

Why the long face? The importance of vertical image structure for biological “barcodes” underlying face recognition

Morgan L. Spence

Perception Lab, School of Psychology,
The University of Queensland,
St. Lucia, Queensland, Australia



Katherine R. Storrs

Perception Lab, School of Psychology,
The University of Queensland,
St. Lucia, Queensland, Australia



Derek H. Arnold

Perception Lab, School of Psychology,
The University of Queensland,
St. Lucia, Queensland, Australia



Humans are experts at face recognition. The mechanisms underlying this complex capacity are not fully understood. Recently, it has been proposed that face recognition is supported by a coarse-scale analysis of visual information contained in horizontal bands of contrast distributed along the vertical image axis—a biological facial “barcode” (Dakin & Watt, 2009). A critical prediction of the facial barcode hypothesis is that the distribution of image contrast along the vertical axis will be more important for face recognition than image distributions along the horizontal axis. Using a novel paradigm involving dynamic image distortions, a series of experiments are presented examining famous face recognition impairments from selectively disrupting image distributions along the vertical or horizontal image axes. Results show that disrupting the image distribution along the vertical image axis is more disruptive for recognition than matched distortions along the horizontal axis. Consistent with the facial barcode hypothesis, these results suggest that human face recognition relies disproportionately on appropriately scaled distributions of image contrast along the vertical image axis.

Introduction

The relative positioning of facial features is important for categorizing and identifying faces (for a review, see Maurer, Le Grand, & Mondloch, 2002). Since all faces share the same first-order relations (two hori-

zontally-aligned eyes above a central nose, above a mouth), recognition of individuals is thought to involve, at least in part, information relating to subtle variations in interfeatural spacing (Diamond & Carey, 1986; Rhodes et al., 1988; for reviews, see Maurer, Le Grand, & Mondloch, 2002; Piepers & Robbins, 2012). This type of information has been referred to as the “second-order” relations, or as the configural properties of a face (Diamond & Carey, 1986; Maurer et al., 2002; Piepers & Robbins, 2012). Perhaps the most often cited and manipulated example of a second-order facial relation is the distance between the two eyes (e.g., Bukach, Grand, Kaiser, Bub, & Tanaka, 2008; Dahl, Logothetis, & Hoffman, 2007; Goffaux, Hault, Michel, Vuong, & Rossion, 2005; Goffaux & Rossion, 2007; Leder & Bruce, 1998; Leder, Candrian, Huber, & Bruce, 2001; Tanaka & Sengco, 1997).

Some studies have conceptualized second-order facial relations and the shape of individual features (such as the nose) as separate constructs. These seem to be, however, inherently intertwined (Hosie, Ellis, & Haig, 1988; Leder & Bruce, 2000; Rhodes, Brake, & Atkinson, 1993). Any facial feature necessarily corresponds with a two-dimensional retinal image, so features can be thought of as having a configuration, even if one chooses to describe this as a shape. If one changes the shape of a given facial feature, this will inevitably impact on distances between this and other facial features (see Leder & Bruce, 2000). Considering this caveat, we use the term “image distribution” to describe the spatial distribution of luminance contrast

Citation: Spence, M. L., Storrs, K. R., & Arnold, D. H. (2014). Why the long face? The importance of vertical image structure for biological “barcodes” underlying face recognition. *Journal of Vision*, 14(8):25, 1–12. <http://www.journalofvision.org/content/14/8/25>, doi: 10.1167/14.8.25.




Figure 1. The image on the left depicts Albert Einstein, a scientist many first year psychology students recognize. This image (not used in the study) has a physical aspect ratio of 1.0 (\pm any printing errors or distortions resulting from the configuration of your display). If you were to view the original image from one side at a horizontal angle of 50° , the aspect ratio of its associated retinal image would be compressed to an extent depicted by the image on the right. Note that in addition to the retinal image compression (in this case, about 64% as wide as it is tall) there would also be a perspective gradient, with regions on the near side corresponding with retinal images 10% larger than matched regions on the far side of the image, assuming central fixation and a viewing distance of 1 m.

changes across a facial image, as this defines both the shapes of individual facial features and their spatial relations.

While second-order relations may be important for face recognition, these must not be defined in *absolute* metric distances, but contextualized relative to the proportions of the retinal image (see Hole, George, Eaves, & Rasek, 2002). The most obvious need for contextualized second-order relations arises from changes in viewing distance: a proximate face will cast a large retinal image, whereas the same face viewed from a distance will cast a smaller retinal image. Since we can recognize individual faces from different distances, any second-order relational information used in face recognition must be calculated in proportion to the size of the facial image on the retina.

A further need to contextualize second-order relations arises because faces are often viewed (and recognized) from different angles. This is perhaps best demonstrated by considering the retinal projection of a picture of a face (see Figure 1). If viewed from directly in front, a square picture will cast an approximately square retinal image. However, if viewed from an angle the same picture will cast a roughly oblong retinal image, foreshortened along the axis of the viewing angle. Perception compensates for this type of retinal image distortion, such that objects in pictures can be recognized, and seem unchanged, when viewed from a

wide variety of angles (see Liu & Ward, 2006; Storrs & Arnold, 2013; Vishwanath, Girshick, & Banks, 2005).

As faces can be recognized in pictures viewed from a variety of angles, any second-order information involved in face recognition must be encoded not only relative to the size of the face's retinal image, but also relative to its aspect ratio. This likely explains why face recognition copes well with linear distortions that preserve an appropriate image distribution given that image's aspect ratio; that is, stretching or compressing a facial image along one of its cardinal axes has little impact on face recognition (Hole et al., 2002). However, facial coding is disrupted by nonlinear distortions, which create an inappropriate distribution relative to the image's aspect ratio, such as by skewing (stretching one half of the image, while compressing the other) an image along its vertical axis (Hole et al., 2002).

Changing the horizontal angle from which a face is viewed will have a substantial impact on the horizontal dimensions of a face's retinal image, but only a minimal impact on its image distribution along the vertical axis. Hence, a greater reliance on vertical image distributions might support viewpoint-invariant face recognition, assuming people view faces (and pictures of faces) from a wider variety of horizontal angles than from different inclinations (Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux & Rossion, 2007). Motivated in part by these considerations, it has been proposed that viewpoint-invariant structural codes for face recognition rely on horizontal "bands" of contrast across the face, and their relative positioning along the vertical axis (Crookes & Hayward, 2012; Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux & Rossion, 2007; Pachai, Sekuler, & Bennett, 2013). This vertical arrangement of horizontal bands of contrast might act like a biological facial "barcode" for recognition, capturing the reflectance of a face as it varies along the vertical axis (Dakin & Watt, 2009; Goffaux & Dakin, 2010).

