Perception, 2000, volume 29, pages 885 – 891

DOI:10.1068/p2851

On the other side of the mean: The perception of dissimilarity in human faces

Volker Blanz, Al.ice J O'Toole¶, Thomas Vetter, Heather A Wild¶
Max-Planck-Institut für biologische Kybernetik, Spemannstrasse 38, D 72076 Tübingen, Germany;
e-mail: volker.blanz@tuebingen.mpg.de; ¶School of Human Development, The University of Texas at Dallas, Richardson, TX 75083-0688, USA
Received 12 October 1998, in revised form 10 September 1999

Abstract. We created a 'face space' using a laser-scan representation of faces. In this space, a caricature can be made by moving a face away from the average face, along the line connecting the particular face to the average face. Here, we move the face along this line in the other direction, proceeding through the mean and 'out the other side'. This results in a face that is 'opposite', in a computational sense, to the original face. We morphed several faces into their anti-faces and sampled the morph trajectory in five discrete steps. We then collected similarity ratings from human participants for all possible pairs of morphed faces to determine how the distances in the 'physical face space' related to the distances in the 'psychological face space'. The data indicate that there is a perceptual discontinuity of face identity as the face crosses over to the 'other side of the mean'. We consider these results in the context of face-space models of human face processing.

1 Introduction

Conceptualizing human representations of faces in a multidimensional space is a common practice in both psychological (Valentine 1991) and computational (Hancock et al 1996; O'Toole et al 1993; Sirovich and Kirby 1987) models of face recognition. This representation has been used widely to account for various empirical findings in the face literature, including the well-known 'other-race effect' (O'Toole et al 1994) and the recognition accuracy advantage for distinctive versus typical faces (eg Light et al 1979). In its abstract form, a face-space representation entails the following: (a) faces are thought of as points in a multidimensional space (or equivalently as 'face vectors' originating from the average face at the origin of the space); (b) the axes of the space represent the feature dimensions with which faces are encoded; and (c) the distance between any two faces is considered a measure of the similarity between the faces.

In its more concrete form, two kinds of face spaces can be created. *Psychological face spaces* can be derived by applying multidimensional scaling (MDS) analysis to face similarity ratings gathered from human observers. This produces a spatial representation of the perceptual 'distances' between faces in a multidimensional space. The interpretation of the axes is usually made by looking at faces that cluster at different ends of the axes and assessing the dimension or feature that is maximally contrasted by faces at opposite ends of the axis (eg Busey 1998).

Physical face spaces can be derived by applying a principal component analysis (ie metric MDS) to physical-similarity data computed on a particular 'physical' encoding of a face (eg an image or three-dimensional surface model). The interpretation or meaning of the axes or principal components (PCs) is again left up to the researcher. However, because the analysis is carried out on a physical representation of the face, it is possible to synthesize faces at opposite ends of the axis in a way that isolates only the single feature captured by the axis. Using both images and the separable components of three-dimensional shape and two-dimensional reflectance maps from laser scanned faces, single PCs have been shown to contrast relatively global features for specifying face gender (O'Toole et al 1993, 1997a).

In addition to the individual feature axes, other trajectories through a physical face space have also proved interesting and psychologically meaningful. For example, automatic caricature generators enhance the distinctiveness of a particular face by 'moving it' away from the average face in a physical face space (Benson and Perrett 1991; Brennan 1985). This move is made along the line connecting the face to the average face. Figure 1a illustrates a schematic face space in which the relative positions of a face and its caricature are indicated. Anti-caricatures, or less distinctive versions of individual faces, can be made in an inverse fashion by 'moving' faces closer to the average face, again along the line connecting the face to the average. This is illustrated in figure 1b. Psychologists have made extensive use of both caricatures and anti-caricatures for studying the effects of face distinctiveness on a variety of tasks (eg Benson and Perrett 1991, 1994; O'Toole et al 1997b; Rhodes et al 1987; Rhodes and McLean 1990; Rhodes and Tremewan 1994, 1996). A detailed review of these psychological findings is beyond the scope of the present paper, but can be found in Rhodes (1996).

