Available online at www.sciencedirect.com
SCIENCE DIRECT®
Vision Research 45 (2005) 1213–1223
Vision Research
www.elsevier.com/locate/visres

The role of surface pigmentation for recognition revealed by contrast reversal in faces and Greebles

Quoc C. Vuong [a,c,1], Jessie J. Peissig [b,*,1], Marianne C. Harrison [c], Michael J. Tarr [c]

[a] Max Planck Institute for Biological Cybernetics, Spemannstr. 38, 72076 Tübingen, Germany
[b] Department of Neuroscience, Brown University, Box 1978, 190 Thayer Street, Providence, RI 02912, USA
[c] Department of Cognitive and Linguistic Sciences, Brown University, 190 Thayer Street, Providence, RI 02912, USA

Received 10 November 2004

Abstract

Faces are difficult to recognize when viewed as negatives [Galper (1970). Recognition of faces in photographic negative. *Psychonomic Science*, *19*, 207]. Here we examined the contribution of surface properties to this contrast effect, and whether it is modulated by object category. We tested observers in a matching task using faces or Greebles, presented with or without pigmentation. When stimulus pairs were shown with mismatched contrast (e.g., positive–negative), there was a decrement in performance. This decrement was larger when the stimuli were shown with pigmentation, and this difference was more pronounced with faces than with Greebles. Overall, contrast reversal disrupts the recognition of both faces and objects to a greater degree in the presence of pigmentation, suggesting that surface properties are important components of the object representation.

Keywords: Object recognition; Face recognition; Shading; Texture

1. Introduction

The visual input contains a rich array of information for recognizing familiar and unfamiliar objects in the environment. However, it is often hypothesized that shape is the critical source of information for recognition (Biederman, 1987; Marr & Nishihara, 1978). At the same time, other visible properties (e.g., motion, shading, texture, and so on) may also provide cues to an object's identity rather than serving strictly as precursors for shape recovery. In this vein, researchers have begun to examine how properties of visible surfaces contribute directly to the recognition process, the underlying assumption being that such properties are integral to the object representation (e.g., Hayward & Williams, 2000; Naor-Raz, Tarr, & Kersten, 2003; Rossion & Pourtois, 2004; Tarr, Kersten, & Bülthoff, 1998; Wurm, Legge, Isenberg, & Luebker, 1993). The difficulty here is that how surface properties appear to an observer depends, in part, on an object's shape and the conditions under which that object is viewed; for example, facial luminance variations will depend jointly on skin pigmentation and shading (which in turn is a product of lighting direction(s), surface curvature, and pose relative to the observer). Moreover, the functional contribution of surface properties to recognition may also depend on factors such as visual similarity among objects (e.g., Price & Humphreys, 1989) or observers' prior familiarity and visual expertise with the category of objects (e.g., Gauthier, Williams, Tarr, & Tanaka, 1998).

Here we examined the contribution of surface pigmentation by reversing the luminance of our stimuli so

* Corresponding author. Tel.: +1 401 863 1196.
E-mail addresses: quoc.vuong@tuebingen.mpg.de (Q.C. Vuong), jessie_peissig@brown.edu (J.J. Peissig).
[1] The first two authors contributed equally.

0042-6989/$ - see front matter © 2004 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2004.11.015

that light regions of an image become dark and vice versa. This manipulation can affect the perception of surface textures and shading, both of which may be used to recover shape or used as cues to an object's identity (e.g., a zebra's stripes). In the face recognition literature, how observers interpret luminance variations in an image (e.g., as surface pigmentation, as shading, as shadows, etc.) has been found to affect their ability to recognize faces (e.g., Bruce & Langton, 1994; Kemp, Pike, White, & Musselman, 1996; Liu, Collin, Burton, & Chaudhuri, 1999). As raised above, the problem from a computational perspective is that this interpretation is confounded by the fact that the luminance gradient of a face (or any object) is the product of many different factors, such as shape, skin coloration, and the lighting conditions under which a face is seen. To address this problem, the goal in the present study was twofold. First, using computer graphics, we attempted to tease apart the contributions of facial shape and facial skin pigmentation to face recognition by presenting the same faces with and without pigmentation (Blanz & Vetter, 1999; see also O'Toole, Vetter, & Blanz, 1999; Troje & Bülthoff, 1996). Second, we examined whether the effects of pigmentation on recognition are restricted to faces by directly comparing faces and novel "Greebles" within the same recognition paradigm (Gauthier & Tarr, 1997; http://www.tarrlab.org/stimuli). A comparison across these two stimulus categories also addresses the degree to which a decrease in recognition performance as a result of contrast reversal is accounted for by familiarity and expertise with the objects used as stimuli.

Galper (1970) initially demonstrated that contrast reversal made faces more difficult to recognize. This *contrast effect* is easily demonstrated by attempting to recognize individuals in photographic negatives. She suggested that the deficit in recognizing photographic-negative faces was due to the inability to accurately perceive facial expressions. However, White (2001) demonstrated that observers exhibit no deficit in identifying facial expressions when faces are shown in reverse contrast, but that these same observers did show a significant performance decrement for matching identities of contrast-reversed faces. Thus, whatever properties of faces are disrupted by contrast reversal, they seem tied to the additional perceptual analysis required to make identity judgments, presumably one of the most difficult object discriminations faced by our visual systems.

One alternative to the idea that it is the difficulty of individual identity judgments that underlies the contrast effect is that this effect is related to the particular category of objects in question. That is, for categories such as faces, the configuration of parts is critical for accurate face recognition (e.g., Tanaka & Farah, 1993); therefore, reversing image contrast may affect observers' ability to recover facial shape and/or parts (Cavanagh & Leclerc, 1989; Kemp et al., 1996). Thus, the configural information available will be different for positive and negative faces (White, 2001). However, configural changes are typically not nearly as disruptive in the recognition of non-face objects. Consequently, reversing the contrast of most non-face objects might not disrupt recognition to the same extent as it does for faces. To test this alternative, Subramaniam and Biederman (1997) used a sequential-matching task to compare the effects of contrast reversal for faces and chairs. Observers determined whether two sequentially presented stimuli were the same or different. Both stimuli had either matched contrast (e.g., positive–positive) or mismatched contrast (e.g., positive–negative). Recognition performance for faces, but not chairs, was worse when face pairs differed in their contrast polarity. This finding supports the face-specificity of the contrast effect.

