Role of ordinal contrast relationships in face encoding

Sharon Gilad[a,1], Ming Meng[b,1], and Pawan Sinha[b,2]

[a]Department of Neurobiology, Weizmann Institute of Science, Rehovot 76100, Israel; and [b]Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139

Communicated by Richard M. Held, New England College of Optometry, Cambridge, MA, December 12, 2008 (received for review September 10, 2008)

What aspects of facial information do we use to recognize individuals? One way to address this fundamental question is to study image transformations that compromise facial recognizability. The goal would be to identify factors that underlie the recognition decrement and, by extension, are likely constituents of facial encoding. To this end, we focus here on the contrast negation transformation. Contrast negated faces are remarkably difficult to recognize for reasons that are currently unclear. The dominant proposals so far are based either on negative faces' seemingly unusual pigmentation, or incorrectly computed 3D shape. Both of these explanations have been challenged by recent results. Here, we propose an alternative account based on 2D ordinal relationships, which encode local contrast polarity between a few regions of the face. Using a novel set of facial stimuli that incorporate both positive and negative contrast, we demonstrate that ordinal relationships around the eyes are major determinants of facial recognizability. Our behavioral studies suggest that destruction of these relationships in negatives likely underlies the observed recognition impairments, and our neuro-imaging data show that these relationships strongly modulate brain responses to facial images. Besides offering a potential explanation for why negative faces are hard to recognize, these results have implications for the representational vocabulary the visual system uses to encode faces.

contrast negation | face perception | fMRI | neural representation | object recognition

In principle, a contrast negated image is exactly as informative as its positive counterpart; negation perfectly preserves an image's 2D geometric and spectral structure. However, as anyone who has had to search through a roll of negatives for a snapshot of a particular person knows, this simple operation has dramatically adverse consequences on our ability to identify faces, as illustrated in Fig. 1 (1–6). Exploring the causes of this phenomenon is important for understanding the broader issue of the nature of information the visual system uses for face identification.

Several researchers have hypothesized that negated face-images are hard to recognize because of the unnatural shading cues in negatives, which compromise shape from shading processes (7–11). The resulting problems in recovering veridical 3D facial shape are believed to impair recognition performance. Although plausible, it is unclear whether this explanation is a sufficient one, especially in light of experimental results showing preserved recognition performance in the absence of shading gradients (12), and theories of face recognition that are based on the use of 2D intensity patterns rather than recovered 3D shapes (13, 14). Another prominent hypothesis is that negation causes faces to have unusual pigmentation (15, 16). However, the adequacy of this "pigmentation hypothesis" has been challenged by data showing that hue negation, which also results in unnatural pigmentation (making the entire face look bluish-green), has little effect on recognition performance (11). It is also unclear whether observers can usefully extract pigmentation information across different illumination conditions (17, 18). We propose an alternative account of negation-induced impairment. Of the infinitely many aspects of facial photometry, our results suggest that the destruction of a small set of stable 2D contrast polarity relationships might underlie negation-induced decrements in face-recognition performance.

In our past work on face representation under variable illumination conditions (19, 20), we have found that polarity of contrast around the eyes is a remarkably stable feature, with the eyes usually being darker than the forehead and the cheeks, as illustrated in Fig. 2. The absolute magnitude of contrast across different regions of a face changes greatly under different imaging/lighting conditions, but the local polarity relationships between the eye regions and their neighborhood are maintained in all but the most unnatural lighting setups (such as lighting a face from below). Dark or oily facial complexion also does not typically disrupt these relationships. Watt (21) and Chen *et al.* (22) too have remarked on the stability of these polarity relationships. The use of ordinal brightness relationships confers significant tolerance for photometric variations in images and may help explain perceptual and neural invariance to illumination changes (23, 24) and their effectiveness for classifying faces when used in machine vision systems (19, 25).

To the extent that the visual system learns object concepts by encoding regularities in the observed data (26–28), it seems likely that stable contrast polarity relationships would come to be incorporated into the facial representation used by the brain. Mismatches between this internal representation and input data will then be expected to lead to decrements in recognition.

Contrast negation of face images leads to the destruction of otherwise highly consistent polarity relations. We hypothesized that this may be a factor leading to the poor recognizability of negated faces. To test this hypothesis, we created a set of "contrast chimeras." These are faces that are photo-negatives everywhere except in the eye region (thus preserving local contrast polarity in that neighborhood). Such faces still have unnatural pigmentation and shading cues over much of their extents and present largely the same problems to shape-from-shading processes as the full negatives do. Explanations based on disruptions due to unnatural pigmentation or incorrect shape-from-shading cues would not, therefore, predict significant improvements in performance with such chimeric faces beyond the improvements derived from the intrinsic recognizability of the eyes. However, if performance in negatives is compromised, at least in part, because of the destruction of local polarity relations between the eyes and their neighborhood, performance with contrast chimeras is expected to be significantly better than that with contrast negatives.

