Recognizing rotated faces and Greebles: What properties drive the face inversion effect?

Alan R. S. Ashworth III

Air Force Research Laboratory

Brooks AFB Texas & Yale University, USA

Quoc C. Vuong*

Department of Cognitive & Computational Psychophysics,

Max Planck Institute for Biological Cybernetics, Germany

Bruno Rossion

Unité Cognition et Développement & Laboratoire de Neurophysiologie,

Université Catholique de Louvain, Belgium

Michael J. Tarr

Department of Cognitive & Linguistic Sciences,

Brown University, USA

Running Head: Face inversion

* Corresponding author: Mailing address: Division of Psychology, Henry Wellcome Building for Neuroecology, Newcastle University, Newcastle upon Tyne, NE2 4HH, UK; phone: +44 (0)191 222-6183; fax: +44 (0)191 222-5622; email: q.c.vuong@ncl.ac.uk

Abstract

The fact that faces are strongly affected by picture-plane inversion has often been cited as evidence for face-specific mechanisms. It is unclear, however, whether this “face inversion effect” is driven by properties shared by faces or whether the effect is specific to faces as a category. To address this issue, we compared the recognition of faces and novel Greebles, which were specifically matched to faces along various stimulus dimensions. In two experiments, participants were required to name individual faces or Greebles following training at either single or multiple orientations. We found that performance systematically decreased with increasing misorientation from either the upright (Experiment 1) or nearest trained orientation (Experiment 2). Importantly, the magnitude of this orientation effect was similar for both faces and Greebles. Taken together, these results suggest that the face inversion effect may be a consequence of the visual homogeneity of the stimulus category, regardless of the category.

Introduction

Images inverted in the picture plane are often more difficult to recognize than their upright counterparts, regardless of whether the images depict faces (Yin, 1969), common non-face objects (Jolicoeur, 1985) or novel 2D shapes (Tarr & Pinker, 1989). However, inversion appears to affect the recognition of faces disproportionately more than the recognition of many other non-face objects. This effect has become known as the *face inversion effect* (e.g., Diamond & Carey, 1986; Scapinello & Yarmey, 1970; Valentine & Bruce, 1986; Yin, 1969). Our purpose in the present study was to test whether this (apparent) disproportionate effect of inversion on recognition performance for faces relative to non-face objects is due, at least in part, to functional properties of faces as a stimulus category. More specifically, faces have a prominent vertical axis and similar parts arranged in the same configuration across individuals, they are often identified at an individual level, and observers are experienced at individuating faces.

One potential difference between the processing of faces and non-face objects is the dominant visual information that is used for recognition purposes. Several researchers have advocated that both the local features, such as the eyes and mouth, and the configuration of these features are critical for face recognition (for a review, see Maurer, Grand, & Mondloch, 2002). Most importantly, researchers have found that inversion in the picture plane impairs how observers performed with facial configurations rather than facial features (Collishaw & Hole, 2002; Freire, Lee, & Symons, 2000; Goffaux & Rossion, in press; Leder, Candrian, Huber, & Bruce, 2000; Tanaka & Farah, 1993). Others have found that inversion simply affects how efficiently facial features are processed (Sekuler, Gasper, Gold, & Bennett, 2004; Riesenhuber, Jarudi, Gilad, & Sinha, 2004; Yovel & Kanwisher, 2004). By comparison, non-face objects may be recognized primarily on the basis of features and coarse spatial relations (Biederman, 1987) without a corresponding reliance on metric configurations of features.

To test the hypothesis that the face inversion effect is partially driven by general properties that faces happen to have, rather than faces *per se*, we compared faces and non-face control stimuli that shared many of these properties that have been shown to be important in face perception. Following numerous other studies, we used Greebles, in particular, because individual Greebles share similar features arranged in similar configurations, thus prompting individual-level discriminations to be based on subtle differences in featural and configural information (Gauthier & Tarr, 1997a). At the same time, Greebles are not faces – indeed, there is behavioral and neurological data to suggest that faces and Greebles categories are processed differently, particularly in Greeble novices (e.g., Gauthier, Behrmann, & Tarr, 2004).¹

To better understand how rotation in the picture plane affects the recognition of faces and Greebles, particularly with respect to the face inversion effect, we presented both stimulus categories at a wide range of orientations in the picture plane. Although earlier studies have compared upright and inverted faces and other homogeneous categories (e.g., houses, airplanes, dogs), the use of only two orientations may fail to reveal possible alignment mechanisms involved in face processing (e.g., Ullman, 1989). In fact, the systematic rotation of images in the picture plane has provided key insights in regards to the encoding and representation of non-face objects (e.g., Tarr & Pinker, 1989). For example, Jolicoeur (1985) demonstrated that familiar common objects rotated in the picture plane become increasingly more difficult to name as the stimulus is rotated further from an upright orientation (typically with respect to gravity), particular for objects which have a dominant orientation. This initial orientation effect diminishes substantially with repeated exposure to the same images, sometimes as soon as the second presentation (Jolicoeur, 1985; Jolicoeur & Milliken, 1989; Murray, Jolicoeur, McMullen, & Ingleton, 1993). Similar decreases in the effect of orientation with practice have been observed for novel 2D shapes (Gauthier & Tarr, 1997b; Shinar &

Owen, 1973; Tarr & Pinker, 1989). Finally, it is worth noting that the level at which objects are recognized (Hamm & McMullen, 1998) and the visual similarity of the set of objects to be recognized (Lawson & Jolicoeur, 1999; Murray, 1998) can affect the magnitude of any orientation effect.

A number of studies of face recognition have likewise explored the recognition of faces rotated in the picture plane. Most notably, Valentine and Bruce (1988) reported several experiments using faces rotated in 45° increments (0°, 45°, 90°, 135°, and 180°). In one experiment, they measured the time subjects required to correctly identify famous faces rotated in the picture plane. In a second experiment, they measured the time subjects required to judge whether a sequence of two faces, a canonical upright face immediately followed by a rotated face, were the same or different. Both experiments revealed a linear relationship between response time and the magnitude of rotation from upright, suggesting that the recognition of rotated faces may involve normalization processes qualitatively similar to those used for the recognition of other objects.

Following Valentine and Bruce (1988), several investigators have also presented faces at multiple picture-plane orientations in conjunction with other manipulations, such as blurring or featural and configural changes (Bruyer, Galvez, & Prairial, 1993; Collishaw & Hole, 2002; Lewis, 2001; Murray, Yong, & Rhodes, 2000; Sjöberg & Windes, 1992; Stürzel & Spillman, 2000). Not surprisingly, these studies consistently find recognition performance that is dependent on the magnitude of the rotation from the upright. That is, they do not observe all-or-none decreases in performance for upside-down faces. In the majority of such studies, as in Valentine and Bruce's initial study, the actual orientation-dependent pattern is linear. Although some studies have reported some non-linearity in the data, this non-linearity seems to be localized to particular orientations (e.g., 180°) and specific tasks (Lewis, 2001; Murray et al., 2000; Stürzel & Spillman, 2000).