Finally, it is possible to invoke category-general object recognition processes to explain the contrast effect. For example, Bruce and Langton (1994) hypothesized that a contributing factor is facial pigmentation. Here pigmentation refers to surface variations that are non-uniformly distributed; that is, changes in coloration arising from physical sources such as the material or surface markings independent of shape (e.g., a mole on a face). In Bruce and Langton's study face images were rendered from 3D head models to remove variations in image luminance due to pigmentation. They found that observers recognized these faces equally well whether they were shown in positive or negative contrast, suggesting that pigmentation contributed to the contrast effect. Kemp et al. (1996), on the other hand, argued that shape-from-shading cues, rather than pigmentation, are disrupted when faces are shown in negative contrast. They found that reversing the luminance while maintaining the hue of color photographs impaired recognition performance, but that reversing the hue while maintaining the luminance did not. Kemp et al. argued that luminance reversal disrupts shape-from-shading whereas hue reversal disrupts pigmentation. Lastly, Liu et al. (1999) proposed that contrast reversal disrupts learned constraints that the visual system uses for object recognition, such as the assumption that objects are typically lit from above (e.g., Ramachandran, 1988). They argued that reversing the contrast is similar to changing the lighting direction from top-lit to bottom-lit. Consistent with this claim, Liu et al. reported an interaction between the lighting direction and the contrast effect. Although each of these explanations may individually account for how contrast reversal affects object recognition, they are not necessarily mutually exclusive. Thus, it is possible that multiple factors (as outlined above) may together determine how observers recognize faces and objects. Here we focused on one such factor: the potential contribution of surface pigmentation to recognition.

Because the aforementioned studies emphasize face recognition and rarely provide a control condition using non-face objects, they do not address the more general issue of the extent to which surface properties per se contribute to recognition, and how such properties may interact with observers' familiarity with the object category. There has been no direct comparison of contrast reversal with pigmented and non-pigmented versions of the same stimuli, with illumination, shading, shadows, and so on, held constant across categories (e.g., Bruce & Langton, 1994; Gauthier et al., 1998; but see Nederhouser, Mangini, Biederman, & Okada, 2003). Thus, the question of what the contrast effect reveals about the underlying computations that are used in face and object recognition remains unanswered.

In the present study we examined the role of pigmentation in recognizing faces, but extended the paradigm by including a homogeneous non-face object category—Greebles. Thus, we are able to address the question of the extent to which pigmentation (and surface properties more generally) interact with category familiarity. We chose Greebles as our control category for two reasons. First, like faces, Greebles belong to a homogeneous category in which individual members have curved surfaces and a similar configuration of parts (Gauthier & Tarr, 1997). Second, although Greebles are arguably "face-like", both behavioral and brain-imaging studies have shown that faces and Greebles are processed differently and by different neural substrates, at least initially (Gauthier & Tarr, 1997; Gauthier et al., 1998; Tarr & Gauthier, 2000). However, following training, faces and Greebles are processed more similarly. For example, naïve Greeble observers do not show an inversion effect (Yin, 1969) with Greebles but do with faces. Trained observers, on the other hand, also show an inversion effect with Greebles (e.g., Gauthier et al., 1998). In addition, similar neural substrates become recruited for both categories following training (e.g., Tarr & Gauthier, 2000).

In our study, observers were shown pairs of faces or pairs of Greebles, with or without surface pigmentation (see Figs. 1 and 2). The image pairs were either contrast matched or contrast mismatched. When pigmentation is present, variations in image luminance are due to a combination of shading and pigmentation. By comparison, if an object has a surface that uniformly reflects light at all points, variations in image luminance are only due to shading, i.e., how the object's shape interacts with the light source and the observers' viewpoint (e.g., Horn, 1975). Because we present the same faces and Greebles in both of these conditions (pigmented and uniform reflectance), we can directly test the extent to which surface pigmentation contributes to the recognition of faces and non-face objects (unconfounded by category) by explicitly controlling other factors such as lighting condition, familiarity, and stimulus repetition. Our hypothesis is that surface properties (e.g., color, texture) are encoded in the object representation, in addition to possible effects of shading and shadows on surfaces (e.g., Tarr et al., 1998). Consequently, contrast reversal will be more detrimental for recognizing pigmented than uniform stimuli, as pigmented stimuli contain an additional cue to identity—a prediction that is consistent with previous results (e.g., Bruce & Langton, 1994). Second, if we observe a contrast effect for Greebles, it suggests that surface properties are also integral components of the representation of unfamiliar, non-face objects, or at a minimum not exclusive to a specific object category, for example, faces.

