Visual search for a socially defined feature: What causes the search asymmetry favoring cross-race faces?

DANIEL T. LEVIN and BONNIE L. ANGELONE
Kent State University, Kent, Ohio

Levin (1996, 2000) reported that white subjects search for black targets more quickly than they search for white targets, suggesting that black faces are perceived as having a feature that is lacking in white faces. Here we test one of the implications of this asymmetry by having subjects search for same-race (SR) and cross-race (CR) faces that are distorted to look less like each other (producing caricatures that enhance race-specifying features), or are distorted to look more like each other (a prototypical distortion expected to reduce the salience of race-specifying features). Experiments 1 and 2 show that caricaturing the feature-positive CR distractors speeds search for the SR face and that prototypical distortion slows this search. The same distortions in SR faces did not affect the search slopes. However, these distortions also eliminated the overall advantage for CR faces. Experiment 3 shows that trial-to-trial variation in the specific distractors in each display can eliminate the asymmetry and suggests that this asymmetry depends on the subjects' ability to set a consistent *a priori* perceptual criterion when searching for a CR target, while the distortion effects emphasize the importance of distractor-rejection processes in determining the form of a serial search asymmetry.

As many of the papers in this issue no doubt make clear, search asymmetries can serve as a useful tool for understanding the way that our visual system defines features relative to each other. First described extensively in Treisman and Souther (1985) and Treisman and Gormican (1988), search asymmetries occur when searching for some target stimulus A among distracting Bs is faster than the converse task in which B is the target hidden among As. They argued that this occurs when one target has a feature that the other lacks and that it is fundamentally easier to locate a feature-positive target among feature-negative distractors than to do the reverse. For example, Treisman and Gormican had subjects complete one block of trials in which a tilted line was the target hidden among straight lines and found that search was faster than in another block in which the target was a straight line hidden among tilted lines. Accordingly, the search result supports the conclusion that the visual system defines "tilt" as a feature that is added to the norm of straightness. Here, we suggest that search asymmetries

can also be used to understand more complex, cognitively defined features, and we test one of the implications of such an asymmetry. Specifically, we explore a search asymmetry based on the race of faces in which subjects locate a cross-race (CR) face among same-race (SR) faces more quickly than they do the reverse. If this asymmetry occurs because subjects code CR faces as having race-specifying features that SR faces simply lack, we predict that increasing the physical salience of race-specifying information in the CR face should affect search slopes more strongly than should similarly increasing the salience of SR race-specifying information. In the situation where white subjects are searching for black (CR) faces and white (SR) faces, distorting the CR faces to make them look more or less "black" should affect slopes more strongly than should similar distortions in the SR faces.

A Search Asymmetry for the Race of Faces

Although Treisman and Gormican (1988) used search asymmetries as a diagnostic for simple primitive feature relationships, it is possible to extend the same reasoning to more complex features that are not primitives and may be subject to cognitive definition. For example, von Grünau and Dubé (1994) found that drawings of boxes shown from below are more easily located among boxes viewed from above than the reverse and argued that this occurs because the unusual viewpoint inherent to the boxes viewed from below served as a feature that stood out from the normative viewpoint. This effect is not limited to abstract drawings and can occur in more natural stimuli as

We thank Julia Noland and Frank Tong for reading and commenting on previous versions of this manuscript. Portions of this work were presented at the 1996 Association for Research in Vision and Ophthalmology conference in Ft. Lauderdale, FL, and were reported as part of D.L.'s doctoral dissertation submitted to Cornell University. Preparation of this report was supported by NIMH Grant 1-R03-MH60137-01. Correspondence should be addressed to D. T. Levin, Department of Psychology, Kent State University, Kent, OH 44242-0001 (e-mail: dlevin@kent.edu).

well. Von Grünau and Anston (1995) found that eyes looking directly at the subjects were more easily located among averted eyes than the reverse. In this case, the feature is presumably an engaging (or challenging) social signal among more neutral faces. Here, we attempt to understand the mechanism underlying another search asymmetry observed by Levin (1996, 2000) in which subjects locate a CR face among SR faces more quickly than they do the reverse (in both cases this refers to white subjects searching for cross-race black faces). In this case, the feature is membership in a contrasting race, whereas faces belonging to one's own race are feature-negative norms. Not only are these features complex, but they are also sensitive to contextual influences, reinforcing other findings that visual search is responsive to learned contrasts (Wolfe, 1998).

It is important to note that stimulus factors inherent to black and white faces do not cause this asymmetry. Rather, it is the relationship between the subjects' social cognitions and the race of the face that causes the asymmetry. Levin (1996) tested native African subjects (living in America for an average of 3 years) and found no asymmetry, and Levin (2000) found that white subjects who recognized black faces accurately (and therefore did not show the typical difficulty in recognizing CR faces) showed a significantly reduced asymmetry. Accordingly, this feature is defined not by simple perceptual characteristics of black and white faces, but by the cognitive context provided by the processing of these faces as social cues. In particular, a wide variety of findings in the social cognition literature converge on the hypothesis that members of social groups distinctly apart from one's own are processed more at the category level than as individuals (e.g., Sedikides & Ostrom, 1993). Therefore, one might expect that category-specifying visual information such as the features specifying membership in another race would be salient, whereas features specifying membership in one's own race would be less salient. The reduced asymmetry in the white subjects who recognized black faces accurately further suggests that it is possible to overcome the effects of this context by focusing on members of the other group as individuals (Li, Dunning, & Malpass, 1994).

In the present experiments, we tested one of the predictions that appears to follow from the basic asymmetry. If CR faces are coded as feature-positive deviations from the feature-negative CR norm, manipulating the salience of the features that specify race in the CR face should affect search slopes more strongly than should equivalent manipulations in the SR face. These manipulations were achieved by first comparing an average black face with an average white face. As is shown in Figure 1, the average for each race was a digital blend of 16 faces from that race (equated for mean distinctiveness and attractiveness). The salience of race-specifying information was increased by caricaturing the faces. This involved distorting their shape in order to enhance the difference between them. Thus, the black caricature is a distortion that

enhances the differences between it and the white face, and the white caricature represents the same transformation of the white face. It is also possible to do the reverse and distort the faces to look more similar to each other. These will be referred to as prototypical distortions.