What is still unknown based on these earlier studies is the degree to which the orientation-dependent recognition functions obtained for faces are equivalent to or different from those obtained for non-face objects. Moreover, there is the larger issue of whether the orientation dependence observed for faces is the product of a face-specific processing system versus a general object recognition mechanism in which observers are simply best at recognizing objects in their most highly-familiar views (Tarr & Pinker, 1989; Palmer, Rosch, & Chase, 1981). In particular, no study has explored how training with non-upright orientations for faces and non-face objects affects orientation-specific object representations, as well as the ability to generalize from novel to familiar orientations (Jolicoeur & Milliken, 1989; Tarr & Pinker, 1989). Because this manipulation in the object recognition literature has provided insights into the mechanisms involved in object recognition, it seems highly pertinent to the face inversion effect, yet studies have not explored how such effects change with practice (e.g., Collishaw & Hole, 2002; Valentine & Bruce, 1988). To address this issue, first we trained observers with both upright and rotated faces and Greebles; second, we tested how observers generalized to new orientations; and third, we tested how further practice with these new orientations affected face recognition.

Another critical issue when comparing orientation effects on faces and non-face objects is the default level of identification at which stimuli are recognized. Faces are typically recognized at the individual level. By comparison, objects are recognized at a basic or entry level, the level at which object shapes are best distinguished from one another (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). Studies have shown that recognizing objects at the subordinate level (e.g., collies vs. poodles) prompts the largest orientation effects, whereas recognizing objects at the basic level (e.g., dog vs. cat) or superordinate level (e.g., animal vs. artifact) incurs little or no orientation effects (Hamm & McMullen, 1998). To

take this factor into account, subjects in the current study named individual faces and Greebles, thereby equating the explicit level of discrimination across the two stimulus sets.

To summarize, the literature reviewed above suggests that there are systematic effects of picture-plane rotations on the recognition of both faces and non-face objects. However, no study has directly compared orientation effects for faces and appropriately matched non-face objects across a large range of orientations (Brooks, Rosielle, & Cooper, 2002, compared priming in faces and common objects). Likewise, no study has compared how such orientation effects change with practice across a range of orientations (Robbins & McKone, 2003, tested learning effects only for inverted faces). It is our view that a finer-grained analysis of orientation effects will help address the larger question of whether the face inversion effect is unique to faces or rather is simply a consequence of properties that faces, as a category, happen to have.

Experiment 1

Method

Subjects

The 72 subjects participating in Experiment 1 were undergraduate students and other members of the Yale University community. The undergraduates participated to fulfill a class requirement while the others were paid for their participation. All subjects had normal or corrected-to-normal vision. All subjects used their dominant hand for responding. None of the subjects participated in the other experiment reported here.

Material and Stimuli

Faces comprised the first type of stimuli in the present study, examples of which are shown in Figure 1. A total of 144 male faces were scanned from full face photographs taken

from a Harvard Business School yearbook. All scans were 256 shades of gray. Faces with distinguishing characteristics or accessories such as facial hair, scars, or glasses were excluded. The faces were normalized so that the eyes were horizontal, and were rescaled to be approximately the same size. A circular region was superimposed on the central portion of each face, thus excluding most of the featural extremities such as ears, hair, and bottom of the chin. The region enclosed by the circle was then removed from the original background and placed into a white circular background within a black surround. Therefore, upon presentation the silhouette of each face appeared round, and the face itself was centered in a white circle.

Figure 2 shows examples of the Greebles which comprised the second type of stimuli. A total of 30 Greebles were created for the present study. Note that a circular region was not superimposed over the external contours of the Greebles, as was done with the faces. This difference in how the stimuli were cropped was the end-result of our attempt to do as much as possible to equate Greeble and face identification. Consider that extant studies of face recognition almost always uses face stimuli with cropped external contours – this is done to prevent subjects from using hair style/shape as diagnostic features disconnected from what are thought of as “face recognition” mechanisms. In using cropped faces, we follow this convention, which is nearly universal in the literature. At the same time, using similar cropping for Greebles would have rendered them almost impossible to recognize – much of the information about part shapes and part relations is carried in protruding parts and their external contours.² These differences actually bring us closer to our goal of equating our two stimulus sets in terms of overall homogeneity, as well as an individuation task that is likely to prompt recruitment of configural representations relating individual parts in more complex spatial arrangements. Faces with hair would have simply made our stimuli less equivalent along these critical dimensions.

In Experiment 1, 18 of the 144 faces were used for all observers. Three sets of stimuli were made such that each set contained six different target faces, with the remaining twelve faces used as distractors. In this manner all faces served as targets and distractors equally across subjects. By comparison, 15 of the 30 Greebles were used. Three sets were made such that each set contained five different Greebles as named targets, with the remaining ten used as distractors. In this manner all Greebles served as targets and distractors equally across subjects. The difference in the number of target and distractor faces and Greebles in this and the subsequent experiment likely made it more challenging to learn and recognize faces. However, increasing the level of difficulty for faces may lead to larger inversion effects as observers may have to make more fine-grain discriminations (i.e., recognizing more faces at the individual level).

----- Insert Figure 1 and 2 about here -----

All stimuli were presented on a high resolution color monitor, driven by an IBM compatible 80486DX-33 computer. Both faces and Greebles were presented at the center of the screen. The stimuli were rotated in the picture plane about the center of the screen. Thus, the image always remained at the center of the screen. Subjects used a chinrest that was 45 cm from the monitor. The stimuli subtended a visual angle of approximately 12°.

Design and Procedure

In Experiment 1, subjects first learned the names of target faces or Greebles. The names were attached to keys on the keyboard, and during the experiment subjects pressed the key with the name that corresponded to the stimuli presented. Thirty-six subjects participated in the face group, while 36 subjects participated in the Greeble group. For both stimulus

types, there was learning phase followed by a testing phase. In the learning phase, all of the target stimuli to be learned were presented in the canonical upright orientation and subjects were instructed to keep their heads vertical at all times. Subjects were explicitly informed that later in the experiment they would be asked to identify each of the target stimuli by pressing the key with its corresponding name. They were not, however, informed that the targets would be rotated during the testing phase.

Learning phase. During the learning phase, the stimuli were only presented in their upright orientation. Within this phase, there were two blocks. The purpose of these blocks was to enable the subjects to learn the name for each stimulus, and map the name of each stimulus onto the correct response key.

For faces, there were five presentations of each of the six target faces with its corresponding name for a total of 30 trials on the first learning block. Subjects were instructed to study each face for the entire 5 sec presentation, and then press the key with the appropriate name after the stimulus cleared the display. The names used were “Bob,” “Dan,” “Jim,” “Lee,” “Ray,” and “Wes.” Next, there were 20 presentations of the six target faces without its name for a total of 120 trials on the second learning block. For these learning trials, subjects were instructed to respond accurately yet quickly. Subjects wore headphones through which they were given accuracy feedback. For a correct response, there was no sound. Incorrect responses were followed by a beep. Presentation order was randomized for each of these learning trial groups for each subject.

