Perception, 2007, volume 36, pages 1368 – 1374

doi:10.1068/p5779

Real-world face recognition: The importance of surface reflectance properties

Richard Russell¶, Pawan Sinha
Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA; ¶ and Department of Psychology, Harvard University, 33 Kirkland Street, #720, Cambridge, MA 02138, USA; e-mail: rrussell@fas.harvard.edu
Received 21 August 2006, in revised form 29 January 2007; published online 8 October 2007

Abstract. The face recognition task we perform most often in everyday experience is the identification of people with whom we are familiar. However, because of logistical challenges, most studies focus on unfamiliar-face recognition, wherein subjects are asked to match or remember images of unfamiliar people's faces. Here we explore the importance of two facial attributes—shape and surface reflectance—in the context of a familiar-face recognition task. In our experiment, subjects were asked to recognise color images of the faces of their friends. The images were manipulated such that only reflectance or only shape information was useful for recognizing any particular face. Subjects were actually better at recognizing their friends' faces from reflectance information than from shape information. This provides evidence that reflectance information is important for face recognition in ecologically relevant contexts.

1 Introduction

Determining which features are used to recognize faces is a central goal for the study of face perception. The visual information available to distinguish faces can be divided into two broad classes—shape and surface reflectance (pigmentation).[1] Recent studies using computer graphics (O'Toole et al 1999b; Russell et al 2006, 2007) and facial transplants conducted on cadavers (Siemionow and Agaoglu 2006) investigated the relative utility of shape and surface reflectance information for recognizing *unfamiliar* faces, and found evidence that they are about equally useful.

The primary reason for the importance of face recognition as a subject of inquiry is the ecological importance of the task. Recognizing another person allows the perceiver to bring to bear memories of how that person has behaved in the past, including how that person interacted with the perceiver and with others. As humans are highly social animals, this kind of knowledge about others and their interpersonal relations is of critical importance. But, despite the fact that face recognition is considered an important research topic in large part because of its ecological relevance, it is typically studied in the laboratory with methods that are far from ecological. In particular, the vast majority of recognition experiments, including all those that have compared the utility of shape and reflectance, use 'unfamiliar-face recognition' tasks, in which the subject does not actually know the people whose faces are presented as stimuli. In contrast, 'familiar-face recognition' tasks involve recognizing the faces of people actually known to the subjects (eg naming a person represented in a photograph).

[1] Here we use the term 'surface reflectance', or simply 'reflectance', to refer to the complete light-transfer function of the surface, including the proportion of incident light that the surface reflects, the proportion of light it reflects as a function of wavelength, subsurface scattering [an important attribute of human skin (Debevec et al 2000)], as well as variation across the surface of these properties. Elsewhere, we have used the term 'pigmentation' to refer to the exact same concept (Russell et al 2006). Though this concept is sometimes referred to as 'albedo', 'color', or 'texture', we believe that these terms are inappropriate in this context because they can also be used to refer to specific subsets of the broader light-transfer function of a surface—the fraction of light emitted by the surface in the case of 'albedo', the fraction of light emitted as a function of wavelength in the case of 'color', and spatially variegated reflectance in the case of 'texture'.

Because subjects performing unfamiliar-face recognition tasks have no other experience with the faces in the images, they are free to make use of all the properties of those images—not only those properties of the face that are stable across different viewpoints, lighting, and facial expressions. Because of this, performance on unfamiliar-face recognition tasks can be driven by image-level features of the particular image in view as well as by face-specific processing. This differs from familiar-face recognition tasks, in which the subject has had experience seeing the test faces under many different conditions. This familiarity with a face allows recognition to proceed despite changes in viewpoint, lighting, and facial expression, which disrupt recognition of unfamiliar faces (Hancock et al 2000). Familiarity also changes the relative utility of different kinds of facial information. The internal facial features (eyes, eyebrows, nose, and mouth) and external facial features (hair, face outline, head shape) are about equally useful for unfamiliar-face recognition, but the internal features are more useful for recognition of familiar faces (Ellis et al 1979; Young et al 1985). It has recently been proposed that unfamiliar faces are processed for identity in a way that is qualitatively different than are familiar faces (Megreya and Burton 2006).

