Journal of Experimental Psychology:
Human Perception and Performance
1999, Vol. 25, No. 4, 1016–1035

0096-1523/99/$3.00

Robust Representations for Faces: Evidence From Visual Search

Frank Tong and Ken Nakayama
Harvard University

We report evidence from visual search that people can develop robust representations for highly overlearned faces. When observers searched for their own face versus the face of an unfamiliar observer, search slopes and intercepts revealed consistently faster processing of self than stranger. These processing advantages persisted even after hundreds of presentations of the unfamiliar face and even for atypical profile and upside-down views. Observers not only showed rapid asymptotic recognition of their own face as the target, but could reject their own face more quickly as the distractor. These findings suggest that robust representations for a highly overlearned face may (a) mediate rapid asymptotic visual processing, (b) require extensive experience to develop, (c) contain abstract or view-invariant information, (d) facilitate a variety of processes such as target recognition and distractor rejection, and (e) demand less attentional resources.

The human visual system appears to follow the well-known adage that "practice makes perfect." Most of us would find it easier to distinguish between different letters of our own alphabet than those of a foreign script. Likewise, most of us would find it easier to learn the complex pattern of a new face than a simple array of random dots. The notion that highly familiar objects may be more easily detected, distinguished, recognized, or recalled rests well with our intuition. From an information-processing standpoint, one might predict a form of perceptual learning such that highly experienced visual objects are more optimally processed.

Face recognition may represent one of the most striking examples of how the human visual system is tuned to its environment. The faces of different people show relatively little variance in their underlying structure, yet the image of a particular face can vary dramatically in visual content by changes in pose, lighting, expression, hairstyle, makeup, age, and so forth. Despite these difficulties, humans show surprising expertise at recognizing faces. There is now a large body of evidence from neurophysiology (e.g., Perrett, Rolls, & Caan, 1982), neuropsychology (Farah, Wilson, Drain, & Tanaka, 1995; Moscovitch, Winocur, & Behrmann, 1997), and human imaging studies (Kanwisher, McDermott, & Chun, 1997; Kanwisher, Tong & Nakayama, 1998; Tong, Nakayama, Moscovitch, Weinrib, & Kanwisher, in press) that shows that face and object recognition reflect different underlying substrates and processes.

People may be expert at recognizing faces in several ways. One type of expertise is the ability to recognize upright faces relative to inverted faces. This topic has been heavily explored and provides perhaps the strongest evidence for face-specific processing. Numerous studies have shown that inversion disrupts face recognition more severely than object recognition (e.g., Yin, 1969; see Valentine, 1988, for a review) and that upright faces may involve more configural processing (A. W. Young, Hellawell, & Hay, 1987), holistic processing (Tanaka & Farah, 1993), or better generalization to novel images (Moses, Ullman, & Edelman, 1996) than inverted faces. Moreover, recent studies of neuropsychological patients with selectively intact or impaired face recognition strongly suggest that upright faces can gain access to a face-specific system but inverted faces cannot (Farah et al., 1995; Moscovitch et al., 1997).

A second type of face expertise may be the ability to recognize extremely well-known faces relative to less familiar faces. Surprisingly few studies have addressed this second type of expertise regarding whether extensively learned faces are represented in a different manner than less familiar faces. Many studies have shown processing advantages for familiar faces—for example, when performing scrambled–intact or male–female judgments (Bruce, 1986), when making age judgments (Bruyer, Lafalize, & Distefano, 1991), or when matching face images across changes in age (Bruck, Cavanagh, & Ceci, 1991). However, these studies typically assess familiarity effects by comparing familiar faces (e.g., famous faces) with faces shown for the first time. Given that observers likely have visual representations for the familiar faces but not for the unfamiliar faces, these studies illustrate how representations can facilitate face processing but fail to address whether face representations, once formed, can differ in their robustness or efficacy.

We introduce the term *robust representation* in this article to describe a concept that the term *familiarity* fails to describe. Whereas familiarity connotes some degree of

Frank Tong and Ken Nakayama, Vision Sciences Laboratory, Department of Psychology, Harvard University.

This research was supported by a postgraduate scholarship to Frank Tong from the Natural Sciences and Engineering Research Council of Canada and Air Force Office of Scientific Research Grant F496209510036. We thank Elinor McKone, Adriane Seiffert, and Patrick Cavanagh for helpful comments on earlier versions of the manuscript and Yuan-Sea Lee for preparing the face images in Experiment 2.

Correspondence concerning this article should be addressed to Frank Tong, Vision Sciences Laboratory, Department of Psychology, Harvard University, 33 Kirkland Street, Cambridge, Massachusetts 02138. Electronic mail may be sent to frank@wjh.harvard.edu.

recognition and perceptual fluency that can be rapidly acquired for upright faces, we believe that robust representations are formed for highly overlearned faces that are encountered under a variety of stimulus conditions and contexts and may mediate optimal visual processing. Admittedly, face recognition can be very accurate after a brief viewing period (e.g., Diamond & Carey, 1986; Yin, 1969). However, everyday experience suggests that this type of rapid learning fails to capture the full knowledge that people can develop for faces. People seem to have robust representations for the faces of close family members or friends. Having frequently viewed these faces across dynamic changes in viewpoint, lighting, and expression, as well as more gradual changes that occur with aging, people may develop robust representations that allow both rapid recognition and effective generalization to novel images.

We believe that robust representations have a number of defining properties. Specifically, they may:

1. mediate rapid asymptotic visual processing;
2. require extensive visual experience to develop;
3. contain some abstract or view-invariant information;
4. facilitate a variety of visual and decisional processes across tasks and contexts; and
5. demand less attentional resources.

This list should not be considered an exhaustive one—there are likely many properties that characterize robust representations. In this article, however, we will focus on evidence for these five properties.

To clarify the meaning of Properties 1 and 2, a plausible relationship between recognition performance and the development of face representations over time is illustrated in Figure 1. Figure 1A shows a typical learning curve in which recognition performance begins sluggishly for an initially unfamiliar face but rapidly improves. This rate of improvement decreases over time such that performance may appear asymptotic within a few dozen presentations of the stimulus. Figure 1B, which represents the true asymptotic performance (Property 1), reveals the deceptiveness of the suggested asymptote in Figure 1A. Although no appreciable learning may be measurable in the latter portion of the curve in Figure 1A, much more extensive visual experience (Property 2) may still lead to very gradual learning such that after several thousand trials, recognition performance is considerably improved (see Figure 1B). Only when extensive additional visual experience leads to negligible improvement can one be confident that true asymptote has been achieved.

Presumably, the learning observed in Figure 1A and 1B reflects actual changes in the visual representation. These changes could take the form of an improved match between stimulus and memory (e.g., a more accurate template or structural description) or the development of a more efficient visual code for a particular face. (The potential role of efficient coding in the development of a robust representation is elaborated in the General Discussion section.) The development of this face representation is schematized in Figure 1C. A single trial is typically sufficient for the initial formation of a face representation (e.g., Paller et al., 1992; Yin, 1969). This representation then develops with additional exposures at a decelerating rate that parallels the learning curves shown in Figure 1A and 1B. Although the representation rapidly develops within the first several dozen presentations such that the stimulus acquires high familiarity, it continues to develop and never fully stabilizes until true asymptotic performance is achieved as shown in Figure 1B. Only at this point might we say that a robust representation has been formed. Therefore, robust representations reflect the endpoint of visual learning, or the most extreme form of familiarity, and are signified by asymptotic performance.

Asymptotic performance and extensive learning requirements are not sufficient criteria for a robust representation however. Many visual learning studies of stereo-depth

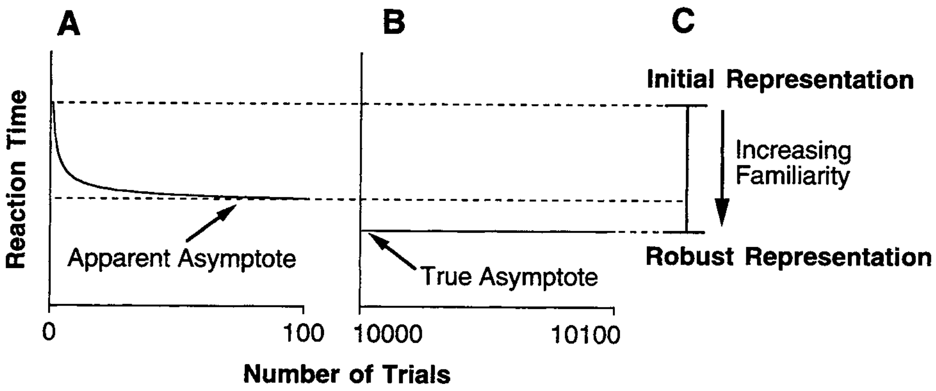

Figure 1. (A) A typical learning curve depicting decelerating improvement in recognition performance for an initially unfamiliar face. Recognition performance appears to reach asymptote within a few dozen presentations. (B) True asymptotic recognition achieved after much more extensive visual experience with the face. Note that additional visual experience leads to negligible learning. (C) Schematic illustration of face representation development and its relationship with recognition performance over time. A fully stabilized or robust representation can be inferred when recognition performance has reached true asymptote.

(Ramachandran & Braddick, 1973), spatial frequency waveform (Fiorentini & Berardi, 1980), vernier (Fahle & Edelman, 1993) and pop-out search (Ahissar & Hochstein, 1996) discrimination reveal remarkable improvements after extensive training with a specific stimulus set. However, this learning typically fails to generalize across very basic changes in stimulus orientation, size, or retinal position.

In contrast, robust representations must be much more flexible, capable of generalizing across changes in stimulus (Property 3), context, or task (Property 4). Well-known faces of family members or friends are frequently seen across dramatic changes in viewpoint, lighting, and expression, yet are easily identified. Perhaps extensive exposure to a face across such image variations leads to representations or visual codes that better capture the view-invariant properties of that face either in terms of distinguishing local features (Penev & Atick, 1996) or global 3-D structure (Atick, Griffin, & Redlich, 1996). If robust representations contain some abstract or view-invariant information (Property 3), one would predict faster processing, not only for highly overlearned views but also for novel or atypical views.

Robust representations should also facilitate a variety of visual and decisional processes, even for unfamiliar tasks or contexts (Property 4). If processing advantages were specific to a particular stimulus–response or stimulus–context association, this would reflect associative learning rather than stimulus-specific learning. On the basis of the assumption that robust representations facilitate a variety of visual and decisional processes, then in the context of visual search, one would predict improvements in both target recognition and distractor rejection despite the rather unusual nature of the latter task.

Finally, one might speculate that robust representations lead to rapid asymptotic processing (Property 1) because they demand less attentional resources (Property 5). According to Norman and Bobrow (1975), learning reflects a decrease in the amount of resources required to achieve a given level of performance. Robust representations might similarly require less attentional resources when processing a relevant face, perhaps by providing a more efficient visual code to mediate such processing. Property 5 leads to the prediction that under conditions of high attentional load, such as large set sizes in visual search, an even greater benefit will be found for robustly represented faces.

We investigated robust representations by using a standard, visual search paradigm in which a particular target is sought among an array of distractors (e.g., Treisman & Gormican, 1988; Treisman & Souther, 1985). Numerous studies have shown that visual search can provide an effective measure of processing efficiency for low-level features such as orientation, size, and color (e.g., Treisman & Gormican, 1988) as well as more complex stimuli such as surfaces (He & Nakayama, 1992), shape from shading (Kleffner & Ramachandran, 1992), gaze direction (von Grünau & Anston, 1995), and facial expressions (Suzuki & Cavanagh, 1995). According to Treisman and Gormican (1988), asymmetries in visual search may reveal what features are strongly coded in the visual system. For example, they found that a curved line "pops out" among straight lines but that a straight line does not pop out among curved lines. They proposed that straight lines may serve as a standard or default value for the visual system, and that standards are difficult to detect as targets, whereas deviations from a standard such as curvature are easy to detect. Similarly, Wang, Cavanagh, and Green (1994) found that a mirror-reversed *Z* would pop out among normal *Z* distractors but that search in the reverse condition was much slower. They concluded that familiarity could support rapid perceptual grouping of distractors and, thereby, lead to pop-out of the novel target. Because our search tasks involved much more difficult discriminations between highly complex stimuli, which resulted in much slower search, our main focus was on the relative search efficiencies among conditions rather than on the parallel versus serial search distinction.