For Greebles, subjects were presented with five repetitions of the five target Greebles with its corresponding name for a total of 25 trials on the first learning block. The names used were “Bob,” “Dan,” “Jim,” “Lee,” and “Ray.” Next, there were 20 repetitions of the five target Greebles for a total for 100 trials on the second learning block.

Test phase. There was a short pause at the end of learning phase to give subjects a break and to present further instructions. Following the two learning blocks, there were two identical test blocks in the test phase, separated by a short break. Each of the six target faces was shown twice at each of the twelve orientations (0° - 330° , in 30° steps) for a total of 144 trials. Each of the 12 distractor faces was shown at four orientations for a total of 48 trials. The 48 distractor trials were distributed equally across the twelve orientations. Thus, there were 192 trials per block in total.

Similarly, each of the five target Greebles was shown twice at each of the twelve orientations for a total of 120 trials. Each of the ten distractor Greebles was shown at four orientations for a total of 40 trials. The 40 distractor trials were distributed as equally as possible across the twelve orientations. Thus, there were 160 trials per block. For both faces and Greebles, 75% of the trials were target trials and 25% of the trials were distractor trials. The presentation order was randomized for each testing block for each subject.

The subjects' task was to accurately and quickly identify each face or Greeble by pressing the key with the appropriate name. There was a key labeled "NA" for "none of the above" that was to be pressed upon presentation of a distractor. Stimulus presentation was preceded by a brief pattern mask that subtended the same visual angle as the stimuli. The mask was a gradient fill from black to white, starting with black in the center and becoming progressively lighter along the radius. The stimulus remained on the screen until subjects responded. Accuracy feedback was provided through headphones, with subjects hearing a beep for incorrect responses.

Results

In both experiments, three common analyses were conducted.³ First, we analyzed the effect of orientation in the picture plane on response times (RTs). For this analysis, we

calculated median RT values from correct target trials (i.e., only the learned stimuli and not the distractors) for each subject. Second, we regressed correct median RTs against orientation to determine the slope of the best-fit line for each subject. The slope provides a common quantitative measure of the effect of orientation across faces and Greebles. Third, we analyzed the effect of orientation on accuracy to ensure that subjects were not trading speed for accuracy. Proportion of incorrect responses per orientation constituted the error data for Experiments 1 and 2.

Response times. The top panels of Figures 3 and 4 show the relationship between RT, orientation, and block, separately for faces and Greebles. The RT data were entered into a mixed-design ANOVA with stimulus type (faces, Greebles) as a between-subjects factor, and block (1-2) and orientation (0° - 180°) as repeated factors. A linear contrast was also computed for orientation. For this ANOVA, RTs were “folded” around 180° with 150° and 210° averaged, 120° and 240° averaged, and so on.

----- Insert Figures 3 and 4 about here -----

All main effects were significant: stimulus type, $F(1,70)=6.28$, $\eta_p^2=0.08$; block, $F(1,70)=156.62$, $\eta_p^2=0.69$; and orientation, $F(6,420)=28.63$, $\eta_p^2=0.29$. However, the effect of stimulus type on RTs in this experiment was small. There was also a significant stimulus type x block interaction, $F(1,70)=8.71$, $\eta_p^2=0.11$. There was a marginally significant but small stimulus type x orientation interaction, $F(6,420)=1.90$, $p=0.08$, $\eta_p^2=0.03$. In addition, the linear contrast of orientation was also significant, and only interacted with block: linear contrast, $F(1,70)=98.63$, $\eta_p^2=0.59$; linear contrast x block interaction, $F(1,70)=11.37$,

$\eta_p^2=0.14$. After computing the linear contrast, there was no significant residual variance associated with the orientation factor.

Recall that, as originally postulated by Yin (1969), inversion should disproportionately affect faces more than non-face objects. Thus, we also analyzed response time at 0° and 180° averaged across blocks for each stimulus type in a mixed-designed ANOVA. For this analysis, there was a significant but small effect of stimulus type, $F(1,70)=7.75$, $\eta_p^2=0.10$ and a significant and relatively large effect of orientation, $F(1,70)=69.42$, $\eta_p^2=0.50$. Importantly, there were no interaction between stimulus type and orientation, $F<1$, $\eta_p^2=0.001$. Thus, although subjects were generally slower with Greebles, rotating the image by 180° was equally detrimental to both faces and Greebles as indicated by the lack of interactions with stimulus type. This conclusion is also confirmed by the slope analysis below.

Slope. Table 1 provides the slopes for faces and Greebles on blocks 1 and 2. The slopes were submitted to a mixed-design ANOVA with stimulus type (faces, Greebles) as a between-subjects factor, and block (1-2) as a within-subjects factor. There was only a significant effect of block, $F(1,70)=11.39$, $\eta_p^2=0.14$; with no significant effect of stimulus type, $F(1,70)=1.20$, $\eta_p^2=0.02$; or significant interaction between stimulus type and block, $F(1,70)=1.34$, $\eta_p^2=0.02$.

----- Insert Table 1 about here -----

Accuracy. The error rates are plotted on the bottom panels of Figures 3 and 4 for faces and Greebles. In the same manner as the response time data, the error data were collapsed across distance from the canonical orientation, and then submitted to the same ANOVA.

There were only main effects of block, $F(1,70)=47.05$, $\eta_p^2=0.40$; and orientation, $F(6,420)=13.46$, $\eta_p^2=0.16$. There was also a significant interaction between block and orientation, $F(6,420)=9.07$, $\eta_p^2=0.12$. As with response times, although the interaction between stimulus type and orientation was marginally significant, its effect size was relatively small, $F(6,420)=1.89$, $p=0.08$, $\eta_p^2=0.03$. The linear contrast for orientation was significant, $F(1,70)=33.05$, $\eta_p^2=0.32$ (with a small residual cubic trend, $F(1,70)=6.86$, $\eta_p^2=0.09$). Lastly, there was a significant block by linear contrast interaction, $F(1,70)=25.98$, $\eta_p^2=0.27$, suggesting that the orientation function was reduced on the second block. There was no evidence of any speed/accuracy tradeoffs.

Discussion

The results of Experiment 1 revealed important similarities in how observers named faces and Greebles across changes in orientation. First, both faces and Greebles yielded linear effects across orientations. Second, subjects showed a decrease in the orientation function (i.e., slope) for both faces and Greebles with practice, as found for other object categories (e.g., Jolicoeur, 1985; Tarr & Pinker, 1989). Finally and critically, we did not find a larger effect of inversion for Greebles than for faces (cf. Yin, 1969). Given that we equated both stimulus types along several relevant dimensions (Gauthier & Tarr, 1997a) and, in particular in terms of the level of identification, it appears that it is not being a face *per se* that produces processing difficulties at the inverted orientation (or any other orientation).