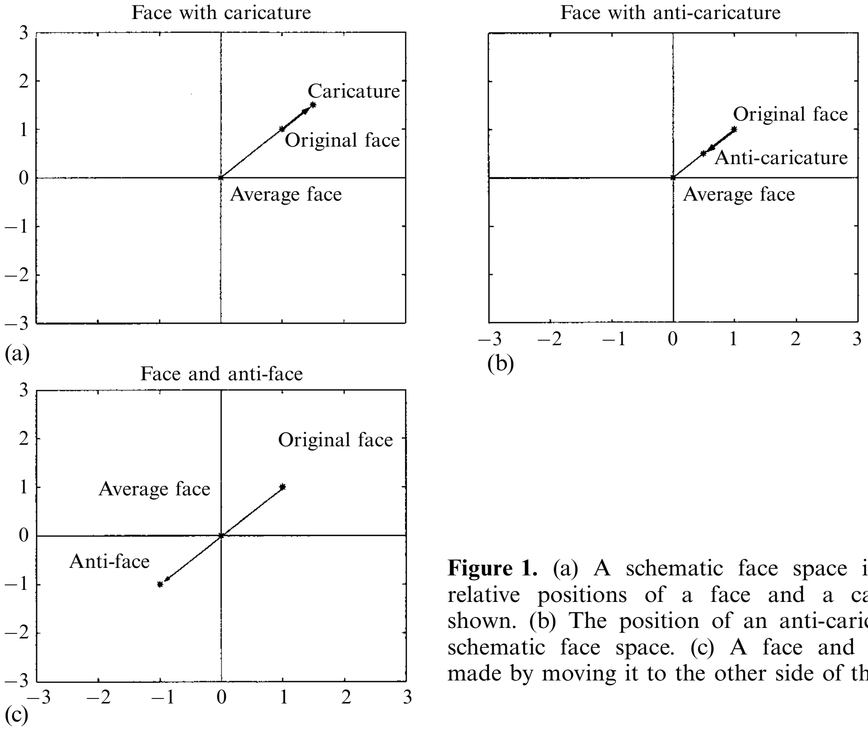

Figure 1. (a) A schematic face space in which the relative positions of a face and a caricature are shown. (b) The position of an anti-caricature in the schematic face space. (c) A face and its anti-face, made by moving it to the other side of the mean.

In summary, distinctiveness in the face-space metaphor is thought to relate to the distance of the face from the average face (Light et al 1979; Valentine 1991), whereas identity is thought to relate to the direction of the face vector in the space (O'Toole et al 1997b). In the present paper, we were curious to see what was on the 'other side of the mean'. We did this by moving an individual face through the mean and out the other side. We call the result of this manipulation an 'anti-face', and wish to distinguish it carefully from an anti-caricature. Although the direction of the move is the same for an anti-face and an anti-caricature, an anti-face is moved to the other side of the mean (see figure 1c). This manipulation should produce the 'opposite' of the face. But, what is the opposite of a face?

In answer to this question, we generated anti-faces for several individual faces by morphing through the line connecting them to the average face, and continuing on to the other side of the mean. We created these faces in two ways, using the general and