Results

Fig. 3*A* shows a few contrast chimeras of the kind used in our experiments. It is interesting to note the large perceptual difference brought about by reinversion of the eyes. We used 24 monochrome celebrity face images spatially normalized to have the same inter-

Author contributions: S.G., M.M., and P.S. designed research; S.G., M.M., and P.S. performed research; S.G., M.M., and P.S. analyzed data; and S.G., M.M., and P.S. wrote the paper.

The authors declare no conflict of interest.

[1]S.G. and M.M. contributed equally to this work.

[2]To whom correspondence should be addressed. E-mail: psinha@mit.edu.

PSYCHOLOGY

NEUROSCIENCE

Fig. 1. Faces in photographic negatives are harder to identify relative to the original images, as illustrated in this negated version of the famous group picture from the 1927 Solvay Conference. Attendees included Schrödinger, Pauli, Heisenberg, Bragg, Dirac, Bohr, Planck, Curie, and Einstein. It is believed that this decrement in recognition is due either to unnatural pigmentation cues across the faces or an impairment of shape from shading processes that provide information about the 3D form of a face. (Photograph by Benjamin Couprie, Institut International de Physique Solvay, Brussels, Belgium.)

pupillary distance. For each original positive image (P), we created three additional variants: a full-negative (N), positive eyes on a head-silhouette (E), and a contrast chimera (C). These variants are shown in Fig. 3*B*. The "eyes on head silhouette" condition was included in our set to allow us to determine whether any observed performance gains in the chimeric condition are due merely to the intrinsic recognizability of the positive eyes on their own.

We conducted 2 experiments. The first used a between-subjects design whereas the second adopted a within-subjects approach for a finer-grain analysis of data. Fifteen subjects participated in the first experiment. They were assigned randomly to 3 equal-sized groups. Subjects in a given group were shown all faces in one of the 3 experimental conditions (full-negative, eyes on head-silhouette or contrast chimera). They were asked to name the individuals shown and were subsequently shown all faces in the full-positive condition to determine which celebrities they were familiar with. Performance for each subject was recorded as the proportion of familiar faces recognized in the experimental condition they had been assigned to. The results are shown in Fig. 3*B*. The mean level of recognition performance we found in the full-negative condition (54.35%) is consistent with results from previous studies that have compared recognizability of negatives relative to positive images (5, 15). Performance with the positive eyes on a head silhouette was poor, averaging 13.37%. Consistent with our hypothesis, however, performance with contrast-chimeras was markedly elevated (mean: 92.32%). All pairwise group differences are statistically significant with $p \ll 0.01$ (with df = 8: $t_{\mathrm{N\ vs.\ C}} = -11.3$; $t_{\mathrm{E\ vs.\ C}} = -19.02$; $t_{\mathrm{E\ vs.\ N}} = -9.3$). These results attest to the superior recognizability of contrast chimeras relative to the full-negatives and the positive eyes by themselves.

To investigate on an individual subject basis how much of the performance boost with chimeras can be accounted for by the intrinsic recognizability of the positive eyes, we conducted a second experiment that followed a within-subjects design. Ten subjects participated in the experiment (none of them had participated in the first experiment). They were first shown all of the negative face images, then the eyes on head-silhouette images, followed by the contrast chimeras and finally the fully positive images. As before, the subjects' task was to identify the individuals shown. From the collected responses for each subject, we determined the set of faces that were not recognized in the full negative condition, but were recognized in the full-positive one. For this set of faces, we calculated the subject's recognition accuracy across the other stimulus conditions.

The results are shown in Fig. 3*C*. As in experiment 1, we found that contrast chimeras were recognized at a significantly higher rate than the full negatives ($P < 0.01$, 1 tailed t test). More importantly, the results showed that the increase in performance could not be explained simply by appealing to the recognizability of the eyes by themselves ($P < 0.01$).

It is important to note that, whereas performance with contrast chimeras is high, it is slightly below that with the full-positives. There could be a few explanations for this. First, we have considered only one set of polarity relationships–those between the eyes and their neighborhood. More such polarity relationships are likely to exist in other regions of the face. The inclusion of these other relationships in the chimeras can conceivably equalize performance between the chimeras and the original positive images. Second, our explanations based on 2D ordinal relationships, and the pigmentation or shape-from-shading based accounts, are not mutually exclusive. It is possible that unusual pigmentation or unnatural shading cues, may lead to a small decrement in performance, enough to account for the differences between the positive and the chimeric conditions.

