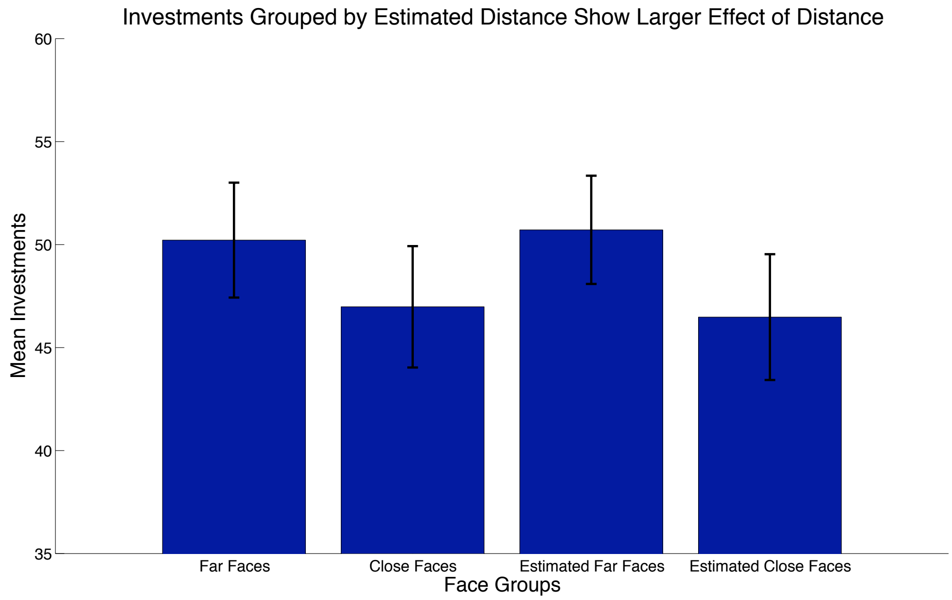

7. Using Estimated Distance to Predict Psychological Responses

The behavioral data described in Chapter 3 showed a dependency on the camera distance. Due to the fact that some faces appear closer than others, even at the same distance, we reason that the psychological responses should depend more heavily on apparent distance than actual distance. This result is verified in the data: by using the classifier estimates of distance, we see a stronger influence on behavior than the actual distance (difference for actual distance = 3.2; for estimated distance = 4.2).

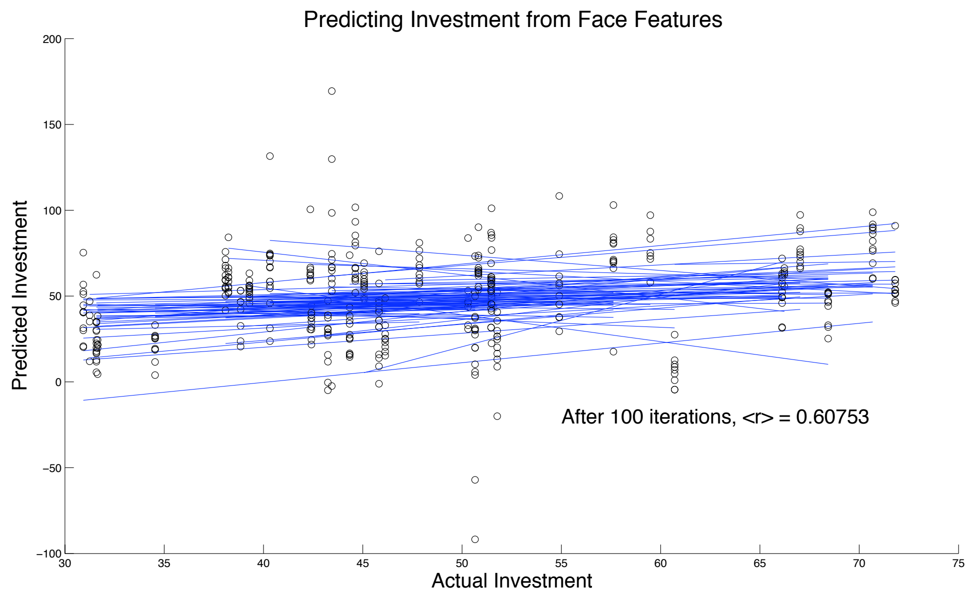

8. Predicting Behavioral Responses

Using the same feature set that classified camera distance, we now train the classifier on the behavioral responses themselves. Iterative resampling of the data is used to produce multiple regression estimates (shown in blue). The average correlation between investment and predicted investment over 100 iterations was r=0.60