A number of recent observations are consistent with the facial barcode hypothesis (Dakin & Watt, 2009). For instance, Goffaux and Rossion (2007) found that image inversion, which has a disproportionate impact on face recognition relative to other objects (Yin, 1969), had a greater impact on sensitivity to vertical relations (e.g., the distance between the nose and mouth) relative to sensitivity to horizontal relations (e.g., the distance between the eyes; see also Crookes & Hayward, 2012). Similarly Goffaux and Dakin (2010) found that filtering images so they only retained information concerning a small range of orientations about vertical disrupted a number of behavioral signatures of face processing (e.g., the face inversion effect, facial identity after-effect, face matching across viewpoint and holistic/interactive processing of face parts), whereas filtering the same visual information about horizontal had a minimal impact on recognition

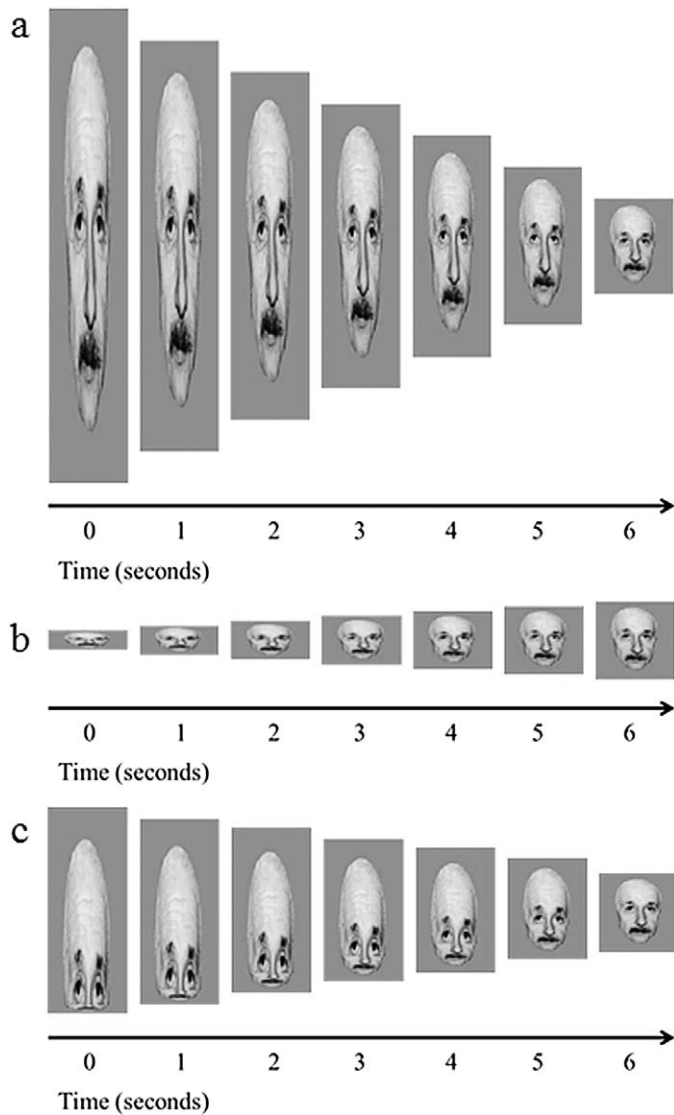


Figure 2. (a) Depiction of a vertically stretched presentation, using a picture of Albert Einstein (not used in study). At the beginning of the presentation the image was set to a ratio of 0.2. The aspect ratio then increased, reaching a value of 1.0 after six seconds. Participants were asked to press a button as soon as they could recognize the face. The button press terminated the image presentation, and prompted the presentation of a test display, wherein the participant chose the correct name from five options. (b) Depiction of a vertically squashed presentation, with an initial image ratio of 4.0, which then decreased. (c) Depiction of a vertically skewed image presentation, wherein the top half of the facial image had an initial aspect ratio of 0.2, while the bottom half had an initial aspect ratio of 4.0.

performance. In another similar manipulation, Pachai and colleagues (2013) added vertical or horizontal Gaussian noise to facial images, and found that the former had a greater adverse impact on recognition. There is thus a growing body of evidence suggesting that face recognition relies disproportionately on

vertical image structure (see also Sekunova & Barton, 2008).

According to the facial barcode hypothesis, vertical image structure is especially important for facial coding, and reliance upon this structure helps people recognize faces from different viewing angles and distances (see Dakin & Watt, 2009). For this strategy to be viable, the barcode associated with a given identity must be expressed in proportion to the height of the facial image on the retina, so recognition should be disrupted by manipulations that alter the vertical distribution of the image. The barcode hypothesis further predicts that the horizontal distribution of the retinal image should be less important, since a facial barcode largely dismisses image variance along the horizontal axis. Alternatively, facial coding might rely on appropriately scaled two-dimensional representations, with spatial interrelationships along the horizontal axis being as important as interrelationships along the vertical axis (Sinha, Balas, Ostrovsky, & Russell, 2006).

Previously, Hole and colleagues (2002) have shown that face recognition is disrupted by nonglobal distortions of an image along its vertical axis (selectively stretching either the top or bottom half of the facial image). This is entirely consistent with the facial barcode hypothesis, as this nonglobal stretching disrupts vertical image structure in proportion to the overall height of the facial image. What is unclear is whether nonglobal distortions along the horizontal image axis will be similarly disruptive, or less disruptive as is predicted by the barcode hypothesis. Here we present a series of experiments that evaluate these possibilities. Using a dynamic distortion paradigm, we measure the impact of stretching, compressing, and skewing facial images along their horizontal or vertical axes. We find that famous face recognition is more impacted by asymmetrical skewing than by linear distortions. More importantly, we find that the detrimental impact of skewing is greater along the vertical than along the horizontal image axis.