We investigated robust representations by having observers search for a very well-known face, their own (self), versus the face of an unfamiliar paired observer (stranger). Although one might intuitively expect faster search for one's own face, such a view depends on the prediction that extensive visual experience (Property 2) is necessary for rapid asymptotic visual processing (Property 1). If asymptotic performance is rapidly achieved within a dozen or so presentations of an initially unfamiliar face, search differences would be negligible over the course of many trials. We were able to test Properties 1 and 2 by studying the time course of face learning for self and stranger (Experiments 2 and 3).

There were a number of theoretical as well as practical advantages in using one's own face to investigate robust representations. First, we were able to test whether robust representations contain any view-invariant information (Property 3). Although people frequently see or even scrutinize their front and three-quarter views in the mirror while grooming, it is impossible to see profile or upside-down views without using multiple mirrors or an image-capturing device such as a camera. By testing generalization to atypical profile (Experiment 2 and 3) and inverted views (Experiments 1 and 2), we were able to assess if representations for one's own face contain any view-invariant information.

Second, observers were most likely inexperienced at searching for their own face. The task of searching for a small gray-scale image of one's own face among an array of distractor faces, especially under profile or inverted conditions, is surely an unfamiliar one (relevant to Property 4). In comparison, searching for the faces of close friends or family members is a much more familiar task. Using one's own face as a probe, therefore, minimized confounding practice effects and emphasized the underlying differences between face representations for self and stranger.

Finally, our method of having *paired* observers search for their own face versus the other observer's face allowed us to control for any effects caused by visual differences between target faces. Such control is crucial in visual search for a repeatedly presented target face, because the presence of a spurious (Purcell, Stewart, & Skov, 1996) or distinctive facial feature (Nothdurft, 1993) can have a profound effect on search efficiency. Other methods that fail to control for such visual differences (e.g., comparing famous versus nonfamous faces) could unintentionally lead to a study of low-level feature search rather than face search.

To tap into face-specific processes, we tried to develop a search task that was as natural as possible. This was done by using real face images, heterogeneous distractors (Experiments 1–3), and heterogeneous targets (Experiments 2 and 3). Heterogeneous distractors minimized the likelihood that the target could be rapidly distinguished by the presence of a low-level image feature (Purcell et al., 1996). Heterogeneous targets ensured that search could not be based on a single template-matching strategy and increased the realism of the task. Unlike most visual search experiments that present a single version of a target, the precise image or viewpoint of a target object can rarely be anticipated when searching in the real world. By using heterogeneous face images, we hoped to mimic a demanding real-world task that was analogous to searching for a friend in a crowd. Heterogeneous images would also discourage the use of feature or conjunction search strategies and encourage the use of face-recognition processes in order to effectively discriminate between the target and visually similar distractors, yet generalize across different target images.

In Experiments 1 and 2, evidence for robust representations was investigated by manipulating target identity and face orientation in a visual search task. Experiment 1 tested whether reliable differences could be found between searching for a single front-view image of self versus stranger among heterogeneous distractor faces and whether these differences would generalize to inverted views. Experiment 2 investigated whether self versus stranger processing advantages would also generalize to atypical profile views and whether recognition performance across trials would suggest asymptotic processing of one's own face. In Experiment 3, observers searched for self or stranger among the opposite distractor to test if observers would not only recognize their own face more quickly as the target but reject their own face more quickly as the distractor.

Experiment 1: Front-View Target as Self or Stranger

Observers searched for a single front-view image of their own face or the face of an unfamiliar paired observer among heterogeneous distractor faces (see Figure 2), with the entire display presented upright or inverted. This method of pairing observers counterbalanced any effects that might be due to differential distinctiveness of target faces.

The two measures of greatest theoretical interest, based on a serial, self-terminating search model, were search slopes and Set Size 1 intercepts. This model seemed appropriate for our data, which always revealed large search slopes and target-absent slopes that were approximately twice the size of target-present slopes, as is diagnostic of serial, self-terminating search (Treisman & Gormican, 1988; Treisman & Souther, 1985). If one assumes that visual attention or eye movements serially shift from distractor to distractor item until the target is found and search is terminated, then search slopes reflect the rate at which distractors can be serially rejected. Search rates can also be considered a measure of the amount of visual attention required to process each additional distractor item (Treisman & Souther, 1985) and, therefore, may reflect the effects of increasing attentional load (relevant to Property 5). In contrast, Set Size 1 intercepts for target-present trials reflect the amount of time required to recognize the target when no distractors are present.[1] Although slopes and intercepts are often positively correlated in search tasks—both tend to increase with task difficulty—theoretically, it should be possible to dissociate distractor-rejection and target-recognition processes (as tested in Experiment 3).

We hypothesized that target recognition at Set Size 1 would be faster for self than stranger and that this difference would persist across numerous trials (an indirect test of Properties 1 and 2). We also expected faster search rates for self than stranger on the basis of the assumption that robust representations demand less attentional resources (Property 5). Although most previous visual search studies have failed to find an advantage for upright schematic faces in various detection and discrimination tasks (Kuehn & Jolicoeur, 1994; Nothdurft, 1993; but also see Suzuki & Cavanagh, 1995), we suspected that more realistic displays consisting of real faces and heterogeneous distractors would reveal an upright versus inverted face advantage for both target-recognition and distractor-rejection processes. It was less clear whether advantages in self versus stranger processing would persist under inverted conditions. Given that face inversion is believed to disrupt configural (A. W. Young et al., 1987) or holistic processing (Tanaka & Farah, 1993), or to entirely disrupt face-specific processing (e.g., Farah et al., 1995), one might expect that a self versus stranger advantage should be abolished by face inversion. That is, seeing an upside-down image of one's own face should be akin to seeing the upside-down image of an unfamiliar face and require learning of that specific view. However, if representations for one's own face contain some view-invariant information (Property 3) or if matches to memory occur after image orientation is normalized, then the self versus stranger advantage should persist for inverted face displays.

Method

Methods that are common to all three experiments have been appropriately noted.

Observers. Observers in all three experiments were Caucasian men between 18 and 31 years of age with normal or corrected-to-normal vision. All observers were naive regarding the purpose of the experiment, and no observer participated in more than one experiment. The observers in Experiment 1 consisted of 4 paid undergraduates and 4 volunteer graduate students or postdoctoral fellows.

Apparatus. All experiments were conducted on a Macintosh Quadra 610 and presented on a 14-inch CRT monitor (832 × 624 pixels) with 256 gray levels.

[1] A limited-capacity, parallel search model might lead to similar patterns of performance. However, the serial model is discussed here because it provides a more intuitive way of conceptualizing these two processes.

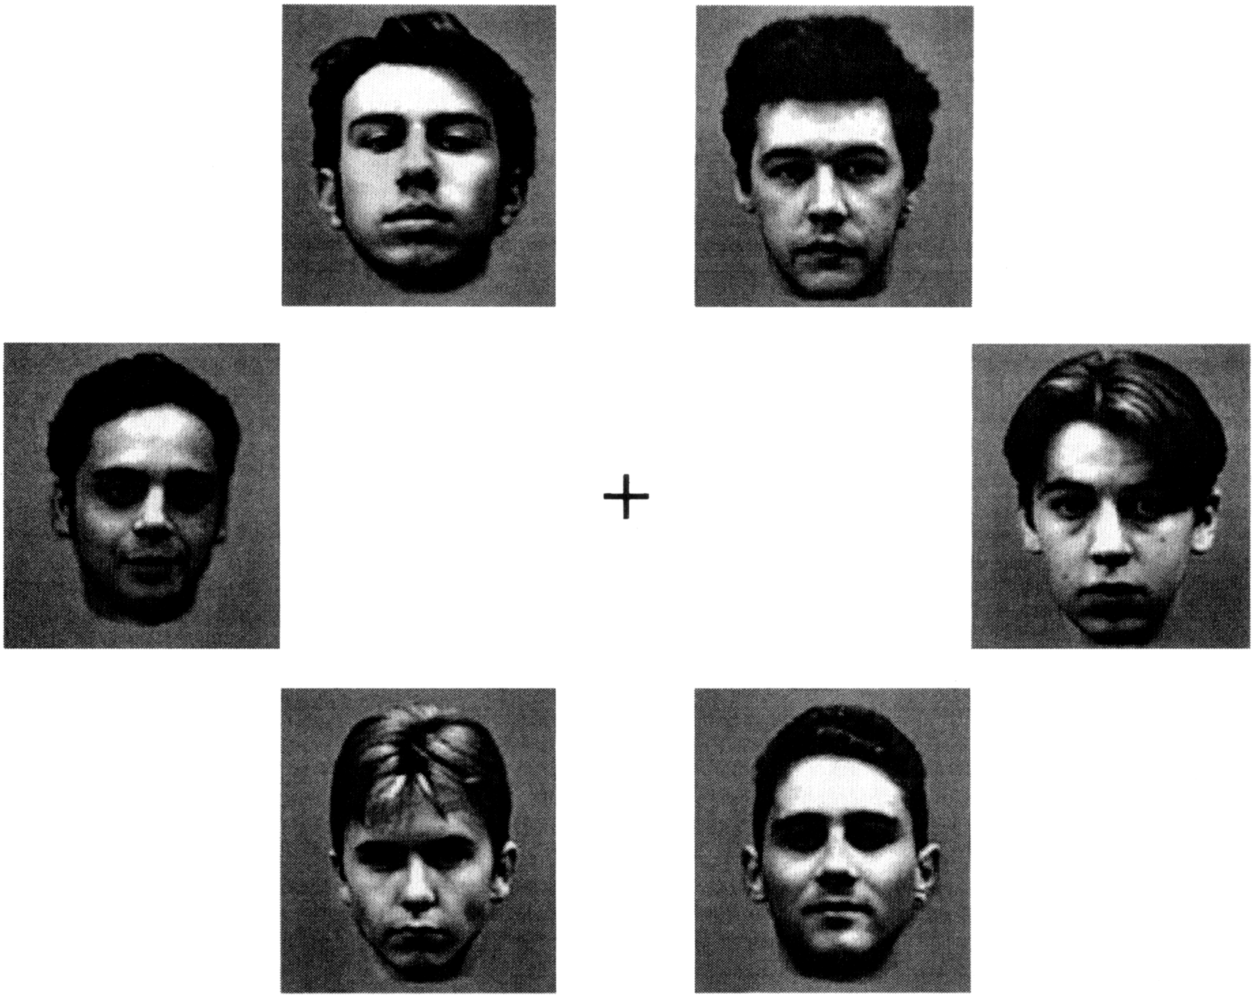

Figure 2. An example of a visual display containing the six distractor face images used in Experiment 1. Observers searched for a front-view image of self or stranger among distractors, with the display presented upright or inverted on four separate blocks of trials.

Stimuli. All face images were of clean-shaven Caucasian men with short hair between the ages of 18 and 31 years. These faces were chosen to minimize recognition that might be based on obvious visual cues related to hair length, facial hair, sex, race, or age.

The unfamiliar distractor faces in Experiment 1 consisted of front-view images of six undergraduate men taken from a publicly accessible face database and are shown in Figure 2. All other face images reported in this article, including those of each observer, were taken in our laboratory on a separate session prior to the experiment. Faces were digitally scanned under controlled room-lighting conditions using a Hi-8 Sony video camera connected to a Macintosh computer. Images were stored with 256 gray levels and cropped so that faces were matched for size and placement within a fixed window size (2.9° × 3.4° or 113 × 126 pixels in Experiment 1). The window background and screen background always differed in luminance to reduce the salience of the external contours of the faces. Gray windows on a white screen were used in Experiments 1 and 3 (Figures 2 and 10), and white windows on a gray screen were used in Experiment 2 (Figures 4 and 5).

The images in Experiment 1 were mildly high-pass filtered to minimize possible differences in low spatial frequency energy between target and distractor faces. Faces were inverted by reversing the image along its horizontal axis.[2]

Procedure. Observers searched for self or stranger with the entire display of faces presented upright or inverted in separate blocks of trials. The order of the four conditions was counterbalanced using a pseudo-Latin square design in which target identity changed between each block and orientation changed between Blocks 2 and 3 (e.g., self–upright, stranger–upright, self–inverted, stranger–inverted). The target image was shown at the beginning of

[2] Although inverting the face image also inverted the light source (i.e., top lighting to bottom lighting), Enns and Shore (1997) have shown that the effects of face orientation and lighting (brow vs. chin lighting) are independent in face-identification tasks. Moreover, they found that inverted chin-lit faces were more poorly recognized than inverted, brow-lit faces. Therefore, maintaining brow-lighting in this experiment would, if anything, underestimate the size of the inversion effect.

each block, followed by 36 practice trials of searching for the target and 150 test trials.