In a recent study, Yovel and Kanwisher (2004) found an inversion effect for faces but not for houses in their behavioural experiment. Observers in their study saw two very brief (250 ms) sequential presentations of faces or houses and had to decide if these were the same or different. Furthermore, upright faces and houses were run in a block before inverted faces

and houses for all observers. We can only speculate that differences in methods are the reasons for the different findings. First, Yovel and Kanwisher's task may not equally engage individual level recognition for both faces and houses. Second, they made featural changes (changing parts) or configural changes (changing the metric relationships between parts) from the same face and house and limited set of parts. Third, our initial training procedure may have given subjects the opportunity to acquire some familiarity with the Greebles. To address this issue, in a second experiment we tried to replicate our results and at the same time explore further possible similarities and differences between these stimulus type by training observers to recognize faces and Greebles at non-upright orientations.

Experiment 2

One of the enduring findings of experiments involving the recognition of stimuli rotated in the picture plane is that practice effects are found as the subjects proceed through the experiment (Eley, 1982; Jolicoeur, 1985; Tarr & Pinker, 1989). That is, subjects are both faster and more accurate across all tested orientations at the end of the experiment than at the beginning. With practice, observers can learn orientation-invariant diagnostic features (e.g., Eley, 1982; Jolicoeur & Milliken, 1989) or orientation-specific templates (e.g., Tarr & Pinker, 1989). In either case, practice effects seem to be specific to the learned stimuli. This high degree of view specificity is often revealed by "surprising" subjects with novel orientations following training (e.g., Tarr & Pinker, 1989). In Experiment 2, we tested whether learning individual faces and Greebles at multiple orientations show similar behavioral patterns across multiple learned views. Thus, Experiment 2 was similar to Experiment 1 except that subjects received two additional hours of training with target faces or Greebles at the upright orientation and other specific orientations prior to testing.

Method

Subjects

The 72 new subjects participating in Experiment 2 were drawn from the same populations as Experiment 1.

Material and Stimuli

A total of 114 of the 144 faces were used in Experiment 2. Of this total, 18 faces were used as targets for all observers, while the remainder served as distractors. The targets were blocked into three groups of six faces, with each group being used equally often across subjects.

Eighteen of the 30 Greebles were used as targets for all observers in Experiment 2, while the remainder served as distractors. The targets were blocked into three groups of six Greebles, with each group being used equally often across subjects. Thus, for both faces and Greebles, each subject learned six target stimuli. The equipment used for stimulus presentation was identical to that used in Experiment 1.

Design and Procedure

Experiment 2 was similar to Experiment 1 in that a naming task was used in which subjects learned to associate names with the target stimuli. The names were attached to keys on the keyboard, and during testing the subjects pressed the key with the name that corresponded to the stimulus currently being presented. Thirty-six subjects participated in the face group, while 36 subjects participated in the Greeble group.

In contrast to the previous experiment, Experiment 2 added two one-hour training sessions on consecutive days prior to the testing session on the third day. During the two training days, subjects practiced naming the target stimulus at three of the twelve possible

orientations used at test. Subjects practiced naming the target at 0°, 150° and 240°. At test, the same target stimuli were presented at all 12 orientations in the picture plane (0°-330° in 30° steps).

Training sessions. For faces, each training session was divided into three blocks of training trials preceded by two blocks of learning trials. For Greebles, the sessions were divided into two blocks of training trials preceded by two blocks of learning trials.

The purpose of the first two learning blocks was to enable subjects to learn the name for each face or Greeble, and map the name of the target stimulus onto the correct response key. During these two blocks, the faces were always presented upright. As in Experiment 1, accuracy feedback was provided through headphones, with subjects hearing a beep for incorrect responses.

For both faces and Greebles, the first learning block contained six presentations of each of the six targets with its corresponding name for a total of 36 trials. For these trials, the subjects were instructed to study each stimulus for the entire 5 sec presentation, then to press the key with the appropriate name after the stimulus was removed. The second learning block contained 12 presentations of each of the six targets without its name for a total of 72 trials. For these trials, the subjects were instructed to respond accurately yet quickly.

During these two learning blocks, all stimuli were presented upright and subjects were instructed to keep their heads vertical. Subjects were explicitly informed that later in the experiment they would be asked to identify each target by pressing the key with its corresponding name. They were not informed that the stimuli would sometimes be rotated in the picture plane. Presentation order was randomized for each of these learning groups for each subject.

Following each of the two learning session subjects ran in either three blocks of training trials for faces or two blocks of training trials for Greebles. During these blocks, three

of the six target stimuli were shown at 0° and 150°, and the other three were shown at 0° and 240°. The first set of items is referred to as the 150-Set, while the latter is referred to as the 240-Set. The distractor stimuli were presented at the same orientations as the target faces. A different set of 12 distractors were used for the different training blocks for faces. By comparison, the same set of 12 Greeble distractors was used given the limited number of these stimuli.

During the training blocks, the three face or Greeble stimuli from the 150-Set were presented 12 times at 0° and 12 times at 150° for a total of 72 trials. Likewise, the three stimuli of each type comprising the 240-Set were presented the same number of times at 0° and 240° for a total of 72 trials. The 12 distractor faces and Greebles were presented twice at 0°, once at 150°, and once at 240° for a total of 48 trials. Therefore each training block contained 192 trials, 75% of which were targets and 25% of which were distractors. Presentation order was randomized for each training block for each subject. During these blocks, the subjects' task was to accurately and quickly identify each face by pressing the key with the appropriate name. There was a key labeled "NA" for "none of the above" that was to be pressed upon presentation of a distractor. Stimulus presentation was preceded and followed by a pattern mask, and the trials were response terminated. Accuracy feedback was provided through headphones, with subjects hearing a beep for incorrect responses. Accuracy and response times were recorded for each response.

Testing session. After the two days of training, subjects returned the next day for the final test session. They were told that this last session was no different from the previous two. Like the two training sessions, the subjects were initially instructed to view the target stimulus in the canonical orientation (0°), and to practice naming the faces or Greebles.

Following next was a "refresher" block of trials and then two blocks of test trials in which the main body of data for Experiment 3 was collected. Other than being abbreviated,

the refresher block of trials was identical to the training trials. The subjects trained with the same 150-Set and 240-Set of stimuli that were used during training. The 150-Set was presented six times at 0° and six times at 150°, and the 240-Set was presented six times at 0° and six times at 240° for a total of 72 target trials. For the twelve distractor faces used, all were shown once at 0°; whereas six were shown at 150° and six were shown at 240° for a total of 24 trials. Of the 96 total trials, the 72 trials during which target faces were shown represents 75% of the total trials.

The final two blocks of trials represented the main body of data for the experiment. These two blocks were identical in all respects except for the mapping of distractors to orientation. In each of these blocks, the same six stimuli with which the subjects had trained were presented. However, unlike the training sessions the stimuli were now shown at all twelve orientations in the picture plane (0°-330° in 30° steps). Therefore, each face and each Greeble was shown in ten novel orientations and in two previously trained orientations (including the upright orientation).