Experiment 1

Methods

Twenty-one psychology students (14 females, $M_{\text{age}} = 23.4$ years, $SD_{\text{age}} = 5.05$, range: 17–36 years) volunteered to participate in Experiment 1. All were naïve as to the purpose of the experiment and had normal or corrected-to-normal visual acuity. Stimuli were generated using Matlab R2012b software (MathWorks, Natick, MA) in conjunction with the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) and displayed on




Figure 3. Top Left—Depiction of the appearance of test images in baseline trials, or after six seconds of animation in other trials. Top Middle—Depiction of the initial appearance of a test image during squashed Horizontal Symmetrical distortion trials. Top Right—Depiction of the initial appearance of test images on stretched Horizontal Symmetrical distortion trials. Bottom Left—Depiction of the initial appearance of a test image on Horizontal Skewed distortion trials, with a squashed left half and a stretched right half. Bottom Right—Depiction of the initial appearance of a test image on Horizontal Skewed distortion trials, with a stretched left half and a squashed right half.

either a 19 in. Samsung SyncMaster 950p (Samsung, Seoul, South Korea), a 19 in. Samsung SyncMaster 900SL (Samsung), or on a 19 in. Sony Trinitron Multiscan G420 monitor (Sony, Tokyo, Japan). All monitors were set to a resolution of 1024×768 with a refresh rate of 75 Hz. Stimuli were viewed from a distance of 57 cm, controlled using a chin rest, while seated in a dark room.

Test images were 96 pictures of Caucasian celebrity faces, with predominantly neutral expressions, sourced from Google Images. Images were converted to grayscale, and adjusted using Adobe Photoshop Elements 10 (Adobe Systems Inc., San Jose, CA) to set the background to grey and eliminate all but the natural facial contours and hairline. The average luminance of all images was then equated using the SHINE toolbox for Matlab (Willenbockel et al., 2010).

Experimental procedure

On each trial, participants made speeded recognition judgments regarding famous faces. Participants were instructed to indicate as quickly as possible, by pressing a mouse button, if they recognized a face, and not to respond if they did not recognize the face. Image presentations could last for up to eight seconds. In half of the trials images were animated for the initial six seconds of the presentation (shrinking or expanding along the vertical or horizontal axis, see below) and were then static for a further two seconds. In other trials, images were static for the full duration. If the participant failed to recognize the face after eight seconds, the face was extinguished, and there was a three second interstimulus interval (ISI) before the next trial.

On trials where the participant reported that they had recognized the face, their response extinguished the facial image presentation and, after a one second pause, a test display was presented. The test display presented five celebrity names, vertically arranged on the display monitor, and the participant chose one by moving a mouse and clicking on the appropriate name. One of the five names was that of the celebrity presented (positioned randomly within the list), and the other four were distractor names selected randomly from among the other 47 gender-matched celebrities (whose facial images were presented on other trials). In some cases, a brief descriptor was included in addition to the name, for example, “Hugh Jackman (Wolverine).” Recognition times for the initial response were only recorded if the participant chose the correct name on the subsequent test display.

Experimental conditions

There were four experimental conditions (see supplementary materials for animated demonstrations). In the Vertical Symmetrical distortion condition, images were initially set to either an aspect ratio of 0.2 (stretched, see Figure 2a) that then linearly increased, or to an initial aspect ratio of 4.0 (squashed, see Figure 2b) that then linearly decreased. In both cases, an aspect ratio of 1.0 was reached after six seconds. The initial width and height of vertically stretched images subtended 4×20 degrees of visual angle (dva) at the retina, and the initial width and height of vertically squashed images subtended 4×1 dva. Half of the trials for this condition involved squashed presentations, the other half stretched pre-

Condition	Baseline	Test
Horizontal Symmetrical	1.11 (<i>SD</i> = 0.64)	2.30 (<i>SD</i> = 0.64)
Vertical Symmetrical	0.97 (<i>SD</i> = 0.31)	2.25 (<i>SD</i> = 0.64)
Horizontal Skewed	1.06 (<i>SD</i> = 0.31)	3.08 (<i>SD</i> = 0.99)
Vertical Skewed	1.06 (<i>SD</i> = 0.42)	4.55 (<i>SD</i> = 0.84)

Table 1. Mean and standard deviation (seconds) of response times for experimental conditions, and associated baseline measures, in Experiment 1.

sentations. In the Vertical Skew distortion condition, the initial aspect ratio was 0.4, with a total initial image width and height of 4×10.5 dva. Vertically skewed images were created by setting the top or bottom half of the image to a height subtending 0.5 dva and the other half of the image to a height of 10 dva. In both cases, the initially squashed half image expanded and the stretched half image contracted until an aspect ratio of 1.0 was achieved after six seconds. Half of the trials for this condition were initially squashed upper-half and stretched lower-half images, and vice versa for the other trials.

Details for the Horizontal Symmetrical distortion condition were as for the Vertical Symmetrical distortion condition, except that images were initially stretched (width 20 dva, height 4 dva) or squashed (width 1 dva, height 4 dva) along the horizontal axis (see Figure 3). Similarly, details for the Horizontal Skew distortion condition were as for the Vertical Skew condition, with the left and right halves of the image being distorted in opposite directions (total initial image width and height of 10.5×4 dva; squashed side: 0.5×4 dva, stretched side: 10×4 dva).

Trial blocks

In total, there were 96 facial images, 48 male and 48 female, each presented twice for a total of 192 individual trials. These were completed in a single block in a random order. For one out of two presentations, each image was static, with an image ratio of 1.0. Performance on this trial served as a baseline measure for this image. For the other presentation, each image was assigned to one of the four experimental conditions, such that each experimental condition included 24 presentations of facial images, 12 female and 12 male. As presentation order was randomized, it was equally probable that the baseline presentation preceded or followed the experimental presentation of the same image.

Results and discussion

The majority of faces were recognized and correctly matched to a name ($M = 86\%$, $SD = 16\%$). Recognition was only slightly poorer for experimental trials ($M = 85\%$, $SD = 16\%$) than baseline trials ($M = 87\%$, $SD =$

15% ; paired $t_{20} = 2.59$, $p = 0.018$). A repeated measures ANOVA on individual accuracy difference scores (experimental trials – baseline) revealed a trend toward a main effect for type of distortion (symmetrical/skewed), such that skewed distortions resulted in more impaired recognition ($M = -2\%$, $SD = 0.05$) than symmetrical distortions, $M = -1\%$, $SD = 5\%$; $F(1, 20) = 3.42$, $p = 0.079$, $\eta_p^2 = 0.15$. There was no main effect for axis of distortion (vertical or horizontal), $F(1, 20) = 1.03$, $p = 0.321$, $\eta_p^2 = 0.05$, nor an interaction between distortion axis and type, $F(1, 20) = 0.45$, $p = 0.834$, $\eta_p^2 = 0.002$. Note that since all trials culminated in a two-second presentation of an undistorted image if the participant had not already recognized the face, only subtle accuracy effects were anticipated. We are primarily interested in tolerance to image distortions, which were assessed by response times. Response times were recorded from trials in which participants reported recognizing the facial image, and subsequently chose the correct name from a list of five alternatives. Mean response times for the four experimental conditions and their respective baselines are presented in Table 1 and depicted in Figure 4.