Observers viewed the monitor from a distance of 70 cm and initiated each trial with a key press. On each trial, a central fixation point appeared followed by an array of 2, 4 or 6 faces that appeared 500 ms later. The distractor faces shown were randomly selected from the set of six distractors, and no face could appear more than once on a given trial. Faces were randomly assigned to one of six possible locations, which formed a hexagon around the fixation point subtending a visual angle of 10.1° × 7.7°. Figure 2 shows an example of the display and the distractor faces used in Experiment 1.

Because of the complexity of the visual stimuli and the difficult discriminations involved, eye movements were allowed in all experiments. The target was present on 50% of the trials, and observers responded "target present" by pressing the "/" key with their right hand and "target absent" by pressing the "z" key with their left hand. Auditory feedback was given for incorrect responses. Observers were instructed to respond as quickly as possible while maintaining high accuracy. All set size by target-type trials occurred with equal frequency and were randomly ordered within a given block. Observers were encouraged to rest between blocks, and each experiment required approximately 1 hr to complete.

Data analyses. For all experiments, data analyses were performed on reaction times for correct responses because processing efficiency was our main interest. However, error rates were checked for potential speed–accuracy trade-offs. To assess general trends in the data, a repeated measures analysis of variance (ANOVA) was first performed on mean reaction times with all factors included.

More focused analyses were conducted on the two measures of theoretical interest: target-present reaction times at Set Size 1 and target-present search rates. These measures reflected the amount of time required for target recognition and serial distractor rejection, respectively. Target-absent slopes were used to assess whether search performance appeared serial (on the basis of an approximate 2:1 absent:present slope ratio) but were not further analyzed because of the much more variable criteria used in deciding target absence (Chun & Wolfe, 1996). Our focused analyses on target-present data provided a more reliable measure of distractor-rejection search rates as well as a method of estimating target-recognition processes at Set Size 1. In cases where Set Size 1 reaction times were not directly measured, values were extrapolated using a linear model.

Results

Figure 3 shows mean reaction times for target-present trials, plotted as a function of set size. Two effects are readily apparent: Search was faster for upright than inverted faces and faster for self than stranger. There was no obvious interaction between these two factors.

Search slopes were much steeper than those typically found in feature or conjunction search, and target-absent slopes were more than twice the size of target-present slopes, suggesting serial, self-terminating search. These large search slopes reflected the difficulty of the task and the fact that observers made eye movements in order to discriminate the target. Overall error rates were extremely low in all search conditions (0.8%–1.3%) and showed no evidence of a speed–accuracy trade-off. A summary of present–absent

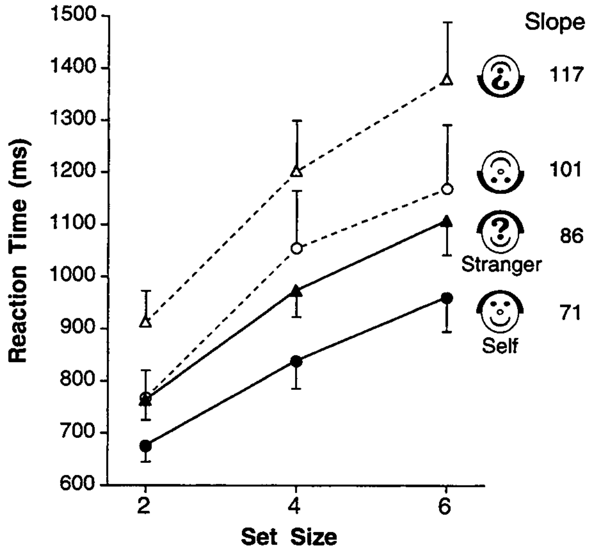

Figure 3. Mean reaction times for target-present trials in Experiment 1. Solid symbols = upright faces; open symbols = inverted faces; circles = self target; triangles = stranger target. Search slopes for each condition are shown in ms/item. Vertical bars represent standard error of the mean (*SEM*). Reaction times are faster for self than stranger and for upright than inverted faces.

search slopes and error rates for all three experiments can be found in the Appendix.

General analyses. A repeated measures ANOVA (Orientation × Target Identity × Target Presence × Set Size) revealed significantly faster reaction times for upright than inverted faces, $F(1, 7) = 7.7, p < .05$, and for target as self than as stranger, $F(1, 7) = 6.7, p < .05$. There was also a significant Orientation × Set Size interaction, $F(2, 14) = 7.0, p < .01$, and a significant Identity × Set Size interaction, $F(2, 14) = 4.0, p < .05$, which suggested a difference in search slopes. Interactions between orientation and identity interaction were not significant.

These general effects were further decomposed into the two measures of interest: search slopes and Set Size 1 intercepts. Focused analyses were performed on target-present data. Given that target presence did not significantly interact with any factors other than set size, our focused analyses on target-present trials should still have been representative of the larger data set.

Distractor-rejection search rates. Although search rates tended to be faster for self than stranger (see Figure 3), the difference between self and stranger was not significant, $F(1, 7) < 1.0$. This appeared to be due to large within-subject and between-subject variability in reaction times and search slopes. Search rates were significantly faster for upright than inverted faces, $F(1, 7) = 11.4, p < .05$.

Target recognition at Set Size 1. Set Size 1 reaction times were extrapolated for the search functions of each observer to estimate the speed of target recognition. Set Size 1 reaction times were significantly faster for self than stranger, $F(1, 7) = 11.2, p < .05$ (self–upright = 611 ms, stranger–upright = 688 ms, self–inverted = 704 ms, stranger–inverted = 827 ms), and faster for upright than inverted faces, $F(1, 7) = 8.8$, $p < .05$. The Orientation × Identity interaction was not significant, $F(1, 7) = 2.5$, $p = .16$.

Discussion

In the upright condition, target recognition was 77 ms faster for one's own face than the stranger's face. (Set Size 1 intercepts for self = 611 ms vs. stranger = 688 ms. Note that intercept values were inflated by saccadic and motor-response latencies to the target item.) This is a dramatic processing advantage, especially when one considers that this difference occurred after 36 practice trials and across 150 test trials. Given that observers made very few errors when searching for the stranger's face and performed hundreds of identification responses to this repeated face image, the stranger's face must have become quite familiar. This target-recognition difference that seemed to endure over time provides some support for the notion that robust representations for a highly overlearned face may lead to rapid processing (Property 1) and may require extensive visual experience to develop (Property 2).

Observers showed the same target-recognition advantage for their own face under inverted conditions (Set Size 1 intercepts for self = 704 ms vs. stranger = 827 ms). This was somewhat surprising given that inversion is known to severely impair configural or holistic processing for faces (Tanaka & Farah, 1993; A. W. Young et al., 1987). Apparently, knowledge of one's own face can generalize to novel picture–plane orientations, despite a recognition cost that is due to inversion. This could either reflect the fact that representations for one's own face contain some view-invariant information (Property 3) or that the inverted images were first normalized for orientation and then matched to memory.

We also found that target-recognition and distractor-rejection processes (as reflected by Set Size 1 intercepts and search slopes, respectively) were significantly faster for upright than inverted faces. These results demonstrate the severity of the face inversion effect. Even though the inverted target consisted of a single, repeatedly presented image, observers could not learn to discriminate this image nearly as well as the upright target. Our findings differ from those of previous visual search studies, which failed to find a face inversion effect (Kuehn & Jolicoeur, 1994; Nothdurft, 1993; but see also Suzuki & Cavanagh, 1995), perhaps because our experiment used a difficult identity-discrimination task and heterogeneous natural images.

Although search rates were generally faster for self than stranger, this difference was not significant. The within-subject data were rather variable because of accumulating practice effects and because the similarity between a given observer's face and the distractor faces could not be directly controlled (one observer's face appeared particularly distinctive). Although our design would have counterbalanced such similarity effects across observer pairs, variations in target–distractor similarity within an observer pair would have introduced considerable within-subject variability. Experiment 2 attempted to address these problems.

Experiment 2: Multiple-View Targets as Self or Stranger

Whereas Experiment 1 demonstrated faster target recognition of self than stranger for canonical front views, Experiment 2 investigated whether representations for one's own face can effectively generalize to atypical depth-plane rotations. This was done by generating heterogeneous target and distractor images for front, three-quarter, and profile views of each face (see Figures 4 and 5). Multiple views also ensured that observers could not guide their search by using a single template-matching strategy, thereby increasing the realism and difficulty of the search task. To reduce the problem of variable target–distractor similarity observed in Experiment 1, we attempted to pair observers who had more visually similar faces and eliminated potentially distinguishing hair cues by using a black cap (see Figure 4). Although covering a salient facial feature such as hair could disturb the matching process between the image and one's internal representation, we suspected that observers would still show an advantage in recognizing the internal features of their own face (cf. A. W. Young, Hay, McWeeny, Flude, & Ellis, 1985).

We were interested in whether representations for one's own face contain view-specific or view-invariant information. If visual representations for highly overlearned faces develop from view-specific encoding, then observers should show a pronounced self versus stranger advantage for canonical front and three-quarter views but a negligible advantage for atypical profile views. In other words, viewing one's own face in profile should be akin to viewing the stranger's face in profile. Both faces should appear unfamiliar and lead to similar rates of processing. However, if representations for one's own face contain some view-invariant information (Property 3), then target-recognition and distractor-rejection processes should be consistently faster for self when compared with stranger as a baseline across all views. The use of stranger as a baseline measure was important because we also expected generally slower search for profiles (even though each view was presented an equal number of times). This was based on the simple fact that capped faces provide less distinguishing visual information from the profile view (cf. front, three-quarter, and profile views in Figure 4; see also Bruce, Valentine, & Baddeley, 1987).

Another key question was whether target recognition for self would improve over repeated presentations or suggest asymptotic processing (see Figure 1A and 1B). If representations for one's own face are truly robust, then additional visual experience should lead to minimal processing benefits (Property 1). By contrast, observers should show considerable improvement for the stranger's face across repeated presentations. We were also interested in whether learning for this initially novel face would proceed gradually or rapidly, and whether recognition for stranger would eventually reach the level of performance seen for self. By using Set Size 1 as a direct measure of target recognition across

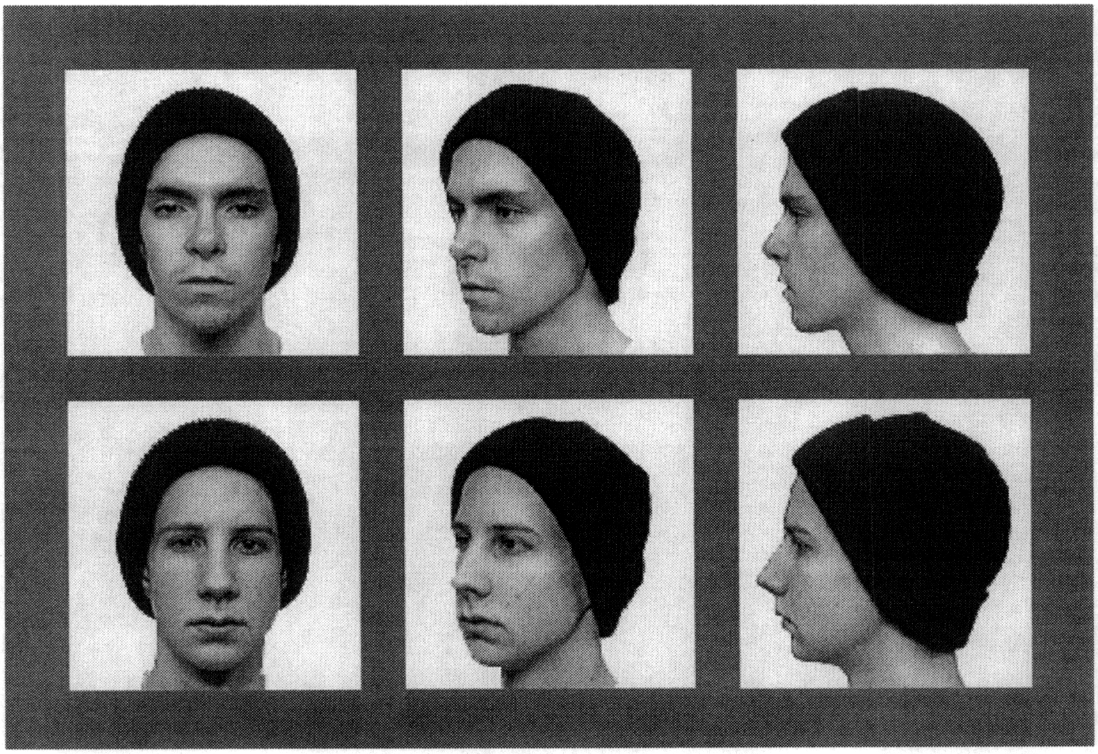

Figure 4. An example of target stimuli in Experiment 2. Front, three-quarter, and profile views of two paired observers are shown.

trials, we were able to study the time course of face learning for self and stranger.