The question is, how will these distortions affect the search process? First, if the CR face is a feature-positive stimulus and the SR face is a feature-negative norm, we might expect that distortions would have effects only in the feature-positive face because these are feature-relevant, whereas the equivalent distortions in feature-negative SR faces will have no effect. This prediction gains plausibility from Levin (2000), who observed that white subjects showed increased perceptual sensitivity to variants of black faces than to white faces. This effect is based on an ABX discrimination task in which subjects discriminated among face variants that lay on continua between black faces and white faces. Thus, the faces varied in terms of race—black faces were distorted to look progressively more like white faces (while still being classified consistently as being black), and white faces were distorted to look more like black faces (while still being classified consistently as being white). White subjects were more accurate at discriminating among black variants, but the effect was limited to white subjects who showed the CR recognition deficit (i.e., who recognized black faces poorly in a separate task). Thus, white subjects who are poor at recognizing black faces are good at discriminating among black faces so long as the stimuli vary in terms of race. This fits well with the hypothesis that white people code black faces as feature-positive stimuli and are therefore sensitive to variations in the salience of that feature, although they are less sensitive to equivalent variations in feature-negative white faces.

A second issue is the direction of the distortion effect. For example, it is possible that caricaturing the CR target will increase the size of the asymmetry. This is particularly likely if the asymmetry rests heavily on the increased salience of the CR target itself, relative to the SR target. However, as Tong and Nakayama (1999) point out, serial search asymmetries can also be caused by differences in the difficulty of distractor rejection as the subject scans the display. In the present case, the white target may be difficult to find because it is difficult to reject black distractors. Accordingly, caricaturing might actually reduce the size of the asymmetry if it makes black distractors easier to reject, thereby speeding the white-target search.

In summary, we are interested in testing the implications of the feature-coding asymmetry in CR and SR faces. In Experiments 1 and 2, the salience of race-specifying information is manipulated in both CR and SR faces. The basic prediction is that this manipulation will affect search for CR faces and not SR faces. The specific form of this effect is also important and will be used to help understand the processes underlying the search asymmetry. If distortions that increase the salience of the feature increase the size of the search asymmetry, a model emphasizing differential target salience would be most appropriate for explaining the

A. Stimuli

B. Sample Trials

Figure 1. Caricatured and undistorted black and white average faces.

asymmetry. If, on the other hand, the manipulation decreases the asymmetry, in particular by speeding the SR-target search, then differential ease of distractor rejection may also play a role in causing the asymmetry. Experiment 3 replicates Experiment 1 and further explores the unexpected finding that mixing trials with normal faces among trials with distorted faces appears to reduce or eliminate the basic search asymmetry that favors black faces.

EXPERIMENT 1

In Experiment 1, each face was distorted to look less like the other via caricaturing. Thus, the white average was caricatured by distorting its shape to look less like the black average, and vice versa for the black face. The subjects in all experiments in this report completed two blocks of trials—one in which they searched for a black face among white faces and one in which they searched for a white face among black faces. In Experiment 1, caricaturing was a between-subjects manipulation so that one group of subjects searched for normal or caricatured

black targets among normal white distractors and normal white targets among normal or caricatured black distractors. The other group searched for normal or caricatured white targets among normal black distractors and normal black targets among normal or caricatured white distractors. The trials with normal versus caricatured targets or distractors were randomly intermixed within the white-target or black-target blocks. On any given trial, the distractors were a homogeneous set of caricatured or normal stimuli (see Figure 1).

Method

Subjects. A total of 18 Cornell University undergraduates completed Experiment 1. Of these, 4 were eliminated for failure to meet the 5% error criterion. Of the remaining 14 subjects, 7 viewed displays with normal and distorted black faces, and 7 viewed displays with normal and distorted white faces. The error rate for those remaining was 2.8%. Of these, 11 were female, 2 were Asian, and 12 were white.

Stimuli. The subjects searched for normal and caricatured faces of each race. The normal faces were the black average and white average faces used in Levin (2000). Each of these faces was a digital blend of 16 faces. The two averages were matched for average bright-

ness and contrast. In addition, they had the same external features (i.e., hair, ears, and jawline), which were created by averaging all 32 faces. Thus, the faces differed only in terms of the shape of the internal features. The caricatures were created by distorting the shape of the faces so that the difference between them was increased. The caricaturing process had two steps. First, facial features were mapped by defining a set of approximately 250 data points that corresponded between the faces. Given these maps, the faces can be distorted by extrapolating the line defined by corresponding map points (Brennan, 1985). In the present case, this produced images of each race that accentuated the differences between them. The white face was distorted by increasing the differences between it and the black face, and vice versa for the black face. The degree of distortion can be expressed in terms of the difference between each undistorted face map, which is simply defined as 100%. The stimuli for Experiment 1 (Figure 1) were 50% caricatures, which means that they were extrapolations away from the opposite averages that are equal to half of the difference between the average faces. Distortions were created using the Morph program and included only shape transformations. Texture maps were not manipulated.

Individual faces were presented with a resolution of 64 (h) $\times$ 90 (v) pixels (approximately 23 $\times$ 32 mm screen size and 2.2° $\times$ 3.1° of visual angle) and were viewed from a distance of approximately 60 cm.

Procedure. Groups of 4–8 subjects were tested on individual computers in the same room.

In the search task, the subjects were instructed to determine whether the target face was present as quickly as possible without sacrificing accuracy. Then, they saw an instruction screen on their computer indicating which face was the target and which was the distractor for the upcoming block of trials. On each trial they saw a stimulus display containing two, four, or eight faces (inclusive of the target) and pressed the “1” key on the computer keyboard if the target was present and the “2” key if it was absent. After giving each response, the subjects received feedback in the form of a “+” for correct responses or a “–” for incorrect responses. For the eight-item displays, faces were arranged in a symmetric, roughly diamond-shaped pattern on the screen. The total size of these displays was approximately 125 (h) $\times$ 110 (v) mm (11.8° $\times$ 10.4° of visual angle). Two- and four-item displays were created by placing faces in symmetrical subsets of the locations used in the eight-item displays.

The subjects completed two blocks of trials—one with the white face as the target and one with the black face as the target. The race of the caricatured face was manipulated between subjects so that each subject’s stimulus set contained undistorted faces of both races in addition to caricatures of one race. Half of the trials included all undistorted stimuli, and half included caricatures. On a given caricature trial, caricatures were either targets or distractors, but never both. For example, if a subject was in the white caricature group, he/she searched four types of displays: white undistorted targets and black undistorted distractors, white caricatured targets and black undistorted distractors, black undistorted targets and white undistorted distractors, and black undistorted targets and white caricatured distractors.

There were 16 trials in each task (black target/white target) $\times$ distortion (caricature/normal) $\times$ display size $\times$ target presence cell of the study making a total of 384 experimental trials for each subject. In addition, the subjects completed 24 practice trials before beginning each search task (black target/white target).