Besides the aforementioned behavioral experiments, another approach for examining the significance of ordinal relationships is to determine whether the inclusion or violation of these relations modulates neural responses to face images. To this end, we recorded brain responses to negative, positive, chimeric and eye

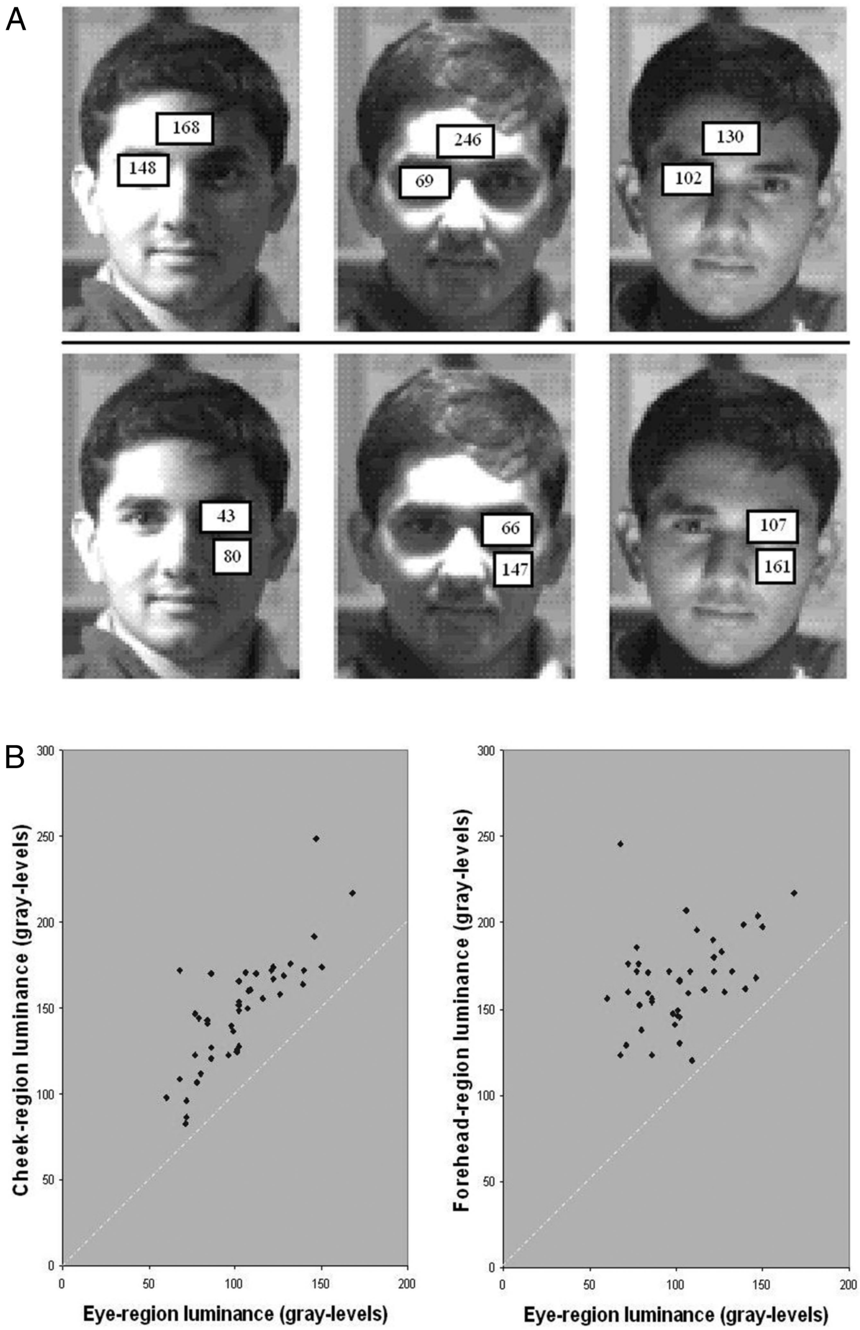

Fig. 2. The ordinal luminance relationships between the eyes and the neighboring regions on a face are stable across a range of imaging conditions. (*A*) Faces are shown under different lighting set-ups. The numbers within the boxes are their average gray-level values. Although the box-averages change significantly across different conditions, their pairwise ordinal relationships do not. For instance, the eye regions are darker than the forehead or cheeks across all conditions shown. (*B*) The scatter-plots show the average cheek and forehead luminance plotted against the average eye region luminance for fifty randomly chosen face images. The stability of these ordinal relationships is evident from the fact that the points in both plots lie on one side of the diagonal. Some of the points in these scatter-plots correspond to faces with dark and/or oily complexions, indicating that these relationships are robust across many face types. We hypothesize that contrast negation impairs recognition because it destroys these highly consistent relationships.

images, using functional magnetic resonance imaging (fMRI). Eight subjects were shown 25 instances each of these 4 classes of face images in M-sequences (29) and their brain activations recorded using a rapid event related design. For any one subject, a particular face appeared in only 1 of the 4 conditions. The condition in which a particular face was shown was counterbalanced across different participants. Subjects were asked to continuously monitor a small fixation cross (0.7° in extent) so as to avoid confounds from potentially differing patterns of scan-paths across the different stimulus conditions. They were not required to make any judgments regarding the faces shown.