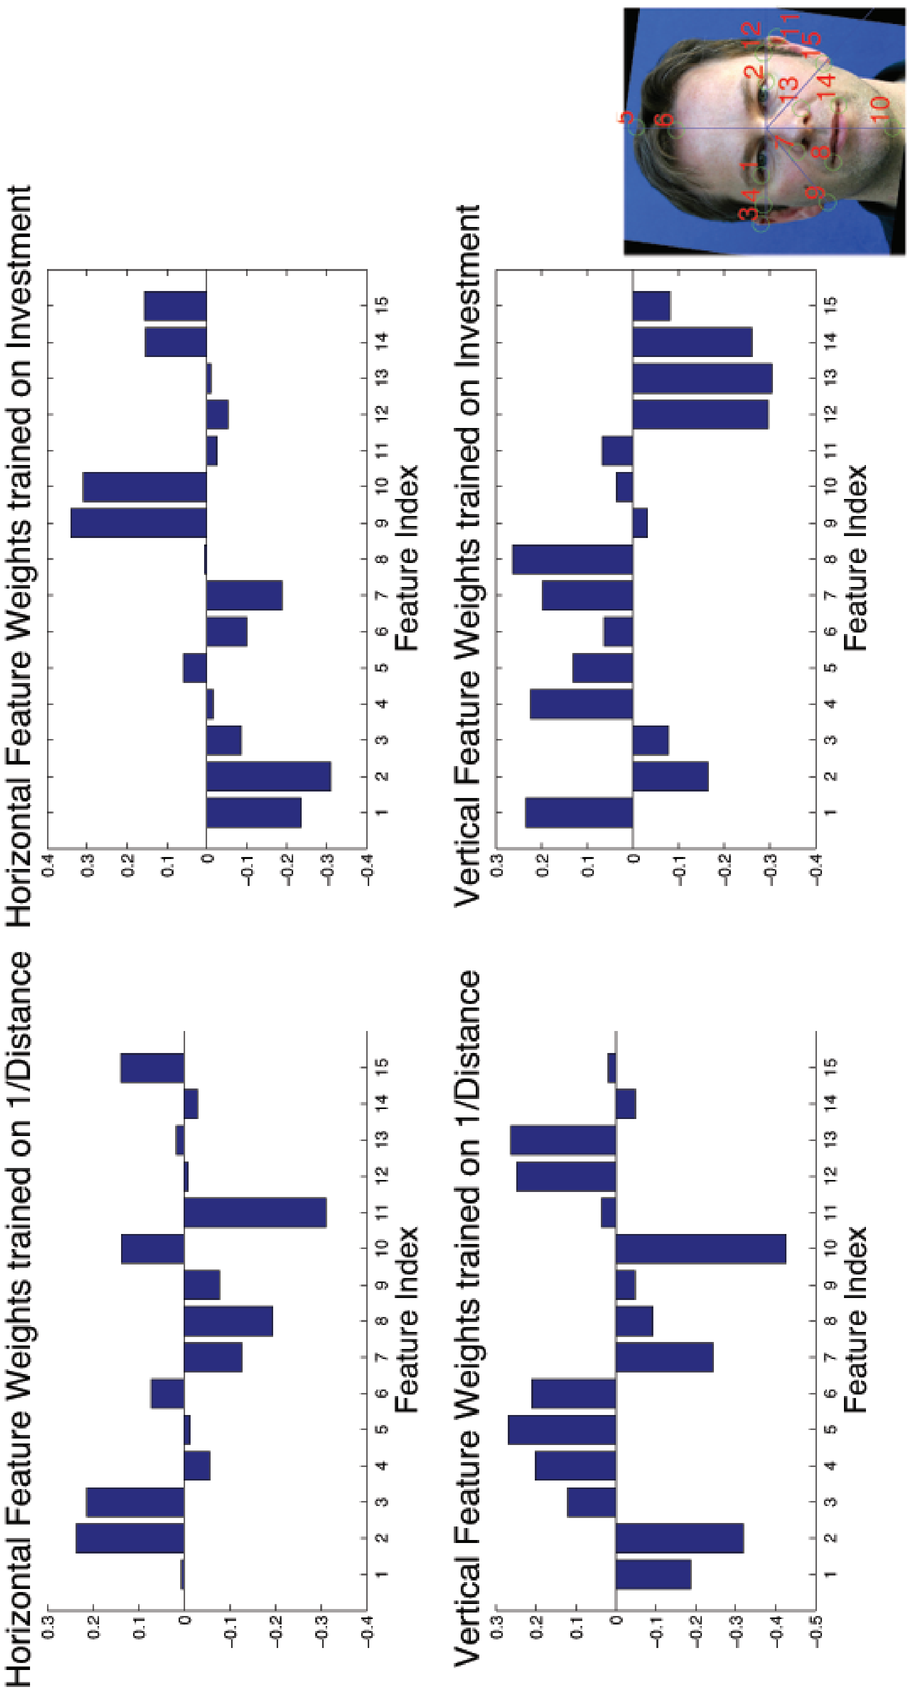
Horizontal Feature Weights trained on 1/Distance
Horizontal Feature Weights trained on Investment
Vertical Feature Weights trained on 1/Distance
Vertical Feature Weights trained on Investment
Feature Index
Feature Index
Feature Index
Feature Index