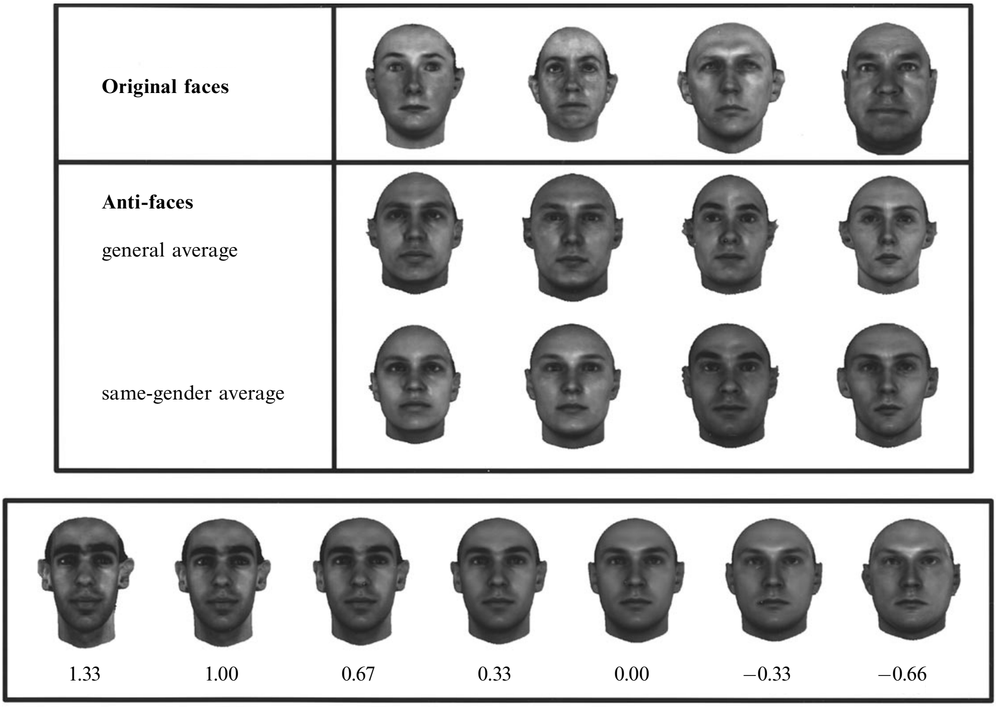

Figure 2. Row 1 contains the pictures of four individuals. Row 2 contains their 'anti-faces' or opposites, computed by using the line between the face and the general average. Row 3 contains the 'anti-faces' or opposites, computed by using the line between the face and the same-gender average. The bottom row illustrates seven equally spaced steps between a caricature and an anti-face. The original face appears second from the left (1.00), the male average face appears third from the right (0.00).

same-sex averages (eg the female average for female faces). Figure 2 contains the original pictures of four individuals (row 1). Anti-faces computed by using the line between the face and the general average, and anti-faces computed by using the line between the face and the same-gender average, appear in rows 2 and 3, respectively. The final row of this figure illustrates seven equally spaced steps between a caricature and an anti-face. The original face appears second from the left, the average face appears third from the right. We arbitrarily labeled the levels of the morph from 1.0 (the original face) to −1.0 (the anti-face).

The faces in figure 2 and their anti-face opposites suggest an interesting perceptual discontinuity of facial identity as the faces cross the origin or average of the space. Specifically, there seems to be little to associate the face and anti-face perceptually—although at a cognitive level, it seems possible to understand why the opposites appear as they do. Thus, for the faces in the first and third column of figure 2, the skin tone and hair change from light to dark. Likewise, for the face in the second column, the small round face becomes large and square-shaped. For the face in the fourth column, the roundish cheeks become gaunt.

The perceptual discontinuity in the identity of the faces as they cross over the mean occurs despite the continuity of the linear trajectory along which the faces fall. This discontinuity seems evident both for anti-faces made by using the general average, which changes gender, and for anti-faces made by using the same-gender average as the reference mean. Before discussing the implications of this observation for the face-space model, we wanted to confirm it formally with an experiment. To do so, we created morphs between a set of faces and their anti-faces, and sampled the morph trajectory in five discrete steps. For each original face, this yielded five images in face space, with the endpoint images equal to the face and its anti-face. Participants rated the similarity of all possible pairs of faces. We used these data to determine if the perceptual distances between equidistant faces were equivalent for equidistant points on the face to anti-face trajectory. If there is a perceptual discontinuity at the mean, we would predict a larger perceptual distance between two faces that straddle the mean than between similarly spaced points elsewhere on the trajectory.

2 Experiment

2.1 *Participants*

Fourteen undergraduate students (thirteen females and one male) from the University of Texas at Dallas volunteered to participate in exchange for research credit in an introductory psychology course.

2.2 *Apparatus*

The experiment was programmed on a Macintosh computer with PsyScope$^{\text{TM}}$ (Cohen et al 1993).

2.3 *Representation and algorithm*

The caricaturing and morphing algorithms used in this study were based on 3-D morphing software developed recently by Blanz and Vetter (see Blanz and Vetter 1999; Vetter and Blanz 1998). This software operates by finding a complete correspondence between individual 3-D laser-scan representations of faces and the average of these. Although this algorithm is similar in principle to most commercial morphing algorithms, three important differences are worth noting. First, whereas most morphing algorithms operate on faces that are put into correspondence by locating a subset of landmark feature matches (eg corners of the eyes, etc), the present algorithm operates on faces that are in complete correspondence. In other words, *all* data points on the head surface were matched. Second, the correspondence problem is solved automatically rather than by the hand-placement/matching of the landmark points. Finally, the correspondence algorithm is applied simultaneously to 3-D laser scans and their reflectance maps, rather than to image data. Combined, these features yield a very-high-resolution representation on which the morphing and caricaturing processes are based. Full details of the processing required to put the heads into full correspondence with the average appear in Vetter and Blanz (1998) and in Blanz and Vetter (1999).