Method

Methods were identical to those described in Experiment 1 except as indicated below.

Observers. Observers consisted of 1 graduate and 15 undergraduate Caucasian men ranging from 18 to 26 years of age. Observers were either paid or received course credit in an introductory psychology course for participation. Data from 3 other observers were discarded because 1 observer experienced computer malfunction, another observer recognized his paired observer, and the pair of the third observer failed to participate.

Stimuli. All target and distractor faces were imaged under identical conditions. Models wore a black cap that concealed their hair, and images were taken of their front, three-quarter, and profile views (0°, 45°, and 90° depth-plane rotations) against a white background. Faces were illuminated by a 40-W incandescent lamp placed just above the camera. The lamp shined 15° downward at the face, and diffuse room lighting also provided a small amount of illumination. Given that the faces were almost front-lit, face inversion should have had a reduced effect on apparent lighting direction compared with Experiment 1 (see also Enns & Shore, 1997).

Six models who did not participate in the experiment served as distractor faces (see Figure 5). An additional model served as the practice target. These faces were unfamiliar to all participating observers. All face images were cropped and matched for size and placement in a 3.4° × 3.5° window.

Target faces were mirror-reversed so that front-view images matched what observers would see in a mirror. After mirror reversal, three-quarter and profile view images faced toward the left. Inverted face images were generated by rotating upright images by 180° in the picture plane. Whenever possible, we attempted to pair observers with visually similar faces. Such pairings attempted to minimize differences in skin, eye, or eyebrow color as well as obvious differences in face shape (e.g., wide or narrow face, rounded or sharp facial features). Figure 4 shows the face images of 2 paired observers.

Procedure. Observers searched for self or stranger with faces presented upright or inverted in four separate blocks. Within each block, randomly mixed trials of front, three-quarter, and profile view faces occurred. The order of the four blocks was counterbalanced using a pseudo-Latin square design. Target identity alternated between each block, and orientation changed between Blocks 2 and 3.

To accurately measure the time course of face learning for self and stranger, it was important to minimize the confounding effect of having to learn the distractor faces in a novel orientation in Blocks 1 and 3. To do this, observers were first familiarized with these novel orientation distractors and trained at searching for the practice target among these distractors prior to the actual test phase that involved searching for a test target (i.e., self or stranger). The distractor-familiarization phase consisted of sequential presentations of front, three-quarter, and profile images of each distractor face for a 2-s viewing period. Observers were instructed to carefully study each face. Observers were then shown front, three-quarter, and profile images of the practice target for free viewing, followed by 56 trials of searching for the practice target

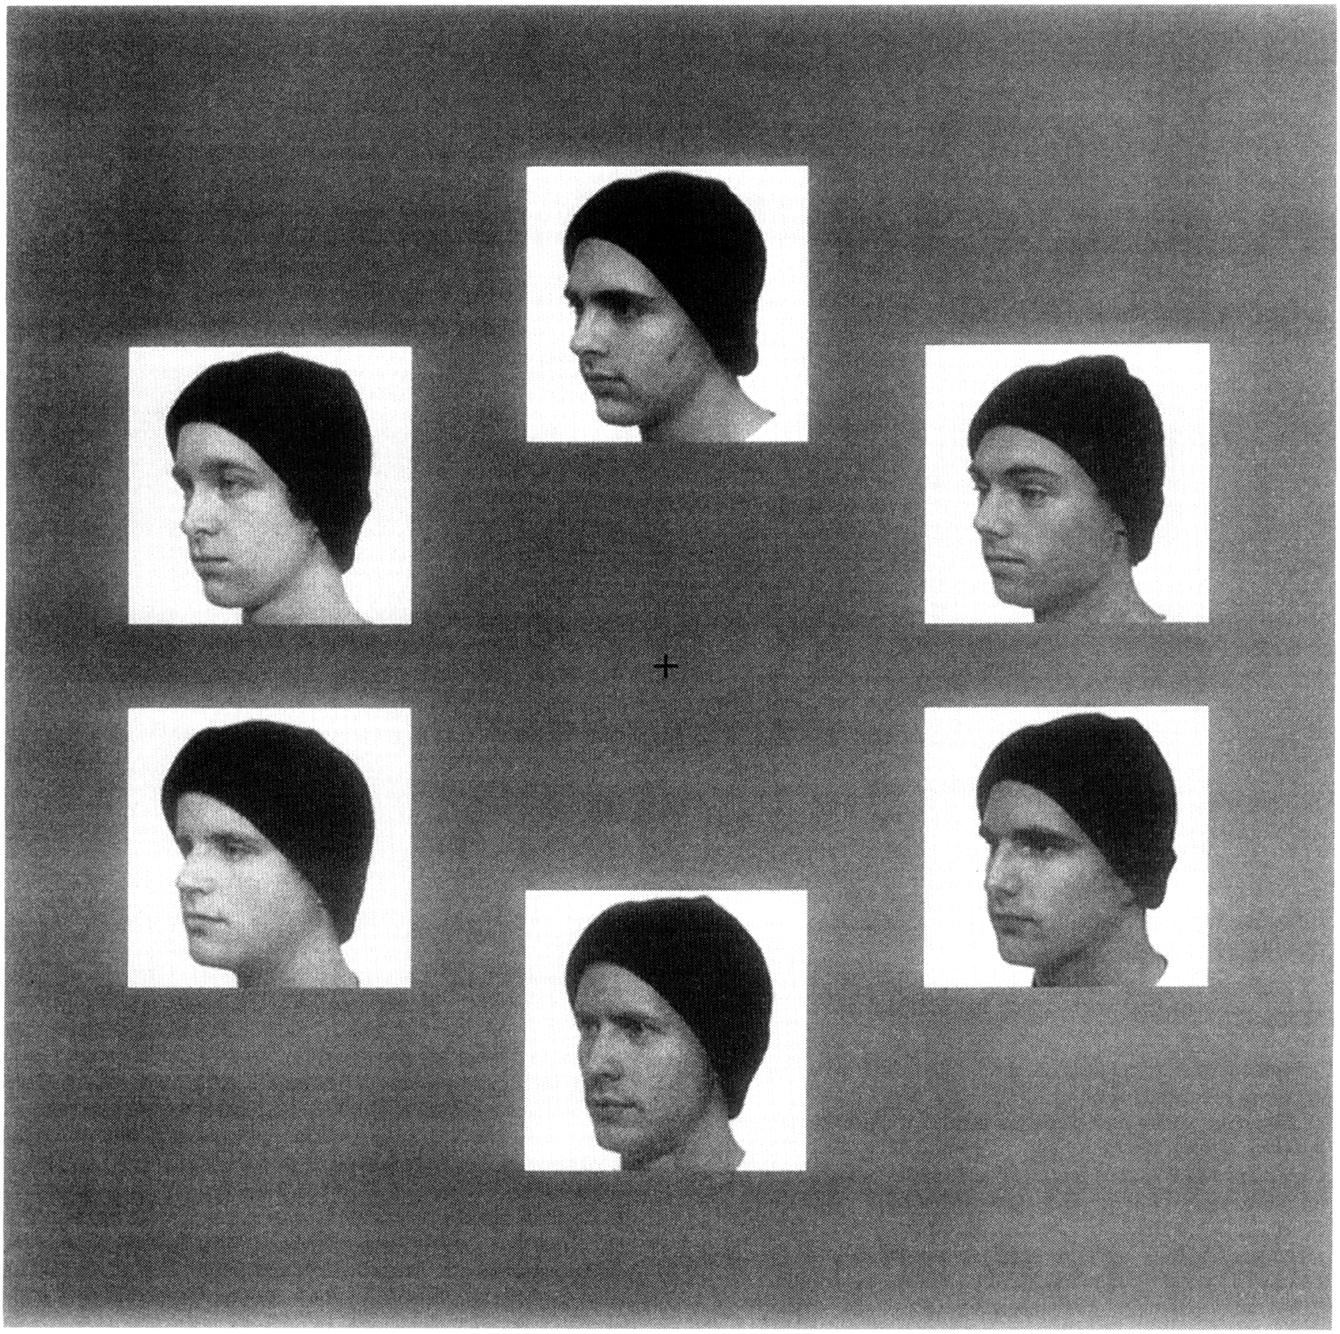

Figure 5. The six distractor faces from Experiment 2 shown in three-quarter view and arranged in a sample visual array. On a given trial, all faces were directed in the same viewpoint and could appear in front, three-quarter, or profile view. Observers searched for self or stranger among distractors, with the display presented upright or inverted on four separate blocks of trials.

under this novel orientation condition. These familiarization and practice target phases were followed by the actual test phase of searching for self or stranger. Observers were allowed to freely view the three images of the test target, followed by 18 practice trials and 144 test trials of searching for the test target.

On a given practice or test trial, one, three, or six faces of differing identities were presented on a gray screen background, all directed in the same viewpoint. An example of a visual display containing three-quarter views of the six distractor faces is shown in Figure 5. Faces were evenly spaced in angular units around the fixation point, and the 12 possible face locations formed an elliptical pattern. Horizontal and vertical axes of the ellipse subtended $11.8° \times 9.2°$ and were approximately scaled for cortical magnification along the principle axes (Rovamo & Virsu, 1979).

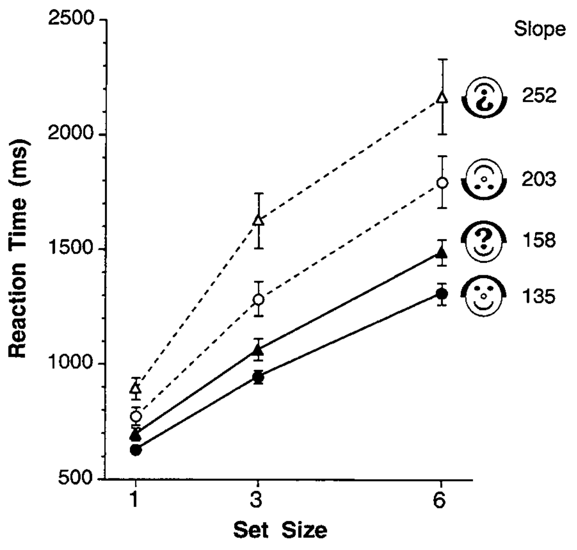

Figure 6. Mean reaction times for target-present trials in Experiment 2. Solid symbols = upright faces; open symbols = inverted faces; circles = self target; triangles = stranger target. Search slopes are listed in ms/item. Plot symbols may exceed the size of vertical bars, which represent ±1 *SEM.* Note that search slopes and Set Size 1 reaction times are consistently faster for self than stranger and for upright than inverted faces.

An intertrial interval of 1,500 ms was used. Within a practice or test block, the target appeared an equal number of times for each View × Set Size condition (8 trials/condition × 3 views × 3 set sizes = 72 target presentations, out of 144 test trials).

Results

Figure 6 shows mean reaction times for target-present trials averaged across viewpoint and plotted as a function of set size. The pattern of results strongly matched those of Experiment 1 except search slopes were generally steeper (cf. Figure 6 and Figure 3; note that ordinate axes differ in scale). Search rates were faster for self than stranger and faster for upright than inverted faces. Response times for Set Size 1 also increased in a similar manner across conditions. Target-absent slopes strongly corresponded with target-present slopes in an approximate 2:1 ratio, as is typically found in serial search (see Appendix). Overall error rates were low in the upright condition (self = 2.5%, stranger = 2.6%) and higher in the inverted condition (self = 7.1%, stranger = 13.6%), with no evidence of a speed–accuracy trade-off (see Appendix).

General analyses. A 5-way ANOVA (Orientation × Target Identity × Target Presence × View × Set Size) revealed that all main effects and two-way interactions involving set size were highly significant at the .001 level, as well as many other significant interactions. More detailed analyses were conducted on Set Size 1 reaction times and search slopes for target-present trials. Given that target presence did not significantly interact with any factors other than set size, our focused analyses on target-present trials should have been representative of the larger data set.

Distractor-rejection search rates. Figure 7 shows search slopes for target-present trials as a function of viewpoint. Search rates were significantly faster for self than stranger, $F(1, 15) = 11.0$, $p < .005$, and for upright than inverted faces, $F(1, 15) = 22.3$, $p < .001$. There was also a main effect of viewpoint, $F(1, 15) = 5.9$, $p < .01$, and a polynomial contrast revealed a significant linear component, $F(1, 15) = 7.8$, $p < .05$, suggesting that search slopes were somewhat larger for profiles. No interactions were significant, including the Identity × Viewpoint interaction ($F < 1.0$). In Figure 7, it is readily apparent that the stranger minus self difference for upright search slopes is as large for atypical profile views (stranger − self = 173 − 139 = 34 ms/item) as for typical front views (147 − 124 = 23 ms/item).