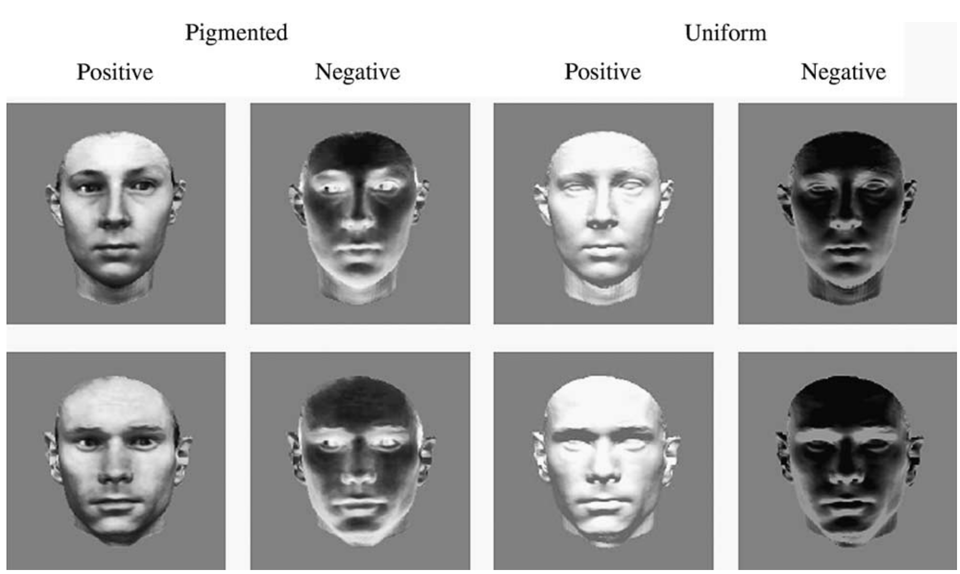

Fig. 1. Examples of face stimuli: pigmented and uniform faces shown with positive and negative contrast.

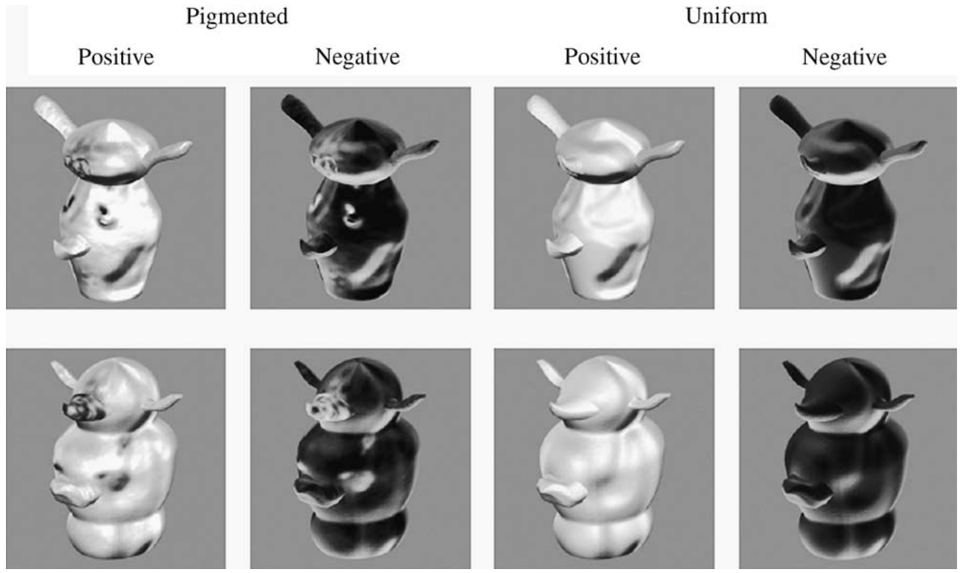

Fig. 2. Examples of Greeble stimuli: pigmented and uniform Greebles shown with positive and negative contrast.

2. Method

2.1. Participants

Forty volunteers from Brown University participated in this experiment for pay. Informed consent was obtained.

2.2. Stimuli

Figs. 1 and 2 present examples of faces and Greebles, respectively, with or without surface pigmentation, and in positive and negative contrast. The face stimuli consisted of 100 male and 100 female full-front faces from the database collected at the Max Planck Institute for Biological Cybernetics (Blanz & Vetter, 1999; http://faces.kyb.tuebingen.mpg.de/). The Greeble stimuli consisted of variants of the original Greebles (Gauthier & Tarr, 1997) modeled in 3D Studio Max Version 4.0 (Discreet; Montreal, CANADA). There were 96 Greebles divided into three families of 32 Greebles. Members of each family shared a common body shape, which was deformed to differing degrees (four possible deformations). Each member of a family also had different parts and configurations of these parts (eight possible configurations). Faces and Greebles were presented as 256-level grayscale images against a mean gray level background, and subtended approximately 8°–10° of visual angle.

There was a *pigmented* and a *uniform-reflectance* version of each stimulus. Both uniform faces and Greebles used a white texture (i.e., pixel value of 255 at all pixel locations). Pigmented faces were rendered with the texture acquired from a 3D laser scanner mapped onto the corresponding 3D head models. The head models were lit from directly in front and above.

Pigmented Greebles had non-uniform textures mapped onto the corresponding largest central component as well as individual parts. Fig. 3 presents examples of the two sets of textures that were used in the present study. One set was mapped to the largest central component of each Greeble, and the second set was mapped to the upper middle part of that Greeble. For the textures mapped to the central component, we copied regions of a few randomly selected faces (mostly forehead and cheek regions) using the Clone tool in Photoshop (Adobe Inc.). We then added different dark "spots" of various sizes and shapes to roughly the same regions across the different textures. For the second set of textures, we created a smaller set of dark gray noise textures. There were 12 textures in the first set, and four in the second set. All textures were initially created in color and then subsequently converted to grayscale. Following the texture mapping, the Greebles were rendered from a three-quarter view. Light sources were arbitrarily positioned in the same locations around each Greeble model.

