Emotion
2006, Vol. 6, No. 2, 193–207

Copyright 2006 by the American Psychological Association
1528-3542/06/$12.00 DOI: 10.1037/1528-3542.6.2.193

Search Asymmetries With Real Faces: Testing the Anger-Superiority Effect

Gernot Horstmann and Andrea Bauland
Bielefeld University

The anger-superiority hypothesis states that angry faces are detected more efficiently than friendly faces. Previously research used schematized stimuli, which minimizes perceptual confounds, but violates ecological validity. The authors argue that a confounding of appearance and meaning is unavoidable and even unproblematic if real faces are presented. Four experiments tested carefully controlled photos in a search-asymmetry design. Experiments 1 and 2 revealed more efficient detection of an angry face among happy faces than vice versa. Experiment 3 indicated that the advantage was due to the mouth, but not to the eyes, and Experiment 4, using upright and inverted thatcherized faces, suggests a perceptual basis. The results are in line with a sensory-bias hypothesis that facial expressions evolved to exploit extant capabilities of the visual system.

Keywords: emotion, expression, attention, evolution, visual search

Starting with the seminal study by Hansen and Hansen (1988), many researchers have pursued the idea that certain facial expressions of emotion (Ekman, 1972) are preattentively available to the cognitive system (e.g., Eastwood, Smilek, & Merikle, 2001; Fox et al., 2000; Fox, Russo, & Dutton, 2002; Hansen & Hansen, 1988; Mogg & Bradley, 1999; Nothdurft, 1993; Öhman, Lundqvist, & Esteves, 2001; White, 1995). This assumption is often framed in an evolutionary or ecological argument that it is of adaptive advantage to respond quickly and with as little conscious preponderance as possible to potentially damaging stimuli (see also, LeDoux, 1998). Therefore, the processing of negatively valenced social and nonsocial stimuli is assumed to have primacy over, for example, the processing of positive or beneficial stimuli. As a specific hypothesis, a processing advantage is assumed to exist for negative or threatening facial expressions over positive or happy facial expressions.

To test this hypothesis, researchers have adopted a frequently used method from vision research, namely the visual search paradigm, where several stimuli are presented simultaneously, and the participants' task is to find a target among distractors (e.g., Treisman & Gelade, 1980; Wolfe, 1998, 2001). If the target is characterized by a feature that is available to the cognitive system prior to its attentional processing, it can be found efficiently with detection latency being independent of set size, which is the total number of stimuli presented. For example, if a red disk can be found among 10 green disks with a latency of 1,000 ms, and among 20 or even 30 disks with a latency of 1,000 ms as well, search is efficient by definition and the feature is assumed to be accessible before attention is directed to its location. In contrast, if the latency of finding a stimulus is positively related to set size, it is concluded that the detection of the stimulus is the result of a serial deployment of attention on each stimulus in succession until the target is detected. For example, if the finding of a T lasts 1,000 ms among 10 Ls, 1,300 ms among 20 Ls, and 1,600 ms among 30 Ls, search is nonefficient by definition, and the detection of the presence of the T is assumed to follow attention rather than precede it.

Practically, regression analysis is used to obtain the linear equation ($y = bx + a$) relating the latency of finding the target (y) to set size (x). If the slope of the function (b) is below 10 ms per item, search is labeled "efficient," whereas search slopes exceeding 20 ms /item are considered as "inefficient" (Wolfe, 1998). In theory (e.g., Treisman & Gelade, 1980; Wolfe, 1994), the nonefficiency of the search of Ts among Ls is because Ts and Ls share basic features (horizontal and vertical lines), and that it is the specific conjunction of these features that defines their identity. A target that is defined by a conjunction of basic features that also appear in the target's surroundings, in turn, normally requires spatial attention to be detected (e.g., Feature Integration Theory, Treisman & Gelade, 1980). Initially it was assumed that feature search is always efficient, while conjunction search is always inefficient (Treisman & Gelade, 1980; Treisman & Souther, 1985). Later evidence revealed that in some cases, conjunctions of features can be found efficiently, which led to modifications of the original theory (e.g., Guided Search 2.0, Wolfe, 1994). According to these later modifications, efficient search is a necessary but not a sufficient condition to qualify a feature as basic in perception, and preattentively available (cf. Wolfe & Horowitz, 2004).

A particular variant strategy in the investigation of preattentive processing is the examination of search asymmetries (e.g., Treisman & Souther, 1985; Treisman & Gormican, 1988; Wolfe, 2001). A search asymmetry amounts to the finding of pop-out and of serial search with only two types of stimuli. For example, search for a "lollipop" (a circle with a vertical stroke dissecting its basis) among circles is efficient, whereas search for a circle among

---

Gernot Horstmann and Andrea Bauland, Department of Psychology, Bielefeld University.