We calculated individual difference scores by subtracting baseline response times for each image from response times for experimental presentations of the same image, with the condition that the participant had correctly identified the face on both trials. In all cases average individual response times were longer for distorted experimental presentations relative to baseline (Vertical Symmetrical distortion, paired $t_{20} = 13.21$, $p < 0.00001$; Vertical Skewed distortion, $t_{20} = 23.04$, $p < 0.00001$; Horizontal Symmetrical distortion, $t_{20} = 9.96$, $p < 0.00001$; Horizontal Skewed distortion, $t_{20} = 11.55$, $p < 0.00001$). These data show that face recognition was disrupted, relative to baseline, across all experimental conditions.

Response time difference scores are depicted in Figure 5. These were subjected to a 2 (Distortion Type: symmetrical/skewed) \times 2 (Distortion Axis: vertical/horizontal) repeated-measures ANOVA. Results revealed a significant main effect of Distortion Type, such that skewed distortions ($M = 2.76$, $SD = 1.05$) were more disruptive than symmetrical distortions, $M = 1.21$, $SD = 0.48$; $F(1, 20) = 194.65$, $p < 0.001$, $\eta_p^2 = 0.91$. There was also a significant main effect of axis, such that vertical distortions ($M = 2.39$, $SD = 1.26$)

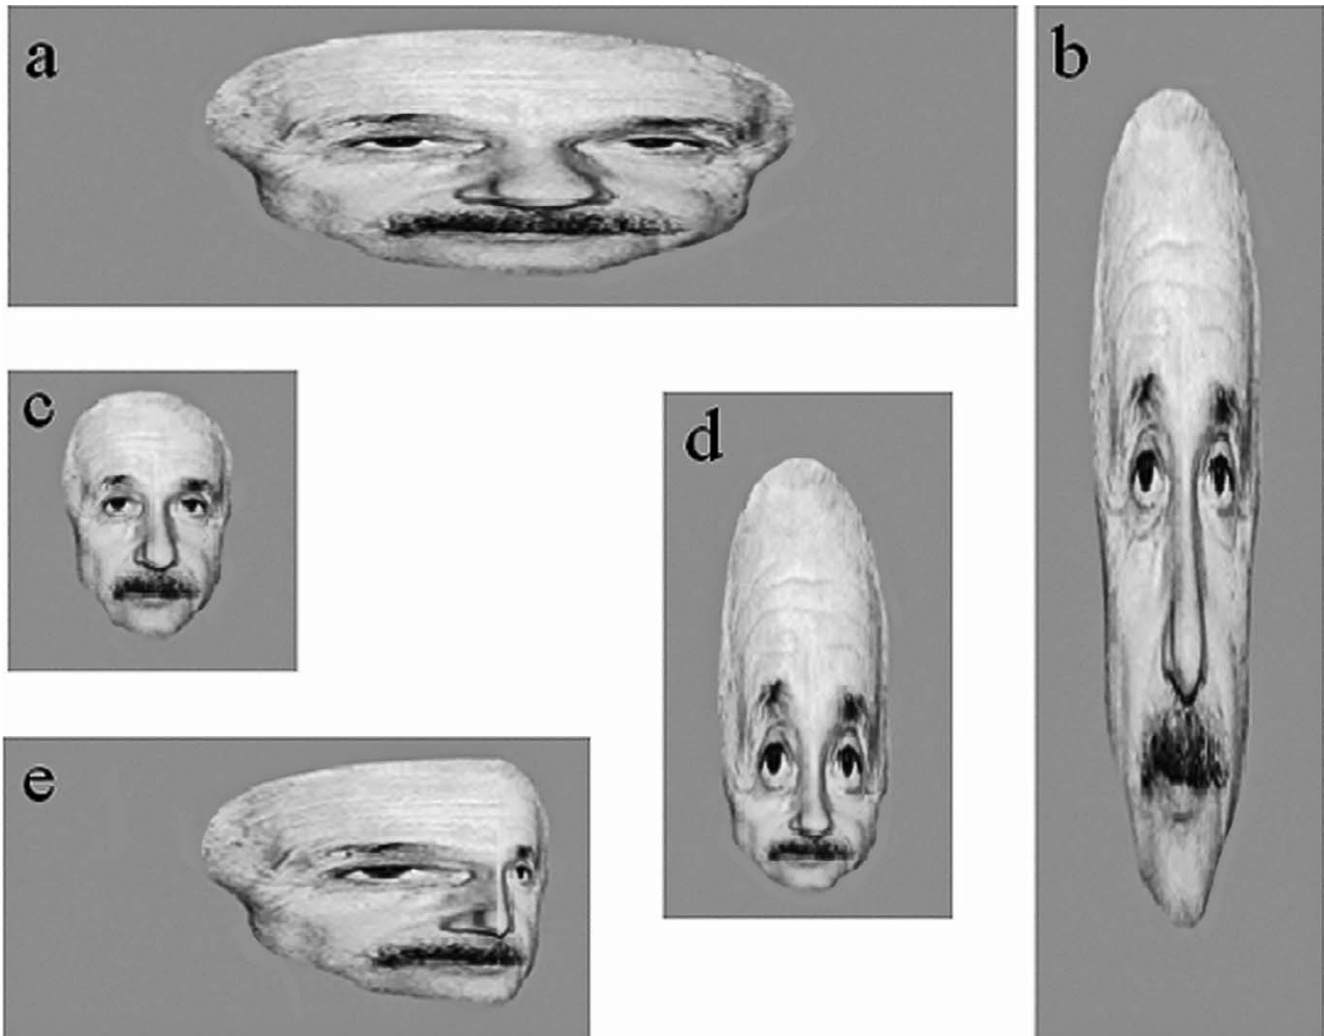


Figure 4. Depiction of average image distortion levels at the time when participants reported recognizing faces in different conditions. Images are shown for (a) Horizontal Symmetrical Distortion—stretched presentation, (b) Vertical Symmetrical distortion—stretched presentation, (e) Horizontal Skewed distortion—left stretched presentation, and (d) Vertical Skewed distortion—top stretched presentation. For comparison, (c) shows the same image as it appeared at baseline, or after six seconds of animation in other conditions. Note, our data suggest images (a), (b), (d), and (e) should be equally recognizable, even though the symmetrical image distortions are greater than the skewed image distortions; compare left portions of (a) and (e), and top portions of (b) and (d).