Target recognition at Set Size 1. Figure 8 shows Set Size 1 recognition times for target-present trials as a function of viewpoint. Target-recognition responses were significantly faster for self than stranger, $F(1, 15) = 8.8$, $p < .001$, and for upright than inverted faces, $F(1, 15) = 43.0$, $p < .001$. Although the self versus stranger difference seemed somewhat smaller for upright conditions (stranger − self = 694 ms − 627 ms = 67 ms) than inverted conditions (stranger − self = 926 ms − 820 ms = 106 ms), this Orientation × Identity interaction was not significant, $F(1, 15) = 2.2$, $p = .16$.

There was also a significant effect of viewpoint, $F(2, 30) = 13.6$, $p < .001$. Polynomial contrasts revealed a highly significant linear component, $F(1, 15) = 36.4$, $p < .001$, which reflected increasing reaction times as the face rotated from front to three-quarter to profile view. As can also be seen in Figure 8, this viewpoint effect was significantly greater for inverted than upright conditions, $F(2, 30) = 4.2$, $p < .05$, because of a stronger linear trend in the inverted condition, $F(1, 15) = 8.9$, $p < .01$.

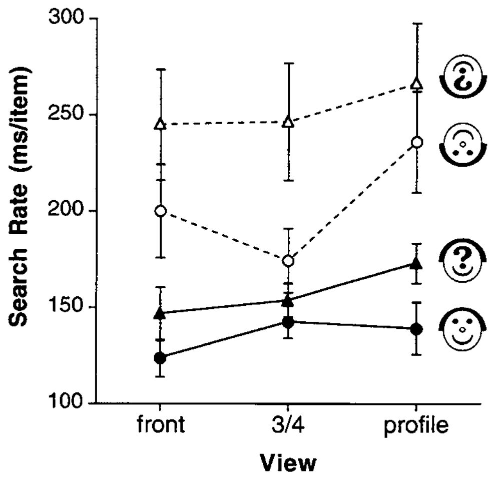

Figure 7. Mean search rates for target-present trials as a function of viewpoint. Solid symbols = upright faces; open symbols = inverted faces; circles = self target; triangles = stranger target. Vertical bars represent ±1 *SEM.* Note that the self versus stranger advantage is evident across all views, even for atypical profile and inverted views.

Surprisingly, observers recognized upright profiles of their own face much more quickly than the stranger's face (stranger − self = 733 ms − 657 ms = 76 ms). This self versus stranger recognition advantage for profiles was as large if not somewhat larger than the advantage seen for canonical upright front views (stranger − self = 670 ms − 625 ms = 45 ms). The lack of a Target Identity × Viewpoint interaction, $F(1, 15) = 2.0$, $p = .16$, suggested that the effects of target identity and viewpoint were independent and additive. That is, the self versus stranger recognition advantage was invariant across the tested views.

Time course analysis. To investigate the time course of face learning, all observed target-present trials for Set Size 1 were sorted in bins of 18 trials for the initial 18 practice trials and following 144 test trials. Mean reaction times and standard errors were then calculated for each trial bin. As can be seen in Figure 9, target recognition in the upright condition was faster for self than stranger across all trial bins, but the time course functions for the two conditions were very different. Recognition times for the stranger dropped sharply between practice trials and the first 18 test trials, but thereafter decreased at a minimal rate of −0.37 ms/trial across test trials, which linear regression revealed as not significant, $F(1, 370) = 1.4$, $p = .24$. In contrast, the time course function for self was almost flat and asymptotic, showing no evidence of decreasing between practice and test trials. Again, the gradual decrease in recognition times across test trials (slope = −0.11 ms/trial) was not significant, $F(1, 373) < 1.0$.

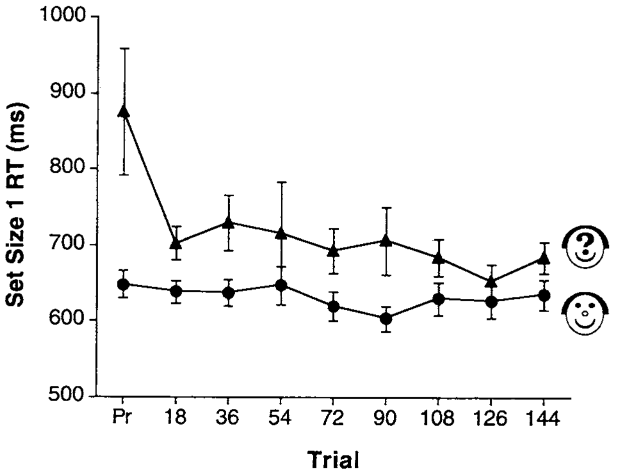

Figure 9. Set Size 1 reaction times for target-present trials as a function of trial position. Observations were averaged in bins of 18 trials. Trial numbers represent the last test trial in each bin, and "Pr" represents the first 18 practice trials. Target was self–upright (circles) or stranger–upright (triangles). Vertical bars depict ±1 *SEM.* Note the rapid, early-learning component for stranger, whereas performance for self appears flat and asymptotic. The consistent self versus stranger difference across test trials further indicates a much slower long-term component in face learning.

Discussion

The pattern of results strongly supported our hypotheses and replicated the general findings of Experiment 1. Distractor-rejection and target-recognition processes, as reflected by search slopes and Set Size 1 reaction times, respectively, were faster for self than stranger and for upright than inverted faces. These results provide some support for the notion that robust representations for one's own face may facilitate a variety of visual processes, including distractor rejection and target recognition (Property 4). If one considers the addition of distractor faces as an increase in attentional load, then faster search rates for self also suggest that robust representations may require less attentional resources (Property 5).

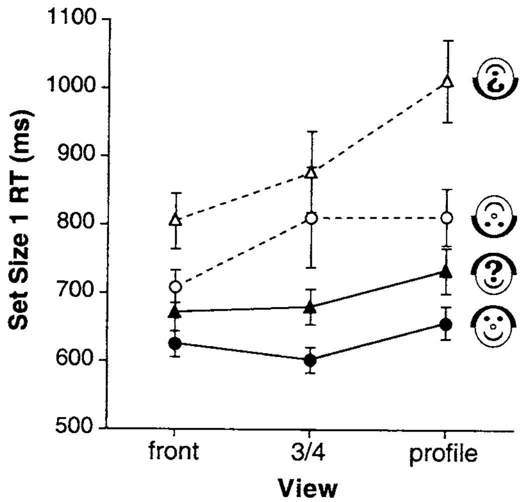

Figure 8. Mean Set Size 1 reaction times for target-present trials as a function of viewpoint. Solid symbols = upright faces; open symbols = inverted faces; circles = self target; triangles = stranger target. Vertical bars represent ±1 *SEM.* Note that target recognition is consistently faster for self than stranger, even for atypical profile and inverted views.

A particularly striking finding was how well observers could visually process their own profile. Reaction times were generally slower for profiles because the capped faces were more visually similar from this vantage point and, therefore, were more difficult to discriminate (as can be seen in Figure 4; see also Bruce et al., 1987). More important, however, the self versus stranger advantages in target recognition and distractor rejection were as large if not larger for profile views as for front views. These results are quite surprising given that observers had equal amounts of visual experience of the stranger's front and profile view but far greater experience of their own front view than their profile.

The fact that self versus stranger processing advantages generalized across all views, including atypical profile and inverted views, strongly suggests that robust representations for one's own face include some view-invariant information (Property 3). At first glance, these results seem to contradict popular notions of face representations as a series of stored two-dimensional views (e.g., Bruce & Young, 1986; Perrett, Mistlin, & Chitty, 1987). However, it is unclear whether observers capitalized on view-invariant featural properties (e.g., skin tone and texture), local shape cues (e.g., size and shape of nose, mouth, etc.), or global 3-D shape to generalize to atypical views. Observers may even have emphasized different visual properties depending on the distinguishing properties of the target face in question. These questions extend beyond the scope of this article but would be interesting to explore in future research.

Another key question was how target-recognition efficiency would change across repeated presentations of a well-known and initially unfamiliar face. Figure 9 suggests the operation of two time components in face learning that resemble the predicted learning functions of Figure 1A and 1B. Observers showed dramatic improvement in recognition times for the stranger's face between practice and test trials but, thereafter, showed only a weak, nonsignificant trend of decreasing recognition times across test trials. In contrast, recognition times for self revealed an almost flat time course function that strongly suggested asymptotic perceptual performance (Property 1). This optimal recognition of self that occurred at the very outset of the task is all the more striking when one considers the novelty of the task—searching for a small gray-scale image of one's own face that may appear in front, three-quarter, or profile view among other similarly capped distractor faces. These data strongly support the notion that robust representations for one's own face may mediate asymptotic visual processing (Property 1) even for highly unfamiliar tasks (Property 4).

The early-learning component for stranger appeared to reflect rapidly acquired familiarity for a novel face. This is consistent with the finding that a single presentation of a novel face is typically sufficient for excellent recognition (e.g., Yin, 1969) or priming (Paller et al., 1992). In contrast, the consistently faster recognition of self than stranger across test trials (627 ms vs. 694 ms, respectively) suggested a second, much slower face-learning component. Time course functions for self and stranger appeared essentially flat and parallel across test trials (self = -0.11 ms/trial, stranger = -0.37 ms/trial). The fact that the self versus stranger recognition advantage persisted across test trials with no evidence of diminishing strongly suggests that extensive visual experience is required for the development of a robust representation (Property 2) and that the time course of this second learning component is very slow indeed.

Experiment 3: Search Asymmetry Between Self as Target Versus Distractor

Experiments 1 and 2 showed that robust representations of target identity can facilitate both target-recognition and distractor-rejection processes. Experiment 3 investigated whether these two processes could be dissociated. This was done by having observers search for self among stranger distractors or stranger among self distractors. Therefore, unlike the previous two experiments, the distractors shown on a given trial consisted of a single person. To encourage discrimination based on facial identity rather than low-level features, we attempted to pair observers with visually similar faces. We also generated heterogeneous targets and distractors by randomly alternating between three-quarter and profile view trials using left and right facing images of self and stranger (see Figure 10). This design led to many more presentations of the stranger's face (1,440 presentations) than in the previous experiments and allowed us to test if greater visual training would abolish this self versus stranger advantage. On the basis of Properties 1 and 2, we suspected that rapid asymptotic processing of an initially unfamiliar face would require much more extensive visual experience than that of a single experimental session.

We were especially interested in whether search would be more rapid with self as target or distractor. Intuitively, one might expect faster search for self targets in a manner analogous to "popout" of a person's own name in the cocktail party effect. A person's own face might be easy to detect as the target but difficult to ignore as the distractor. This view resembles Treisman and Gormican's (1988) theory that rapid search is determined by a distinguishing property of the target that is absent in the distractors. Such target-salience models of visual search predict that search slopes and intercepts should both increase as the target becomes more difficult to detect. Target-present search functions should therefore "fan apart" as illustrated in Figure 11A. Most search-asymmetry studies do yield this signature pattern, supporting the target-salience model (e.g., Kleffner & Ramachandran, 1992; Treisman & Gormican, 1988; Treisman & Souther, 1985). However, researchers have noted that target-salience models may not apply to serial search situations in which the target fails to pop out (e.g., Treisman & Souther, 1985).