Results

Caricaturing black CR faces affected search slopes primarily because caricaturing black distractors speeded search for white SR targets, whereas the same manipulation in white faces did not affect slopes.

Search slopes were computed and were entered into four separate within-subjects two-factor stimulus (normal/caricature) $\times$ target presence (present/absent) analyses of variance (ANOVAs)—one for each task (black search/white search) and each stimulus manipulation (white caricature/black caricature).

Caricature effects. Caricaturing the white face failed to affect search slopes both when it was a target and when it was a distractor (all F s < 1), whereas caricaturing the black face did affect slopes. Caricaturing speeded search nonsignificantly when the black face was the target [$F(1,6) = 3.507$, $MS_e = 199$, $p = .110$; note that this comparison includes only target-present data because there is no caricatured stimulus in this condition for target-absent trials], but it did significantly affect search slopes when the black face was the distractor [$F(1,6) = 6.530$, $MS_e = 882$, $p = .0432$]. As can be seen in Figure 2, caricaturing black distractors reduced search slopes for the white face (target-present normal, 74 msec/item; caricatured, 49 msec/item; target-absent normal, 149 msec/item; caricatured, 116 msec/item), whereas it had a smaller effect on black-target-present search trials (target-present normal, 60 msec/item; caricatured, 46 msec/item), although the interaction between the stimulus caricatured and search slopes was nonsignificant [$F(1,6) = 1.03$, $MS_e = 179$, $p = .3493$]. Distractor-caricaturing effects were marginally nonsignificant considering target-present and target-absent trials alone [target-present, $F(1,6) = 4.279$, $MS_e = 487$, $p = .084$; target-absent, $F(1,6) = 4.696$, $MS_e = 810$, $p = .073$].

Search asymmetry. Data from no-distortion trials were tested for a search asymmetry favoring detection of black faces by entering slopes into a three-factor group (black distortion/white distortion) $\times$ task (black search/white search) $\times$ target presence (target-present/target-absent) mixed-factors ANOVA with group as the between-subjects factor. The task effect revealed a nonsignificant trend for the black search to be faster than the white search [$F(1,12) = 2.799$, $MS_e = 1,330$, $p = .1201$], which was unmodified by any group $\times$ task interaction ($F < 1$; see Figure 2). The trend was similarly marginal for both the target-present data [black-target slope = 58 msec/item, white-target slope = 70 msec/item; $F(1,12) = 1.808$, $MS_e = 603$, $p = .204$] and the target-absent data [black-target slope = 127 msec/item, white-target slope = 147 msec/item; $F(1,12) = 2.598$, $MS_e = 1,093$, $p = .133$].

Two-item RTs. RTs for two-item displays were entered into a mixed factors group $\times$ task $\times$ distortion $\times$ target presence ANOVA with group as the between-subjects factor. This analysis revealed a significant main effect for task [black-target search, 733 msec; white-target search, 760 msec; $F(1,12) = 10.414$, $MS_e = 1,994$, $p = .0073$] and a significant group $\times$ task interaction whereby the task effect was stronger for the black distortion group [$F(1,12) = 6.185$, $MS_e = 1,994$, $p = .0286$]. Finally, the target-presence effect was highly significant [707 msec for target-present trials and 785 msec for target-absent

Figure 2. Results of Experiment 1. Target-present and target-absent search slopes for trials with undistorted faces and caricatures.

trials; $F(1,12) = 43.508$, $MS_e = 3,896$, $p < .0001$]. No other main effects or interactions were significant.

Discussion

The primary prediction that distorting the CR face would affect search slopes was confirmed. Caricaturing black targets nonsignificantly speeded the black-target search, and caricaturing black distractors significantly speeded the white-target search. The other primary finding was that distorting targets had no effect on search slopes and only marginally decreased overall search times for black targets. This finding is reasonable if one assumes that search is either a serial scan of individual objects, or a serial scan of small groups of objects within

which search is parallel (e.g., Pashler, 1987; Treisman & Gormican, 1988). In either case, the target will not be detected until some number of distractors have been searched first. Therefore, it is likely that ease of distractor rejection does contribute to the asymmetry. However, before settling on this conclusion, it is important to test an alternative explanation for the form of the caricaturing effect. It may be that caricaturing obscures race-specifying information instead of enhancing it. The distortion was relatively extreme, and, given the fact that the caricature advantage for identifying grayscale faces quickly reversed itself with anything but the smallest distortion (Benson & Perrett, 1991), it is plausible that the resulting face did not match the subjects' representations of black faces. If

a slow SR search was caused by the difficulty inherent to rejecting a feature, ruining that feature could speed the search. The other possibility is that caricaturing the CR distractors made them more effectively different from the SR target, thereby speeding the SR search. This is particularly plausible if the subjects are simply more sensitive to race-relevant variations in CR faces than they are to the same variations in SR faces.

These two hypotheses make contrasting predictions regarding the effects of prototypical distortion in which the faces are made to look more like each other. If the SR search was speeded because the distorted CR distractors obscured race-related features, then an equally extreme prototypical distortion should also speed the SR search because it too should obscure the race feature. If, on the other hand, increased speed is due to increased target-distractor discriminability, prototypically distorting the black face should slow the white search because black distractors would be effectively more similar to the white target.

In Experiment 2, we therefore included both prototypical and caricatured distortions. The prototypical distortions were created by the same process that was used to create the caricatures. In this case, however, distortions were opposite to those used in Experiment 1. Instead of increasing the difference between the races, they reduced the difference by distorting the shape of each face so that it would be more similar to the other. Otherwise, the procedure was similar to that in Experiment 1 with the exception that the entire design was within subjects and only distractors were distorted. This latter change was necessary to limit the length of the task.

EXPERIMENT 2

Method

Subjects. A total of 19 Cornell University undergraduates completed Experiment 2. Six were removed from the analysis due to high error rates. The average error rate of those who remained was 2.5%. Of the remaining 13 subjects, 6 were male, 2 were Asian, 1 was Filipino, and 10 were white.

Stimuli. In addition to searching for caricatured faces, the subjects searched for prototypical distortions created with the same pro-

cess that was used for caricaturing, except that the faces were made to look more similar by changing the shape of each face so that each was the midpoint between the two. Thus, the configuration of the faces was essentially the same. The texture information remained different enough, however, to allow for easy determination of the race of each face (see Figure 3).