The negative-contrast stimuli were created from their positive-contrast counterparts by subtracting the 8-bit grayscale value of each pixel of the positive stimulus from 255 (the maximum value). Although this subtraction is a common method for creating negatives, it reduces the luminance of the negative stimulus due to the nonlinear relationship between the 8-bit pixel value and the display luminance. To control for this confound, Liu et al. (1999) re-created their negative-contrast stimuli in the luminance domain so that both positive and negative faces had the same (mean) lumi-

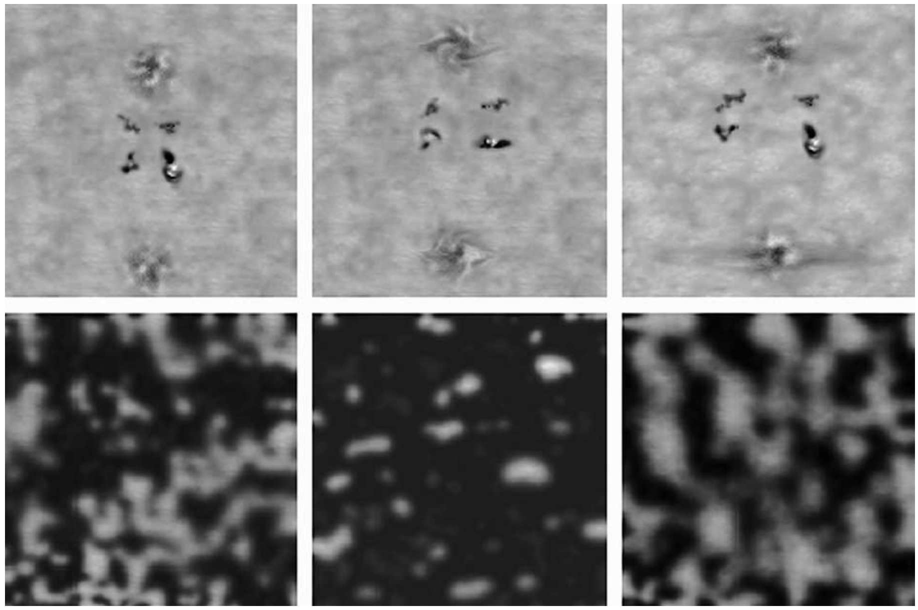

Fig. 3. Example textures used on the novel Greeble stimuli. The top row illustrates different textures used for the largest central component of the Greebles (12 different textures total). The bottom row illustrates the textures used for the upper middle part of the Greebles (four different textures total).

nance. Critically, they found similar results regardless of the method used to create the negative stimuli.

Masks for faces and Greebles consisted of rectangular pieces cropped from a few sample stimuli. The cropped pieces were pasted to create an 8°–10° square (the size of the mask was matched to the size of the preceding stimulus). All masks were created from stimuli with the same contrast polarity as the image that preceded it.

2.3. Design

Stimulus (faces, Greebles) was tested between participants: 20 subjects were tested with faces and 20 were tested with Greebles. For each stimulus type there were three within-participants factors: Trial Type (same, different), Surface (pigmented, uniform), and Contrast (matched, mismatched). For matched contrast, both stimuli had either positive or negative contrasts. For mismatched contrast, either the first or second stimulus in the sequence was positive.

For participants presented with faces, 72 of the 100 faces of each sex were used on experimental trials; the remainder was used on practice trials. The selection and assignment of faces to conditions was randomly determined for each participant. There were 32 practice trials in which feedback was provided. There were 96 experimental trials (48 with male pairs, 48 with female pairs) in which no feedback was provided. This block of trials was presented three times in a different random order (288 trials total). Likewise, for participants presented with Greebles, 72 of the 96 Greebles (24 from each family) were randomly selected and used on experimental trials; the remainder was used on practice trials. These stimuli were randomly assigned to each of the eight possible conditions with the constraint that, on *different* trials, the two Greebles were members of the same family. Participants practiced with three blocks of 16 trials in which feedback was provided. There were 48 experimental trials repeated in a different random order for six blocks (288 trials total). No feedback was provided on these trials.

2.4. Procedure

A sequential-matching task was used in which participants judged whether two stimuli (faces or Greebles) depicted the same individuals. Each trial sequence for faces proceeded as follows: a central fixation cross for 500 ms, the first face for 150 ms, a mask for 500 ms, the second face for 150 ms, and a second mask for 500 ms. For Greebles, each trial sequence proceeded as follows: a fixation cross for 500 ms, the first Greeble for 250 ms, a mask for 250 ms, the second Greeble for 250 ms, a second mask for 250 ms. The presentation time for the Greebles was increased to avoid floor effects. In addition to masking, two further manipulations were used to prevent image matching: First, both images and their masks were spatially shifted randomly by up to 50 pixels horizontally and vertically, and second, the first or second stimulus was reduced in size by 15%.

Participants were instructed to respond as quickly and as accurately as possible following the presentation of the second stimulus by pressing the "same" or "different" key. They were instructed to ignore changes to position, size, and "color" when making their responses. There were practice trials to familiarize participants with the procedure. Feedback was provided during practice to ensure that participants understood which changes to disregard. The stimuli were presented on an iMac CRT monitor with a 1024 pixel × 768 pixel resolution. No correction was used to linearize the pixel luminance of the monitor output, although Macintosh OS9's built-in color tools and default color-table were used to set the response of the iMac CRT. Participants sat approximately 50 cm from the monitor. The experiment was programmed using MATLAB Release 5.0 (Mathworks, Natick, MA) and PsychToolbox (http://www.psychtoolbox.org/; Brainard, 1997; Pelli, 1997).

3. Results

The analyses were based on sensitivity (d'), and correct response times (RTs) from *same* trials. Sensitivity was used instead of accuracy to account for response biases. For computing d', hits were defined as responding "same" on *same* trials, and false alarms were defined as responding "same" on *different* trials. We also controlled for outliers and anticipatory responses by removing RTs greater than 2000 ms and less than 300 ms. The sensitivity and RT data were then submitted to a mixed-design Analysis of Variance (ANOVA) with Stimulus (faces, Greebles) as a between-participant factor; and Surface (pigmented, uniform) and Contrast (matched, mismatched) as within-participant factors. The significance level for all analyses reported was set to $p = 0.05$.

Figs. 4 and 5 plot mean d' and RTs, respectively, for faces across the remaining conditions. Similarly, Figs. 6 and 7 plot mean d' and RTs, respectively, for Greebles. Tables 1 and 2 show means and standard errors for proportion correct responses and RT on *same* and *different* trials for faces and Greebles, respectively.