Previous studies have implicated regions in the fusiform gyri [fusiform face areas (FFAs)], especially in the right cerebral hemisphere, in face perception (30). We identified the right FFA via separate localizer runs for each subject and then examined activations in this region corresponding to the 4 stimulus conditions. Fig. 4 shows the results.

Consistent with past reports (2), fully positive faces led to significantly increased brain activity in the right fusiform face area relative to fully negative faces ($F_{1,\,40} = 4.8$, $P < 0.05$). Interestingly, responses evoked by contrast chimeras are as high as, and statistically indistinguishable from, those corresponding to fully positive faces ($F_{1,\,40} = 0.11$, $P = 0.74$). By contrast, eyes embedded in the head-silhouette lead to minimal activation in the fusiform face areas, much less than positives and chimeras (P vs. E: $F_{1,\,40} = 13.76$, $P = 0.001$; C vs. E: $F_{1,\,40} = 7.18$, $P = 0.01$), and similar to full negatives ($F_{1,\,40} = 0.89$, $P = 0.35$). These results suggest that the regular polarity of contrast around the eyes is important for eliciting significant neural activation to faces. A similar trend is observed in the left fusiform face area, although the level of statistical signifi-

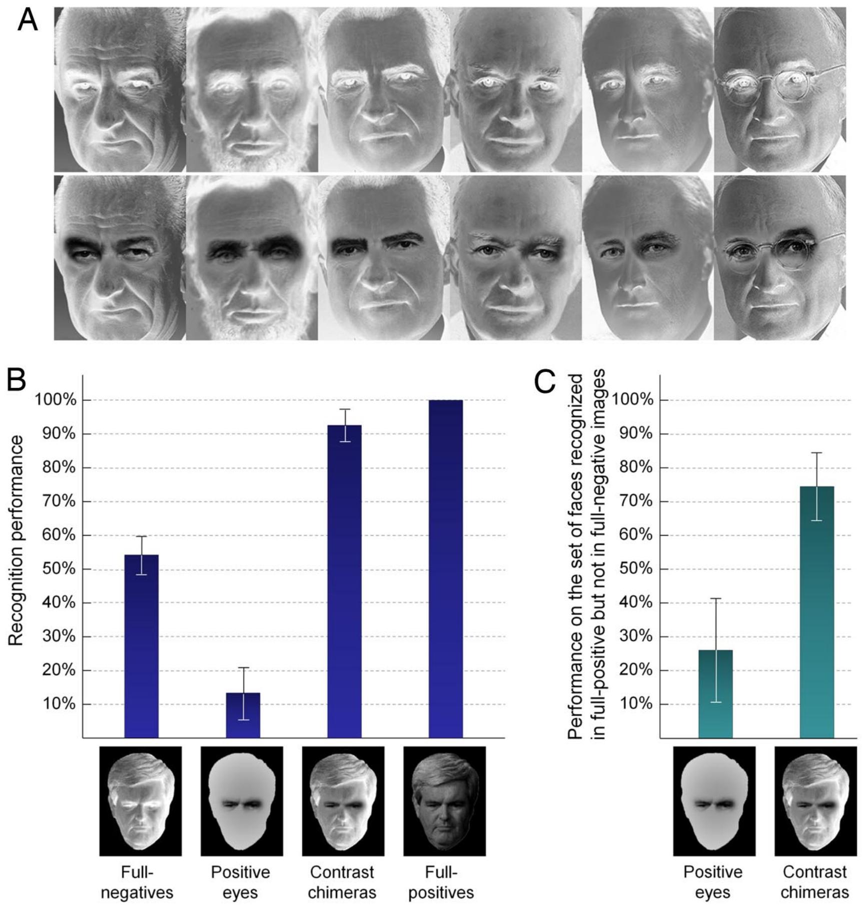

Fig. 3. Examples of contrast chimeras and behavioral recognition results. (*A*) A few negative faces (*Upper*) and the corresponding contrast chimeras (*Lower*). The only difference between these sets of images is that the chimeras are renegated in the eye-regions. They are thus a composite of negative faces and positive eyes. This local transformation significantly restores facial recognizability, as data in *B* show. The examples included here are representative of the kinds of images used in our experiments. Copyright considerations prevent us from displaying actual stimuli, which comprised images of contemporary celebrities. (*B*) Recognition performance as a function of stimulus type. Each of the 3 experimental conditions (negatives, eyes, and chimeras) was shown to separate groups of subjects, followed by the fully positive stimuli. (*C*) Recognition results from 10 subjects. Data are reported for the faces that were not recognized in the photographic negatives but were recognized in the original positives (referred to as set A in the following). This normalizes data so that performance with the negatives is 0% and with the positives, 100%. Performance with contrast chimeras is much improved relative to the full negative condition. This cannot be explained simply in terms of the intrinsic recognizability of the eyes. Although >70% of the faces in set A were recognized in the chimeric condition, <27% of set A faces were recognized based on the positive eyes embedded in head silhouettes. Error bars in *B* and *C* represent ±1 standard deviation. (Images in *A* from the White House web site, www.whitehouse.gov.)

cance is weaker because not all subjects exhibited a distinct left FFA ($p_{C\ vs.\ N} = 0.09$; $p_{C\ vs.\ P} = 0.41$; $p_{C\ vs.\ E} = 0.02$; $p_{N\ vs.\ E} = 0.30$; $p_{N\ vs.\ P} = 0.47$; $p_{E\ vs.\ P} = 0.07$).