were significantly more disruptive to recognition than horizontal distortions, $M = 1.58$, $SD = 0.81$; $F(1, 20) = 79.62$, $p < 0.001$, $\eta_p^2 = 0.80$. There was also a significant Distortion Type \times Distortion Axis interaction, $F(1, 20) = 79.53$, $p < 0.001$, $\eta_p^2 = 0.80$, such that the effect of skew differed depending on axis of distortion. Skewing along the vertical axis ($M = 3.49$, $SD = 0.69$) was more disruptive than skewing along the horizontal axis ($M = 2.03$, $SD = 0.80$; $t_{20} = 10.182$, $p < 0.001$).

The key finding of Experiment 1 was that image skews along the vertical axis were more disruptive than horizontal skews. Previously it has been shown that face recognition can be disrupted by faces adopting different emotional expressions (e.g., Wang, Fu, Johnston, & Yan, 2013; see Calder & Young, 2005 for a review). For this reason, as far as possible, we selected

facial images adopting neutral expressions in an attempt to control for this potential interaction. However, if our image distortions had systematically biased facial expression and this effect was greater than any direct deleterious impact on face recognition, and this effect was exaggerated for image skewing relative to linear distortion, and this effect was greater along the vertical than the horizontal image axis, then our key finding could have resulted indirectly from the impact of our image distortions on facial expression. We believe this combination of contingencies is unlikely, so we attribute our key finding to image distortions having a direct impact on face recognition processes.

Our data are consistent with the facial barcode hypothesis (see Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux & Rossion, 2007), suggesting that face




Figure 5. Line graph depicting differences in response times (seconds) for baseline and experimental presentations of the same facial images. Data are shown for images distorted along either their vertical or horizontal axes, using Symmetrical or Skewed image distortions (see main text for an explanation of terms). Error bars depict ± 1 SEM.

recognition relies disproportionately on appropriately scaled distributions of image contrast along the vertical image axis. However, an alternative possibility is that identity could be estimated relatively independently from either side of a horizontally skewed facial image, but not from either side of a vertically skewed image (e.g., Vetter, Poggio, & Bühlhoff, 1994). Our reasoning for this hypothesis is as follows: skewed images were created by independently distorting each image half, squashing one and stretching the other. As horizontally skewed faces were split along their line of symmetry, they had redundant identity information on either side of the split and this information was internally consistent given the overall dimensions for that side of the face. This was not true for vertically skewed images, where feature shapes and positioning on either side of the face were internally inconsistent. We assessed whether our results were likely driven by a confound between axis of skew and axis of symmetry by repeating Experiment 1, this time with test stimuli consisting of half facial images (doing away with symmetry) skewed along axes centered on the eye/cheek region (see Figures 6 and 7).

Experiment 2

Methods

Methodological details for Experiment 2 were as for Experiment 1, with the following exceptions. There were twenty-one participants (14 females, 7 males, M_{age}




Figure 6. Depictions of test appearances in Experiment 2. Top Left—Appearance of test images in baseline trials, or after six seconds of animation in other trials. Top Middle—Initial test appearance during squashed Horizontal Symmetrical distortion trials. Top Right—Initial test appearance on stretched Horizontal Symmetrical distortion trials. Bottom Left—Initial test appearance on Horizontal Skewed distortion trials, with a squashed right half and a stretched left half. Bottom Right—Initial test appearance on Horizontal Skewed distortion trials, with a stretched right half and a squashed left half.

= 19.6 years, $SD_{\text{age}} = 3.62$, age range: 17–33 years), all naïve as to the purpose of the experiment with normal or corrected-to-normal visual acuity. Test images consisted of the same 96 pictures of celebrity faces, cropped to depict just the right-hand side of the face, extending from the facial midline to the cheek.

For linearly stretched presentations, images were initially extended in width (Horizontal Symmetric condition) or height (Vertical Symmetric condition) to four times the undistorted image width or height. Likewise, for linearly squashed presentations, images were initially compressed in width or height to one-fourth the width or height of undistorted images. For skewed presentations, images were distorted in opposite directions to either side of an axis passing through the middle of the half-facial image (i.e., a point lying on the right cheek of the face). On Vertical Skew trials, the top half of the image was squashed to one-fourth of its original height and the bottom stretched to four times




Figure 7. Depictions of test appearances in Experiment 2. Left—Initial test appearance on stretched Vertical Symmetrical distortion trials. Top Center—Appearance of tests in baseline trials, or after six seconds of animation in other trials. Top Right—Initial test appearance during squashed Vertical Symmetrical distortion trials. Bottom Center—Initial test appearance on Vertical Skewed distortion trials, with a squashed top half and a stretched bottom half. Bottom Right—Initial test appearance on Vertical Skewed distortion trials, with a stretched top half and a squashed bottom half.

its original height, or vice versa. Likewise, on Horizontal Skew trials, the left half of the image was squashed by a factor of four, while the right half was stretched by the same amount, or vice versa. In neither

condition did the axis of skew now correspond with an axis of symmetry.

Results and discussion

Fewer faces were recognized in Experiment 2 ($M = 65\%$, $SD = 20\%$) than in Experiment 1 ($M = 86\%$, $SD = 16\%$; $t_{40} = 4.00$, $p < 0.001$). This reflects the detrimental impact of eliminating half the face. Recognition was only slightly poorer in experimental trials ($M = 63\%$, $SD = 19\%$) relative to baseline ($M = 67\%$, $SD = 20\%$; repeated measures $t_{20} = 3.27$, $p = 0.004$). Only subtle differences in recognition, for experimental versus baseline presentations, were anticipated as experimental trials culminated in a two second presentation of an undistorted face, if it had not already been recognized. As in Experiment 1, our primary measure was therefore the *speed* of recognition. Mean recognition response times and associated standard deviations for the four experimental conditions and their respective baselines are presented in Table 2. These times only pertain to faces that were reported to be recognized, and subsequently successfully matched to a name, in both the baseline and experimental presentation of that image.