9. Comparison of Learned Weights for Distance and Investment

The same features that predict distance also can be used to predict investment, but the feature weightings are very different.

10a. Increasing the Trustworthiness of an Example Stimulus

10b. Decreasing the Trustworthiness of an Example Stimulus

Using the learned association between feature position and trustworthiness, we generated a warp to manipulate the appearance of the face in a way that our model predicts observers will find more trustworthy.

Chapter 5. Discussion

Summary

In this thesis, I have examined the role perspective projection and vertical viewing angle of facial images play in influencing decisions about interpersonal trust. I devised two methods of simultaneous photography for eliminating the confound of facial expression in the acquisition of the stimuli necessary to perform the experiments, designed and conducted a series of studies to investigate which perceptions and behaviors are affected by the manipulations, and have suggested a statistical method for how the brain or a computer might quickly process the photographic distance of facial images. The result is an end-to-end framework for investigations into a new field of psychological research that links behavioral economics to the sociology of interpersonal spacing via cognitive science of face perception and emotional processing.

In Chapter 2, I discussed two novel photographic methods for simultaneously acquiring images from multiple locations. The challenge for simultaneous capture of two images along the same optical axis is that the closer camera will necessarily occlude the view of the farther camera. I solved this problem by placing a half-silvered mirror at a 45 degree angle between the portrait subject and one of the cameras. The other camera was then positioned to capture the reflected image from the mirror. By reflecting, resizing, and normalizing the color, I produced images that differ only in the geometric configuration of the facial features due to perspective distortion.

Similarly, to address the role of vertical spacing in determining perceptions of face images, I also devised a method with the help of SURF student Pranay Kothari for acquiring multiple images along a vertical axis. The challenge for this method was in avoiding the confounding factor of direct vs averted gaze. Previous studies that addressed this issue used stimuli with eyes closed, but we felt that it was worth the effort to use open-eye stimuli to match the viewing conditions of the horizontal spacing experiment. Therefore, we employed digital manipulation techniques to transfer the direct gaze from one photograph onto another. This produced a set of images that displayed the exact same emotional expression and direct gaze, but differed in the vertical viewing angle.

Both of these techniques create the illusion of time being frozen to allow for camera movement and produce images that previously were not physically possible. By isolating emotional expression as a visual cue, we were able to investigate the role of vertical and horizontal displacement on the perception of facial images.