Discussion

The imaging results demonstrate marked neural response facilitation by the reinversion of eyes in the chimeric condition. They serve to complement the behavioral findings of significant restoration of recognizability for such facial images. Taken together, these results cast new light on the long-standing question of why photographic negatives are hard to recognize. They suggest that the difficulty in analyzing negative facial images may be driven in large part by the destruction of 2D contrast polarity relations between a few key regions of the face. The special significance of the eye neighborhood is borne out by data from an additional experiment we conducted. The stimuli used in this experiment were contrast chimeras that had the mouth region (instead of the eye region) reinverted. We found that these chimeras did not have a statistically significant facilitatory effect on recognition performance. Nine subjects, different from the ones in the eye-chimera study, participated in this study. Relative to performance with full-positive images, performance with full negatives was 66.7% whereas that with mouth-positive chimeras was 71.0%. For a 2-tailed *t* test, the computed P value was 0.54. This result indicates that the increase in recognition performance depends on the specific region of the face that is reinverted, with the eyes being particularly important. Future studies might examine more exhaustively the relative perceptual significance of additional chimeric variants.

The reason why ordinal relationships around the eyes might be especially significant could be based in the statistics of facial-feature variability and diagnosticity. The eye neighborhood appears to embody a very stable set of relations across many different faces and illumination conditions. This is borne out in analyses that Viola and Jones (25) undertook toward designing a computational face-detection system. They considered many thousands of relationships across various image regions and evaluated each in terms of its ability to indicate whether the underlying image was a face or not. The top two features from this analysis are shown in Fig. 5. It is interesting to note that they both comprise contrast relationships in the eye neighborhood.

This importance of ordinal contrast relationships around the eyes may also explain a few other well-known observations regarding face recognition. A primary one is the difficulty people experience in recognizing faces that are lit from below (making the eye region brighter than the local neighborhood, thereby upsetting the polarity relationships) (10). An interesting prediction that the ordinal-relation hypothesis makes is that for bottom-lit faces, negatives should be easier to recognize than positives, because the ordinal relationships are restored by this manipulation. This counterintuitive prediction is indeed supported by empirical data. Liu and colleagues (31) found precisely such an effect in their studies of lighting direction and negation. (As an aside, it is worth noting that negatives of faces lit from above are not the same as positive images of faces lit from below.)

Another finding that the ordinal-relations hypothesis can help explain is the difficulty of recognizing unpigmented face depictions (32), such as stone busts. It is unlikely that the difficulty arises from disrupted shape-from-shading processes, because such busts present ideal conditions for the operation of these processes (33). However, they do not provide the contrast polarity relationships between the eyes and their neighborhood that real faces do. Interestingly, Bruce and Young have noted that painting dark irises on such busts makes them look much more "life-like" (7, pages

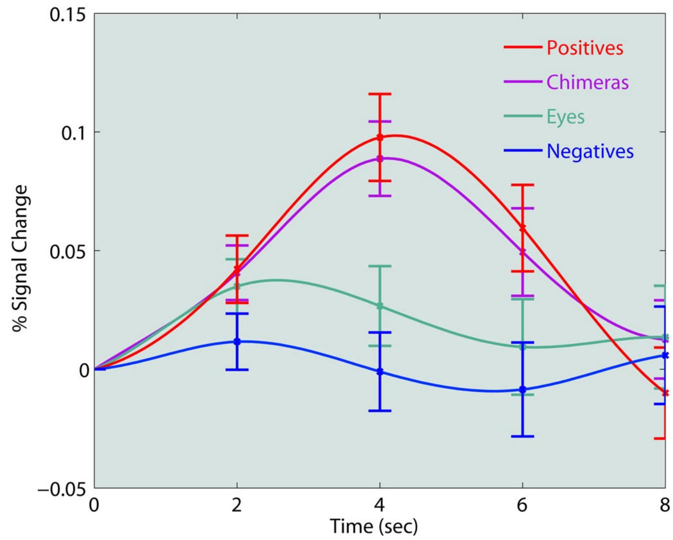

Fig. 4. Averaged activations in the right fusiform face area corresponding to 4 different facial image types: full negatives, positive eyes on a head silhouette, contrast chimeras and full positives. Percentage MR signal change was calculated by using activation at $t = 0$ seconds of each trial as the baseline. Activations corresponding to full-positives and chimeras are statistically indistinguishable from each other, but are significantly higher than those corresponding to the full negatives and the eyes on head silhouettes. Error bars indicate standard error.