Procedure. The procedure for Experiment 2 was identical to that in Experiment 1 with the following exceptions. First, the design was entirely within subjects so that each subject searched displays containing undistorted faces, caricatures, and prototypical distortions. Second, only the distractors were distorted. Targets were always undistorted faces.

Results

In Experiment 2, prototypical distortions of CR distractors slowed the search for SR targets, whereas caricaturing CR distractors again facilitated the search (for target-absent trials).

Distortion effects. Search slopes were entered into a single three-factor task (black search/white search) $\times$ distortion level (none/caricature/prototype) $\times$ target presence (present/absent) within-subjects ANOVA.

The task $\times$ distortion level interaction was significant [$F(2,24) = 5.539, MS_e = 241, p = .0105$], which, along with the pattern of means in Figure 4, shows that distortions of black distractors affected search slopes more than did distortions of white distractors. A two-factor distortion level $\times$ target-presence ANOVA limited to the white-target task (and therefore to distortions of the black distractors) confirmed that there was a strong effect for distortions of black distractors [$F(2,24) = 9.494, MS_e = 357, p = .0009$], and simple effects analysis showed that the distortion effect was significant for both target-present trials [$F(2,24) = 3.710, MS_e = 259, p = .039$] and target-absent trials [$F(2,24) = 12.393, MS_e = 261, p < .001$]. Pairwise comparisons for target-present trials showed that white-target slopes were steeper with prototypically distorted black distractors (60 msec/item) than with undistorted faces (44 msec/item, $p < .01$, Duncan test), but that there was no difference between slopes for caricatured and normal stimuli (48 msec/item vs. 44 msec/item). For target-absent trials, prototypical distractors slowed search (116 msec/item for prototypical distortions vs. 100 msec/item for undistorted faces; $p < .01$, Duncan

Figure 3. Prototypical distortions added in Experiment 2.

Figure 4. Search slopes from Experiment 2.

test), whereas caricatured distractors speeded it (85 msec/item, $p < .01$, Duncan test).

Search asymmetry. To test for a search asymmetry, slopes from trials with undistorted stimuli were entered into a two-factor task (black search/white search) $\times$ target presence (present/absent) ANOVA. The analysis revealed no search asymmetry (task effect, $F < 1$). Simple effects of task were similarly nonsignificant for target-present (black search = 45 msec/item, white search = 44 msec/item; $F < 1$) and target-absent trials [black search = 89 msec/item, white search = 99 msec/item; $F(1,12) = 1.095$, $MS_e = 661$, $p = .316$].

Two-item reaction times. RTs for the two-item displays were entered into a three-factor task $\times$ distortion $\times$ target presence ANOVA. This analysis revealed a significant main effect of distortion [undistorted: 694 msec, caricature: 693 msec, prototype: 716 msec; $F(2,24) = 8.974$, $MS_e = 1,027$, $p = .0012$] and task [black target: 682 msec, white target: 721 msec; $F(1,12) = 5.469$, $MS_e = 10,745$, $p = .0375$]. Finally, the target-presence effect was highly significant [664 msec for target-present trials and 738 msec for target-absent trials; $F(1,12) = 127.303$, $MS_e = 1,680$, $p < .0001$]. No other main effects or interactions were significant ($F_s < 1$).

Discussion

Experiment 2 replicated the finding that distortion affects search slopes when only the feature-positive face is manipulated. In particular, search slopes for the SR target increased when the CR distractors were prototypically distorted. Conversely, slopes decreased when the CR distractors were caricatured. Combined, these findings suggest that a discriminability function is superimposed on the asymmetry. CR faces are easier to detect among SR faces because they are positively defined. When the positively defined stimulus is the target, distractor variation does little to affect search slopes. In the reverse case,

when the feature-negative stimulus is the target, similarity between the target and feature-positive distractors does affect search slope.

EXPERIMENT 3

One distinctive element common to both Experiments 1 and 2 is that the basic search asymmetry favoring detection of black faces is small to nonexistent. Given that Levin (1996, 2000) observed this asymmetry in three different experiments, understanding the exact difference between Experiments 1 and 2 and previously reported experiments may help clarify the conditions that are necessary to observe the asymmetry. The clearest difference between the previously observed search asymmetries and the present experiments is the presence of trial-to-trial variability in the specific faces used as stimuli. In previous successful demonstrations of this asymmetry, subjects searched displays that were consistent within a given block—that is, there was only one version of the target and one of the distractors across trials. In Experiments 1 and 2, on the other hand, the faces were slightly different across trials, causing variation in target–nontarget similarity between trials. Accordingly, the search asymmetry may depend on the subjects' ability to set a consistent criterion for rejecting SR distractors when looking for a CR target. The asymmetry might then occur because subjects learn that the relative salience of the CR face allows them to search more quickly through SR distractors in anticipation of a relatively strong target-present signal once the target is fixated. Such a strategy would probably depend on good a priori knowledge of the exact faces used as targets and distractors, so that if they were different between trials, the asymmetry might disappear because subjects would abandon the attempt to boost search speed.

In Experiment 3, we tested the possibility that trial-to-trial stimulus variability could eliminate the asymmetry.

The subjects were assigned to one of two different conditions. In the consistent-first condition, the subjects completed two blocks (one searching for the black target and one searching for the white target) of consistent trials in which targets were always the undistorted version of one face among undistorted versions of the other face, followed by two blocks of variable trials in which distractors were undistorted faces on some trials and caricatures on others. In the mixed-first condition, a different group of subjects first completed two blocks of trials in which distractors varied, then completed two blocks of consistent trials. If the search asymmetry depends on the subjects' ability to set a consistent rejection criterion, we would expect it to disappear when the consistent blocks are preceded by blocks that confound a simple criterion.

Method

Subjects. Thirty-one introductory psychology students from Kent State University completed the experiment. Of these, 19 failed to meet the 5% error criterion used in Experiments 1 and 2. Because the primary purpose of Experiment 3 was to explore an unexpected finding from those initial experiments, the initial analysis did not include these subjects in order to equate performance across experiments. Four participants remained in the mixed-first condition, and 8 remained in the consistent-first condition. The mean error rate for those who remained was 3.0%. See below for further discussion.

Procedure. The stimuli and procedures were similar to those used in Experiments 1 and 2, with the following exceptions. Instead of completing two blocks of mixed trials, the subjects completed four blocks of trials, two in which distractors were mixed from trial to trial, and two in which they were always undistorted faces. In mixed blocks, distractors were either caricatures or undistorted stimuli. The subjects in the mixed-first condition completed a block of mixed black-target trials and a block of mixed white-target trials (order of black-target and white-target trials was counterbalanced across subjects), followed by two blocks of consistent trials. The subjects in

the consistent-first condition completed the consistent blocks first, followed by the mixed blocks.