In Chapter 3, I present two papers that use the stimuli described in Chapter 2 to investigate how interpersonal displacements influence the trait attributions made from face stimuli. The growing literature on automatic trait attribution suggests that face perception can influence real world behavior such as voting and financial decision-making. Since faces are such a ubiquitous stimulus class, it is important to characterize all the factors that affect how our behavior is influenced by their appearance. Previous accounts focused on properties of the face itself such as the expression or the shape of the head. These approaches are completely valid, but are limited to drawing inferences from the stimulus itself and so neglect the relationship of the observer to the stimulus.

We found both horizontal and vertical manipulations of viewing position influenced attributions of trust, but not other attributes. Images taken within personal space were seen as less trustworthy, as were images taken from above eye-level. These findings were consistently found across six experiments in different subject groups and form compelling corroboration of the idea that social spacing not only matters, but its perception is activated through manipulation of the accompanying visual transformation. Furthermore, most of our participants were unaware of the fact that the faces differed in appearance from trial to trial, indicating that not only does this manipulation matters, but also it often escapes conscious awareness.

An interesting facet of this discovery is that it didn't have to be the case. Since both vertical elevation and horizontal spacing are easily described with more overt cues such as position and size, it may not have been necessary at all for our brains to learn the more subtle differences in appearance. One could imagine doing quite well in the world judging distance to faces just based on their size, or judging the height of a conspecific just based on the vertical position of their head. However, whether one considers it necessary or not, the fact of the matter is that facial appearance does systematically change with social spacing and so it forms an available channel of information from which to make inferences about the three dimensional position of a person relative to oneself. These percepts have been 100% correlated for the vast majority of the history of our perceptual faculties. However, since the advent of photography and magnification optics, human brains have been increasingly exposed to images that disentangle the intertwined visual cues of size and perspective projection. As video and photography continue to dominate the visual landscape of our environment, it is only more likely that the manipulations discussed in this thesis will play a role in shaping human experience.

In Chapter 4, I provide an initial framework for the automated labeling of face images according to viewing distance. By analyzing the two-dimensional pattern of displacement of important anatomical locations as they vary with distance, we can predict the distance to a face from a single image. Without knowledge of the person's three-dimensional head shape, this method is likely to produce errors of over-estimation or under-estimation. Future research can correct this by choosing an appropriate training set for each individual face to be tested. Chapter 4 also presented a method for adjusting a face's feature position based on the learned weights to the investment that face received and validated the method with a pilot experiment showing increased valence ratings compared to the face shifted in the opposite direction.

I suggest the use of these tools not only to control images used in psychological research, but also for characterizing the perspective projections that are commonly used in visual media. It may be the case that our subconscious reaction to images viewed from different distances has shaped the selection of viewing distances commonly used. I anticipate that advertisers and political campaigns will use tools such as these to optimize their presentation to the public.

Further Investigation

To better understand the relevance of viewing distance as it might be applied in the world, it is important to further the psychological research started here. This thesis presents studies that are intentionally very narrow in scope: Caucasian males viewed frontally wearing neutral emotional expressions and gazing directly into the camera. I have conducted pilot studies showing that cross-cultural factors may influence the role viewing distance plays. It seems likely that as race, gender, age, and other demographic traits are varied across both the participants and the stimulus classes, it is likely that significant interactions will be found between these factors.

Since a proposed psychological mechanism for trust reduction due to physical proximity is threat assessment, further investigation should vary parameters likely to influence the perception of threat. Two factors immediately come to mind: familiarity and emotional expression. It is known that as faces become more familiar, the neural representation changes and the behavioral attitudes toward them changes as well. Therefore, if participants become accustomed to the faces in the study, it is possible that the magnitude of the effect will decrease. Similarly, if participants are personally familiar with the faces, it is also possible that the magnitude of the effect will decrease. These effects, however, might be reversed, depending on the relationship to the viewer. If participants receive negative information or associations, then it is possible that the adverse reaction to a personal space invasion would increase. These effects are then likely to be mediated by the emotional expression and gaze direction of the stimuli as well.