Given these substantial differences between familiar-face recognition and unfamiliar-face recognition tasks, we cannot assume that the relative utility of shape and reflectance found with unfamiliar-face recognition tasks is representative of the situation with familiar-face recognition. Because most real-world situations involve recognizing people known to an observer, familiar-face recognition tasks are justifiably considered the 'gold standard' for face recognition by humans. For these reasons, we conducted an experiment in which we investigated the relative utility of shape and reflectance cues for *familiar*-face recognition. Our approach was to manipulate photographs of real people to create novel face images that had either the same shape or the same reflectance as the original photos. The people photographed were a group of students living in the same dormitory. They were chosen because they knew each other, and could be subsequently asked to perform a task in which they recognized one another in photographs. In this way we conducted a familiar-face recognition experiment in which the subjects were shown faces that had either the distinctive shape or the distinctive reflectance of their friends' faces, and were asked to name the friend whose face had been altered to produce the image. This allowed us to assess the subjects' ability to recognize their friends by using only shape cues or only reflectance cues.

2 Methods

2.1 *Subjects*

Twenty-six subjects participated in the experiment (eighteen females); they ranged in age from 18 to 34 years, with an average age of 20 years. All subjects were among the thirty individuals whose photographs were used as stimuli (described below). Eight of the subjects were Asian females, ten were Caucasian females, and eight were Caucasian males. Three individuals were photographed but were not included as subjects in the experiment because they took part in a pilot version, and a fourth individual was unavailable for testing.

2.2 *Stimuli*

Color frontal photographs were taken of thirty faces (ten male). The individuals photographed ranged in age from 18 to 34 years, with an average age of 20 years. All ten males were Caucasian, ten of the females were Caucasian, and ten were East Asian.

It was important that variation in lightness of the face be the result of variation in reflectance rather than variation in illumination. For this reason, we paid very careful attention to photographing all the faces under the same illumination. Toward this end, every photo was taken in the same room, in which all light came from two studio lamps with diffusing heads. These two lamps, the tripod and camera, and the chair on which the photographic subjects sat were kept in locations that were fixed with respect to each other and the room. Similarly, the heights of the tripod and lamps were fixed. Finally, the height of the chair on which the subjects sat was adjusted such that the subjects' heads were all placed at the same height. The lights were centered at 0° elevation to eliminate cast shadows and to minimize the effects of shading cues (Liu et al 1999).

The photographs were then manipulated with Morph Man 3.0 (Stoik Imaging). There were three groups of faces: Asian females, Caucasian females, and Caucasian males, each comprising ten faces. For each group, the ten original faces were morphed together to produce an average face. We created the stimuli for the shape condition by warping this average face into the shape of each of the original ten faces. This produced ten new faces that each had the same shape as one of the original faces but the reflectance of the average face. Similarly, we created the stimuli for the reflectance condition by warping each of the ten original faces into the shape of the average face. This produced ten new faces that each had the same reflectance as one of the original faces but the shape of the average face. Thirty new faces were thus created for the shape condition and thirty for the reflectance condition. By 'morphing' we refer both to moving pixels in the image plane and to the averaging of the pixel intensities of different images, while by 'warping' we refer only to moving pixels in the image plane. In order to achieve very high fidelity between images (to ensure, for example, that the shape of the images in the reflectance set was as identical as possible), we used approximately 250 reference points per face to perform the morphing and warping. This method has been described elsewhere in much greater detail (Russell et al 2007), including a discussion of its pros and cons, and a thorough description of warping and morphing in the appendix. Examples of the stimuli are shown in figure 1.

2.3 *Procedure*

Subjects first viewed a list of the names of the thirty faces that would be presented and were asked to indicate by checking boxes whether they knew or did not know each individual. Eighteen subjects knew all the people listed, six subjects knew all but one, one subject knew all but three, and another knew all but four. Subjects were told that they would be shown images produced by manipulating photographs of those thirty individuals who had just been listed, and that their task with each presented face was to name the individual whose face had been manipulated to produce that image. The subjects were not informed how the images had been manipulated.

Each subject was then shown twenty-five faces—ten Asian females, ten Caucasian females, and five Caucasian males (ten Caucasian male stimuli were produced, but because of experimenter error only five were actually presented to the subjects—always the same five). Half of these faces were taken from the shape condition and half from the reflectance condition. This was counterbalanced across two different versions of the experiment, with a reflectance version of a particular individual's face being shown in one version and a shape version being shown in the other version. There were no significant differences in performance between the two versions of the experiment. Faces were presented on a computer screen, and subjects responded by writing the name of the individual on a sheet of paper.

After display of the manipulated faces, the subjects were shown the twenty-five original (unaltered) photographs and asked to name the individuals depicted. This part of the experiment served as a baseline. If a subject was unable to recognize the unaltered photograph of an individual (which occurred on 6% of these trials), his/her response to the manipulated image corresponding to that individual was removed from subsequent analysis.