In addition, the subjects searched only two- and eight-item displays. This was done simply to shorten the experiment.

Results

The search asymmetry favoring CR faces reappeared for consistent blocks in the consistent-first condition and was not present for consistent blocks that were preceded by mixed blocks.

Search asymmetry. Slopes from consistent blocks were first entered into a three-factor group (consistent first/mixed first) $\times$ task (black target/white target) $\times$ target presence (present/absent) mixed-factors ANOVA with group as the between-subjects factor. This analysis revealed a trend for a significant search asymmetry [task effect, $F(1,10) = 2.956$, $MS_e = 1,549$, $p = .1163$], but no task $\times$ group interaction (see Figure 5). However, this appears to be due to unusually steep target-absent slopes in the consistent-first black-target trials (target-present: target-absent ratio = 1:3.28), so the analysis was repeated for the target-present trials only. This analysis revealed a search asymmetry overall [task effect, $F(1,10) = 7.949$, $MS_e = 238$, $p = .0182$] and a task $\times$ group interaction [$F(1,10) = 5.0784$, $MS_e = 238$, $p = .0480$]. Simple effects analysis revealed a strong asymmetry favoring black targets in the consistent-first condition [black-target slope = 43 msec/item, white-target slope = 77 msec/item; $F(1,10) = 19.29$, $MS_e = 238$, $p = .001$] and no asymmetry in the mixed-first condition (black-target slope = 74 msec/item; white-target slope = 78 msec/item; $F < 1$).

Distortion effects: SR face. Data from mixed blocks were entered into a three-factor group (consistent first/mixed first) $\times$ distortion (normal/caricature) $\times$ target

Figure 5. Search slopes for consistent blocks in Experiment 3. In the mixed-first condition, the consistent blocks were preceded by blocks in which distractors were normal or caricatured on any given trial.

presence (present/absent) ANOVA. This revealed no significant effects of distortion ($F < 1$) or group [$F(1,10) = 2.582$, $MS_e = 7,992$, $p = .1392$], and the distortion $\times$ group interaction was also nonsignificant ($F < 1$). Another ANOVA on only target-present data revealed no significant main effects [group effect, $F(1,10) = 2.940$, $MS_e = 3,390$, $p = .1172$; distortion effect, $F(1,10) = 1.268$, $MS_e = 46$, $p = .2864$], but a near-significant group $\times$ distortion interaction [$F(1,10) = 4.198$, $MS_e = 46$, $p = .0676$]. Simple effects analysis showed that in the consistent-first condition search slopes were significantly shallower when the white distractors were undistorted (39 msec/item) than when caricatured [48 msec/item; $F(1,10) = 7.56$, $MS_e = 46$, $p = .020$]. The effect of distortion in the mixed-first condition was nonsignificant ($F < 1$).

Distortion effects: CR face. The same analysis of distortion for black faces revealed a significant main effect for distortion in the three-factor group $\times$ distortion $\times$ target presence ANOVA [$F(1,10) = 7.110$, $MS_e = 821$, $p = .0236$]. An analysis focusing on target-present data yielded similar results with a significant distortion effect [$F(1,10) = 11.499$, $MS_e = 430$, $p = .0069$] and nonsignificant group $\times$ distortion interaction [$F(1,10) = 1.954$, $MS_e = 430$, $p = .1924$]. Simple effects analysis of target-present data showed that the mixed-first group search was significantly faster with black caricatured distractors (56 msec/item) than with black undistorted distractors [99 msec/item; $F(1,10) = 8.600$, $MS_e = 430$, $p = .015$], whereas the effect of distortion was a nonsignificant trend in the consistent-first condition [caricatures = 44 msec/item, undistorted distractors = 61 msec/item; $F(1,10) = 2.979$, $MS_e = 430$, $p = .115$].

Two-item RTs. RTs for two-item displays were entered into two different ANOVAs. A mixed factors group $\times$ task $\times$ target presence ANOVA that included data from consistent blocks revealed a significant main effect for target presence [857 msec for target-present trials and 919 msec for target-absent trials; $F(1,10) = 10.098$, $MS_e = 4,046$, $p = .0099$], but not for task (black target, 873 msec; white target, 904 msec; $F < 1$). The group $\times$ target presence interaction was nearly significant [$F(1,10) = 2.209$, $MS_e = 10,007$, $p = .1044$]. The target-present advantage was larger in the distortion-first trials.

A mixed factors group $\times$ task $\times$ distortion $\times$ target presence ANOVA was run on data from the mixed blocks, revealing significant main effects for group [mixed-first, 932 msec; consistent-first, 803 msec; $F(1,10) = 5.239$, $MS_e = 67,402$, $p = .0451$] and target presence [target-present, 840 msec; target-absent, 896 msec; $F(1,10) = 11.543$, $MS_e = 5,874$, $p = .0068$]. The group $\times$ target presence interaction was significant [$F(1,10) = 6.979$, $MS_e = 5,875$, $p = .0247$], and the pattern of means suggests that the target-presence effect was stronger in the mixed-first condition than in the consistent-first condition. The task $\times$ target presence interaction was also significant in that target-absent RTs were relatively slower in searching for the white target [$F(1,10) = 13.308$, $MS_e = 1,131$, $p = .0045$]. Finally, this interaction was modified

by a significant group $\times$ task $\times$ target presence interaction [$F(1,10) = 6.736$, $MS_e = 1,131$, $p = .0267$].

Analysis of errors. Given the large number of subjects who did not satisfy the error criterion, an analysis of the pattern of errors that included all 31 subjects was helpful in understanding why they did so poorly. First, the correlation between each subject's mean search RT (across all trials) and their error rates was significant both for target-present and target-absent trials (target-present, $r = -.58$, $p < .001$; target-absent, $r = -.56$, $p < .001$), with high error rates associated with fast responses, suggesting a speed-accuracy tradeoff.