Any feature that is found to be psychologically relevant should be accompanied by a statistical technique for assessing the trait in images. Many such classifiers for gender, expression, and other face traits are already available. These may be combined with the technique described here to define a factorial model of how trust is contingent on demographics, familiarity, and expression.

Another important area in which to extend this research is into a wider range of distances. The distances chosen for this thesis are motivated by Hall's sociological theory of personal space, but Hall also defined other threshold distances, including social space and public space. I chose to restrict this research into just the threshold between social space and personal space because this is the threshold with the clearest psychological predictions and has the greatest affect on the appearance of the face. Therefore, it was a suitable candidate for establishing proof of principle and best practices for further research. As camera technology improves, it is likely that images will be taken from farther and farther away, producing faces images that approach the fully orthographic projection. It should be investigated how the effect we describe in this thesis varies with longer camera distances.

This thesis investigated a binary distinction between close and far images that were taken across the threshold of personal space, but future research should investigate distance as a continuous quantity. By using the morphing technique described in Chapter 3, MDP Experiment 2, it is possible to adjust the parameter of viewing distance by interpolating between long and short distances. By interleaving these images in a psychophysical staircase design, we should be able to find the ideal viewing distance for each individual face. Once each face's ideal distance is determined, this information can be fed into a learning algorithm to try to predict the optimal viewing distance for new faces.

This thesis used only static, grayscale images to test trait attributions, but in reality, approach/avoidance behavior is highly dynamic and involves feedback. Therefore, a worthwhile extension of this research would be to show videos of faces approaching or retreating, with and without the added cue of perspective distortion. It would be our hypothesis that the size manipulation alone will not as strongly activate representations about personal space violation.

Other extensions to this work would include non-face stimuli, including whole body representations, objects, people in context with background, and a host of other manipulations to increase the ecological relevance of the cue we isolated here.

Conclusion

It is my hope that the work described in this thesis will just be a starting point for future research into the relationship between interpersonal spacing, perception, and decision-making. I have suggested some specific projects above, but there are many others to be undertaken. The field provides a test-bed of experiments about evolutionary theory, development, and neural representations. Are animals sensitive to this kind of perspective cue? When do children develop sensitivity? How are face representations in the brain dependent on viewing distance?

Applications of this work are varied, from the analysis of marketing and campaign materials, the selection of camera positions for film and television, to the consultation for optimal photographs for personal ads. People are intensely social creates and are naturally drawn toward each other. Hopefully this research sheds a new light on the extent to which our brains have developed to facilitate this process.

References

Adams, L. & Zuckerman D. (1991) The Effect of Lighting Conditions on Personal Space Requirements. *The Journal of General Psychology,* 118(4): 335-340.

Adolphs, R. (2004). Emotional Vision. *Nature*, 7(11): 1167-1168.

Adolphs, R., & Tranel, D. (2003). Amygdala damage impairs emotion recognition from scenes only when they contain facial expressions. *Neuropsychologia*, 41, 1281–1289.

Adolphs, R., Tranel, D., Damasio, A.R. (1998). The human amygdala in social judgment. *Nature*, 393, 470–74.

Adolphs, R., Tranel, D., Damasio, H., & Damasio, A. R. (1995). Fear and the human amygdala. *Journal of Neuroscience*, 15, 5879–5891.

Antonakis, J. & Dalgas, O. (2009) Predicting Elections: Child's Play. *Science*, 323 (27): 1183.

Argyle, M., & Dean, J. (1965). Eye-Contact, Distance, and Affiliation. *Sociometry*, 28 (3): 289 – 304.

Atkinson, M., Enos, R., and Hill, S. (2009) Candidate Faces and Election Outcomes: Is the Face-Vote Correlation Caused by Candidate Selection? *Quarterly Journal of Political Science*, 4: 229-249.

Bar, M., Neta, M., and Linz, H. (2006). Very First Impressions. *Emotion*, 6 (2): 269 – 278.