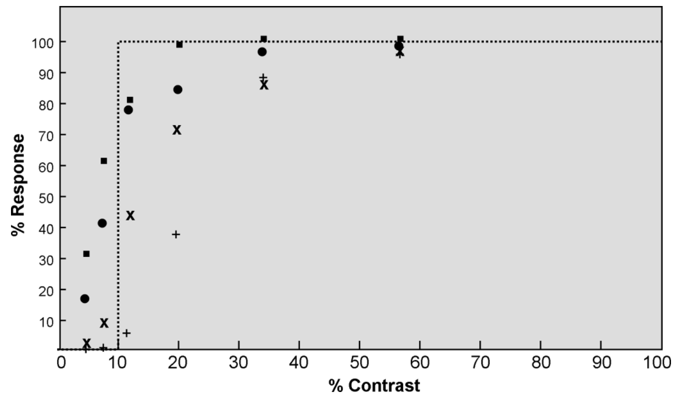

Fig. 6. Contrast response profiles of 4 cells in the primary visual cortex. It is interesting to note how rapidly responses saturate as a function of contrast. Such cells can serve as the neural substrate for extracting ordinal contrast relationships from images (neural data from ref. 39). The dotted line indicates the response profile of an ordinal comparator which responds maximally when contrast exceeds a low threshold value, and is silent otherwise.

163–165), further suggesting the role of ordinal contrast relationships in the representation of faces.

How can these findings regarding the perceptual significance of photometric relationships be reconciled with our ability to recognize line-drawings of faces? Intuitively, line-drawings appear to contain primarily contour information and very little photometric information over which to define the luminance relations. However, experimental data suggest otherwise. Studies (34, 35) have found that line-drawings that contain exclusively contour information are very difficult to recognize (specifically, they found that subjects could recognize only 47% of the line-drawings compared with 90% of the original photographs). How can we resolve such findings with the observed recognizability of line-drawings in everyday experience? Bruce and colleagues (7, 36) have convincingly argued that such depictions do in fact contain significant photometric cues, i.e., the density and weight of the lines change the relative intensity of different facial regions. In essence, the contours included in such depictions by artists correspond not merely to a low-level edge-map, but in fact embody a face's photometric structure. It is the skillful inclusion of these photometric cues that is believed to make human generated line-drawings more recognizable than computer generated ones (37). If the recognizability of line-drawings is governed in part by the photometric cues they contain, then it stands to reason that they would be susceptible to contrast negation, just as continuous-tone images are. This prediction is indeed supported by experimental data (38).

Fig. 5. The top two features selected by a procedure for determining which image measurements are most effective for classifying an image as a face. Both features comprise relationships between eyes and their local neighborhood, and suggest that the statistics of their consistency might help explain why such features are perceptually significant (after 25).

These results point to several interesting questions for future research. We mention three, and offer tentative answers for two. First, is the ordinal encoding strategy a physiologically plausible one? In other words, how can ordinal measurements be extracted by the human visual system? A potential answer might be found in the contrast response profiles of neurons in the primary visual cortex. As Fig. 6 shows, many neurons in mammalian V1 exhibit rapidly saturating responses as a function of contrast. These profiles are approximations of step-functions that would characterize an ordinal comparator. Thus, neurons in the early stages of the mammalian visual pathway provide a plausible substrate for extracting ordinal image relationships.

A second important question for future research concerns the specificity of these results to the domain of faces. Negation effects are known to be more pronounced for faces than for other classes of objects (1). Does this mean that ordinal contrast relationships are involved only in encoding faces? Not necessarily. Faces constitute a particularly homogenous class, with a very consistent set of photometric relationships within them. A more heterogeneous collection of objects, even though comprising a single category such as "houses" or "cars," might not preserve such consistency. The windows of a car can, for instance, be brighter or darker than the rest of the body. Given this variability, an ordinal code would not privilege one direction over the other. Rather, it will simply include both directions as equally valid, effectively indicating that a difference relationship ought to be present in a given image location, without mandating what the sign of the relationship must be. Such a code will not be affected by negation. Even though this seems to be too permissive an encoding strategy, it has been demonstrated to be quite effective for machine vision systems designed to recognize real-world objects (40).

A third question pertains to the relevance of these results to the study of neurological disorders such as autism, which are believed to be associated with face processing abnormalities. Specifically, individuals with autism are known to avoid eye-contact and tend to focus more on the mouth region (41, 42). If this tendency changes the relative saliencies of image regions and relationships to be weighted more toward the mouth, we would expect that mouth-chimeras would be more facilitatory for recognition performance of autistic observers rather than eye-chimeras, in contrast to the patterns of results we have described above with neuro-typical observers. Such a study would provide clues about the nature of facial representation in autism.

To summarize, our findings provide a potential answer to the long-standing question of why faces are hard to recognize in negative contrast images. More generally, they suggest that contrast polarity relations between face regions in the vicinity of the eyes might be embodied in the visual system's facial representations and serve as strong determinants of recognition performance.