Each of the six targets was shown three times at each of the twelve orientations for a total of 216 trials per block. The twelve distractors were divided into two groups of six. In the first test block, one group of distractors was shown six times at each of the following orientations: 0°, 60°, 120°, 180°, 240°, and 300°. The other group of distractors was shown six times at each of the following orientations: 30°, 90°, 150°, 210°, 270°, and 330°. Therefore, there were 72 distractor trials in all. For the second test block, the orientations were switched between groups of distractors. Of the 288 trials per block, the 216 trials during which target faces were shown represents 75% of the total trials.

For each of the two test blocks, the subjects' task was to accurately and quickly identify each target by pressing the key with the appropriate name. There was a key labeled "NA" for "none of the above" that was to be pressed upon presentation of a distractor.

Stimulus presentation was preceded and followed by the same pattern mask used in Experiment 1, and the trials were response terminated.

Results

Response times. Naming times as function of orientation, block, and stimulus set for faces and Greebles are plotted in Figures 5 and 6, respectively.

----- Insert Figures 5 and 6 about here -----

In Experiment 2, subjects learned faces or Greebles at two orientations (0°, and either 150° or 240°). For this experiment, we were primarily interested in naming times for specific targets during the test trials as a function of the angular distance to their two trained orientations. Therefore, the data were collapsed across distance from familiar orientation. For example, to calculate the value that represents 30° from a familiar orientation, the following points were averaged: the points on the 150-Set function that were 30° from 150° (120° and 180°); the points on the 240-Set function that were 30° from 240° (210° and 270°); the points on either function that were 30° from 0° (30° and 330° from both the 150-Set and the 240-Set).

The data were entered into a mixed-design ANOVA with stimulus type (faces, Greebles) as a between-subjects factor; and block (1-2) and orientation (0°-120°) as repeated factors. As in Experiment 1, a linear contrast was computed for the orientation factor. There was a significant effect of stimulus type, $F(1,70)=92.39$, $\eta_p^2=0.57$; block, $F(1,70)=11.59$, $\eta_p^2=0.14$; and orientation, $F(4,280)=27.54$, $\eta_p^2=0.28$. There was a significant interaction between stimulus type and block, $F(4,70)=9.15$, $\eta_p^2=0.12$; and a significant but small block x

orientation interaction, $F(4,280)=2.73$, $\eta_p^2=0.038$. There was also a significant effect for the linear contrast computed for the orientation factor, $F(1,70)=47.13$, $\eta_p^2=0.40$, and a significant linear contrast x block interaction, $F(1,70)=5.13$, $\eta_p^2=0.07$, indicating a small difference in the slope of the orientation functions from Block 1 to Block 2 (see Table 1). After computing the linear contrast, there was no significant residual variance associated with the orientation factor.

Our primary analysis in Experiment 2 tests recognition performance relative to any of the learned orientations (i.e., 0° , 150° , and 240°). In a second analysis, we also tested the 150-Set and 240-Set separately to determine whether learning these orientations generalizes to other orientations, as found by Jolicoeur and Milliken (1989) for common animals and objects (see also Tarr & Pinker, 1989). Figures 7 and 8 plot naming times as a function of trained orientation and stimulus set for faces and Greebles, respectively. For these analyses, we only analyzed RTs from the 0° , 150° and 240° test orientations in two separate mixed-design ANOVAs with stimulus type (faces, Greebles) as a between-subjects factor, and block (1-2) and orientation (0° , 150° , 240°) as within-subjects factors.

----- Insert Figures 7 and 8 about here -----

For the 150-Set stimuli, there were only significant main effects of stimulus type, $F(1,70)=87.51$, $\eta_p^2=0.56$; and orientation, $F(2,140)=25.90$, $\eta_p^2=0.27$. The linear contrast for orientation was also significant, $F(1,70)=43.80$, $\eta_p^2=0.39$, with no residual quadratic trends. Importantly, there were no interactions with stimulus type. For the 240-Set stimuli, there were only main effects of stimulus type, $F(1,70)=83.33$, $\eta_p^2=0.54$; and orientation, $F(2,140)=28.73$, $\eta_p^2=0.29$. There were also significant linear, $F(1,70)=11.47$, $\eta_p^2=0.14$; and

quadratic trends in the data, $F(1,70)=42.31$, $\eta_p^2=0.38$, with the quadratic contrast explaining a larger proportion of the variance (38% vs. 14%, respectively). Again, there were no interactions between stimulus type and orientation. Overall, these analyses show that subjects responded more quickly with trained orientations (0° and 150° for the 150-Set and 0° and 240° for the 240-Set) than with novel orientations for faces and Greebles. Based on the trend analysis and on previous data (e.g., Tarr & Pinker, 1989), we confirmed this finding by a planned contrast comparing trained versus non-trained orientations ($F(1,35)=36.34$ for faces, and $F(1,35)=43.75$ for Greebles). For both sets, stimuli presented at the 0° orientation was also correctly named the quickest. This finding was confirmed by a planned contrast comparing the 0° versus the non-upright orientations ($F(1,35)=24.14$ for faces, and $F(1,35)=47.20$ for Greebles). Finally, although faces were named more quickly than Greebles, overall rotations in the picture plane appeared to be equally detrimental to both stimulus types (i.e., no interactions between stimulus type and orientation).

Slope. The collapsed RT data, per stimulus type and block, were regressed against orientation to determine the orientation function. Table 1 also presents the slopes for faces and Greebles for the first and second block from Experiment 2. A mixed-design ANOVA with stimulus type (faces, Greebles) as a between-subjects factor and block (1-2) as a within-subjects factor revealed that the slopes were not significantly different between faces and Greebles, $F(1,70)=0.96$, $\eta_p^2=0.01$; but the rates of rotation were significantly different across blocks, $F(1,70)=5.11$, $\eta_p^2=0.07$. However, the block effect was relatively small. Lastly, there was no significant interaction between stimulus type and block, $F<1$, $\eta_p^2=0.006$.

Accuracy. The error rates for faces and Greebles as a function of orientation, block, and stimulus set are plotted in Figures 9 and 10. As with RTs, the error data were submitted to a mixed-design ANOVA with stimulus type (faces, Greebles) as a between-subjects factor,

and with block (1-2) and orientation (0° - 120°) as repeated factors. There were significant main effects of block, $F(1,70)=10.39$, $\eta_p^2=0.13$; and orientation, $F(4,280)=11.45$, $\eta_p^2=0.14$. The main effect of stimulus type was marginally significant and small, $F(1,70)=3.62$, $p=0.06$, $\eta_p^2=0.05$. There was also a small interaction between stimulus type and block, $F(1,70)=6.21$, $\eta_p^2=0.08$. Overall, there was no evidence of a speed/accuracy tradeoff.