For sensitivity, there was a significant main effect of Contrast, $F(1,38) = 54.95$, $p < 0.001$. In addition, the interactions between Contrast and Surface, $F(1,38) = 7.24$, $p < 0.05$; and Contrast and Stimulus, $F(1,38) = 5.94$, $p < 0.05$, were significant. For RTs, there was a significant main effect of Surface, $F(1,38) = 14.15$, $p < 0.001$; and of Contrast, $F(1,38) = 61.09$, $p < 0.001$. The interaction between Contrast and Stimulus was significant, $F(1,38) = 4.93$, $p < 0.05$; and the interaction between Contrast and Surface was marginally significant, $F(1,38) = 3.23$, $p = 0.08$. Consistent with previous studies, there was a robust contrast effect when recognizing unfamiliar faces; that is, observers were faster and more accurate when the two faces had the same

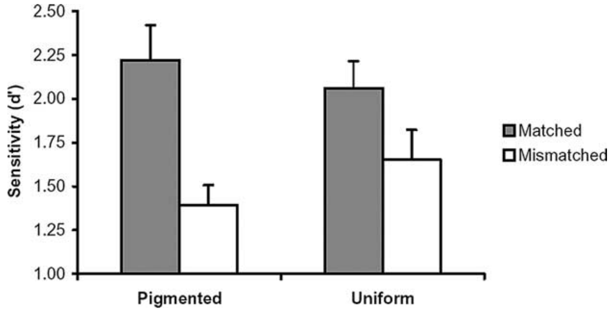

Fig. 4. Mean d' scores for faces as a function of Surface (pigmented, uniform) and Contrast (matched, mismatched). Error bars are +1 standard error of subject means.

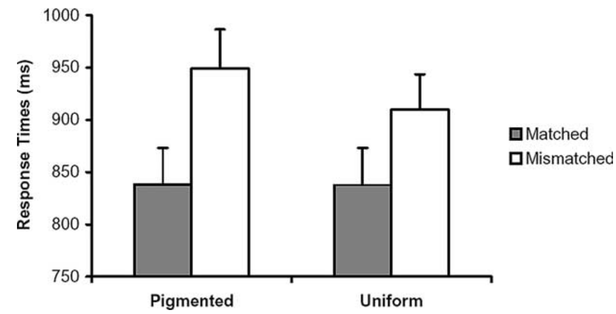

Fig. 5. Mean correct response times on *same* trials for faces as a function of Surface (pigmented, uniform) and Contrast (matched, mismatched). Error bars are +1 standard error of subject means.

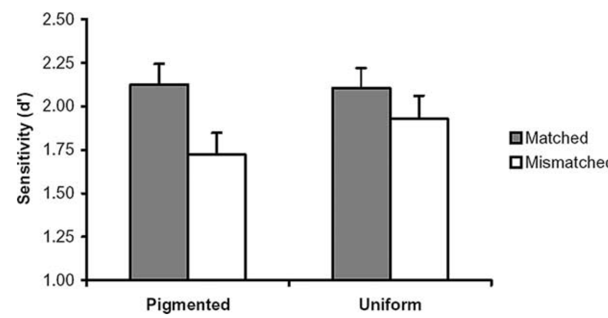

Fig. 6. Mean d' scores for Greebles as a function of Surface (pigmented, uniform) and Contrast (matched, mismatched). Error bars are +1 standard error of subject means.

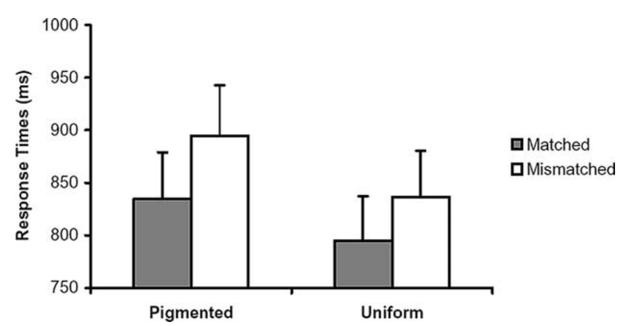

Fig. 7. Mean correct response times on *same* trials for Greebles as a function of Surface (pigmented, uniform) and Contrast (matched, mismatched). Error bars are +1 standard error of subject means.

Table 1
Mean (standard errors of the means) proportion correct and correct response times for faces on *same* and *different* trials

		Same		Different	
		Proportion correct	RT (ms)	Proportion correct	RT (ms)
Pigmented	Matched	0.90 (0.01)	837 (23)	0.77 (0.03)	879 (30)
	Mismatched	0.68 (0.03)	951 (26)	0.76 (0.03)	916 (30)
Uniform	Matched	0.92 (0.01)	834 (24)	0.68 (0.03)	884 (29)
	Mismatched	0.76 (0.02)	912 (23)	0.75 (0.03)	903 (29)

Table 2
Mean (standard errors of the means) proportion correct and correct response times for Greebles on *same* and *different* trials

		Same		Different	
		Proportion correct	RT (ms)	Proportion correct	RT (ms)
Pigmented	Matched	0.94 (0.01)	835 (20)	0.66 (0.04)	849 (44)
	Mismatched	0.83 (0.02)	899 (23)	0.75 (0.04)	847 (40)
Uniform	Matched	0.94 (0.01)	795 (19)	0.65 (0.03)	868 (39)
	Mismatched	0.91 (0.01)	839 (20)	0.60 (0.04)	876 (43)

contrast polarity (e.g., Bruce & Langton, 1994; Galper, 1970; Liu et al., 1999; Subramaniam & Biederman, 1997). At the same time, to the best of our knowledge, this is the first study to report a similar contrast effect for recognizing novel and unfamiliar objects, at least under some conditions (see also Gauthier et al., 1998). Note that on *different* trials there is little performance difference between contrast-matched and contrast-mismatched trials (Tables 1 and 2); note also that observers might have used any of a large number of different features to perform the task on these trials. Thus, it appears that the effects of contrast reversal on sensitivity and response times for both faces and Greebles are largely driven by *same* trials. Moreover, for sensitivity, observers performed nearly at ceiling (accuracy of 90% or better) on contrast-matched *same* trials, with performance being much worse on contrast-mismatched *same* trials. These findings suggest that contrast reversal may lower observers' decision thresholds for matching the identity of two stimuli that do not physically match (recall that, in addition to contrast, the size and position of the two stimuli differed on contrast-mismatched *same* trials). That said, relative to the questions we are addressing, it should be pointed out that the presence of surface pigmentation significantly affects this decision threshold for *both* categories of objects tested.