By contrast, a strictly serial model of visual search predicts a crossover search-asymmetry pattern for target-present search functions as illustrated in Figure 11B. This prediction is based on two assumptions: (a) Target-recognition times and distractor-rejection rates reflect independent processes, and (b) the "easy" stimulus (as labeled in Figure 11B) is processed more quickly than the "hard" stimulus, irrespective of whether it appears as target or distractor. It follows that target recognition at Set Size 1 should be faster for the easy-target/hard-distractors condition but that distractor-rejection search rates should be faster for the hard-target/easy-distractors condition. As the number of distractors increases, reaction times for the easy-target/hard-distractors should eventually exceed reaction times for the hard-target/easy-distractors, yielding a crossover interaction.[3]

[3] It is important to note that some studies of slow serial search fail to conform to this strictly serial model. For example, Levin (1996) found that searching for a cross-race target face among same-race distractor faces was faster than the opposite condition for all set sizes including Set Size 1, indicating that observers were detecting a salient feature in the cross-race faces rather than rejecting same-race distractor faces more rapidly. Similarly, von Grünau and Anston (1995) found consistently faster detection of straight-gaze targets among averted-gaze distractors across all set sizes, suggesting that observers found the presence of eyes looking at them to be more salient than eyes looking away. Clearly, certain

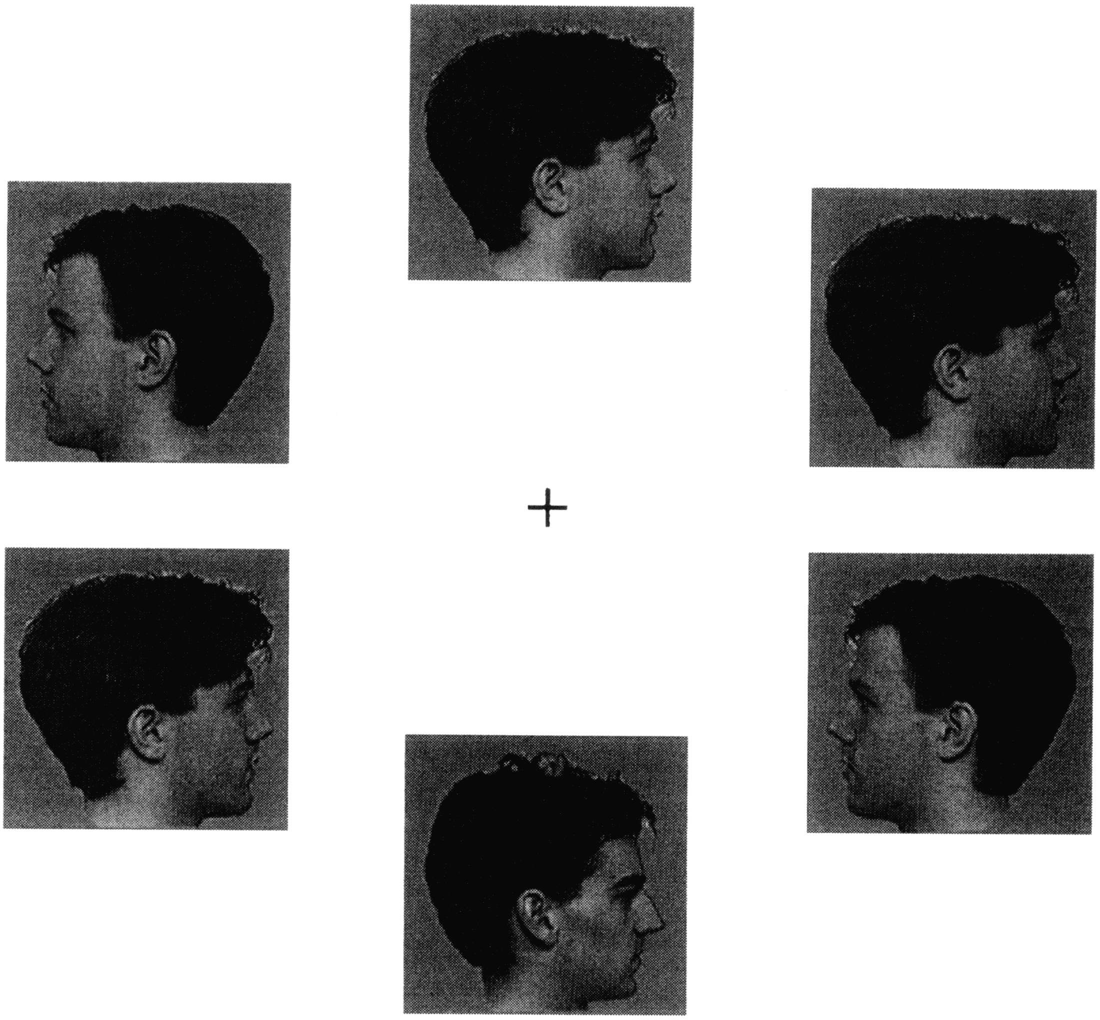

Figure 10. An example of a visual display containing two paired observers from Experiment 3. Observers searched for self target among stranger distractors or stranger target among self distractors on separate blocks. On a given trial, either three-quarter view (not shown) or profile view faces were presented (upright only), and both left and right facing images could appear. A right-facing profile target is shown in the bottommost position.

Therefore, in Experiment 3, we tested whether extensive visual experience for one's own face specifically increases target saliency or more generally facilitates a variety of visual and decisional processes (Property 4), such as target recognition and distractor rejection. Although the task of recognizing self as target may be quite familiar, the task of having to reject arrays of three-quarter and profile views of self distractors to find a target face is certainly novel. On the basis of Property 4, however, we predicted that robust representations for one's own face should facilitate both target-recognition and distractor-rejection processes and lead to a crossover interaction. Such a crossover interaction would not only dissociate these two processes and validate our serial search model, it would further rule out motivational explanations regarding faster search for self (e.g., observers trying harder when searching for their own face as the target). We also predicted that these crossover effects would generalize to atypical profile views on the basis that robust representations contain view-invariant information (Property 3).

facial features that yield slow serial searches may still be more salient as a target.

Method

Methods were identical to those described in Experiment 1, except as indicated below.

Observers. Observers were 10 Caucasian undergraduate men between 19 and 21 years of age. Observers either received payment or credit in an introductory psychology course for participation.

Stimuli. Left and right facing profile and three-quarter view images of each observer were scanned under controlled diffuse room lighting conditions. Face images were matched for size and placement on a standard gray background of size 3.4° × 3.5°. When possible, we attempted to pair observers with visually similar faces to minimize obvious differences in the hair color, hair style, skin tone, and aspect ratio of the faces (i.e., wide vs. narrow face shape). Figure 10 shows profile images of 2 paired observers arranged in a sample visual array.

Procedure. Observers searched for their own face or the face of an unfamiliar paired observer among opposite identity distractors. The experiment consisted of four blocks that followed an A-B-B-A design, with A-B-B-A and B-A-A-B orders counterbalanced across observers. At the beginning of each block, the four images of the target were shown, followed by 32 practice trials and 160 test trials. Observers received a total of 768 practice and test trials.

The display consisted of two, three, four, or six face images presented equally spaced around a fixation point on a uniform white background (see Figure 10). As in Experiment 2, there were 12 possible image locations that formed an 11.8° × 9.2° elliptical array. On a given trial, either three-quarter or profile view faces were shown, although both left and right facing distractors could appear. An intertrial interval of 1,500 ms was used.

Results

Figure 12 shows search functions for target-present trials combined across viewpoints. Reaction times for three-quarter and profile views were combined because there was no main effect of viewpoint, $F(1, 9) = 1.4$, $p = .27$, and no significant interactions involving viewpoint ($F < 1.0$, for all interactions involving viewpoint). Target-absent search slopes were approximately twice the size of target-present slopes, suggesting serial search (see Appendix for details).[4] Overall error rates were low in each Identity × Viewpoint condition (2.5% to 3.9%; details in Appendix) and were not further analyzed.

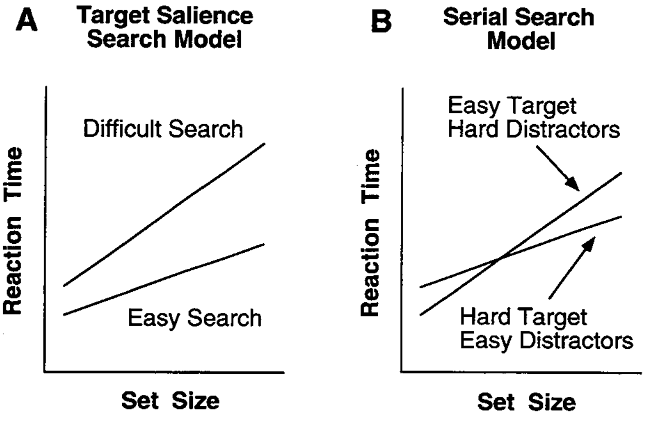

Figure 11. (A) Search asymmetry pattern for target-present trials predicted by a target salience search model. Both search slopes and Set Size 1 intercepts increase as the target becomes more difficult to detect (i.e., less salient). (B) Crossover search asymmetry pattern for target-present trials predicted by a serial search model. Target recognition at Set Size 1 is faster for the easy-target/hard-distractors condition, but distractor-rejection search rates are faster for the hard-target/easy-distractors condition.

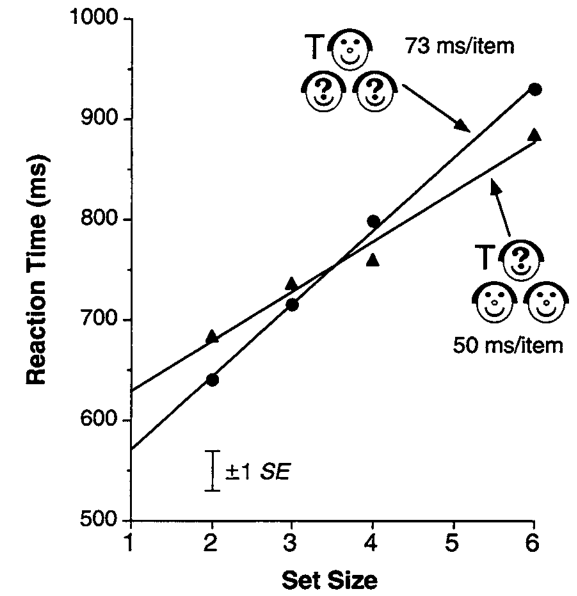

Figure 12. Mean reaction times for target-present trials in Experiment 3. Search conditions were self-target/stranger-distractors (circles) and stranger-target/self-distractors (triangles). Lines of best fit are plotted. Search slopes are listed in ms/item. The vertical bar represents ±1 *SE* of the within-subject difference between conditions. Note the unusual crossover effect. Search rates are faster for self distractors, whereas Set Size 1 intercepts are faster for self targets.

Figure 12 revealed an atypical crossover interaction that supported our serial search model and suggested independent target-recognition and distractor-rejection effects. This was supported by a significant Target–Distractor Identity × Set Size interaction, $F(3, 27) = 5.2$, $p < .01$. On target-present trials, search slopes were faster for self distractors (50 ms/item) than stranger distractors (73 ms/item). In contrast, the intercept difference at Set Size 1 suggested more rapid recognition of self targets (571 ms) than stranger targets (629 ms). Because these results deviated from the predictions of the target-salience model, we investigated the consistency of these search slope and intercept differences across observers by plotting frequency histograms.

Distractor-rejection search rates. A histogram of search slope differences (self-target/stranger-distractors − stranger-target/self-distractors) revealed consistently faster search rates for self distractors (see Figure 13A). This effect, shown by 9 out of 10 observers, was highly significant, $t(9) = 4.17$, $p < .005$. Search rates for self distractors were at least as fast

[4] Search rates were faster for stranger-target/self-distractors than self-target/stranger-distractors. This effect was clearly evident for target-present trials but appeared less evident in the target-absent data. As previously discussed, target-absent slopes provide less reliable measures of search efficiency because of variable criteria used to decide target absence (Chun & Wolfe, 1996). The larger absent:present slope ratio in the stranger-target/self-distractor condition likely reflected the fact that observers adopted a more conservative criterion for deciding the absence of the stranger target.

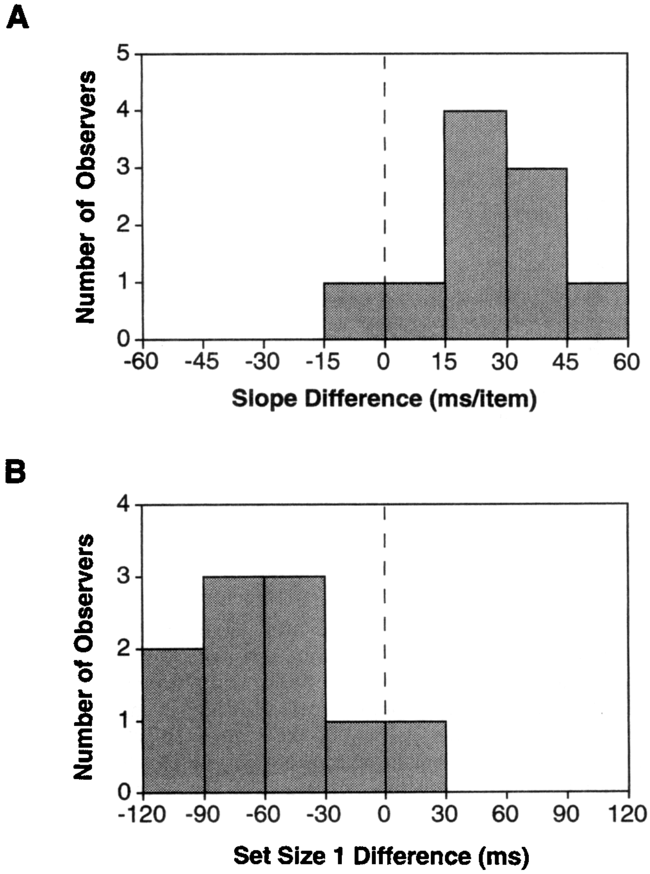

Figure 13. Frequency histograms of search slope differences (Panel A) and Set Size 1 intercept differences (Panel B) for self-target/stranger-distractors minus stranger-target/self-distractors. Note that 9 of 10 observers showed faster search rates for self distractors, whereas 9 of 10 observers showed faster recognition for self targets.

for atypical profile views (46 ms/item) as for canonical three-quarter views (53 ms/item). Moreover, the search rate advantage for self distractors was as large for profile views (stranger distractors − self distractors = 74 − 46 = 28 ms/item) as for three-quarter views (stranger distractors − self distractors = 71 − 53 = 18 ms/item), if not slightly larger.