In all other respects, the algorithm we used for making anti-faces was a standard caricature algorithm, which we applied in the opposite or anti-caricature direction. For a more detailed description of this algorithm, see O'Toole et al (1997b).

2.4 *Stimuli.*

The face stimuli were made from laser scans (Cyberware$^{\text{TM}}$) consisting of a 3-D head model and its associated surface reflectance map. The general average was computed on the basis of two hundred individuals (hundred male and hundred female) and the gender averages were computed on the hundred male and hundred female faces, separately. The stimuli used in the present study consisted of four 3-D images (as described above) of the faces of adult females aged 24, 29, 32, and 37 years. Although this may seem to be a small number of faces, given that we wished to create 5 images from each original face (20 images) and collect similarity ratings between all possible pairs of images (400 comparisons), we worried about subject fatigue and concentration, and thought it best to limit the number of faces.

The original faces were morphed into their anti-faces with the 3-D-morphing software described in Blanz and Vetter (1999). Morphing was done through the same-sex mean so that subject similarity judgments would not be based on perceived gender differences between the faces. The morph trajectory was divided into five equal segments, as follows.

The original face (1.00) and anti-face (−1.00) were the endpoints of the trajectory. The average was at 0.00. In between the average and the face (anti-face), two additional faces were sampled at 0.33 and 0.66 (−0.33 and −0.66). From these seven images, we collected data on the −0.66, −0.33, 0.33, 0.66, and 1.00 faces, as follows. The average was eliminated because it would be the center point for all pairs of faces and anti-faces, and would thus appear in the experiment much more frequently than the other face images. The anti-face was defined here as the face at −0.66. This is because minor artifacts in the correspondence procedure seemed to be amplified in the −1.00 anti-faces, causing an unnatural appearance for some faces.

2.5 *Procedure*

Participants rated all possible pairs of faces according to the following scale: 1: same person, 2: similar person, and 3: dissimilar person. Faces appeared two at a time side-by-side on the computer screen and remained visible until participants responded using the computer keyboard. Face pair order was randomized for all participants.

2.6 *Results*

To determine if there was a perceptual discontinuity in the space across the mean, we proceeded as follows. For each subject, we compiled a matrix of the mean similarity ratings assigned to all pairs of faces. For this analysis, we considered only the similarity ratings among pairs of faces along the trajectory between a face and its anti-face.(1) We traced this trajectory from the anti-face to the face, averaging the ratings for the anti-faces paired with each of their co-linear morphs. Our original question concerned the perceptual 'step size' between the co-linear neighbors, and so we computed the step sizes by taking the differences between the ratings received by neighboring pairs. If no discontinuity exists, we would expect the step sizes to be equal for all equidistant pairs. In the present case, because we did not collect ratings on the faces in combination with the average face, the continuity hypothesis predicts equal-sized perceptual steps at all places on the trajectory, except across the average. Here we would expect the step to be exactly twice the size of the other steps (ie the pair −0.33 and 0.33 involves two steps).

These step-size data were computed for each subject across the four faces and were submitted to a repeated-measures analysis of variance (ANOVA), with the location of the pair on the trajectory taken as the independent variable. As expected, owing to the double step size across the average, there was a significant effect of the pair: $F_{3,39} = 49.95$, $p < 0.001$. More relevant for our hypothesis is the pattern of data, which appears in figure 3 and supports the perceptual-discontinuity hypothesis. Planned comparisons were carried out on all contiguous pairs. These indicated statistically equivalent step sizes (ie no significant differences) at all places on the trajectory except for points around the average. As can be seen by the means and standard error bars in figure 3, the step size for the face pair that crosses the average is about five times greater than the largest step size found elsewhere on the trajectory (1.42 versus 0.25). This is much greater than the two times predicted by the perceptual continuity hypothesis.