Bohnet, I. & Zeckhauser, R. (2004) Trust, Betrayal, and Risk. *Journal of Economic Behavior & Organization*, 55(4): 467-484.

Bryan, R., Perona, P., & Adolphs, R. (in prep). A novel cue for social judgments of faces: perspective distortion from interpersonal distance.

Beaulieu, C. (2004) "Intercultural Study of Personal Space: A Case Study." *Journal of Applied Social Psychology*, 34.4: 794-805. DOI: 10.1111/j.1559-1816.2004.tb02571.x

Berg, J., Dickhaut, J. & McCabe, K (1995). Trust, reciprocity and social history. *Games and Economic Behavior*, 10, 122-142.

Camerer, C.F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press, Princeton, NJ.

Chiao, J., Iidaka, T. Gordon, H., Nogawa, J., Bar, M., Aminoff, E., Sadato, N., and Ambady, N. (2008) Cultural specificity in amygdala response to fear faces. *Journal of Cognitive Neuroscience*, 20(12), 2167-2174.

Davis, M., Whalen, P.J. (2001). The amygdala: vigilance and emotion. *Molecular Psychiatry*, 6, 13–34.

Diamond, R. & Carey, S. (1986) Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, 115(2): 107-117. doi: 10.1037/0096-3445.115.2.107

Dosey, M. and Meisels, M. (1969). Personal space and self-Protection. *Journal of Personality and Social Psychology*, 11(2), 93-97.

Duarte, J., Seigel, S., & Young, L. (2009). Trust and Credit. *AFA 2010 Atlanta Meetings Paper*. Available at SSRN: http://ssrn.com/abstract=1343275

Engell, A., Haxby, J., Todorov, A. (2007). Implicit trustworthiness decisions: automatic coding of face properties in human amygdala. *Journal of Cognitive Neuroscience,* 19, 1508–19.

Engell, A, Todorov, A., and Haxby, J. (2010). Common neural mechanisms for the evaluation of facial trustworthiness and emotional expressions as revealed by behavioral adaptation. *Perception*, 39: 931 – 941.

Evans, G., Wener, R. (2007) Crowding and personal space invasion on the train: Please don't make me sit in the middle. *Journal of Environmental Psychology*, 27(1): 90-94.

Evans, G., & Howard, R. (1973) Personal Space. *Psychological Bulletin,* 80(4): 334-344.

Freiwald, W., Tsao, D., and Livingstone, M. (2009). A face feature space in the macaque temporal lobe. *Nature Neuroscience*, 12 (9): 1187 – 1198.

Gauthier, I., Skudlarski, P., Gore, J., & Anderson, A. (2000) Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, 3: 191-197.

Gobbini, I., Koralek, A., Bryan, R., Montgomery, K., and Haxby, J. (2007). Two Takes on the Social Brain: A Comparison of Theory of Mind Tasks. *Journal of Cognitive Neuroscience*, 19 (11): 1803 – 1814.

Goodall, C. (1991). Procrustes Methods in the Statistical Analysis of Shape. *Journal of the Royal Statistical Society. Series B (Methodological)* , Vol. 53, No. 2 (1991), pp. 285-339

Gross, C. & Schonen, S. (1992) Representation of Visual Stimuli in Inferior Temporal Cortex. *Phil. Trans. R. Soc. Lond. B* 29, 335(1273): 3-10.

Hall, E. *The Hidden Dimension*. Garden City, NY, (Doubleday, 1966).

Hansen, P. C. (1987). The truncated SVD as a method for regularization. *BIT Numerical Mathematics,* 27(4): 534–553. doi:10.1007/BF01937276.

Hansen, J. & Schuldt, J. (1982). Physical distance, sex, and intimacy in self-disclosure. *Psychological Reports*, 51 (1): 3 – 6.

Hassin R, & Trope Y. (2000) Facing faces: studies on the cognitive aspects of physiognomy. *Journal of Personality and Social Psychology,* 78(5):837-52.

Haxby, J., Bryan, R., and Gobbini, I. (2006). The representation of mammalian faces in human cortex. *Journal of Vision*, 6 (6): article 86.