To test patterns of errors across conditions, we entered the error data into two four-factor condition (consistent first/mixed first) $\times$ task (black target/white target) $\times$ distortion (normal in consistent block/normal in mixed block/caricature in mixed block) $\times$ display size (2/8) mixed-factors ANOVAs with condition as the between-subjects factor. One of these included target-present data, and one included target-absent data. Errors on target-absent trials were less frequent than errors on target-present trials [4.6% vs. 9.5%; $F(1,29) = 29.290$, $MS_e = 0.016$, $p < .0001$]. For target-present data, error rates were higher for the white-target search than for the black-target search [11.0% vs. 8.0%; $F(1,29) = 8.154$, $MS_e = 0.011$, $p = .0079$], and error rates increased across display sizes more for the white-target search (6.3% and 15.9% for two- and eight-item displays, respectively) than for the black-target search [6.1% and 9.9% for two- and eight-item displays, respectively; task $\times$ display size interaction, $F(1,29) = 7.073$, $MS_e = 0.011$, $p = .0126$], suggesting a search asymmetry in errors as well as in RTs. However, this asymmetry was not modified by the condition effect ($F < 1$). Finally, the task $\times$ distortion interaction was significant [$F(1,29) = 3.735$, $MS_e = 0.006$, $p = .0298$]. The effect of distortion was small for the white distractors in the black-target search (9.2% errors for undistorted distractors, 7.5% errors for undistorted distractors in mixed trials, and 7.3% errors for caricatured distractors in mixed trials) and larger for the black distractors in the white-target search (10.6% errors for undistorted distractors, 9.2% errors for undistorted distractors in mixed trials, and 13.4% errors for caricatured distractors in mixed trials). For target-absent data, no main effects were significant, although the condition $\times$ distortion, condition $\times$ display size, and distortion $\times$ display size interactions were significant [$F(2,58) = 5.152$, $MS_e = 0.005$, $p = .0087$; $F(1,29) = 7.419$, $MS_e = 0.006$, $p = .0108$; $F(2,58) = 4.880$, $MS_e = 0.004$, $p = .0110$, respectively].

Further analysis suggests that the error asymmetry was strongest in mixed blocks of trials. Although the distortion $\times$ task $\times$ display size interaction was nonsignificant [$F(2,58) = 2.116$, $MS_e = 0.005$, $p = .1297$], the asymmetry was nonsignificant in an ANOVA that included only consistent blocks of trials ($F < 1$), whereas it was quite strong in another ANOVA that included only mixed blocks [$F(1,29) = 7.742$, $MS_e = .013$, $p = .0094$].

A similar analysis of target-present error rates that included data from only those subjects who were elimi-

nated from the RT analysis revealed a strong asymmetry favoring the black-target search. In the black-target search, the error rate for large displays (12.49%) was 4.73% greater than the rate for small displays (7.67%), whereas it increased by 14.73% across display sizes (22.38% for large displays vs. 7.56% for small displays) for the white-target search [task $\times$ display size interaction, $F(1,17) = 10.758$, $MS_e = 0.013$, $p = .0044$]. This effect was unmodified by a significant condition $\times$ task $\times$ display size interaction [$F(1,17) = 1.052$, $MS_e = 0.013$, $p = .3195$].

In addition to testing correlations between overall RTs and error rates we did a more fine-grained analysis of RTs preceding and following error trials. Chun and Wolfe (1996) argued that subjects optimize the balance between speed and accuracy by progressively responding more quickly across trials until they make an error, then becoming more cautious and slowing down for the next trial in order to be assured of a correct response. Following this, they speed up until they make another error. If this cyclic process constitutes the on-line control over errors, it is possible that subjects who make too many errors are failing to appropriately slow their responses after an error occurs. If this is true, the overall error rate for each subject should be correlated with the degree to which that subject's post-error RTs are greater than their pre-error RTs. This correlation was, indeed, significant for target-present trials ($r = -.49$, $p = .006$), but not for target-absent trials ($r = -.23$, $p = .214$; in both cases, the five trials before and after each error were averaged). It therefore appears that many of the eliminated subjects maintained a poor balance between speed and accuracy, causing artificial deflation of RTs.

Discussion

The results of Experiment 3 confirm the hypothesis that trial-to-trial variability can eliminate the search asymmetry. When the subjects first completed sets of consistent trials in which targets and distractors were the same face across trials, they showed a strong search asymmetry in target-present trials. However, when they completed the same trials after having experienced trial-to-trial variability in the mixed trials, the asymmetry disappeared. Thus Experiment 3 replicated the search asymmetries observed by Levin (1996, 2000), while revealing an important limit to the phenomenon. In addition, the distortion effects were again observed for the CR faces more strongly than for the SR faces. As in Experiment 1 (and to some degree Experiment 2), caricaturing the CR distractors speeded search for the SR face. The pattern of results across the three experiments suggests that the distortion effects and the asymmetry may have the same underlying cause (i.e., asymmetrically defined features) but are independent manifestations of the feature definitions. The distortion effects are consistent across experiments and, in addition, are stronger in distractors, whereas the asymmetry is sensitive to the trial-to-trial variability in the salience of search-relevant features.

One surprising result of Experiment 3 was that caricaturing SR distractors significantly slowed the CR-target

search in the consistent-first condition. The cause of this is not clear and difficult to conceptualize. Perhaps the distortion in some cases attracted attention to the distractors, thus slowing the search. However, it is important to note that the effect was isolated among a large number of negative findings for SR distortions and is smaller in magnitude (9 msec/item) than all of the CR-distortion effects (16–43 msec/item in Experiments 1–3).

The major difficulty with this experiment is clearly that so many subjects had to be dropped from the analysis. The problem appears to be that the subject population in Experiment 3 was unable to control error rates as effectively as those in Experiments 1 and 2. Because the goal of Experiment 3 was to explore a finding from the first two experiments, it was important to maintain a consistent error cutoff so that we could test the hypotheses about the first experiments using subjects who performed the task in a similar manner. Although this forced many subjects out of the initial analysis, it is important to point out that the results of this experiment closely replicate those of other experiments. First, the distortion effects in both conditions were similar to those in the other two experiments. In all cases, distortion affected search for the CR face more strongly, and the mean slopes were similar to those observed in other experiments. Target-present SR-target slopes with caricatured CR distractors were 56 msec/item in the mixed-first condition of Experiment 3 as compared with 49 msec/item in Experiment 1 and 99 msec/item versus 74 msec/item, respectively, for the undistorted conditions. Slopes in the consistent-first condition were also similar to those reported in Levin (1996), in which no subjects had been eliminated, and in Levin (2000) Experiment 1, in which 4 out of 25 subjects had been eliminated. Target-present CR search slopes ranged between 53 and 65 msec/item in the previous experiments as compared with 43 msec/item here, whereas slopes in the SR search were between 75 and 88 msec/item previously as compared with 77 msec/item in the present case. Finally, to ensure that the main interactions were not idiosyncratic, they were rerun with three additional mixed-first subjects who only barely missed the criterion (by less than 1%). This analysis was very similar to the one above, and, in particular, the interaction between group and task for target-present trials was again significant [$F(1,13) = 6.637$, $MS_e = 271$, $p = .0230$], and the means for the CR-target and SR-target searches in the mixed-first condition were almost identical (mixed-first CR search = 77 msec/item, mixed-first SR search = 80 msec/item).