As in Experiment 1, face recognition was disrupted, relative to baseline, in all experimental conditions (Vertical Symmetrical distortion paired $t_{20} = 12.57$, $p < 0.00001$; Vertical Skewed distortion paired $t_{20} = 33.15$, $p < 0.00001$; Horizontal Symmetrical distortion paired $t_{20} = 11.21$, $p < 0.00001$; Horizontal Skewed distortion paired $t_{20} = 24.39$, $p < 0.00001$).

The response time difference scores, which were calculated for each participant, were subjected to a 2 (Distortion Type: symmetrical/skewed) \times 2 (Distortion Axis: vertical/horizontal) repeated-measures ANOVA. Results revealed a significant main effect of Distortion Type, such that skewed distortions ($M = 3.62$, $SD = 0.72$) were more disruptive than symmetrical distortions, $M = 1.78$, $SD = 0.87$; $F(1, 20) = 266.59$, $p < 0.001$, $\eta_p^2 = 0.93$. There was also a significant main effect of Distortion Axis, such that vertical distortions ($M = 2.94$, $SD = 1.38$) were more disruptive than horizontal distortions, $M = 2.44$, $SD = 0.99$; $F(1, 20) = 14.37$, $p = 0.001$, $\eta_p^2 = 0.42$. The Distortion Type \times Distortion Axis interaction was also significant, such that the effect of distortion type (skewed relative to symmetrical) differed depending on the axis of distortion, $F(1, 20) = 16.66$, $p = 0.001$, $\eta_p^2 = 0.45$. Specifically, skewing along the vertical axis ($M = 4.05$, $SD = 0.56$) was more disruptive than skewing along the horizontal axis ($M = 3.18$, $SD = 0.60$; paired $t_{20} = 6.74$, $p < 0.001$), whereas there was no difference in disruption along either axis for symmetrical distortions (vertical $M = 1.70$, $SD = 0.62$; horizontal $M = 1.70$, $SD = 0.69$; paired $t_{20} = 0.001$,

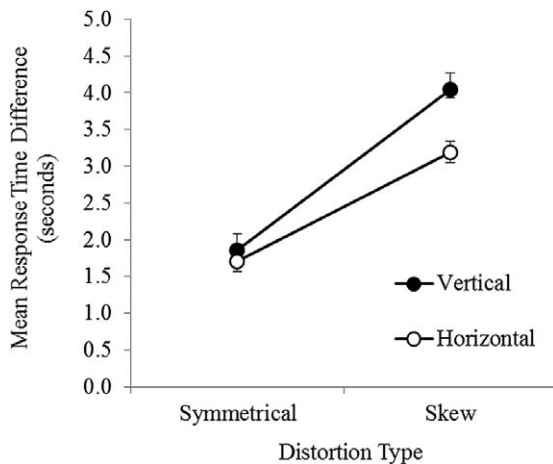


Figure 8. Differences in response times during baseline and experimental presentations of the same images during Experiment 2. Data are shown for images distorted along either their vertical or horizontal axes, using Symmetrical or Skewed image distortions (see main text for an explanation of terms). Error bars depict ± 1 SEM.

$p = 0.999$; see Figure 8). Consistent with the findings from Experiment 1, these results show that even after controlling for bilateral facial symmetry using half facial images, vertical skewing was still more disruptive than horizontal skewing.

Relative impact of horizontal and vertical skews

While image skewing along the vertical axis was more disruptive than skewing along the horizontal, the magnitude of this difference for half facial images in Experiment 2 ($M = 0.87$, $SD = 0.59$) was mitigated relative to full facial images in Experiment 1 ($M = 1.47$, $SD = 0.66$). To examine this differential effect of skewing along the horizontal and vertical axes for whole and half faces, difference scores for horizontally and vertically distorted images were subjected to two separate mixed ANOVAs, with type of distortion (Symmetrical/Skewed) being a within-subjects factor and type of facial image (Half/Whole) a between-subjects factor. Results for horizontally distorted images revealed a significant interaction between type of distortion and type of facial image, such that skewed horizontal distortions were more disruptive for half facial image presentations than for full facial image

presentations, $F(1, 40) = 9.84$, $p = 0.003$, $\eta_p^2 = 0.20$ (see Figure 9a). For vertically distorted images, there was no evidence for an interaction involving type of distortion and type of facial image, so skewed vertical distortions were equally disruptive for half facial image presentations in Experiment 2 as for full facial image presentations in Experiment 1, $F(1, 40) = 0.35$, $p = 0.557$, $\eta_p^2 = 0.009$ (see Figure 9b). These results show that, while it took longer to recognize half than full facial images, vertical image distortions were equally disruptive for both types of image.

If the horizontal image structure was not implicated in face recognition, it would be expected that there should be no differential impact of horizontal skewing for half compared to full facial images. Contrary to this, an image skew centered on the eye/cheek region of a half facial image is more disruptive than a skew centered on the nose of a full facial image. This indicates the differential importance of the image distribution along the horizontal image axis for recognition. The implication of this observation is that facial coding should not be conceptualized as a purely one-dimensional code. Instead, our results suggest a differentially weighted use of information concerning vertical and horizontal image structure for facial coding. Hence, while an appropriate scaling of the vertical information is more important than an appropriate scaling of horizontal information, a proportionate scaling of horizontal image structure is still important. This observation does not contradict the facial barcode hypothesis. While the theory predicts a disproportionate reliance on vertical image structure, it does not discount the possibility that horizontal image structure might play some role in recognition (Goffaux & Dakin, 2010; Goffaux & Rossion, 2007).

General discussion

Our data are consistent with face recognition relying disproportionately on appropriately scaled variance along the vertical, relative to the horizontal, image axis (see also Crookes & Hayward, 2012; Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux & Rossion, 2007; Sekunova & Barton, 2008). This was apparent since image skewing was more detrimental than linearly

Condition	Baseline	Test
Horizontal Symmetrical	1.39 ($SD = 0.36$)	3.09 ($SD = 0.90$)
Vertical Symmetrical	1.52 ($SD = 0.43$)	3.22 ($SD = 0.82$)
Horizontal Skewed	1.47 ($SD = 0.34$)	4.65 ($SD = 0.64$)
Vertical Skewed	1.48 ($SD = 0.58$)	5.53 ($SD = 0.54$)

Table 2. Means and standard deviations (seconds) of response times for the four experimental conditions and their respective baselines in Experiment 2.