Haxby, J., Hoffman, E., and Gobbini, I. (2002). Human Neural Systems for Face Recognition and Social Communication. *Biological Psychiatry*, 51: 59 – 67.

Haxby, J., Gobbini, M., Furey, M., Ishai, A., Schouten, J., & Pietrini, P. (2001) Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science* 28(293): 2425-2430. DOI: 10.1126/science.106373

Hayduk, L. (1983). Personal Space: where we now stand. *Psychological Bulletin*, 94, 293-335.

Hebert, PR, Rich-Edwards, JW, Manson, JE, Ridker, PM, Cook, NR, O'Connor, GT, Buring, JE, & Hennekens, CH (1993). Height and incidence of cardiovascular disease in male physicians. *Circulation* 88:1437-1443.

Horne, N. (2006) "The influence of biological sex, sexuality and gender role on interpersonal distance." *The British journal of social psychology* 45.3: 579-597.

Houser, D., Schunk, D., & Winter, J. (2006): Trust Games Measure Trust. *Discussion Papers in Economics* 38; http://epub.ub.uni-muenchen.de/1350.

Johnson, C., & Dabbs, J. (1976). Self-Disclosure in Dyads as a Function of Distance and the Subject-Experimenter Relationship. *Sociometry*, 39 (3): 257 – 263.

Jourard, SM. & Friedman, R (1970). Experimenter-subject "distance" and self-disclosure. *Journal of Personality and Social Psychology*, Vol 15(3): 278-282. doi: 10.1037/h0029388

Judge, TA, & Cable, DM. (2004) The Effect of Physical Height on Workplace Success and Income: Preliminary Test of a Theoretical Model. *Journal of Applied Psychology*

89(3): 428-441.

Kanwisher, N., McDermott, J., & Chun, M. (1997) The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *The Journal of Neuroscience*, 17(11): 4302-4311.

Kennedy, D. P., Gläscher, J., Tyszka, J. M., & Adolphs, R. (2009). Personal space regulation by the human amygdala. *Nature Neuroscience*, 12: 1226-1227.

Kramer RM. (2006) The great intimidators. *Harvard Business Review,* 84(2):88-96, 164.

Li, S. (2007) How far is far enough?: A measure of information privacy in terms of interpersonal distance. *Environment and Behavior,* 39(3): 317-331.

Mattes, Spezio, Kim, Todorov, Adolphs, and Alvarez, M. (2010) Predicting Election Outcomes from Positive and Negative Trait Assessments of Candidate Images. *Political Psychology*, 31 (1): 41-57.

Marr, D. (1982). *Vision. A Computational Investigation into the Human Representation and Processing of Visual Information.* W.H. Freeman and Company.

McCann, SJ. (2001) Height, societal threat, and the victory margin in presidential elections (1824-1992). *Psychol Rep* 88(3): 741-2.

Mignault, A & Chaudhuri, A. (2003) The many faces of a neutral face: head tilt and perception of dominance and emotion. *Journal of Nonverbal Behavior* 27(2):111-132.

Moghaddam, B. & Yang, M. (2002). Learning Gender with Support Faces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24 (5): 707 – 711.

Oosterhof, N. & Todorov, T. (2008) The functional basis of face evaluation. *Proceedings of the National Academy of Sciences*, 105 (32): 11087 – 11092.

Oosterhof, N., & Todorov, A. (2009). Shared perceptual basis of emotional expressions and trustworthiness impressions from faces. *Emotion*, 9, 128-133.

Patterson, M. (1976). An arousal model of interpersonal intimacy. *Psychological Review*, 83(3): 235-245. doi: 10.1037/0033-295X.83.3.235

Pawlowski, B, Dunbart, RM, & Lipowicz, A. (2000). Tall men have more reproductive success. *Nature* 403: 156.

Perona, P. (2007). A new Perspective on Portraiture. *Journal of Vision,* 7(9): article 992.

Phung, S. & Bouzerdom, A. (2007). A Pyramidal Neural Network for Visual Pattern Recognition. *IEEE Transactions on Neural Networks*, 18 (2): 329 – 343.