Analyses that included all subjects produced no significant task or group effects and no significant interactions between the two. There was a 13-msec asymmetry favoring detection of black faces in the consistent-first condition and a 1-msec asymmetry in the mixed-first condition.

Combining these analyses with the previous error analysis gives the impression that the subjects who were dropped in the initial analysis traded speeded responses on the slower white-target search for errors, at least for mixed

blocks of trials. Thus, the search asymmetry in the error data reflects the increased difficulty of the white-target search, and the subjects who made many errors overall made a particularly large number of errors in the white-target search, especially in the 8-item displays for which we observed error rates as high as 22.38%. Given that the error asymmetry was strongest in the mixed blocks, it is also possible that the absence of the RT asymmetry in Experiments 1 and 2 occurred in part because the asymmetry was reflected in the error data of those experiments. In addition, the analysis suggests that the subjects who did not achieve the error criterion did not sufficiently slow their responses after error trials. However, the error asymmetry was nonsignificant for consistent trials, which suggests that a speed–accuracy tradeoff cannot be assumed to underlie all negative RT findings with this search asymmetry.

Before going on to the General Discussion section, it is important to address one final methodological issue. Throughout this report, we have assumed that differential search rates for black and white faces and differential distortion effects are due to feature definitions based on concepts about social groups. The alternative is that these effects are due to stimulus-based factors inherent to black faces and white faces. Although this is possible to a degree, it is unlikely in the present case. First, as mentioned above, the basic asymmetry and closely related effects, such as facilitated classification of CR faces and improved discrimination for variants of CR faces, have each been repeatedly related to social-cognitive variables, such as group membership and the presence of the CR recognition deficit in a number of reports (Levin, 1996, 2000; Valentine & Endo, 1992).

In addition, the stimuli themselves were carefully controlled so that a number of potential stimulus-based differences, such as luminance, contrast, and external features, were equalized between the races. Also, the distortions involved only internal feature shape and not luminance or contrast. Therefore, only a few of the manipulated features were substantively different between the faces. The primary difference was the increased size of some internal features in the black face. However, slopes were sensitive to the difference between the undistorted black face and the 50% prototypical distortion of that face, whereas they were not sensitive to the difference between the 50% white distortion (which was very similar in basic shape parameters to the prototypically distorted black face) and the undistorted white face. Therefore, to sustain a stimulus artifact hypothesis, one would have to hypothesize a cutoff in basic feature-shape discriminability almost exactly in the middle of the black–white size continuum, which, for each feature, actually covers a relatively small range of absolute sizes to begin with.

GENERAL DISCUSSION

The experiments reported here converge on three basic findings. First, distortions in the feature-positive CR face affected search slopes, whereas the same distortions in

the SR face did not. Second, distortions that increased the prominence of feature-specifying information affected slopes by effectively reducing the size of the search asymmetry, especially because caricaturing CR distractors speeded search for the SR target. Finally, these experiments show that trial-to-trial variability in the specific faces used as stimuli can eliminate the basic search asymmetry favoring detection of CR faces. In Experiments 1 and 2 and in the mixed-first condition of Experiment 3, the search asymmetry was small or nonexistent, whereas in the consistent-first condition of Experiment 3, it was present and quite strong in target-present trials.

Combined, these results suggest two primary constraints on explaining the search asymmetry for race. First, it appears difficult to understand the asymmetry without considering the salience of the target and ease of rejecting distractors independently as suggested by Tong and Nakayama (1999). The second issue to discuss is the need to understand how task contexts can modulate the asymmetry. On the basis of these findings, the search asymmetry for race is apparently predicated on subjects' ability to set a consistent criterion when rejecting SR distractors.

Increase the Feature and Decrease the Asymmetry: Modeling a Serial Search Asymmetry

In a sense, it is paradoxical that increasing the salience of the feature that drives the search asymmetry will end up eliminating it. Given that it is difficult to adapt Treisman and Gormican's (1988) pooled activation model to this relatively slow serial search, it might be useful to consider the explanation for search asymmetries described in Duncan and Humphreys's (1989) model of visual search and stimulus similarity. Under this model preattentively generated descriptions of all objects in the display compete for access to visual short term memory (VSTM) on the basis of their similarity to a template describing the target (i.e., a representation of the target). Duncan and Humphreys explain search asymmetries by assuming that target templates are necessarily more complex for feature-negative targets. For example, where the target is a feature-positive tilted line, the template simply contains the feature "tilt," whereas if the target is a feature-negative straight line, the template must contain both a description of the standard and a proviso that the feature must not be present. It therefore contains the complex conjunction "line *and* no tilt." Search for a feature-negative target is therefore more difficult both because it is complex and also because it implicitly contains a description of the feature-positive distractors. In the line+tilt example, the "no-tilt" specification contains reference to the tilted distractors. This might make the distractors difficult to reject and, furthermore, could cause increased sensitivity to feature-relevant variation in the feature-positive distractors.

Using this framework, then, search for a black target is faster because the template describing it is simpler: The black-target template would be "black," whereas the white-target template would be "face *and* not black." Although this might help explain facilitated detection of

black targets, the model seems less suited to explaining the effects of distortion on rejecting black distractors. If the white-target template is "face *and* not black," it partly matches black distractors that might make rejecting them difficult, but one would predict that increasing the salience of race-specifying information would make the black distractors even harder to reject, which is the opposite of what actually occurred.

The above discussion suggests that it is possible to explain the effect of feature-positive target salience but more difficult to tie this closely with the effects of caricaturing distractors. Therefore, it might be useful to consider the asymmetry as jointly constrained by independent contributions from factors that make targets salient and those that make distractors easy to reject (Tong & Nakayama, 1999). In the case of the present experiments, the feature-positive status of the black face affects these processes differently: It makes the black target easier to detect, perhaps because the black template is simpler, whereas increasing the salience of the CR feature in distractors is more effective in making them distinct from targets. Clearly, the overall pattern of slopes and mean RTs in small displays fits with Tong and Nakayama's prediction that an asymmetry based on target salience will show shallower slopes when the feature-positive stimulus is the target in addition to faster RTs for the feature-positive stimulus in small displays. (In the present case, RTs for 2-item feature-positive target displays were faster than 2-item feature-negative target displays. In addition, Levin [1996] found that individual CR faces are detected and classified more quickly than SR faces.) On the other hand, the distortion effects suggest that the same asymmetrically defined feature can also affect the degree to which distractor-rejection processes contribute to the asymmetry.