We are indebted to Lily-Maria Silny for her assistance in manuscript preparation, and to Christine Broermann, Wulf-Uwe Meyer, and 3 anonymous reviewers for their helpful comments on earlier versions of the manuscript.

Correspondence concerning this article should be addressed to Gernot Horstmann, Department of Psychology, Bielefeld University, POB 100 131, 33 501, Bielefeld, Germany. E-mail: gernot.horstmann@uni-bielefeld.de

lollipops is inefficient. This pattern is considered an important diagnostic of a feature being basic in perception and thus available prior to attention (cf. Wolfe & Horowitz, 2004). More precisely, it indicates that the two stimuli can be discriminated on at least one dimension, which is present in the pop-out stimulus but absent in the serial-search stimulus (Treisman & Souther, 1985), or which is owned in large quantities in the pop-out stimulus, but only in little quantities in the serial-search stimulus (Treisman & Gormican, 1988; see also Wolfe, 2001; Wolfe & Horowitz, 2004).

Several studies have tested the prediction that angry or negative faces are detected preattentively (and can thus be used to guide attention to its position in a crowd) and with priority over positive faces by making use of the search-asymmetry paradigm. In particular, these studies have compared the latency to find an angry face (being negative or threatening; cf. Horstmann, 2003) in a friendly crowd (lacking the characteristic of being negative or threatening) to the latency to find a happy face in an angry crowd. In the following, the pertinent findings bearing on the preattentively available facial threat information will be shortly reviewed. As will be shown, the evidence for a preattentive discrimination of facial threat in the published studies is mixed. Then we will point out some theoretical problems, and we will suggest a possible solution to the diagnosed problems. After that, we will present four new experiments bearing on the threat advantage-hypothesis.

Literature Review

Several experiments, using different methods, have examined the preattentive discrimination of affect (e.g., Eastwood et al., 2001; Fox et al., 2000; Hansen & Hansen, 1988; Horstmann, submitted; Horstmann, Borgstedt, & Heumann, in press; Nothdurft, 1993; White, 1995; Öhman et al., 2001; Fenske & Eastwood, 2003). The present review, however, focuses exclusively on the experiments using the visual search paradigm and is even more restricted to experiments that examined a possible search asymmetry between angry and happy faces. Recently, a number of studies have used a different approach by comparing the efficiency of a search for emotionally valenced target faces (e.g., positive vs. negative faces) among emotionally neutral distractor faces (e.g., Eastwood et al., 2001; Williams et al., 2005). We will discuss the pros and cons of this constant distractor paradigm in the final discussion. For the present, we note that the search-asymmetry design, by using only two contrasting stimuli assumed to differ on a single dimension (in this case: threat), avoids possible problems that result from the use of a third "neutral" stimulus.

Photographic Stimuli

Hansen and Hansen (1988, Experiment 3) were the first to conduct a visual search study with the aim of testing a possible threat-advantage for angry faces. Hansen and Hansen presented angry faces in happy crowds and happy faces in angry crowds with varying set sizes of 4 and 9 faces, arranged in a 2 × 2 and a 3 × 3 matrix, respectively. The faces were photographic stimuli taken from the Ekman and Friesen (1976) set of facial expressions of emotion (see also the publication of Purcell, Stewart & Skov, 1996, for reproductions of the stimuli used in Hansen & Hansen, 1988). Photos of two stimulus persons were used in the experiments, but each participant saw the face of only one stimulus person. In half the trials (target absent trials), no target was presented, with angry and happy crowds being presented equally often. In the remaining trials (target present trials), a happy face or an angry face was presented in a crowd of the other facial expression. The participant's task was to indicate the presence of a discrepant face with a key press.

Hansen and Hansen (1988) found that in target present trials, angry faces were detected efficiently with a slope (b) of the linear function (i.e., $y = bx + a$) relating RT (y) to the set size (x) being 8 ms per face, while friendly faces were detected nonefficiently, with a slope of 52 ms/ face. Moreover, in target absent trials, the slope for happy faces was 90 ms /face—that is, the ratio of the target present slopes and the target absent slopes was about 1:2 (i.e., twice as high in the target absent trials than in the target present trials), indicating serial self-terminating search (cf. Treisman & Souther, 1985). A 1:2 slope is predicted on the basis of serial self-terminating search because in target present trials, the target is, on the average, found after serially scanning half the items in the display, while a target absent judgment requires the scanning of all items in the display. In contrast, the target present to target absent ratio was 1:1.5 with the angry faces, implying that even in target absent trials, participants did not scan all items in the display. Hansen and Hansen (1988) thus obtained a search asymmetry for angry and happy faces, implying that the stimuli they used can be compared to each other by their value of a common feature, with angry faces possessing this feature and with the happy faces lacking it. Thus, the result supports the hypothesis that facial threat is a preattentively available feature that can be used to guide attention to a threatening face, and that preattentive processing is unique to threatening (or negative) faces.