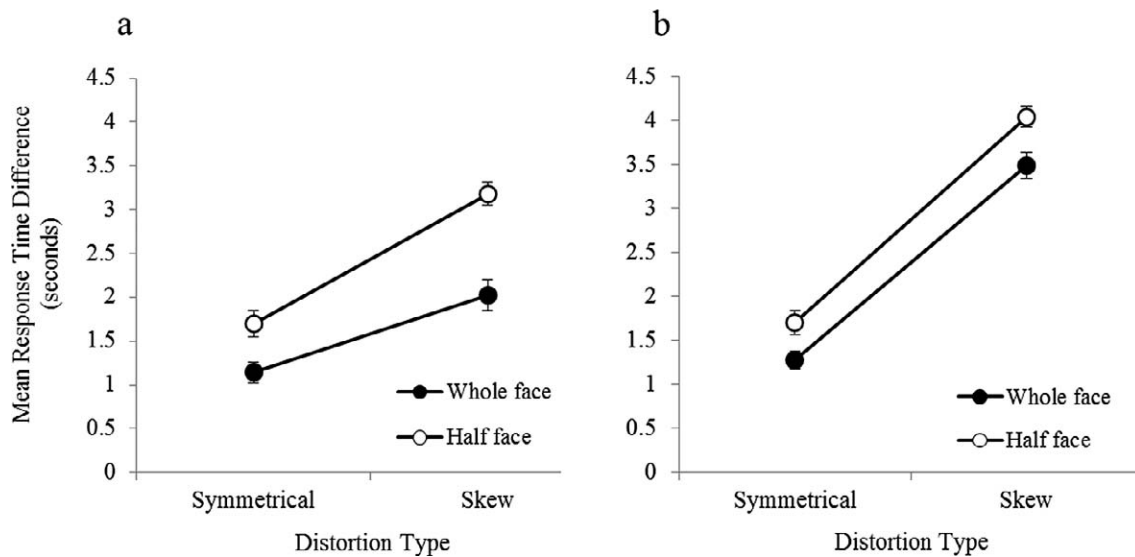


Figure 9. Differences in response times during baseline and experimental presentations of the same images for (a) Horizontal, and (b) Vertical, Symmetrical, and Skewed distortion trials in Experiments 1 and 2. Error bars depict ± 1 SEM.

stretching (or squashing) an image, and this effect was most pronounced for distortions along the vertical image axis. This could not be attributed to the bilateral symmetry of the face conferring an advantage on skews centered on the horizontal (relative to the vertical) facial midline, as the disproportionate impact of skewing along the vertical image axis remained constant for half facial images, which have no axis of symmetry.

Our manipulations were motivated by a consideration of the types of retinal image distortion that ensue when looking at faces, and images of faces, from different distances and viewing angles in a typical visual diet. The changes in retinal image size with viewing distance, and of retinal image shape with viewing angle, demand that any spatial code for face recognition must be contextualized relative to the overall dimensions of the facial image on the retina (see Hole et al., 2002). One way to achieve this would be to place a greater reliance on an appropriate scaling of vertical facial image structure, as the vertical image distribution changes less than the horizontal image structure with naturalistic changes in viewing angle (see Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux & Rossion, 2007). Our data further substantiate this theoretical prediction.

Caveats

Horizontal image structure matters

While our data support previous findings demonstrating a critical role of vertical image structure in face recognition (see Crookes & Hayward, 2012; Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux &

Rossion, 2007), they also evidence a role for horizontal structure. Skewing half facial images about their vertical midline (eye/cheek) was more disruptive than equivalent skewing of full facial images, so it would seem that face recognition can be impacted by differences in horizontal image structure. This observation is not contradictory to the facial barcode hypothesis. While the theory predicts a disproportionate reliance on vertical image structure, it does not discount a contribution from horizontal image structure (e.g., Dakin & Watt, 2009; Goffaux & Dakin, 2010; Goffaux & Rossion, 2007).

We have distorted both image configuration and feature shapes

Facial configuration is often used as a phrase to refer specifically to the relative placement of nameable features (e.g., the nose, the eyes, the mouth), whereas feature shape is described as belonging solely to a nameable feature (e.g., the mouth). Distorting the shape of facial images, as we have done, inevitably impacts both. Consequently, one could say it is ambiguous which distortion, of feature shape or facial configuration, has disrupted face recognition in our study. However, any distortion of a nameable feature shape will impact the distances between parts of that and other facial features, so this conceptual division has been questioned (see Hosie et al., 1988; Leder & Bruce, 2000; Rhodes et al., 1993). This consideration guided our use of the term “image distribution” to describe the spatial distribution of luminance contrast changes within a facial image, as we wanted to avoid any suggestion that we were uniquely manipulating the relative placement of nameable facial features (and not

their shapes). Regardless of this caveat, our data show that it is more important that facial image distributions scale with the overall dimension of the vertical, as opposed to the horizontal, image axis. We are agnostic as to whether this implies that it is important that facial features be appropriately positioned, or appropriately shaped, given the dimensions of the image axis.

Our data relate to recognition of faces in pictures

Like most research on face perception, these data relate directly to recognition of portraiture. The overall retinal image of a picture is foreshortened when viewed from an angle (see Figure 1). As we can evidently recognize portraiture from a variety of viewing angles, face perception must be robust to the retinal image distortions from oblique angle viewing. But to what extent are these observations pertinent when viewing faces in daily life? Foreshortening as a function of viewing angle will also change the dimensions of retinal images of a real person's face; however, but different parts of a three-dimensional face will be differently foreshortened, and there is the additional problem of occlusion. Viewing a face from different angles will result in some parts becoming obstructed (for instance, the cheek by the nose) and in other parts becoming visible (e.g., one side of the head). Consequently, in these circumstances, interfeature distances would not be referenced relative to the absolute aspect ratio of a face's retinal image. Instead, interfeature distances within retinal images of real faces might be judged relative to prominent landmarks (e.g., the edge of the eyebrows, and the jaw- and hairlines).

Conclusion

Our data show that it is more important that facial image distributions be appropriately scaled along their vertical axis, relative to their horizontal axis. These data are consistent with the proposition that facial coding relies on a barcode-like neural representation, which captures how the reflectance of a face varies along the vertical image axis (Dakin & Watt, 2009; Goffaux & Dakin, 2010).

Keywords: face perception, face recognition, barcodes

Acknowledgment

This research was supported by an Australian Research Council Discovery project grant to DHA (DP0878140). The authors have no competing financial interests.

Commercial relationships: none.

Corresponding author: Morgan L. Spence.