----- Insert Figures 9 and 10 about here -----

Discussion

As in Experiment 1, there were little differences in how faces and Greebles were recognized when they were rotated in the picture plane. First, naming times increased in a linear manner as the stimuli were rotated away from *any* of the trained orientations (i.e., 0° , 150° , and 240°). Second, additional practice at novel orientations resulted in an attenuation of the orientation effect so that subjects responded more quickly when the stimuli were presented again at these orientations on the second block of test trials (e.g., Jolicoeur, 1985; Murray et al., 1993; Tarr & Pinker, 1989). However, practice effects were small for both faces and Greebles. Third, the results of Experiment 2 demonstrate the orientation-specificity of learning: following training at specific orientations in the picture plane, subjects named stimuli quickest at their trained orientations. Furthermore, subjects responded fastest for the upright orientation.

As in Tarr and Pinker (1989; see also Jolicoeur & Milliken, 1989, Experiment 1), the speed benefit associated with training with a face or Greeble at a specific orientation does not generalize to other stimuli of the same class, or to trained stimuli shown at novel orientations. Only stimuli trained at 150° showed a benefit at 150° during test, and only stimuli trained at 240° showed a benefit at 240° during test. That the upright orientation is still named more

quickly than the non-upright trained orientation may partly be due to the training procedure: Subjects were trained more often with the upright orientation than either the 150° or 240° orientation. Overall, these findings are consistent with Tarr and Pinker's assertion that object representations are both orientation and object specific. Faces and Greebles do not appear to be exceptions to this rule.

We note, however, that Jolicoeur and Milliken (1989) found evidence of orientation generalization with familiar objects. In their second experiment, subjects named familiar objects that were shown upright in the context of other objects shown rotated in the picture plane. In a surprise block, the upright-only objects were also shown at non-upright orientations. They found that subjects were not affected by orientation for these objects, suggesting that subjects could generalize to novel orientations. In Jolicoeur and Milliken's study, observers named visually distinctive objects at the basic level (see also Hamm & McMullen, 1998). The orientation and stimulus specificity found in this experiment may therefore depend on the visual homogeneity of the stimulus class (as with the homogenous 2D figures used by Tarr & Pinker, 1989) and subordinate-level recognition (e.g., specificity may be increased by prompting subjects to attend to subtle or more complex features).

Lastly, the results of Experiment 2 extend the findings of Experiment 1 and those of previous studies on the recognition of misoriented faces (e.g., Bruyer et al., 1993; Collishaw & Hole, 2002; Lewis, 2001; Murray et al., 2000; Valentine & Bruce, 1988). In particular, we find that representations of faces include specific orientation information about those faces, and that this information can generalize, in an orientation-sensitive manner, to a limited extent to novel orientations. One possibility is that the inclusion of orientation-specific information may be related to subjects' reliance on metric spatial relations between facial features (e.g., the distance between the eyes), which have been shown to be a critical factor for recognizing individual faces (e.g., Collishaw & Hole, 2002).

General Discussion

The majority of studies on the *face inversion effect* have compared the recognition of upright and inverted faces and non-face objects (e.g., Yin, 1969). This comparison leaves unanswered potential similarities and differences in how faces and non-face objects are processed, for example, the types of visual features that may be critical for recognition. Although investigators have systematically examined how faces at other picture-plane orientations are recognized (Bruyer et al., 1993; Collishaw & Hole, 2002; Lewis, 2001; Murray et al., 2000; Sjöberg & Windes, 1992; Stürzel & Spillman, 2000; Valentine & Bruce, 1988), they have not directly compared faces and non-face objects that were matched to faces along various stimulus dimensions. Conversely, although investigators have compared faces and matched controls (e.g., Yin, 1969; Yovel & Kanwisher, 2004), they have not tested both at multiple orientations in the picture plane. Here we systematically rotated both faces and Greebles in the picture plane and measured how well subjects could recognize them across this continuum.

Our main findings are as follows. First, we found a strong and largely linear dependence of response times on orientation for naming faces. That is, there was no qualitative shift (e.g., a step function) in recognition performance for upright and inverted faces (e.g., Collishaw & Hole, 2002; Valentine & Bruce, 1988). Second, we found that observers were generally faster and more accurate with faces than Greebles across all orientations, although faces were potentially more difficult to recognize (e.g., circular outline, and more targets and distractors). That is, there was a baseline difference in performance between the two stimulus types. Third, despite this baseline difference, we consistently found similar orientation effects for faces and Greebles across different learning and testing conditions. Fourth, we found that this orientation effect diminished with practice at familiar

orientations for both faces and Greebles. These practice effects are larger when faces and Greebles are only trained on the upright orientation (Experiment 1), possibly because training at multiple orientations (Experiment 2) benefits subsequent generalization to untrained orientations (e.g., Gauthier & Tarr, 1997b). Finally, we found that recognition was highly specific to learned orientations and did not transfer to exemplars of the same categories for either stimulus type.

Taken as a whole, the data reported here suggest that the *face inversion effect* by itself does not provide a strong argument in favor of separable mechanisms for faces and for non-face objects. Rather, in conjunction with existing data from both the face and object recognition literature, we think that the inversion effect is largely driven by properties of the stimulus categories and not the stimulus category *per se*. Specifically, observers need to discriminate between highly similar features and configurations of features to recognize both faces and Greebles. Whether this high degree of discrimination is carried out by a single mechanism (e.g., Tarr & Cheng, 2003; Valentine & Bruce, 1988) or by qualitatively different mechanisms for faces and non-face objects alike (e.g., Yin, 1969; Yovel & Kanwisher, 2004) is an issue to be addressed by future work.

The present study helps to target properties that may be critical for this line of research. In the face recognition literature, researchers have shown that inversion disrupts configural processing (e.g., Collishaw & Hole, 2002; Goffaux & Rossion, in press; Freire et al., 2000; Leder et al., 2000; Tanaka & Farah, 1993). Consequently, observers may be more likely to rely on individual local features for inverted faces. However, the recognition of some non-face upright objects may also rely on the same configural processes as the recognition of faces. As our data suggest, discriminating between individuals within a homogeneous stimulus class with a dominant single orientation may be sufficient to recruit configural processes and, hence, be equally disrupted by inversion.