Rule, N., Freeman, J., Moran, J., Gabrieli, J., Adams, R., and Ambady, N. (2010). Voting behavior is reflected in amygdala response across cultures. *Social Cognitive Affective Neuroscience*, 5: 349 – 355.

Said, C. P., Dotsch, R., & Todorov, A. (2010). The amygdala and FFA track both social and non-social face dimensions. *Neuropsychologia, 48,* 3596-3605.

Said, C., Sebe, N., & Todorov, A. (2009). Structural resemblance to emotional expressions predicts evaluation of emotionally neutral faces. Emotion, 9, 260-264.

Schonemann, P.H. (1966), "A generalized solution of the orthogonal Procrustes problem", *Psychometrika* **31**: 1–10.

Serre, T., L. Wolf, S. Bileschi, M. Riesenhuber and T. Poggio. (2007) Robust Object Recognition with Cortex-like Mechanisms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29, 3, 411-426, 2007.

Skotko, V. & Langmeyer, D. (1977). The Effects of Interation Distance and Gender on Self-Disclosure in the Dyad. *Sociometry*, 40 (2): 178 – 182.

Spezio, Rangel, Alvarez, O'Doherty, Mattes, Todorov, Kim, and Adolphs, (2008) Stirrat, M., & Perrett, D. (2010). Valid facial cues to cooperation and trust: male facial width and trustworthiness. *Psychological Science*, 21, 349-354.

Tajadura-Jiménez A, Pantelidou G, Rebacz P, Västfjäll D, Tsakiris M. (2011) I-space: the effects of emotional valence and source of music on interpersonal distance. *PLoS One*, 6(10):e26083.

Tsao, D., Freiwald, W., Tootell, R., and Livingstone, M. (2006). A Cortical Region Consisting Entirely of Face-Selective Cells. *Science*, 311 (5761): 670 – 674.

Tsao, D., Moeller, S., and Freiwald, W. (2008). Comparing face patch systems in macaques and humans. *Proceedings of the National Academy of Sciences*, 105 (49): 19514 – 19519.

Turk, M. & Pentland, A. (1991). Eigenfaces for Recognition. *Journal of Cognitive Neuroscience*, 3 (1): 71 – 86.

Viola, P. & Jones, M. (2004) Robust Real-Time Face Detection. *International Journal of Computer Vision*, 57 (2): 137 – 154.

Vranic, A (2003) Personal Space in Physically Abused Children. *Environment and Behavior*, 35(4): 550-565. doi: 10.1177/0013916503035004006

Wheeless, L. & Grotz, J. (1977). The measurement of trust and its relationship to self-disclosure. *Human Communication Research*, 3 (3): 250 – 257.

Winston, J., Strange, B., O'Doherty, J., and Dolan, R. (2002). Automatic and intentional brain responses during the evaluation of trustworthiness of faces. *Nature Neuroscience*, 5 (3): 277 – 283.

Wiskott, L., Fellous, J., Kruger, N., von der Malsburg, C. (1997). Face recognition by elastic bunch graph matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19 (7): 775 – 779.

Williams, T., Pepitone, M., Christensen, S., Cooke, B., Huberman, A., Breedlove, N., Breedlove, T., Jordan, C., Breedlove, S (2000) Finger-length ratios and sexual orientation. *Nature*, (404)6777: 455-456.

Wong, E., Ormiston, M., and Haselhuhn, M. (2011). A Face Only and Investor Could Love: CEOs' Facial Structure Predicts Their Firms' Financial Performance. *Psychological Science*, 22 (12): 1478 – 1483.

Zebrowitz, L. (1997). *Reading faces: window to the soul?* Boulder, CO: Westview Press.

Zebrowitz, L., & Montpare, J. (2008). Social Psychological Face Perception: Why Appearance Matters. *Social and Personality Psychology Compass*, 2 (3): 1497 – 1517.

Zebrowitz, L., & Montpare, J. (1992). Impressions of babyfaced individuals across the lifespan. *Developmental Psychology*, 28 (6): 1143 – 1152.