Finally, we wished also to verify the perceptual non-equivalence of the two-unit steps taken between −0.33 and 0.33, and two-unit steps taken elsewhere on the trajectory. In the present case, we had data on one other two-unit step, which occurred between 0.33 and 1.00. The null hypothesis would predict that the perceptual step from −0.33 to 0.33 should be the same as the step from 0.33 to 1.00. An ANOVA was carried out to compare individual subject rating differences for these two steps, by using the step

---

(1) When we began collecting these data, we had originally intended to present an MDS analysis of the full similarity matrix. We carried out this analysis, and found it consistent with the inferential analyses made of the pure 'distances'. Our reason for excluding the MDS here is that 3–4 dimensions are needed to explain a substantial proportion of variance, which makes a visual interpretation of the data difficult.

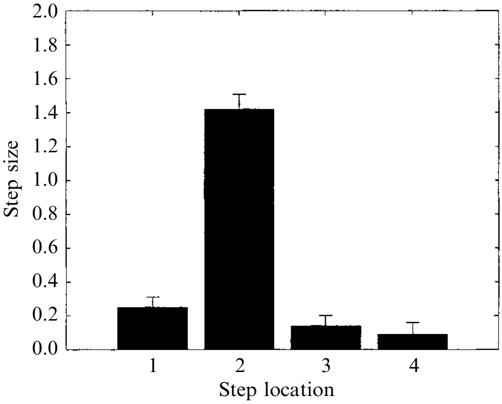

Figure 3. Perceptual step size as a function of the location of the face pair on the trajectory. 1: (−0.66 and −0.33); 2: (−0.33 and 0.33); 3: (0.33 and 0.66); 4: (0.66 and 1.00, the original face). Step sizes suggest a perceptual discontinuity at the average.

location as a within-subjects independent variable. This proved highly significant: $F_{1,13} = 58.83$, $p < 0.001$. The mean for the two-unit step across the mean was 1.42 and the mean for the two-unit step from 0.33 to 1.00 was 0.23.

3 Discussion

The data, therefore, support the existence of a perceptual discontinuity as the face trajectory crosses over the mean. More generally, there seems to be a break in the perceptual continuity of any kind. Not only do the faces and their anti-faces seem strikingly different, human observers do not find it easy to perceptually relate any of the images to their co-linear counterparts on the other side of the mean.

It is interesting also to consider more analytically the difference between using the general versus same-gender average as the reference mean for the face to anti-face transition. Anti-faces made by using the same-gender retain the gender of the original face. Anti-faces made by using the general average change gender. Figure 4 illustrates the relationship between these transitions using a face-space schematic. As can be seen, the anti-face transition amounts to a direct manipulation of the gender information in the face, independent of the identity information. More precisely, the line connecting the two anti-faces is parallel to the line connecting the male and female averages.

In summary, in perceptual terms a line in the physical face space has two identities, one on each side of the mean. For opposites made by using the same-gender averages,

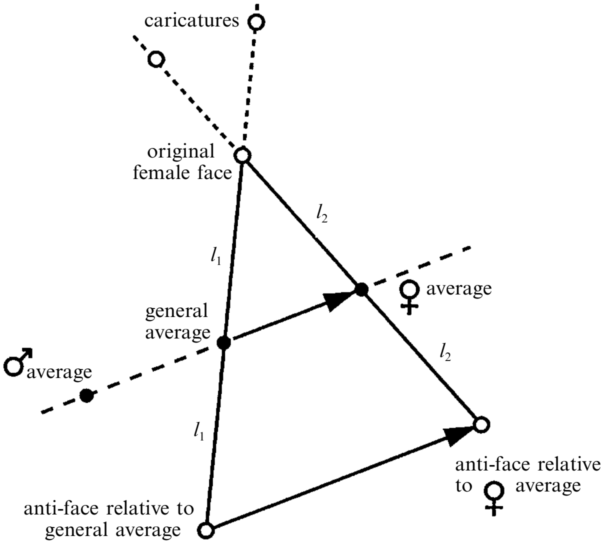

Figure 4. A demonstration that the relationship between an anti-face made by using the general average as the reference and an anti-face made by using the same-gender average as the reference amounts to a direct manipulation of the gender information in the face, independent of the identity information. This is shown by the fact that the line connecting the two anti-faces is parallel to the line connecting the male and female averages.

the two identities share the same gender. For the opposites made by using the general average, the identities do not share a common gender. We do not wish to argue that this demonstration invalidates the usefulness of the 'face-space' framework for modeling human perception and memory for faces. Rather, we think that as a general approach to understanding human face representations, computational models that enable the synthesis of a face at an arbitrary place in a physical face space may give an indication of the physical dimensions that underlie psychological face spaces. The validity of individual physical face spaces as models of human perception can be tested by comparing the properties of physical and perceptual face spaces.