Email: morgan.spence@uqconnect.edu.au.

Address: School of Psychology, The University of Queensland, St. Lucia, Queensland, Australia

Correspondence and requests for materials should be addressed to M. L. S.

References

- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*, 433–436. Retrieved from <http://color.psych.upenn.edu/brainard/PsychToolbox.pdf>
- Bukach, C. M., Grand, R., Kaiser, M. D., Bub, D. N., & Tanaka, J. W. (2008). Preservation of mouth region processing in two cases of prosopagnosia. *Journal of Neuropsychology, 2*(1), 227–244, doi:10.1348/174866407X231010.
- Calder, A. J., & Young, A.W. (2005). Understanding the recognition of facial identity and facial expression. *Nature Reviews, 6*, 641–651, doi:10.1038/nrn1724.
- Crookes, K., & Hayward, W. G. (2012). Face inversion disproportionately disrupts sensitivity to vertical over horizontal changes in eye position. *Journal of Experimental Psychology: Human Perception and Performance, 38*(6), 1428–1437, doi:10.1037/a0027943.
- Dahl, C. D., Logothetis, N. K., & Hoffman, K. L. (2007). Individuation and holistic processing of faces in rhesus monkeys. *Proceedings of the Royal Society B: Biological Sciences, 274*(1622), 2069–2076, doi:10.1016/0001-6918(86)90085-5.
- Dakin, S. C., & Watt, R. J. (2009). Biological “bar codes” in human faces. *Journal of Vision, 9*(4):2, 1–10, <http://www.journalofvision.org/content/9/4/2>, doi:10.1167/9.4.2.
- 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.
- Goffaux, V., & Dakin, S.C. (2010). Horizontal information drives the behavioral signatures of face processing. *Frontiers in Psychology, Perception Science, 1*, 43. doi:10.3389/fpsyg.2010.00143.
- Goffaux, V., Hault, B., Michel, C., Vuong, Q. C., & Rossion, B. (2005). The respective role of low and high spatial frequencies in supporting configural and featural processing of faces. *Perception, 34*(1), 77–86, doi:10.1068/p5370.

- Goffaux, V., & Rossion, B. (2007). Face inversion disproportionately impairs the perception of vertical but not horizontal relations between features. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(4), 995–1002, doi:10.1037/0096-1523.33.4.995.
- Hole, G. J., George, P. A., Eaves, K., & Rasek, A. (2002). Effects of geometric distortions on face recognition performance. *Perception*, *31*, 1221–1240, doi:10.1068/p3252.
- Hosie, J. A., Ellis, H. D., & Haig, N. D. (1988). The effect of feature displacement on the perception of well-known faces. *Perception*, *17*(4), 461–474.
- Leder, H., & Bruce, V. (1998). Local and relational aspects of face distinctiveness. *The Quarterly Journal of Experimental Psychology Section A*, *51*(3), 449–473, doi:10.1080/713755777.
- Leder, H., & Bruce, V. (2000). When inverted faces are recognized: The role of configural information in face recognition. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology*, *53*(2), 513–536, doi:10.1080/027249800390583.
- Leder, H., Candrian, G., Huber, O., & Bruce, V. (2001). Configural features in the context of upright and inverted faces. *Perception*, *30*(1), 73–83, doi:10.1068/p2911.
- Liu, C. H., & Ward, J. (2006). Face recognition in pictures is affected by perspective transformation but not by the centre of projection. *Perception*, *35*, 1637–1650.
- Maurer, D., Le Grand, R., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, *6*, 255–260.
- Pachai, M. V., Sekuler, A. B., & Bennett, P. J. (2013). Sensitivity to information conveyed by horizontal contours is correlated with face identification accuracy. *Frontiers in Psychology*, *4*(74), 1–9, doi: 0.3389/fpsyg.2013.00074.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442. Retrieved from <http://www.psych.nyu.edu/pelli/pubs/pelli1997videotoolbox.pdf>
- Piepers, D. W., & Robbins, R. A. (2012). A review and clarification of the terms “holistic,” “configural,” and “relational” in the face perception literature. *Frontiers in Psychology*, *3*, 559, doi:10.3389/fpsyg.2012.00559.
- Rhodes, G. (1988). Looking at faces: First-order and second-order features as determinants of facial appearance. *Perception*, *17*(1), 43–63, doi:10.1068/p170043.
- Rhodes, G., Brake, S., & Atkinson, A. (1993). What’s lost in inverted faces? *Cognition*, *47*(1), 25–57, doi: 10.1016/0010-0277(93)90061-Y.
- Sekunova, A., & Barton, J. J. (2008). The effects of face inversion on the perception of long-range and local spatial relations in eye and mouth configuration. *Journal of Experimental Psychology: Human Perception and Performance*, *34*(5), 1129–1135, doi:10.1037/0096-1523.34.5.1129.
- Sinha, P., Balas, B., Ostrovsky, Y., & Russell, R. (2006). Face recognition by humans: Nineteen results all computer vision researchers should know about. *Proceedings of the IEEE*, *94*(11), 1948–1962, doi:10.1109/JPROC.2006.884093.
- Storrs, K. R., & Arnold, D. H. (2013). Shape aftereffects reflect shape constancy operations: Appearance matters. *Journal of Experimental Psychology: Human Perception and Performance*, *39*, 616–622, doi:10.1037/a0032240.
- Tanaka, J. W., & Sengco, J. A. (1997). Features and their configuration in face recognition. *Memory and Cognition*, *25*, 1–10.
- Vetter, T., Poggio, T., & Bühlhoff, H. H. (1994). The importance of symmetry and virtual views in three-dimensional object recognition. *Current Biology*, *4*, 18–23.
- Vishwanath, D., Girshick, A. R., & Banks, M. S. (2005). Why pictures look right when viewed from the wrong place. *Nature Neuroscience*, *8*, 1401–1410, doi:10.1038/nn1553.
- Wang, Y., Fu, X., Johnston, R. A., & Yan, Z. (2013). Discriminability effect on Garner interference: Evidence from recognition of facial identity and expression. *Frontiers in Psychology*, *4*, 1–10, doi:10.3389/fpsyg.2013.00943.
- Willenbockel, V., Sadr, J., Fiset, D., Horne, G. O., Gosselin, F., & Tanaka, J. W. (2010). Controlling low-level image properties: The SHINE toolbox. *Behavior Research Methods*, *42*, 671–684.
- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81*, 141–145.