It is interesting to note that, if there is no limit to the benefit that the white-target search can receive from caricatured black distractors, then it might be possible to force the white-target search to become faster than a black-target search (with similarly caricatured white distractors). Accordingly, it is important to consider the possibility that the ultimate outcome of the feature-positive/feature-negative relationship could be a search asymmetry favoring the feature-positive target or the feature-negative target. In the latter case, if feature-positive distractors become sufficiently easy to reject, it will overwhelm the advantage provided by the salience of a feature-positive target. This means that we might need to reconsider the automatic assumption that the feature-positive object is associated with shallower slopes in serial search asymmetries and search for converging evidence to constrain the feature relationship such as that provided by the asymmetric caricaturing effects.

The Importance of Setting a Response Criterion in a Serial Search Asymmetry

Given that the search asymmetry depended on the consistency of stimuli across trials, the present data also emphasize the importance of setting a well-calibrated

criterion for search termination to allow for this search asymmetry. In the present experiments, trial-to-trial variation in the specific degree of feature salience may have prevented the subjects from adopting the optimal criterion for distinguishing target-present and target-absent trials. It is important to note that this effect seems to have occurred not only within a block of trials, but can carry over from one block to another, as illustrated in the mixed-first condition in Experiment 3, in which a consistent block of trials revealed no asymmetry when it was preceded by a mixed block.

Although little is currently known about the role of these criterion-setting processes in the mechanics of visual search, it is possible to specify some of the issues that will be important for understanding these criteria. First, the data from Experiment 3 suggest that the asymmetry disappeared because search for the feature-positive target was slowed in the mixed-first condition. Thus, the criterion effect appears to be a matter of lost opportunity in that the normally faster (40–50 msec/item for target-present trials) feature-positive search is slowed to the same rate as the feature-negative search. Second, it is important to understand what events are critical for setting this criterion. Chun and Wolfe (1996) discussed a similar issue when they analyzed search terminations in target-absent trials. They suggest that subjects may pool information across target-absent displays and respond more and more quickly to nonoccurrences of targets until they make an error, at which point they slow their responses on the next trial, then begin speeding up again. This (assuming immediate feedback), errors are hypothesized to be critical events that determine the response criterion. In this case, it does appear that the error rate in the variable blocks (2.54%) is less than in the consistent blocks [4.23%; $F(1,11) = 5.452, MS_e = 3.153, p = .0395$]. Thus, it is plausible that an error-based control over criterion is used more in the consistent trials, whereas some other mechanism for controlling responses is used in the variable block, and that this mechanism can carry over from block to block. However, there appears to be no interaction between task order (mixed-first or consistent-first) and the relative error rate, so this hypothesis must be considered speculative. Another suggestive bit of data comes from mean RTs to the 2-item displays. In the mixed-first condition, these were 100 msec slower for target-absent displays than for target-present displays, whereas they were only 13 msec slower for the target-absent displays in the consistent-first condition. Again, this suggests that the subjects are less conservative and more able to use a subtle perceptual criterion when the stimuli are consistent across trials.

Summary and Conclusions

The present experiments confirm predictions based on the assumption that white subjects code black faces as being feature-positive for race, whereas they code white faces as feature-negative norms. The search asymmetry favoring detection of CR faces was modulated both by manipulations of the salience of race-specifying informa-

tion in the feature-positive target and by the consistency of the stimuli across trials. When black distractors were caricatured, search for the white target was facilitated, whereas it was slowed when the black distractors were prototypically distorted. There were no equivalent effects for white faces, and the direction of the effects suggests that analyses of serial search asymmetries need to explicitly account for both target-detection and distractor-rejection processes. In addition, it is important also to understand the degree to which subjects are able to consistently use a set target-detection (or distractor-rejection) criterion across search trials. Thus, locating the subtle, but highly meaningful, features that specify race appears to depend on a carefully balanced search for feature-positive objects that vary both in terms of perceptual feature-presence and the complexity of their working-memory representation.

REFERENCES

- BENSON, P. J., & PERRETT, D. I. (1991). Perception and recognition of photographic quality facial caricatures: Implications and the recognition of natural images. *European Journal of Cognitive Psychology*, **3**, 105-135.
- BRENNAN, S. E. (1985). Caricature generator: The dynamic exaggeration of faces by computer. *Leonardo*, **18**, 170-178.
- 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.
- DUNCAN, J., & HUMPHREYS, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, **96**, 433-458.
- LEVIN, D. T. (1996). Classifying faces by race: The structure of face categories. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **22**, 1364-1382.
- LEVIN, D. T. (2000). Race as a visual feature: Using visual search and perceptual discrimination tasks to understand face categories and the cross race recognition deficit. *Journal of Experimental Psychology: General*, **129**, 559-574.
- LI, J. C., DUNNING, D., & MALPASS, R. S. (1994). *Cross-racial identification among European-Americans: Basketball fandom and the contact hypothesis*. Unpublished manuscript.
- PASHLER, H. (1987). Detecting conjunctions of color and form: Reassessing the serial search hypothesis. *Perception & Psychophysics*, **41**, 191-201.
- SEDIKIDES, C., & OSTROM, T. M. (1993). Perceptions of group variability: Moving from an uncertain crawl to a purposeful stride. *Social Cognition*, **11**, 165-174.
- TONG, F., & NAKAYAMA, K. (1999). Robust representations for faces: Evidence from visual search. *Journal of Experimental Psychology: Human Perception & Performance*, **25**, 1016-1035.
- 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., & ENDO, M. (1992). Toward an exemplar model of face processing: The effects of race and distinctiveness. *The Quarterly Journal of Experimental Psychology*, **44A**, 161-204.
- VON GRÜNAU, M., & ANSTON, C. (1995). The detection of gaze direction: A stare-in-the-crowd effect. *Perception*, **24**, 1297-1313.
- VON GRÜNAU, M., & DUBÉ, S. (1994). Visual search asymmetry for viewing direction. *Perception & Psychophysics*, **56**, 211-220.
- WOLFE, J. M. (1998). Visual search. In H. Pashler (Ed.), *Attention* (pp. 13-73). Hove, U.K.: Psychology Press.

(Manuscript received March 7, 2000;
revision accepted for publication December 9, 2000.)