Target recognition at Set Size 1. A histogram of Set Size 1 intercept differences (self-target/stranger-distractors − stranger-target/self-distractors) revealed consistently faster target recognition for self targets (see Figure 13B). This effect, shown by 9 out of 10 observers, was highly significant, $t(9) = 5.13, p < .001$. Target recognition was equally rapid for atypical profile views (571 ms) and canonical three-quarter views (571 ms) of self but appeared slightly slower for profile views (645 ms) than three-quarter views (614 ms) of stranger.

Time course analysis. We further compared slope and intercept differences between Blocks 1 and 2 versus Blocks 3 and 4 to see if these effects showed any evidence of diminishing over time. The intercept difference showed no evidence of decreasing across blocks, $F(1, 9) < 1.0$. Target recognition at Set Size 1 was faster for self than stranger by 56 ms in Blocks 1 and 2 (self = 586 ms, stranger = 642 ms) and by 62 ms on Blocks 3 and 4 (self = 555 ms, stranger = 617 ms). The apparent trend towards faster overall target-recognition times across blocks was not significant, $F(1, 9) = 1.8, p = .22$.

The search rate advantage for self distractors showed a weak trend toward decreasing across blocks, with a difference of 29 ms on Blocks 1 and 2 (self distractors = 50 ms/item, stranger distractors = 79 ms/item) and a difference of 17 ms/item on Blocks 3 and 4 (self distractors = 49 ms/item, stranger distractors = 66 ms/item). However, this trend was not significant, $F(1, 9) = 1.0, p = .34$, and only 6 out of 10 observers showed an effect in this direction.

Discussion

As predicted, searching for self or stranger among the opposite distractor yielded an atypical crossover interaction that could not be explained by differences in target salience or observer motivation (e.g., observers trying harder when searching for self as target). Instead, the crossover interaction revealed that target recognition and distractor rejection were dissociable processes and thus validated our use of the serial search model. Observers were able to recognize their own face more rapidly as the target yet reject their face more rapidly as the distractor. Given the unusual nature of having to search through and rapidly reject multiple distractor images of one's own face, these results support the notion that robust representations may facilitate a variety of visual and decisional processes including distractor rejection (Property 4).

Observers showed equivalent benefits in self versus stranger processing for profile and three-quarter views, similar to Experiment 2. In fact, target-recognition and distractor-rejection processes for one's own face were as rapid for atypical profile views as for canonical, three-quarter views. Given that observers have very limited visual experience of their own profile, these results and those of Experiment 2 strongly suggest that robust representations contain some view-invariant information (Property 3) that can generalize to atypical depth–plane rotations. Regarding why reaction times were slower for profiles in Experiment 2 but not in Experiment 3, one might speculate that the use of a single identity distractor or the inclusion of hair in the face images may have enhanced target–distractor discriminability and thereby minimized the effects of depth–plane rotation.

Another striking finding was that the self versus stranger processing advantages in target recognition and distractor rejection showed no evidence of diminishing by Blocks 3 and 4. By this time, the stranger's face had already appeared 96 times as the target and 624 times as the distractor (practice trials included). The fact that these effects persisted across hundreds of presentations provides strong support for the notion that robust representations mediate rapid asymptotic processing (Property 1) and require extensive visual experience to develop (Property 2).

Finally, manipulating distractor identity appeared to have a strong influence on search slopes compared with manipulations of target identity in the previous two experiments. Search rates were consistently faster for self distractors (50 ms/item) than stranger distractors (73 ms/item) by a considerable proportion. (Note that this search-rate difference actually underestimates the advantage of having self as distractor because the comparison condition, self-target/stranger-distractor, would likely yield faster search rates than if both target and distractor consisted of unfamiliar faces.) These search rates were too rapid to be explained by serial eye movements alone and instead appeared to reflect the time required for attentional selection of the target. Faster search rates for self distractors therefore suggested that robust representations for one's own face may demand less attentional resources (Property 5).

General Discussion

In all three experiments, processing was consistently faster for self than stranger, irrespective of whether the face appeared in front, three-quarter or profile view, upright or upside-down, with or without hair, as target or distractor. Target recognition for one's own face strongly suggested rapid asymptotic performance (Property 1) and was substantially faster than recognition for stranger by 58 ms to 77 ms across experiments. Such processing advantages even persisted after hundreds of presentations of the stranger's face, indicating that far more visual experience than that of a single experimental session is needed to develop a robust representation (Property 2). To our knowledge, this is the first reported evidence of such a slow, long-term component in face learning. Our results reveal the inadequacy of studying the effects of face representations by simply comparing familiar versus novel faces (e.g., Bruce, 1986; Bruck et al., 1991; Bruyer et al., 1991; A. W. Young et al., 1985), given that face representations, once formed, can widely vary in the extent of their development or robustness.

Another important finding was that observers showed the same self versus stranger processing advantages for atypical inverted and profile views of their own face as for canonical, upright front and three-quarter views. In fact, overall processing rates for front, three-quarter, and profile views of self were generally comparable. These results contrast the highly specific learning effects found in visual discrimination studies that typically fail to generalize across simple changes in orientation, size, or retinal position (e.g., Ahissar & Hochstein, 1996; Fahle & Edelman, 1993; Fiorentini & Berardi, 1980; Ramachandran & Braddick, 1973). The fact that our observers could generalize to highly atypical views of their own face strongly supports the notion that robust representations contain some abstract or view-invariant information (Property 3). Apparently, extensive experience of a face leads to more sophisticated visual learning than simple view-based or image-based learning.

When observers searched for self or stranger among unfamiliar distractor faces, both search slopes and intercepts were faster for self than stranger. This suggests that searching for a robustly represented target not only facilitates the familiar task of target recognition but also facilitates the unfamiliar task of serial distractor rejection (Property 4). When self and stranger alternated as target and distractor in Experiment 3, recognition at Set Size 1 was faster for self targets; however, search rates were faster for self distractors. This indicates that the unusual task of distractor rejection is most greatly facilitated when the distractors rather than the targets are robustly represented. To the extent that search slopes are indicative of the amount of visual attention required to process each additional distractor in a visual array (e.g., Treisman & Gormican, 1988), these results further suggest that for a given array of faces, replacing each unfamiliar face with a robustly represented face serves to reduce the overall amount of attention required to process that array. In other words, robust representations demand less attentional resources (Property 5).

Theoretical Considerations

The above discussion provides several lines of evidence that support the notion that people can develop robust representations for highly overlearned faces. We have discussed in detail how the various defining properties of robust representations lead to different predicted patterns of visual search performance when compared to less developed representations. However, certain questions remain to be addressed.

One question is whether observers would show similar evidence of robust representations for other highly overlearned faces, such as those of close friends or family members. Given that people typically have more visual experience of such faces than their own and that robust representations most likely arise from extensive visual experience, we have every reason to expect similar results for other highly overlearned faces. Pragmatically speaking, robust representations would serve a more useful role in rapidly identifying the faces of important others than one's own face.

Another question is whether specific types of visual experience are necessary for the development of a robust representation. For example, is it necessary to sample a much larger number of views or image variations than were shown here, and do these images need to be linked dynamically? Recent studies have shown that learning a front-view face across variations in lighting, rather than front-lighting only, facilitates generalization to novel, front-lit views (Tong & Nakayama, 1997). Similarly, multiple-face views are better learned if they are linked by coherent motion rather than incoherent or no motion (Hill, Schyns, & Akamatsu, 1997). Providing additional 3-D cues in the form of additional lighting (i.e., shape-from-shading) information or structure-from-motion information may therefore lead to a more view-invariant face representation. Although recognition benefits from additional 3-D information usually disappear after a few presentations of an initially novel face view (Tong & Nakayama, 1997), unlike the much more enduring effects found here, one might suspect that extensive visual experience of a face across a variety of changes leads to the development of a robust representation that captures the underlying view-invariant properties of a face (Property 3). Indirect support of this notion comes from visual learning studies. When observers are extensively trained with a specific set of stimuli in psychophysical discrimination tasks, visual learning is dramatic but highly specific and typically fails to generalize to simple stimulus changes in size, orientation, or position (e.g., Ahissar & Hochstein, 1996; Fahle & Edelman, 1993; Fiorentini & Berardi, 1980; Ramachandran & Braddick, 1973). We therefore suspect that deriving some form of object constancy across stimulus change may be crucial for the development of a robust representation.

One might also ask whether robust representations are specific to faces or whether they might also apply to other visual categories. We believe that robust representations arise from extensive experience at discriminating a particular exemplar from other similar exemplars across a variety of viewing conditions and contexts. Given the expertise people have for recognizing faces in general and the extent to which certain faces are highly overlearned, we considered faces to be the most promising category for investigating robust representations. However, similar effects might possibly occur for other overlearned visual items such as letters, words, or visual objects for which one has developed expertise (e.g., Diamond & Carey's, 1986, dog expert study). Our results suggest that visual search can provide a powerful method of investigating evidence for robust representations in other such domains. The data patterns described here, such as the asymptotic recognition performance (Figure 9), crossover search asymmetry for target-present trials (Figure 12), and generalization of processing advantages across viewpoint change (e.g., Figures 7 and 8), may serve as diagnostic measures of robust representations for exemplars in other visual categories.

Robust Representations and Efficient Coding

The final and perhaps the most difficult question is the following: How does a robust representation differ from other representations, not in its effects on visual processing, which have already been thoroughly discussed, but in the nature of its visual code? Given that little is currently known about the visual code for faces (or for any other complex stimulus for that matter), any suggestions on our part must remain speculative. On the basis of theoretical ideas on sensory coding (Barlow, 1961) and converging evidence from neurophysiology (e.g., Yamane et al., 1988) and computer vision (e.g., Penev & Atick, 1996; Sirovich & Kirby, 1987), we would like to tentatively suggest that robust face representations arise from the development of an efficient visual code.

According to Barlow (1961), the goal of the visual system is to develop an efficient method of coding the enormous amount of information that falls on our retinae. For example, the oriented Gabor-like filtering properties of V1 neurons may reflect an optimized code for representing the higher order statistical structure of natural images, which contain a great deal of redundant information (i.e., correlations between the luminance intensities of neighboring pixels; Field, 1987). Similarly, binocular V1 neurons can efficiently code a stimulus projected onto the two retinae as well as the disparity of that stimulus (Barlow, Blakemore, & Pettigrew, 1967). It is important to note that in absence of proper visual experience, neurons tuned to orientation (Blakemore & Cooper, 1970) or binocular disparity (Hubel & Wiesel, 1965) may fail to develop. These results suggest that efficient codes (a) reduce the redundancy of incoming visual input and thereby reduce the number of active neurons required to code a stimulus, (b) extract or explicitly code important information (e.g., oriented edges, relative depth), and (c) are shaped by the statistics of an organism's visual experience.

Efficient codes may also be used in high-level vision to describe the statistical structure of complex objects and thereby mediate visual recognition (Barlow, 1961). In computer vision, compact codes for representing face images are most often derived by using principal-components analysis. Although algorithms differ in whether they analyze local features (Penev & Atick, 1996), the global 2-D face image (e.g., Sirovich & Kirby, 1987), or the derived 3-D shape of a face by analyzing shape-from-shading information (Atick et al., 1996), the end result is the same. Rather than requiring, say, 100,000 intensities to code each pixel of a face image, only 50–100 principal components are needed to accurately code such an image. Some of these principal components may even resemble plausible face recognition units. For example, Penev and Atick's local feature analysis model yields units tuned to particular facial features that bear some resemblance to face cells in the temporal cortex of macaques (Yamane, Kaji, & Kawano, 1988). These neurons show differential responses to face identities that are based on their tunings for particular facial features or configurations. Although each face cell appears to code a restricted facial dimension, a sparse population of active cells could theoretically provide an efficient code for identifying individual faces (M. P. Young & Yamane, 1992). Supporting this face-encoding hypothesis, Rolls, Baylis, Hasselmo, and Nalwa (1989) found that face cells were most likely to show changes in response selectivity for an initially novel face within the first 1 or 2 presentations, suggesting that these neurons are involved in encoding novel faces.