Hansen and Hansen's (1988) use of photographic stimuli has been proven to be problematic. In particular, Purcell et al., (1996) substantiated that the original result is because of a confound that occurred during the digital image processing (gray-scale pictures were converted into black-and-white pictures), resulting in conspicuous black spots that pertained only to the angry faces and not to the happy faces. Apparently, the participants detected this confound and used it to discriminate between target present and target absent trials with happy crowds. In the replication of Purcell et al. (1996), only those participants who reported the confound revealed efficient search for angry target faces. Moreover, when the original gray scale pictures were used instead of the digitally processed high contrast derivates, the search asymmetry was not obtained. The search-asymmetry studies reviewed hereafter all used schematized line drawings of facial expressions instead of realistic photographic stimuli, evidently because experimental control over line drawings is much more precise than with photographs.

Schematic Stimuli

Partly using the pioneering study of Hansen and Hansen (1988) as the starting point, four published studies used the search asymmetry design using schematic stimuli (in chronological order: Nothdurft, 1993; White, 1995; Fox et al., 2000; Öhman et al., 2001). The results of these studies are quite heterogeneous. Nothdurft (1993, Series 5) used drawings of faces composed of a circle as the face's outline, covered with hair, with dots as eyes, a "‸" as the nose and a curved line as the mouth (see Figure 1a). The only

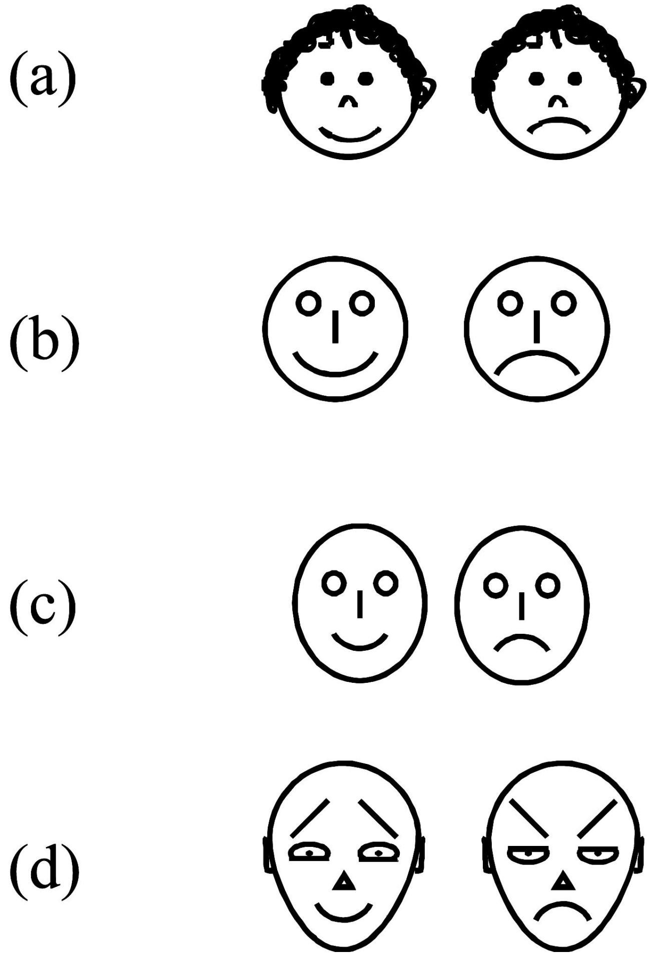

Figure 1. Overview over the stimuli used in the previous research. The stimuli are replicas of stimuli used by (a) Nothdurft (1993), (b) White (1995), (c) Fox et al. (2000), and (d) Öhman et al. (2001).

discriminating feature was the orientation of the curve, indicating a smile or a frown. Consistent with his other research on visual perception (e.g., Nothdurft, 1991), he used a large variation of set sizes, with more than 40 faces for his largest crowds, presented in an irregular rectangular matrix. Search was very inefficient with target present slopes of 62 ms per item, and he apparently did not find differences in search efficiency with positive and negative targets (Nothdurft, 1991, Figure 4). White (1995, Experiment 1) used positive and negative schematic faces similar to those presented in Figure 1b, composed of a circle as the face's outline, a stroke as a nose, and small circles as eyes. The stimuli were presented on the circumference of an imaginary circle, that is, equidistantly from fixation. Three set sizes were used with 2, 4, or 6 stimuli. The experiment revealed practically flat search functions for target present trials, with a 0 ms slope/item for happy targets and a -3 ms/item slope for negative targets. Target absent slopes were about 40 ms/item. In addition, there was a main effect for crowd, with RTs to positive crowds being about 100 ms faster than to angry crowds. For the sake of clarity, it should be emphasized that a crowd effect is not predicted on the hypothesis of a preattentive detection of threat or negative affect: it is the slope (*b*) of the function relating RT to set size ($y = bx + a$), that defines search efficiency, not the intercept term (*a*) of the function. Rather, the intercept may reveal processes that occur before the beginning of the search or between the termination of the search and the production of the response, but not during search.