There are two related limitations of the present results that need to be addressed. First, Greebles are arguably “face-like”, which may explain why there was no difference in the inversion effect between these two stimulus categories. Greebles certainly have three smaller parts attached in a symmetric manner to a larger central part. They do have a biological appearance. However, they do not appear to be faces in the absolute: the shapes of Greeble parts have little in common with most faces parts; the surface texture and patterning on Greebles is nothing like those found on faces; and they have a different 3D structure with all smaller parts protruding from the larger part. In the present study, we used Greebles as our matched stimulus controls because individual Greebles share similar features arranged in similar configurations, thus forcing discrimination among them to be based on subtle differences among configural and featural information. Our view is that such image geometry coupled with individual-level recognition is critical in driving the face inversion effect. Although Greebles share these properties with faces by design, there is converging evidence indicating that Greebles and faces are processed differently by observers who are unfamiliar with these novel objects (Gauthier & Tarr, 1997a). Put another way, Greebles are not “face-like” by default in Greeble novices. For example, the visual presentation of Greebles does not activate the middle fusiform gyrus as measured by fMRI in naïve observers (Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999). That is, the middle fusiform has been implicated in face processing (Kanwisher, McDermott, & Chun, 1997), but it is only *after* observers have become Greeble experts that they exhibit category-selective neural responses in this functionally-defined area. Reinforcing the idea that the critical variable in any *similarity* between faces and Greebles is one of expertise – automatized processing at the individual level – and not image geometry, two behavioral studies have found that some configural effects are obtained with Greeble experts, but not Greeble novices (Gauthier & Tarr, 1997a; Gauthier, Williams, Tarr, & Tanaka, 1998). Finally, in a neuropsychology study designed

specifically to address the issue of the “face-like” nature of Greebles, CK, an agnosic patient with preserved face recognition abilities but impaired object recognition abilities, performed poorly at several tasks with Greebles, suggesting that he was unable to extend his spared face-specific abilities to Greebles (Gauthier et al., 2004). Similarly, there is a developmental prosopagnosic patient who is impaired with face recognition but learns Greebles as well as control observers (Duchaine, Dingle, Butterworth, & Nakayama, 2004).

That being said, whether Greebles are or are not “face-like” may remain an unanswerable question. There are also data that argue against the ideas that factors other than image geometry play a critical role in obtaining configural effects (e.g., Yovel & Kanwisher, 2004) or that expertise is critical for category-selective responses in the middle fusiform gyrus (Rhodes, Byatt, Miche, & Puce, 2004). In the end, we should emphasize that it was not the purpose of our study to test whether Greebles are face-like or not. They are not *de novo* processed in a manner equivalent to faces, thus, for our present purposes they provide an appropriate control for examining the effect of misorientation in the picture-plane on individual-level object recognition.

The second limitation of our interpretation of the present results is that the category of faces may be the predominant one in which exemplars are mono-oriented and share similar features arranged in a similar configuration. As we used artificial non-face objects as our controls, the ecological validity of our results may be limited. We would like to emphasize that the results point to properties that seem to be important for driving the face inversion effect. Faces may indeed be “special” in this sense but it is important to clarify what this means more precisely (see also Gauthier et al., 2004).

It is important to note that one critical dimension that often varies between faces and non-face objects is the level of expertise (Gauthier et al., 1999; Gauthier & Logothetis, 2000; Diamond & Carey, 1986; Tanaka & Farah, 1993). Studies have shown that with the

development of expertise, observers can become sensitive to configural information that may be disrupted by rotations in the picture plane. Not surprisingly, expertise is often seen within homogenous object domains such as dogs and cars (Bukach, Gauthier, & Tarr, 2006). For example, Diamond and Carey found a large inversion effect using human faces for both the dog experts and novices but only the dog experts showed a large inversion effect for pictures of dogs (but see Robbins & McKone, *in press*, who did not replicate these results). Given the limited number of stimuli in Experiment 1 and the additional training in Experiment 2, observers may have gained some degree of expertise with the novel Greebles tested. These factors may lead to similar inversion effects found for faces and Greebles in the present study. Future work is needed in this area to determine the extent that expertise may contribute to inversion effects for faces and non-face objects (Gauthier & Tarr, 1997a).

In sum, the degree to which misorientation in the picture plane affects recognition depends on at least the visual homogeneity of the stimulus class (i.e., similar features and configuration of features), whether exemplars have a single dominant orientation, and the level at which exemplars are recognized. Our results indicate that these factors magnify in a systematic and quantitative manner the difficulties associated with recognizing misoriented faces and non-face objects alike.

Acknowledgments

Funding for this research was provided to Alan Ashworth by the Air Force Research Laboratory, Brooks Air Force Base, San Antonio, Texas., and to Michael J. Tarr by the Air Force Office of Scientific Research Grant #F49620-91-J-0169. We thank Scott Yu for creating the novel objects used as control stimuli, Isabel Gauthier for her help in generating these stimuli, and 3 anonymous reviewers for their comments and suggestions.

References

- Biederman I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Brooks, B. E., Rosielle, L. J., & Cooper, E. E. (2002). The priming of face recognition after metric transformations. *Perception*, 31, 297-313.
- Bruyer, R., Galvez, C., & Prairial, C. (1993). Effect of disorientation on visual analysis, familiarity decision and semantic decision on faces. *British Journal of Psychology*, 84, 433-441.
- Bukach, C., Gauthier, I., & Tarr, M. J. (2006). Beyond faces and modularity: The power of an expertise framework. *Trends in Cognitive Sciences*, 10, 48-63.
- Collishaw, S. M., & Hole, G. J. (2002). Is there a linear or a nonlinear relationship between rotation and configural processing of faces? *Perception*, 31, 287-296.
- Diamond, R., & Carey, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, 115, 107-117.
- Duchaine, B. C., Dingle, K., Butterworth, E., & Nakayama, K. (2004). Normal Greeble learning in a severe case of developmental prosopagnosia. *Neuron*, 43, 469-473.
- Eley, M. G. (1982). Identifying rotated letter-like symbols. *Memory & Cognition*, 10, 25-32.
- Freire, A., Lee, K., & Symons, L. A. (2000). The face-inversion effect as a deficit in the encoding of configural information: Direct evidence. *Perception*, 29, 159-170.
- Gauthier, I., Behrmann, M., & Tarr, M. J. (2004). Are Greebles like faces? Using the neuropsychological exception to test the rule. *Neuropsychologia*, 42, 1961-1970.
- Gauthier, I., & Logothetis, N. (2000). Is face recognition not so unique after all? *Cognitive Neuropsychology*, 17, 125-142.

Gauthier, I., & Tarr, M. J. (1997a). Becoming a “Greeble” expert: Exploring the face recognition mechanism. *Vision Research*, 37, 1673-1682.

Gauthier, I., & Tarr, M. J. (1997b). Orientation priming of novel shapes in the context of viewpoint-dependent recognition. *Perception*, 26, 51-73.

Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform “face area” increases with expertise in recognizing novel objects. *Nature Neuroscience*, 2, 568-573.

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.

Goffaux, V., & Rossion, B. (in press). Face inversion disproportionately impairs the perception of vertical but not horizontal relations between features. *Journal of Experimental Psychology: Human Perception & Performance*.

Hamm, J. P., & McMullen, P. A. (1998). Effects of orientation on the identification of rotated objects depend on the level of identity. *Journal of Experimental Psychology: Human Perception & Performance*, 24, 413-426.

Jolicoeur, P. (1985). The time to name disoriented natural objects. *Memory & Cognition*, 13, 289-303.

Jolicoeur, P., & Milliken, B. (1989). Identification of disoriented objects: Effects of context of prior presentation. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 15, 200-210.

Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*, 17, 4302-4311.

Lawson, R., & Jolicoeur, P. (1999). The effect of prior experience on recognition thresholds for plane-disoriented pictures of familiar objects. *Memory & Cognition*, 27, 751-758.

Leder, H., Candrian, G., Huber, O., & Bruce, V. (2001). Configural features in the context of upright and inverted faces. *Perception*, 30, 73-83.

Lewis, M. B. (2001). The Lady's not for turning: Rotation of the Thatcher illusion. *Perception*, 30, 769-774.

Maurer D., Grand, R. L., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, 6, 255-260.

Murray, J. E. (1998). Is entry-level recognition viewpoint invariant or viewpoint dependent? *Psychonomic Bulletin & Review*, 5, 300-304.

Murray, J. E., Jolicoeur, P., McMullen, P. A., Ingleton, M. (1993). Orientation-invariant transfer of training in the identification of rotated natural objects. *Memory & Cognition*, 21, 604-610.

Murray, J. E., Yong, E., & Rhodes, G. (2000). Revisiting the perception of upside-down faces. *Psychological Science*, 11, 492-496.

Palmer, S., Rosch, E., & Chase, P. (1981). Canonical perspective and the perception of objects. In J. Long & A. Baddeley (Eds.), *Attention and Performance IX* (pp. 135-151). Hillsdale, NJ: Lawrence Erlbaum.

Rhodes, G., Byatt, G., Miche, P. T., & Puce, A. (2004). Is the fusiform face area specialized for faces, individuation, or expert individuation? *Journal of Cognitive Neuroscience*, 16, 189-203.

Riesenhuber, M., Jarudi, I., Gilad, S., & Sinha, P. (2004). Face processing in humans is compatible with a simple shape-based model of vision. *Proceedings of the Royal Society: Biological Sciences*, 271, S448-S450.

Robbins, R., & McKone, E. (2003). Can holistic processing be learned for inverted faces? *Cognition*, 88, 79-107.

Robbins, R., & McKone, E. (in press). No face-like processing for objects-of-expertise in three behavioural tasks. *Cognition*.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology*, 8, 382-439.

Scapinello, K. F., & Yarmey, A. D. (1970). The role of familiarity and orientation in immediate and delayed recognition of pictorial stimuli. *Psychonomic Science*, 21, 329-330.

Sekuler, A. B., Gaspar, C. M., Gold, J. M., & Bennett, P. J. (2004). Inversion leads to quantitative, not qualitative, changes in face processing. *Current Biology*, 14, 391-396.

Shinar, D., & Owen, D. H. (1973). Effects of form rotation on the speed of classification: the development of shape constancy. *Perception & Psychophysics*, 14, 149-154.

Sjöberg, W., & Windes, J. (1992). Recognition times for rotated normal and "Thatcher" faces. *Perceptual & Motor Skills*, 75, 1176-1178.

Stürzel, F., & Spillman, L. (2000). Thatcher illusion: Dependence on angle of rotation. *Perception*, 29, 937-942.

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

Tarr, M. J., & Cheng, Y. D. (2003). Learning to see faces and objects. *Trends in Cognitive Sciences*, 7, 23-30.

Tarr, M. J., & Pinker, S. (1989). Mental rotation and orientation-dependence in shape recognition. *Cognitive Psychology*, 21, 233-282.

Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition*, 32, 193-254.

Valentine, T., & Bruce, V. (1986). The effect of race, inversion and encoding activity upon face recognition. *Acta Psychologica*, 61, 259-273.

Valentine, T., & Bruce, V. (1988). Mental rotation of faces. *Memory & Cognition*, 16, 556-566.

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

Yovel, G., & Kanwisher, N. (2004). Face perception: Domain specific, not process specific. *Neuron*, 44, 889-898.

Figure Captions

Figure 1. Examples of the cropped face stimuli used in Experiments 1 and 2.

Figure 2. Examples of the Greeble stimuli used in Experiments 1 and 2. Each row shows Greebles from a different family (same body shape). Each column shows exemplars from a family (different parts).

Figure 3. Mean response times and error rates averaged across subjects in Experiment 1 for faces as a function of orientation and block. Error bars in this and subsequent figures are standard errors of the mean (SEM).

Figure 4. Mean response times and error rates averaged across subjects in Experiment 1 for Greebles as a function of orientation and block.

Figure 5. Mean response times averaged across subjects in Experiment 2 for faces as a function of orientation, block and training set. Faces from the 150-Set were trained at 0° and 150°, whereas those from the 240-Set were trained at 0° and 240°.

Figure 6. Mean response times averaged across subjects in Experiment 2 for Greebles as a function of orientation, block and training set. Greebles from the 150-Set were trained at 0° and 150°, whereas those from the 240-Set were trained at 0° and 240°.

Figure 7. Mean response times averaged across subjects in Experiment 2 for faces for the 0°, 150° and 240° orientation as a function of block and training set. Note that response times were fastest for trained orientations (either 0° and 150° in the top panel, or 0° and 240° in the bottom panel).

Figure 8. Mean response times averaged across subjects in Experiment 2 for Greebles for the 0°, 150° and 240° orientation as a function of block and training set. Note that response times were fastest for trained orientations (either 0° and 150° in the top panel, or 0° and 240° in the bottom panel).

Figure 9. Mean error rates averaged across subjects in Experiment 2 for faces as a function of orientation, block and training set. Faces from the 150-Set were trained at 0° and 150°, whereas those from the 240-Set were trained at 0° and 240°.

Figure 10. Mean error rates averaged across subjects in Experiment 2 for Greebles as a function of orientation, block and training set. Greebles from the 150-Set were trained at 0° and 150°, whereas those from the 240-Set were trained at 0° and 240°.

Table 1. Mean slopes (SEM) of Experiments 1 and 2 for faces and Greebles on Blocks 1 and 2.

	Experiment 1		Experiment 2	
	Faces	Greebles	Faces	Greebles
Block 1	1.29	0.93	0.98	1.37
	(0.20)	(0.20)	(0.28)	(0.28)
Block 2	0.58	0.58	0.71	0.88
	(0.12)	(0.12)	(0.17)	(0.17)

Figure 1

Figure 2

Figure 3

Figure 4

Figure 5

Figure 6

Figure 7

Figure 8

Figure 9

Figure 10

¹We will return to the more subtle issue of whether Greebles are “face-like” in the General Discussion.

²Indeed, this is another dimension along which Greebles are not “face-like”.

³For all analyses, we report η_p^2 (partial eta-squared) as a measure of effect size or the strength of the association between an experimental factor and a dependent variable. More precisely, η_p^2 is the proportion variance associated with the experimental factor, partialing out other factors, that is:

$$\eta_p^2 = SS_{factor} / (SS_{factor} + SS_{error}).$$

For example, $\eta_p^2=0.30$ would mean that a factor accounted for 30% of the variance. Finally, for all analyses, an $\alpha=0.05$ was adopted to test for any significant effects.