The critical results in the present experiment, however, are the interactions between contrast and surface properties for faces and Greebles. For faces, Tukey's Honestly Significant Difference (HSD) test revealed a significant contrast effect for both pigmented and uniform faces in sensitivity and RTs. We also computed the absolute mean difference between *mismatched–matched* contrast separately for pigmented and uniform surfaces, which measures the magnitude of the contrast effect for these two surface types. A significant difference in the magnitude was observed in both sensitivity, $t(19) = 2.39$, $p < 0.01$ ($M_{\text{pigmented}} = 0.84$, $SE_{\text{pigmented}} = 0.13$, $M_{\text{uniform}} = 0.51$, $SE_{\text{uniform}} = 0.06$), and response times, $t(19) = 1.63$, $p < 0.05$ ($M_{\text{pigmented}} = 110$ ms, $SE_{\text{pigmented}} = 24$ ms, $M_{\text{uniform}} = 72$ ms, $SE_{\text{uniform}} = 12$ ms).

For Greebles, Tukey's HSD test revealed a significant contrast effect for pigmented but not uniform Greebles in sensitivity, and a significant contrast effect for both pigmented and uniform Greebles in response times. The computed absolute mean difference between *mismatched–matched* contrast was marginally larger for pigmented Greebles than for uniform Greebles in sensitivity, $t(19) = 1.33$, $p = 0.10$ ($M_{\text{pigmented}} = 0.65$, $SE_{\text{pigmented}} = 0.10$, $M_{\text{uniform}} = 0.46$, $SE_{\text{uniform}} = 0.08$), but not in RTs, $t(19) = 0.86$ ($M_{\text{pigmented}} = 60$ ms, $SE_{\text{pigmented}} = 19$ ms, $M_{\text{uniform}} = 42$ ms, $SE_{\text{uniform}} = 10$ ms). Overall, these findings stand in contrast to those reported by Nederhouser et al. (2003). Using a two-alternative forced-choice matching task, they found no contrast effect for their novel "blob" objects with or without pigments. This null effect was found even after participants became "experts" with those stimuli.

One possible explanation for the observed interaction between contrast and surface properties is the presence of luminance changes on mismatched-contrast trials; that is, luminance differences may be larger for pigmented than for uniform stimuli. To test this possibility, we computed the correlation between the difference in *mean luminance* of the two images presented on each trial and observers' mean response times and accuracy on that trial. As shown in Tables 3 and 4, there were no significant correlations between either response time or accuracy and mean luminance difference. Thus, it appears unlikely that observers relied on this low-level feature to perform the discrimination task.

Table 3
Correlation between the difference in mean luminance of the two face images that were presented on each trial and observers' response time and accuracy on that trial

		Same		Different	
		Proportion correct	RT (ms)	Proportion correct	RT (ms)
Pigmented	Matched	0.06	0.00	0.05	−0.06
	Mismatched	0.08	−0.09	0.02	0.00
Uniform	Matched	0.05	0.01	−0.05	−0.03
	Mismatched	−0.02	0.01	0.03	−0.01

Table 4
Correlation between the difference in mean luminance of the two Greeble images that were presented on each trial and observers' response time and accuracy on that trial

		Same		Different	
		Proportion correct	RT (ms)	Proportion correct	RT (ms)
Pigmented	Matched	0.05	−0.03	0.09	0.04
	Mismatched	−0.05	0.02	0.06	−0.06
Uniform	Matched	0.05	−0.03	−0.08	−0.01
	Mismatched	0.06	0.07	0.19	−0.05

4. Discussion

In this study we explored the role of surface properties in visual object recognition. Because inverting the contrast affects the perceptual salience of some surface properties and not others (Galper, 1970), we relied on this manipulation to compare the recognition of faces and Greebles that had pigmented or uniform surfaces. We observed a contrast effect for both pigmented and uniform face stimuli, with a significantly larger contrast effect for pigmented faces. By comparison, we found a contrast effect for pigmented Greebles but a non-significant contrast effect for uniform Greebles in sensitivity. Overall however, as evident in Figs. 4 and 6, there are qualitative similarities between the results for faces and Greebles. That is, we obtained a larger contrast effect for pigmented versions of both categories of objects. Thus, it seems that similar mechanisms may be involved in the analysis of surface properties for both categories of objects.

As alluded to earlier, one issue with negating images by subtraction is that the negative images are darker than their positive counterparts. There are three arguments suggesting that this confound does not materially affect our results. First, as described in the Methods, Liu et al. (1999) found similar contrast effects with the subtraction method and with a 180° phase shift of the Fourier transformation of measured luminance values (the latter not producing this confound). Second, our critical comparison is whether the contrast between the two images matched or mismatched on a given trial for pigmented and uniform stimuli. Thus, contrast-matched trials included both positive-matched and negative-matched trials and we averaged over these two trial types. Consequently, this confound actually works against our hypothesis. Finally, we did not find any correlation between mean luminance changes and observers' performance on contrast-matched and contrast-mismatched trials for both pigmented and uniform stimuli. Although this does not constitute definitive evidence against the confound between contrast polarity and luminance, it is consistent with the hypothesis that observers are *not* using differences in mean luminance to perform the task. To sum up, in combination with previous findings (e.g., Bruce & Langton, 1994; Liu et al., 1999), our present results indicate that surface pigmentation plays an important role in the recognition of faces and non-face objects.