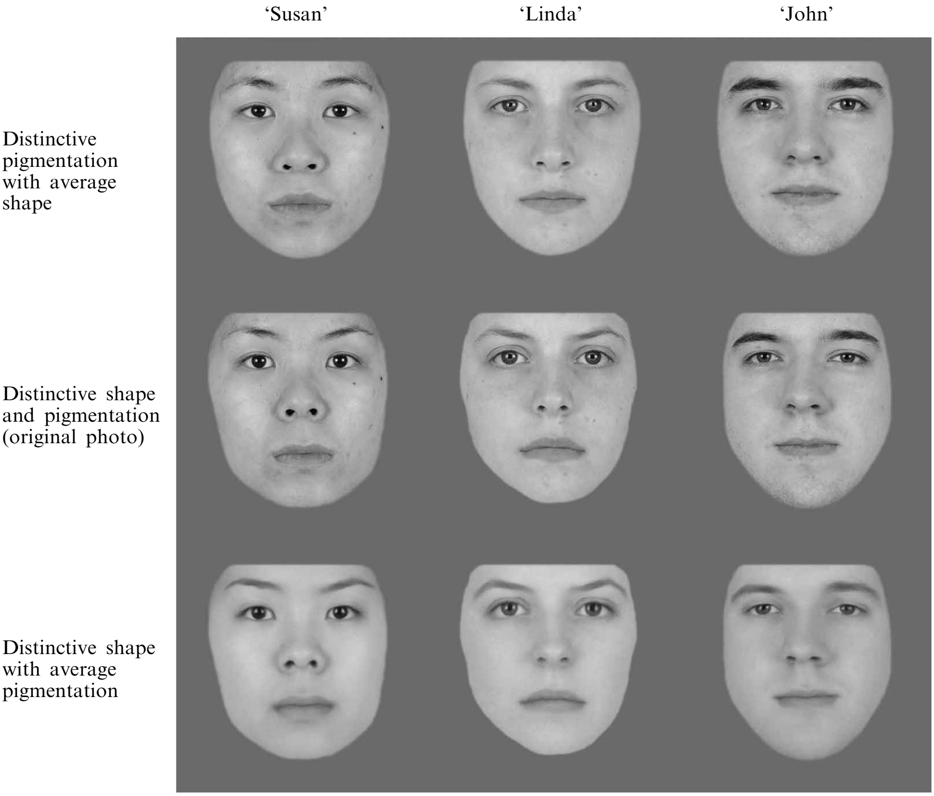

Figure 1. Examples of stimuli. Color images were used in the experiment. The faces in the middle row are cropped photographs of actual people. The faces in the top row (taken from the reflectance condition) have the same shape as the averaged faces, and the reflectance of the actual faces below them. Thus, to recognize the upper left face as corresponding to 'Susan', the subject must use reflectance information. The faces in the bottom row (taken from the shape condition) have the same shape as the original faces above them, but the reflectance of the average faces. To recognize the lower right face as 'John', the subject must use shape information.

3 Results

The percentage of correct responses by condition and face type is shown in figure 2. We ran an ANOVA with condition (shape, reflectance), face type (Asian female, Caucasian female, Caucasian male), and subject group (Asian female, Caucasian female, Caucasian male) as fixed factors, and the percentage of correct responses as the dependent variable. There was a significant main effect of condition ($F_{1,23} = 8.0$, $p = 0.01$), with higher performance overall in the reflectance condition. There was also a significant main effect of face type ($F_{2,46} = 13.3$, $p < 0.001$), and a significant interaction between condition and face type ($F_{2,46} = 6.7$, $p < 0.01$).

It appears from figure 2 that the difference between performance with shape cues and performance with reflectance cues was caused primarily by the Caucasian female faces. This was confirmed with a posteriori paired-samples t-tests, which found significant difference in performance between the shape and reflectance conditions for the Caucasian females ($t_{25} = 5.5$, $p < 0.001$) but not for Asian females or Caucasian males. There was no significant effect of subject group ($F_{2,23} < 1$) or any other significant interactions, including the interaction between subject group and face type ($F_{4,46} < 1$). This indicates that there were no significant differences in how the subjects performed when viewing faces from their own or other race and gender groups.