Materials and Methods

Stimuli. Chimeric face images were generated by digitally compositing positive eye-regions on negative faces. The eye-region was approximately spectacle shaped, extending from the top of the eyebrows to the lower margin of the eyes and from the left to the right tips of the eyebrows. To minimize the creation of edge artifacts in the compositing process, we first changed the overall intensity of the face image so as to make the skin be midgray, and feathered the interface between positive and negative regions.

fMRI Experiment. Eight adults volunteered in the fMRI experiment. Scanning was performed on a 3.0-Tesla Siemens scanner, using a standard head coil at the Martinos Imaging Center at MIT. A high-resolution T1-weighted 3D-MPRAGE anatomical scan was acquired for each participant (FOV 256 $\times$ 256, 1 mm^3 resolution). To measure BOLD contrast, 33 slices parallel to the AC/PC line were acquired using standard T2*-weighted gradient-echo echoplanar imaging (TR 2000 ms, TE 30 ms, flip angle 90°, slice thickness 5 mm, in-plane resolution 3 $\times$ 3 mm). Six experimental scan runs and 3 ROI localization runs were performed with each participant. Stimuli were rear-projected onto a screen in the scanner bore. Experimental stimuli consisted of 100 face images in negative, positive, chimeric and eyes-only conditions. In each experimental run, these 100 faces mixed with 24 fixation-only trials were presented in a different M-sequence (29). Each trial was 2 s long. Each face image was presented for 300 ms followed by 1700 ms ISI. In addition, each run began and ended with 12 s of fixation-rest. For each participant, a particular face was presented in only 1 of the 4 conditions with 25 faces in each of the conditions. The condition in which a particular face was shown was counterbalanced across different participants. In summary, each participant saw the same 6 sequences of the same 100 faces but the faces were shown in different counterbalanced conditions. During the experiments, a green dot was randomly presented onto a random location of the screen. The fixation cross was also presented randomly in red or green. Participants were instructed to continuously monitor these changes and when the fixation was green report whether the previous display had a green dot.

fMRI Analysis. All fMRI data underwent 3-D motion correction, slice time correction and analysis, using SPM2 (www.fil.ion.ucl.ac.uk/spm/software/spm2) and custom-routines in Matlab. Slow drifts in signal intensity were removed by linear detrending, but no other temporal smoothing was applied. fMRI data of each participant were normalized into SPM-MNI space and were spatially smoothed (FWHM $=$ 6 mm). ROI localization was accomplished by 3 separate block-designed runs of full color faces vs. objects. These faces were different from the face images used in the experimental runs. ROI was defined as the set of contiguous voxels in fusiform gyrus that showed significantly stronger activation ($P <$ 10−4, uncorrected) to faces than to objects [fusiform face area (FFA) (30)]. Right FFA was reliably located in all 8 participants, whereas left FFA was found in 6 of the 8 participants. Event-related average of fMRI activity in these ROIs was sorted according to the 4 experimental conditions. Percentage MR signal change was calculated by using activation at $t = 0$ s of each trial as the baseline. Peak amplitude of 4 to 6 seconds after stimulus onset was used to perform statistics.

ACKNOWLEDGMENTS. We thank Irving Biederman, Daniel Kersten, Ami Klin, Thomas Papathomas, and Roger Tootell for their helpful comments on this work. This work was supported by an Alfred P. Sloan Foundation Fellowship in Neuroscience (to P.S.) and by the Simons Foundation.