Although face cells may rapidly encode novel faces, our data on robust representations suggest that much more extensive visual experience is needed to develop an optimal, compact visual code. The gradual development of a robust face representation can perhaps be conceptualized as a reduction in the number of units or principal components needed to accurately describe a face. Over time, the remaining units become more decorrelated (and therefore less redundant), yet more selectively tuned to the face in question. Such compact face codes may resemble higher order recognition units in other domains such as word recognition. For example, the trigram C–A–T can be readily encoded as a single familiar unit of information, whereas the unfamiliar trigram C–G–T cannot. A robustly represented face might also be considered a more coherent chunk or unit of information.

The concept of efficient coding for robust representations may help explain several findings, such as why robust representations demand less attentional resources (Property 5) and lead to faster search rates. If one assumes that the visual system has a maximal processing rate, then the system should read the output of a compact face code more quickly. As the number of faces is multiplied in visual search, processing times for self and stranger are likewise multiplied. This would lead to the predicted crossover search asymmetry found in Experiment 3. Efficient coding might also explain the asymptotic recognition performance for self (Property 1). According to information theory, there is always a maximum amount of redundancy reduction possible, beyond which further abbreviation of a code will lead to a loss of information (Attneave, 1959). Asymptotic recognition for robustly represented faces may therefore reflect an optimal code that cannot be further compacted. Such optimized codes would lead to general rather than specific facilitation effects, speeding any process that accessed this code irrespective of task or context (Property 4). Coding efficiency might also explain the view-invariant processing advantages found for self versus stranger (Property 3). Extended visual experience across changes in viewpoint, lighting, and expression may ultimately lead to a code that efficiently captures the view-invariant properties of a face (cf. Hill et al., 1997; Tong & Nakayama, 1997) rather than the image-based properties of a specific viewing condition.

Finally, the concept of coding efficiency may help explain a more general finding regarding why monkeys (e.g., Perrett et al., 1982, 1987) and humans (Kanwisher et al., 1997, 1998; Tong, Nakayama, Moscovitch, Weinrib, & Kanwisher, in press) appear to have specialized neural regions for face processing. Although the notion of an all-purpose recognition system might seem parsimonious, the use of a common code for describing subtle variations between faces and coarse variations between dissimilar object categories would likely prove very inefficient, given that faces and objects share little overlap in terms of their higher order image structure. Rather than develop a single, inefficient coding scheme, the visual system may opt to develop a separate code optimized for the highly specific yet important task of face recognition.

Concluding Remarks

Face recognition is a complex task that people perform with remarkable ease. Typically, a single presentation is sufficient for accurate face recognition. However, the experiments discussed here demonstrate that face learning continues over a much more extended time period. Visual learning does eventually reach an endpoint, and we propose that this endpoint reflects the development of a robust representation. We found several lines of evidence that robust representations for a highly overlearned face such as one's own may (a) mediate rapid asymptotic processing, (b) require extensive visual experience to develop, (c) contain some view-invariant information for generalization to atypical inverted and profile views, (d) facilitate a variety of visual processes such as target recognition and distractor rejection, and (e) demand less attentional resources. One possible mechanism that may mediate the development of a robust representation is the development of an efficient visual code. Our findings strongly support the notion that extensive visual experience can shape the way people process complex visual stimuli in their environment.

References

Ahissar, M., & Hochstein, S. (1996). Learning pop-out detection: Specificities to stimulus characteristics. *Vision Research, 36,* 3487–3500.

Atick, J., Griffin, P. A., & Redlich, A. N. (1996). Statistical approach to shape from shading: Reconstruction of three-dimensional face surfaces from single two-dimensional images. *Neural Computation, 8,* 1321–1340.

Attneave, F. (1959). *Applications of information theory to psychology.* New York: Holt, Rinehart & Winston.

Barlow, H. B. (1961). The coding of sensory messages. In W. H. Thorpe & O. L. Zangwill (Eds.), *Current problems in animal behavior* (pp. 331–360). Cambridge, England: Cambridge University Press.

Barlow, H. B., Blakemore, C., & Pettigrew, J. D. (1967). The neural mechanism of binocular depth discrimination. *Journal of Physiology, 193,* 327–342.

Blakemore, C., & Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature, 228,* 477–478.

Bruce, V. (1986). Influences of familiarity on the processing of faces. *Perception, 15,* 387–397.

Bruce, V., Valentine, T., & Baddeley, A. (1987). The basis of the ¾ view advantage in face recognition. *Applied Cognitive Psychology, 1,* 109–120.

Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology, 77,* 305–327.

Bruck, M., Cavanagh, P., & Ceci, S. J. (1991). Fortysomething: Recognizing faces at one's 25th reunion. *Memory and Cognition, 19,* 221–228.

Bruyer, R., Lafalize, A., & Distefano, M. (1991). Age decisions on familiar and unfamiliar faces. *Behavioural Processes, 24,* 21–35.

Chun, M. M., & Wolfe, J. M. (1996). Just say no: How are visual searches terminated when there is no target present? *Cognitive Psychology, 30,* 39–78.

Diamond, R., & Carey, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General, 115,* 107–117.

Enns, J. T., & Shore, D. I. (1997). Separate influences of orientation and lighting in the inverted-face effect. *Perception and Psychophysics, 59,* 23–31.

Fahle, M., & Edelman, S. (1993). Long term learning in vernier acuity: Effects of stimulus orientation, range and of feedback. *Vision Research, 33,* 397–412.

Farah, M. J., Wilson, K. D., Drain, H. M., & Tanaka, J. R. (1995). The inverted face inversion effect in prosopagnosia: Evidence for mandatory, face-specific perceptual mechanisms. *Vision Research, 35,* 2089–2093.

Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *Journal of the Optical Society of America, 4,* 2379–2394.

Fiorentini, A., & Berardi, N. (1980). Learning specific for orientation and spatial frequency. *Nature, 287,* 43–44.

He, Z. J., & Nakayama, K. (1992). Surfaces versus features in visual search. *Nature, 359,* 231–233.

Hill, H., Schyns, P. G., & Akamatsu, S. (1997). Information and viewpoint dependence in face recognition. *Cognition, 62,* 201–222.

Hubel, D. H., & Wiesel, T. N. (1965). Binocular interaction in

striate cortex of kittens reared with artificial squint. *Journal of Neurophysiology, 28,* 1041–1059.
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.
Kanwisher, N., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. *Cognition, 68,* B1–B11.
Kleffner, D. A., & Ramachandran, V. S. (1992). On the perception of shape from shading. *Perception & Psychophysics, 52,* 18–36.
Kuehn, S. M., & Jolicoeur, P. (1994). Impact of quality of image, orientation, and similarity of the stimuli on visual search for faces. *Perception, 23,* 95–122.
Levin, D. T. (1996). Classifying faces by race: The structure of face categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 22,* 1364–1382.
Moscovitch, M., Winocur, G., & Behrmann, M. (1997). What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. *Journal of Cognitive Neuroscience, 9,* 555–604.
Moses, Y., Ullman, S., & Edelman, S. (1996). Generalization to novel images in upright and inverted faces. *Perception, 25,* 443–461.
Norman, D. A., & Bobrow, D. G. (1975). On data-limited and resource-limited processes. *Cognitive Psychology, 7,* 44–64.
Nothdurft, H.-C. (1993). Faces and facial expressions do not pop out. *Perception, 22,* 1287–1298.
Paller, K. A., Mayes, A. R., Thompson, K. M., Young, A. W., Roberts, J., & Meudell, P. R. (1992). Priming of face matching in amnesia. *Brain and Cognition, 18,* 46–59.
Penev, P. S., & Atick, J. J. (1996). Local feature analysis: a general statistical theory for object representation. *Network: Computation in Neural Systems, 7,* 477–500.
Perrett, D. I., Mistlin, A. J., & Chitty, A. J. (1987). Visual neurones responsive to faces. *Trends in Neurosciences, 10,* 358–364.
Perrett, D. I., Rolls, E. T., & Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. *Experimental Brain Research, 47,* 329–342.
Purcell, D. G., Stewart, A. L., & Skov, R. B. (1996). It takes a confounded face to pop out of a crowd. *Perception, 25,* 1091–1108.
Ramachandran, V. S., & Braddick, O. (1973). Orientation specific learning in stereopsis. *Perception, 2,* 371–376.
Rolls, E. T., Baylis, G. C., Hasselmo, M. E., & Nalwa, V. (1989). The effect of learning on the face selective responses of neurons in the cortex in the superior temporal sulcus of the monkey. *Experimental Brain Research, 76,* 153–164.
Rovamo, J., & Virsu, V. (1979). An estimation and application of the human cortical magnification factor. *Experimental Brain Research, 37,* 495–510.
Sirovich, L., & Kirby, M. (1987). Low-dimensional procedure for the characterization of human faces. *Journal of the Optical Society of America, 4,* 519–524.
Suzuki, S., & Cavanagh, P. (1995). Facial organization blocks access to low-level features: An object inferiority effect. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 901–913.
Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology, 46A,* 225–245.
Tong, F., & Nakayama, K. (1997). Multiple lightings facilitate face recognition for novel views. *Investigative Ophthalmology & Visual Science (ARVO Annual Supplement), 38,* 4658.
Tong, F., Nakayama, K., Moscovitch, M., Weinrib, O., & Kanwisher, N. (in press). Response properties of the human fusiform face area. *Cognitive Neuropsychology.*
Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review, 95,* 15–48.
Treisman, A., & Souther, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General, 114,* 285–310.
Valentine, T. (1988). Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology, 79,* 471–491.
von Grünau, M., & Anston, C. (1995). The detection of gaze direction: A stare-in-the-crowd effect. *Perception, 24,* 1297–1313.
Wang, Q., Cavanagh, P., & Green, M. (1994). Familiarity and pop-out in visual search. *Perception and Psychophysics, 56,* 495–500.
Yamane, S., Kaji, S., & Kawano, K. (1988). What facial features activate face neurons in the inferotemporal cortex of the monkey. *Experimental Brain Research, 73,* 209–214.
Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology, 81,* 141–145.
Young, A. W., Hay, D. C., McWeeny, K. H., Flude, B. M., & Ellis, A. W. (1985). Matching familiar and unfamiliar faces on internal and external features. *Perception, 14,* 737–746.
Young, A. W., Hellawell, D., & Hay, D. C. (1987). Configural information in face perception. *Perception, 16,* 747–759.
Young, M. P., & Yamane, S. (1992). Sparse population coding of faces in the inferotemporal cortex. *Science, 256,* 1327–1330.

Appendix

Summary of Search Experiments: Search Slopes and Error Rates

Display orientation	Target identity	Target viewpoint	Search slopes: Present	Search slopes: Absent	% errors by set size			
Experiment 1					2	4	6	
Upright	Self	Front	71	180	1.3	2.4	3.8	
	Stranger	Front	86	239	2.0	2.9	4.0	
Inverted	Self	Front	101	302	1.8	3.6	3.3	
	Stranger	Front	117	335	2.2	2.6	3.8	
Experiment 2					1	3	6	
Upright	Self	Front	124	231	0.8	2.0	2.3	
		3/4	142	252	2.3	2.7	2.7	
		Profile	139	266	3.5	2.7	3.5	
	Stranger	Front	147	292	2.0	1.6	2.0	
		3/4	154	305	2.0	3.1	4.3	
		Profile	173	322	3.1	2.3	2.7	
Inverted	Self	Front	200	345	1.6	5.5	4.7	
		3/4	174	398	4.7	8.6	10.6	
		Profile	236	410	8.2	6.6	13.3	
	Stranger	Front	245	362	8.6	6.6	14.5	
		3/4	247	441	7.4	16.0	18.4	
		Profile	267	530	11.7	16.8	22.3	
Experiment 3					2	3	4	6
Upright	Self	3/4	71	148	2.3	3.0	2.0	2.5
		Profile	74	132	4.8	3.0	2.3	5.5
	Stranger	3/4	53	128	3.5	3.5	3.3	3.0
		Profile	46	121	3.0	3.0	4.3	4.8

Received July 29, 1997
Revision received May 13, 1998
Accepted July 1, 1998 ■