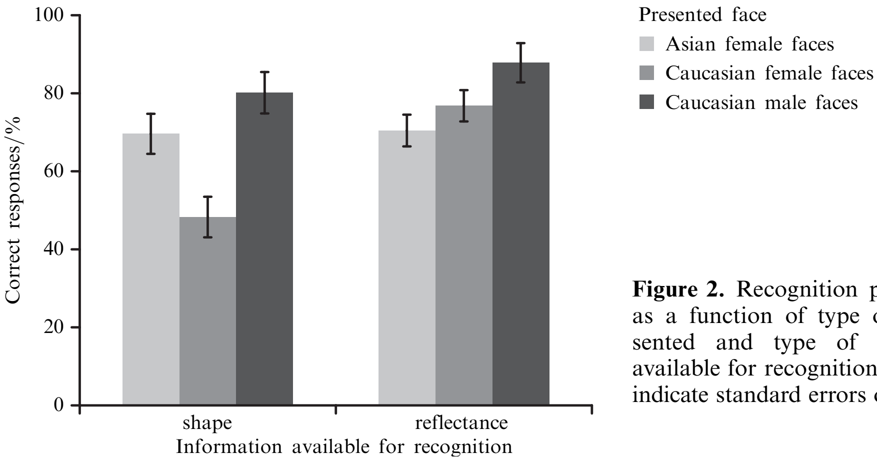

Figure 2. Recognition performance as a function of type of face presented and type of information available for recognition. Error bars indicate standard errors of the mean.

4 Discussion

The primary finding of this experiment was that subjects were capable of recognizing images of their friends' faces even when only reflectance cues were available. This provides direct evidence that reflectance cues are useful for the recognition of familiar faces. Indeed, with all three groups of faces, subjects' performance when they used reflectance cues alone was as good as or better than performance with shape cues alone. This suggests that reflectance cues are actually as or more important than shape cues for familiar-face recognition. This is consistent with previous findings that shape and reflectance cues are about equally important for unfamiliar-face recognition, and extends them to the more ecologically relevant domain of familiar-face recognition.

Considering only the Caucasian faces, there is a sex difference in the relative utility of shape and reflectance. For the male faces, both cues were about equally useful, while for the female faces shape was much less useful. This is consistent with another study in which it was found that familiar male faces were better recognized than familiar female faces when only shape information was available (Bruce et al 1991). This difference in the ability to recognize male and female faces by using only shape cues could be a consequence of male faces having greater shape variation than female faces. An extensive literature dating back to Darwin suggests that there is greater variation in male than in female faces (reviewed by Ramsey et al 2005), and anthropometric measurement of the faces shows that there is indeed greater shape variation in male faces than in female faces (Farkas and Munro 1987). Behavioral evidence suggests that this physical difference is perceived, with more female faces than male faces being rated as typical (Vokey and Read 1988) and with typical faces harder to recognize than atypical faces (Light et al 1979). Although there is evidence that male faces have greater shape variability than female faces, there is no comparable evidence with regard to reflectance, and it is likely that male faces are not more variable in terms of reflectance than female faces. This would lead to shape, but not reflectance, being more useful for recognizing male than female faces, consistent with the current results.

The hypothesis of greater shape variability in male faces would not explain why shape and reflectance were equally useful for Asian female faces. However, much less is known about variation among Asian faces, and we did not have Asian male faces with which to compare the Asian female faces. It is possible that the sex difference in shape variation of Caucasian faces does not exist in Asian faces. Studies with Asian and other non-Caucasian faces would be interesting in this regard.

A few caveats are in order. Because only five Caucasian male faces were presented but ten each of the other two face types, differences between performance with Caucasian male faces and the other face types could be due in part to the comparison being unbalanced. Also, the separation of shape and reflectance cues is not perfect. The images in the reflectance condition vary primarily in terms of reflectance, but also vary slightly in terms of shape, because these images vary in terms of shading (though this variation was minimized through the use of diffuse, frontal illumination). Finally, the current method cannot address the question of which regions of the face, or which aspects of reflectance or shape, are important for recognition.

A growing body of literature is beginning to highlight the important role that reflectance information plays in several aspects of face perception, such as in determining the apparent sex of an individual (Hill et al 1995; O'Toole et al 1998; Tarr et al 2001) and the attractiveness of a face (Jones et al 2004; Little and Hancock 2002; O'Toole et al 1999a; Russell 2003). It has recently been proposed that not only is reflectance information useful for face perception but also that the very faculty of color vision in primates might have evolved for the purpose of discriminating reflectance changes in the faces and rumps of conspecifics (Changizi et al 2006). The experiment we reported here adds to this body of work by demonstrating the contribution of reflectance information to the ecologically crucial task of familiar-face recognition.