1. Galper RE (1970) Recognition of faces in photographic negative. *Psychonomic Sci* 19:207–208.
2. George N, *et al.* (1999) Contrast polarity and face recognition in the human fusiform gyrus. *Nat Neurosci* 2:574–580.
3. Hayes T, Morrone MC, Burr DC (1986) Recognition of positive and negative bandpass-filtered images. *Perception* 15:595–602.
4. Nederhouser M, Yue X, Mangini MC, Biederman I (2007) The deleterious effect of contrast reversal on recognition is unique to faces, not objects. *Vision Res* 47:2134–2142.
5. Phillips RJ (1972) Why are faces hard to recognize in photographic negative. *Perception Psychophys* 12:425–426.
6. White M (2001) Effect of photographic negation on matching the expressions and identities of faces. *Perception* 30:969–981.
7. Bruce V, Young A (1998) *In the Eye of the Beholder—The Science of Face Perception* (Oxford Univ Press).
8. Cavanagh P, Leclerc YG (1989) Shape from shadows. *J Exp Psychol Hum Percept Perform* 15:3–27.
9. Hill H, Bruce V (1996) Effects of lighting on the perception of facial surfaces. *J Exp Psychol Hum Percept Perform* 22:986–1004.
10. Johnston A, Hill H, Carman N (1992) Recognising faces: Effects of lighting direction, inversion, and brightness reversal. *Perception* 21:365–375.
11. Kemp R, Pike G, White P, Musselman A (1996) Perception and recognition of normal and negative faces: The role of shape from shading and pigmentation cues. *Perception* 25:37–52.
12. Moscovitch M, Winocur G, Behrmann M (1997) What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. *J Cognit Neurosci* 9:555–604.
13. Beymer D, Poggio T (1996) Image representations for visual learning. *Science* 272:1905–1909.
14. Turk M, Pentland A (1991) Eigenfaces for Recognition. *J Cognit Neurosci* 3:71–86.
15. Bruce V, Langton S (1994) The use of pigmentation and shading information in recognising the sex and identities of faces. *Perception* 23:803–822.
16. Vuong QC, Peissig JJ, Harrison MC, Tarr MJ (2005) The role of surface pigmentation for recognition revealed by contrast reversal in faces and Greebles. *Vision Res* 45:1213–1223.
17. Braje WL (2003) Illumination encoding in face recognition: Effect of position shift. *J Vision*3:161–170.
18. O'Toole AJ, *et al.* (2007) Face recognition algorithms surpass humans matching faces over changes in illumination. *IEEE Trans Pattern Anal Mach Intell* 29:1642–1646.
19. Balas BJ, Sinha P (2006) Region-based representations for face recognition. *ACM Trans Appl Perception* 3:354–375.
20. Sinha P (2002) Qualitative representations for recognition. *Lecture Notes in Computer Science*, eds Wallraven C, Buelthoff HH (Springer, Berlin), Vol 2525.
21. Watt R (1994) A computational examination of image segmentation and the initial stages of human vision. *Perception* 23:383–398.
22. Chen HF, Belhumeur P, Jacobs DW (2000) In search of illumination invariants. *Proc IEEE Conference Comp VisionPattern Recognition* 1:254–261.
23. Braje WL, Legge GE, Kersten D (2000) Invariant recognition of natural objects in the presence of shadows. *Perception* 29:383–398.
24. Vogels R, Biederman I (2002) Effects of illumination intensity and direction on object coding in macaque inferior temporal cortex. *Cereb Cortex* 12:756–766.
25. Viola P, Jones MJ (2004) Robust real-time face detection. *Int J Comp Vision*57:137–154.
26. Fiser J, Aslin RN (2002) Statistical learning of new visual feature combinations by infants. *Proc Natl Acad Sci USA* 99:15822–15826.
27. Purves D & Lotto RB (2003) *Why We See What We Do: An Empirical Theory of Vision* (Sinauer, Sunderland, MA).
28. Wallis G, Bulthoff H (1999) Learning to recognize objects. *Trends Cognit Sci* 3:22–31.
29. Buracas GT, Boynton GM (2002) Efficient design of event-related fMRI experiments using M-sequences. *NeuroImage* 16(3 Pt 1):801–813.
30. Kanwisher N, McDermott J, Chun MM (1997) The fusiform face area: A module in human extrastriate cortex specialized for face perception. *J Neurosci* 17:4302–4311.
31. Liu CH, Collin CA, Burton AM, Chaudhuri A (1999) Lighting direction affects recognition of untextured faces in photographic positive and negative. *Vision Res* 39:4003–4009.
32. Bruce V, *et al.* (1991) Recognising facial surfaces. *Perception* 20:755–769.
33. Horn BKP, Brooks MJ (1989) *Shape from Shading* (MIT Press, Cambridge, MA).
34. Davies G, Ellis H, Shepherd J (1978) Face recognition accuracy as a function of mode of representation. *J Appl Psychol* 63:180–187.
35. Rhodes G, Brennan S, Carey S (1987) Identification and ratings of caricatures—Implications for mental representations of faces. *Cognit Psychol* 19:473–497.
36. Bruce V, Hanna E, Dench N, Healey P, Burton M (1992) The importance of ''mass'' in line-drawings of faces. *Appl Cognit Psychol* 6:619–628.
37. Pearson DE, Robinson JA (1985) Visual communication at very low data rates. *Proc IEEE* 73:795–812.
38. Pearson DE, Hanna E, Martinez K (1990) *Images and Understanding*, eds Barlow H, Blakemore C, Weston-Smith M (Cambridge Univ Press), pp 46–60.
39. Albrecht DG, Hamilton DB (1982) Striate cortex of monkey and cat: Contrast response function. *J Neurophysiol* 48:217–237.
40. Oren M, Papageorgiou C, Sinha P, Osuna E, Poggio T (1997) Pedestrian detection using wavelet templates. *Proc IEEE Conference Comp Vision Pattern Recognition*, pp 193–199.
41. American Psychiatric Association (2000) *Diagnostic and Statistical Manual of Mental Disorders* (American Psychiatric Association, Arlington, VA) 4th Ed.
42. Klin A, Jones W, Schultz R, Volkmar F, Cohen D (2002) Defining and quantifying the social phenotype in autism. *Am J Psychiatry* 159:895–908.