A wide variety of results in addition to our present study also indicate that skin pigmentation and possibly high-contrast regions such as the eyes contribute significantly to face recognition (Thoresz, Lipson, & Sinha, 2002). Although Bruce and Langton (1994), among others, have made similar claims, there were several methodological limitations in their study that we addressed: (1) Both types of faces were rendered from the same head models under similar lighting conditions so that the effects of pigmentation (independent of shape and illumination) were measured directly; (2) We used a large set of male and female faces that enabled us to present new faces in each condition without repeating any individual, at least within a block; [2] (3) We controlled for low-level cues that could be used in the sequential-matching task by randomly displacing the two faces, changing their relative size, and presenting

[2] A similar pattern of result was found across blocks so they were averaged to yield more statistical power in the overall ANOVA. Recall that blocks were repetitions of the same trial conditions.

image masks. These methodological differences did lead to some differences in results: in comparison to Bruce and Langton, we found a significant contrast effect for uniform faces. Another important difference in our study and theirs is that Bruce and Langton tested observers who were already familiar with the individuals used as stimuli. Consequently, the participants may have had more robust knowledge of the shape of individual faces, thereby rendering them better able to compensate for the absence of pigmentation cues.

Alternatively, Kemp et al. (1996) claimed that reversing contrast disrupts the recovery of shape-from-shading. They argued that with grayscale images both skin pigmentation and shading are encoded as luminance changes. To disentangle these two sources of image-luminance variation, they used color images and compared luminance reversals to hue reversals, since shape-from-shading processes are insensitive to hue. As expected, they found a significant contrast effect for luminance, but not hue reversals. However, their claim that shape-from-shading is disrupted does not necessarily follow from this finding. That is, reversing the contrast of a color (or grayscale) image does not eliminate non-uniformly distributed light and dark regions due to skin pigmentation, which complicates recovering shape from luminance variations in the image. Indeed, most shape-from-shading algorithms assume uniform reflectance, i.e., no textures or markings (Horn, 1975). If the surface does not have uniform reflectance, as is the case with faces, a change in luminance may be caused by an object's shape or by a "spot" on its surface. Here we directly eliminated pigmentation from the faces. That we still found a contrast effect with uniform faces strongly argues against Kemp et al.'s hypothesis.

Contrast reversal does affect the local surface orientation (i.e., defined as the normal to the tangent plane of a point on the visible surface) recovered from luminance variations in the image (Horn, 1975). If, however, we integrate across these local orientation estimates, the *3 D shape* would be the same whether the estimates were derived from positive or negative images. The results by Bruce and Langton (1994) and Liu et al. (1999) suggest that observers can compensate for these orientation changes to a certain degree (e.g., for highly familiar faces or by assuming a prior lighting direction). By comparison, for uniform faces, we found that contrast reversal can be detrimental to face recognition when these faces are not highly familiar to the observers and with lighting direction fixed (but unknown). Another related way to conceptualize the information available in the luminance variations in images is in terms of lines drawn through regions of constant luminance (i.e., isophotes). These isophotes are not affected by contrast reversal, regardless of whether the surface has uniform or non-uniform (i.e., pigmented) reflectance. Again, however, observers in our study showed differential performance for both pigmented and uniform faces and Greebles. Thus, our results suggest that *both* pigmentation and shading affect object recognition.

A second conclusion of our present study is that the contrast effect is not specific to faces, nor even to categories for which we have visual expertise. In addition to the results reported for faces, we also found that novice observers were less accurate and responded more slowly when pigmented Greebles have mismatched contrast. A similar contrast effect was also observed in response times for uniform Greebles. That said, based on Gauthier and her colleagues' work (Gauthier & Tarr, 1997; Gauthier et al., 1998; Tarr & Gauthier, 2000), it is likely that the role of pigmentation—or any other surface property—may become more pronounced with the onset of expertise. For example, Gauthier et al. (1998) found that Greeble experts responded more slowly than novices in naming negative-contrast Greebles. More recently, using a sequential-matching task, Gauthier (personal communication) also found that experts responded less accurately than novices when the two Greebles had mismatched contrasts. Nederhouser et al. (2003) also recently addressed this issue and found that the effects of contrast reversal did not increase with experience on a matching task. However, it is unclear whether their observers were truly experts as Nederhouser et al.'s criterion was simply a large number of trials on a matching task, rather than achieving the specific performance criterion used to define perceptual expertise (Tanaka & Taylor, 1991).

Our results provide some evidence for the hypothesis that visual expertise may modulate the degree to which observers rely on surface properties in recognition. In particular, faces produced a significant sensitivity and response-time cost when contrast was reversed even with uniform faces (cf. Bruce & Langton, 1994). By comparison, uniform Greebles did not result in a significant cost in sensitivity with contrast reversal. That said, as with faces, the contrast effect was still more pronounced for pigmented than for uniform Greebles. There was also a cost in response times, suggesting that contrast reversal may affect observers' decision threshold for responding "same" (see Section 3). It remains a matter of future research to determine how visual expertise may modulate the contrast effect, particularly for uniform surfaces. Another issue for future research is what visual information used by experts is perturbed by contrast reversal. For example, the increased sensitivity to configural information that comes with expertise may lead observers to rely on subtle shape differences between individual faces. Contrast reversal may adversely affect these shape differences even in the absence of pigmentation.