Acknowledgments. We are grateful to A Breckenkamp, M Recalde, and especially C Zhao for help with data acquisition and stimulus preparation, to Matthew Bronstad, Ken Nakayama, and Mary Potter for helpful comments, and to Nancy Kanwisher for suggesting the experiment. The study was supported by a grant to PS from the DARPA Human ID Program administered by the Information Awareness Office.

References

Bruce V, Healey P, Burton M, Doyle T, Coombes A, Linney A, 1991 "Recognising facial surfaces" *Perception* **20** 755 – 769

Changizi M A, Zhang Q, Shimojo S, 2006 "Bare skin, blood and the evolution of primate colour vision" *Biology Letters* **2** 217 – 221

Debevec P, Hawkins T, Tchou C, Duiker H-P, Sarokin W, Sagar M, 2000 "Acquiring the reflectance field of a human face" *SIGGRAPH 2000 Conference Proceedings* (ACM Press/Addison-Wesley, New York) pp 145 – 156

Ellis H D, Shepherd J W, Davies G M, 1979 "Identification of familiar and unfamiliar faces from internal and external features: some implications for theories of face recognition" *Perception* **8** 431 – 439

Farkas L G, Munro I R, 1987 *Anthropometric Facial Proportions in Medicine* (Springfield, IL: Charles C Thomas)

Hancock P J B, Bruce V, Burton A M, 2000 "Recognition of unfamiliar faces" *Trends in Cognitive Sciences* **4** 330 – 337

Hill H, Bruce V, Akamatsu S, 1995 "Perceiving the sex and race of faces: the role of shape and colour" *Proceedings of the Royal Society of London B* **261** 367 – 373

Jones B C, Little A C, Burt D M, Perrett D I, 2004 "When facial attractiveness is only skin deep" *Perception* **33** 569 – 576

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

Little A C, Hancock P J B, 2002 "The role of masculinity and distinctiveness in judgments of human male facial attractiveness" *British Journal of Psychology* **93** 451 – 464

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

Megreya A M, Burton A M, 2006 "Unfamiliar faces aren't faces: Evidence from a matching task" *Memory & Cognition* **34** 865 – 876

O'Toole A J, Deffenbacher K A, Valentin D, McKee K, Huff D, Abdi H, 1998 "The perception of face gender: The role of stimulus structure in recognition and classification" *Memory & Cognition* **26** 146 – 160

O'Toole A J, Price T, Vetter T, Bartlett J C, Blanz V, 1999a "3D shape and 2D surface textures of human faces: the role of 'averages' in attractiveness and age" *Image and Vision Computing* **18** 9 – 19

O'Toole A J, Vetter T, Blanz V, 1999b "Three-dimensional shape and two-dimensional surface reflectance contributions to face recognition: an application of three-dimensional morphing" *Vision Research* **39** 3145 – 3155
Ramsey J L, Langlois J H, Marti C N, 2005 "Infant categorization of faces: Ladies first" *Developmental Review* **25** 212 – 246
Russell R, 2003 "Sex, beauty, and the relative luminance of facial features" *Perception* **32** 1093 – 1107
Russell R, Biederman I, Nederhouser M, Sinha P, 2007 "The utility of surface reflectance for the recognition of upright and inverted faces" *Vision Research* **47** 157 – 165
Russell R, Sinha P, Biederman I, Nederhouser M, 2006 "Is pigmentation important for face recognition? Evidence from contrast negation" *Perception* **35** 749 – 759
Siemionow M, Agaoglu G, 2006 "The issue of 'facial appearance and identity transfer' after mock transplantation: A cadaver study in preparation for facial allograft transplantation in humans" *Journal of Reconstructive Microsurgery* **22** 239 – 334
Tarr M J, Kersten D, Cheng Y, Rossion B, 2001 "It's Pat! Sexing faces using only red and green" *Journal of Vision* **1** 337
Vokey J R, Read J D, 1988 "Typicality, familiarity and the recognition of male and female faces" *Canadian Journal of Psychology* **42** 489 – 495
Young A W, Hay D C, McWeeny K H, Flude B M, Ellis H D, 1985 "Matching familiar and unfamiliar faces on internal and external features" *Perception* **14** 737 – 746

© 2007 a Pion publication

ISSN 0301-0066 (print) ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 36 2007

www.perceptionweb.com

Conditions of use. This article may be downloaded from the Perception website for personal research by members of subscribing organisations. Authors are entitled to distribute their own article (in printed form or by e-mail) to up to 50 people. This PDF may not be placed on any website (or other online distribution system) without permission of the publisher.