Acknowledgements. Alice O'Toole gratefully acknowledges support from NIMH grant 1R29MH 5176501A4 and the Alexander von Humboldt Foundation. The authors are grateful to Isabelle Bülthoff and Yi Cheng for helpful comments on this manuscript.

References

Benson P J B, Perrett D, 1991 "Synthesising continuous-tone caricatures" *Image and Vision Computing* **9** 123 – 129

Benson P J, Perrett D I, 1994 "Visual processing of facial distinctiveness" *Perception* **23** 75 – 93

Blanz V, Vetter T, 1999 "A morphable model for the synthesis of 3D faces", in *1999 Symposium on Interactive 3D Graphics—Proceedings of SIGGRAPH'99* (New York: ACM Press) pp 187 – 194

Brennan S E, 1985 "The caricature generator" *Leonardo* **18** 170 – 178

Busey T A, 1998 "Physical and psychological representations of faces: Evidence from morphing" *Psychological Science* **9** 476 – 483

Cohen J D, McWhinney B, Flatt M, Provost J, 1993 "PsyScope: A new graphic interactive environment for designing psychology experiments" *Behavior Research Methods, Instruments & Computers* **25** 257 – 271

Hancock P J B, Burton A M, Bruce V, 1996 "Face processing: Human perception and principal components analysis" *Memory and Cognition* **24** 26 – 40

Light L, Kayra-Stuart F, Hollander S, 1979 "Recognition memory for typical and unusual faces" *Journal of Experimental Psychology: Human Learning and Memory* **5** 212 – 228

Light L, Hollander S, Kayra-Stuart F, 1981 "Why attractive people are harder to remember" *Personality and Social Psychology Bulletin* **7** 269 – 276

O'Toole A J, Abdi H, Deffenbacher K A, Valentin D, 1993 "Low dimensional representations of faces in high dimensions of the space" *Journal of the Optical Society of America A* **10** 405 – 410

O'Toole A J, Deffenbacher K A, Valentin D, Abdi H, 1994 "Structural aspects of face recognition and the other-race effect" *Memory and Cognition* **22** 208 – 224

O'Toole A J, Vetter T, Troje N F, Bülthoff H H, 1997a "Sex classification is better with head structure than with image intensity information" *Perception* **26** 75 – 84

O'Toole A J, Vetter T, Volz H, Salter E M, 1997b "Three-dimensional caricatures of human heads: Distinctiveness and the perception of facial age" *Perception* **26** 719 – 732

Rhodes G, 1996 *Superportraits* (Hove, East Sussex: Psychology Press)

Rhodes G, Brennan S E, Carey S, 1987 "Identification and ratings of caricatures: Implications for mental representations of faces" *Cognitive Psychology* **19** 473 – 497

Rhodes G, McLean I G, 1990 "Distinctiveness and expertise effects with homogeneous stimuli: Towards a model of configural coding" *Perception* **19** 773 – 794

Rhodes G, Tremewan T, 1994 "Understanding face recognition: Caricature effects, inversion, and the homogeneity problem" *Visual Cognition* **1** 275 – 311

Rhodes G, Tremewan T, 1996 "Averageness, exaggeration, and facial attractiveness" *Psychological Science* **7** 105 – 110

Sirovich L, Kirby M, 1987 "Low-dimensional procedure for the characterization of human faces" *Journal of the Optical Society of America A* **4** 519 – 554

Valentine T, 1991 "A unified account of the effects of distinctiveness, inversion, and race in face recognition" *Quarterly Journal of Experimental Psychology A* **43** 161 – 204

Vetter T, Blanz V, 1998 "Estimating coloured 3D face models from single images: An example based approach", in *Proceedings of the Fifth European Conference on Computer Vision (ECCV '98), Freiburg, Germany* Eds H Burkhardt, B Neumann, Springer Lecture Notes in Computer Science 1407 (Berlin: Springer) pp 499 – 513