Taking a step back, the contrast effect has often been tested in the context of face recognition, yet as we discussed above, most current accounts of this effect are not specific to faces. For example, Bruce and Langton

(1994) attribute this effect to pigmentation, which our data support. Liu et al. (1999) hypothesized that contrast reversals may also violate the lit-from-above assumption of human vision. Other researchers have suggested that contrast reversal may adversely affect aspects of shape recovery, which depend on both surface textures and shading patterns (e.g., Kemp et al., 1996; see also O'Toole et al., 1999; Troje & Bülthoff, 1996). Although our present results do not allow us to definitively speak to whether these various explanations are contributing to the contrast effect across categories, what is clear is that the presence of pigmentation independent of other factors is sufficient to induce recognition costs with contrast reversal for both categories of objects. This direct comparison between pigmented and uniform faces and Greebles forms the critical set of factors in our study. Additional research is needed to address whether this effect is driven by this single factor or multiple factors; for example, expertise and familiarity may interact with the perception and representation of surface properties.

In conclusion, the contrast effect we report here signals an important role for surface properties in both face and object recognition. Critical to this claim is our systematic control of the presence or absence of pigmentation for both faces and homogeneous non-face objects. The difference in the contrast effects across these conditions indicates that people are sensitive to an object's surface properties in recognition, irrespective of what such properties tell us about shape or illumination conditions. That is, both intrinsic surface properties such as color (e.g., Naor-Raz et al., 2003; Price & Humphreys, 1989; Rossion & Pourtois, 2004; Wurm et al., 1993) and surface properties arising from shape–illumination interactions (e.g., shading and shadows; Tarr et al., 1998) appear to be integral components of human object representation and recognition.

Acknowledgements

This research was funded by the Perceptual Expertise Network (PEN), number 15573-S6, a collaborative award from the James S. McDonnell Foundation (I. Gauthier, PI), and a Natural Sciences and Engineering Research Council of Canada (NSERC) postgraduate scholarship to QCV. We would like to thank Barbara Knappmeyer for assistance in creating the face stimuli, Bruno Rossion for advice on experimental design, Anjula Joshi and Giulia Righi for help in data collection, and three anonymous reviewers for comments.

References

Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological Review, 94*, 115–147.

Blanz, V., & Vetter, T. (1999). A morphable model for the synthesis of 3D faces. In *Proceedings of SIGGRAPH 99*, Los Angeles (pp. 187–194).

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*, 433–436.

Bruce, V., & Langton, S. (1994). The use of pigmentation and shading information in recognising the sex and identities of faces. *Perception, 23*, 803–822.

Cavanagh, P., & Leclerc, Y. G. (1989). Shape from shadows. *Journal of Experimental Psychology: Human Perception and Performance, 15*, 3–27.

Galper, R. E. (1970). Recognition of faces in photographic negative. *Psychonomic Science, 19*, 207–208.

Gauthier, I., & Tarr, M. J. (1997). Becoming a "Greeble" expert: exploring mechanisms for face recognition. *Vision Research, 37*, 1673–1682.

Gauthier, I., Williams, P., Tarr, M. J., & Tanaka, J. (1998). Training "Greeble" experts: a framework for studying expert object recognition processes. *Vision Research, 38*, 2401–2428.

Hayward, W. G., & Williams, P. (2000). Viewpoint dependence and object discriminability. *Psychological Science, 11*, 7–12.

Horn, B. K. P. (1975). Obtaining shape from shading information. In P. H. Winston (Ed.), *Psychology of computer vision* (pp. 115–155). New York: McGraw-Hill.

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.

Liu, C. H., Collin, C. A., Burton, A. M., & Chaudhuri, A. (1999). Lighting direction affects recognition of untextured faces in photographic positive and negative. *Vision Research, 39*, 4003–4009.

Marr, D., & Nishihara, H. K. (1978). Representation and recognition of three dimensional shapes. *Proceedings of the Royal Society of London, Series B, 200*, 269–294.

Naor-Raz, G., Tarr, M. J., & Kersten, D. (2003). Is color an intrinsic property of object representation? *Perception, 32*, 667–680.

Nederhouser, M., Mangini, M. C., Biederman, I., & Okada, K. (2003). Invariance to contrast inversion when matching objects with face-like surface structure and pigmentation [Abstract]. *Journal of Vision, 3*, 93a.

O'Toole, A. J., Vetter, T., & Blanz, V. (1999). Three-dimensional shape and two-dimensional reflectance contributions to face recognition: an application of three-dimensional morphing. *Vision Research, 39*, 3145–3155.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision, 10*, 437–442.

Price, C. J., & Humphreys, G. W. (1989). The effects of surface detail on object categorization and naming. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 41*, 797–827.

Ramachandran, V. S. (1988). Perception of shape from shading. *Nature, 331*, 163–166.

Rossion, B., & Pourtois, G. (2004). Revisiting Snodgrass and Vanderwart's object pictorial set: the role of surface detail in basic level object recognition. *Perception, 33*, 217–236.

Subramaniam, S., & Biederman, I. (1997). Does contrast reversal affect object recognition? *Investigative Ophthalmology and Visual Science, 38*, 998.

Tanaka, J., & Farah, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 46A*, 225–245.

Tanaka, J. W., & Taylor, M. (1991). Object categories and expertise: is the basic level in the eye of the beholder? *Cognitive Psychology, 23*, 457–482.

Tarr, M. J., & Gauthier, I. (2000). FFA: a flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience, 3*, 764–769.

Tarr, M. J., Kersten, D., & Bülthoff, H. H. (1998). Why the visual recognition system might encode the effects of illumination. *Vision Research, 38*, 2259–2275.

Thoresz, K., Lipson, P., & Sinha, P. (2002). Common representations for scenes and objects [Abstract]. *Journal of Vision, 2*, 409a.

Troje, N., & Bülthoff, H. H. (1996). Face recognition under varying poses: the role of texture and shape. *Vision Research, 36*, 1761–1771.

White, M. (2001). Effect of photographic negation on matching the expressions and identities of faces. *Perception, 30*, 969–981.

Wurm, L. H., Legge, G. E., Isenberg, L. M., & Luebker, A. (1993). Color improves object recognition in normal and low vision. *Journal of Experimental Psychology: Human Perception and Performance, 19*, 899–911.

Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology, 81*, 141–145.