Fox et al. (2000, Experiment 5) presented schematic faces that were similar to those of White (1995) except that the outline of White's faces was a circle whereas that of Fox et al. (2000) was an oval (Figure 1c). In some of the experiments reported in Fox et al., the stimuli had brows that made the negative face look more unambiguously angry than sad, which were however omitted in the only experiment in which two set sizes (4 vs. 8 faces) were used and in which exclusively positive and negative faces were presented. The faces were presented on the imaginary outline of a circle, similar to White's experiment. The authors found shallower search slopes for the angry targets (16 ms) than for the happy targets (29 ms). We will call this results pattern a relative search asymmetry henceforth, to distinguish it from a search asymmetry proper, with efficient search for one stimulus and inefficient search for the other stimulus. The target absent slopes were quite steep amounting to 82 ms per face. The high target present/absent ratio of about 1:3.6 is a somewhat odd feature of the results pattern because an extreme slope ratio may indicate that the participants were highly uncertain that no target is present in target absent trials (cf. Treisman & Souther, 1985), while the relatively small slopes in target present trials indicate that the target, if present, was found very quickly. However, a threat-advantage was also evident in the error rates, which did not increase with set size for angry faces but did with happy faces as targets.

Öhman et al. (2001, Experiment 3) presented quite elaborate schematic facial expressions, where angry and happy faces differed by the shape of eyes, mouth, and brows (Figure 1d). In particular, the faces were constructed such that when the orientation of the eyes, mouth, and brows in happy faces was considered as a 0°, each feature had a 180° orientation in the angry face. Öhman et al. (2001) used 3 set sizes of 4, 9, and 16 faces arranged in regular matrices of 2×2, 3×3, and 4×4, respectively. Öhman et al. (2001) did not find a search asymmetry but rather inefficient search for both angry and happy target faces, with search slopes of approximately 35 ms/face in the target present trials and about 75 ms/in the target absent files (the latter figure is derived from the RTs in Figure 5 of Öhman et al.). The target present to target absent ratio was 1:2.1, being compatible with serial self-terminating search. Although the slopes for both faces were clearly inefficient, it is possible that a speed–accuracy tradeoff masked a relative search asymmetry because the error rates revealed a set size effect with angry crowds and friendly targets, but not in the reversed condition. Öhman et al. (2001) also obtained a crowd effect, with responses to happy crowds being faster and error rates being lower, at least for target present trials. In sum, a search asymmetry with pop-out for angry or threatening faces was not found, search was clearly inefficient, and only the error data may be interpreted as revealing an advantage of angry targets in happy crowds.

Some Potential Problems

The review of the experimental literature that have tested the hypothesized search asymmetry for a threat or a negative valence dimension for facial stimuli has revealed a rather heterogeneous set of search slopes. The target present slopes ranged between 0 ms (White, 1995) and 62 ms (Nothdurft, 1993), and some, but not all studies found a relative search asymmetry. (Actually, one of us [Horstmann, submitted] has tested replicas of Fox et al.'s (2000), White's (1995), and Öhman et al.'s (2001) stimuli with the same visual search procedure as used in the present experiments, and found relative search asymmetries with all stimuli, replicating, however, a wide range of search slopes, and pronounced differences in the sizes of the asymmetries).

A possible implication is that the experiments, which aimed at comparing stimuli of the categories of facial threat versus facial friendliness (or alternatively, positive vs. negative facial expressions) actually revealed that different instantiations of the categories produce quite different results. This is problematic with respect to internal and external validity. First, given that efficient search was found in some studies (Hansen & Hansen, 1988; White, 1995), but not in others (Fox et al., 2000; Nothdurft, 1993; Öhman et al., 2001), what should be concluded concerning the hypothesized preattentive processing of facial affect (i.e., what internal validity do the studies have)? Second, because different versions of schematic faces were used, it is hard to know which of the stimulus pairs—if any—is ecologically relevant (i.e., what external validity the studies have).

Ecological relevance is an important issue here because the evolutionary threat-advantage hypothesis is about real faces, not about schematic faces. Just to mention two particular problems: First, the most favorable evidence for a threat advantage comes from studies that use relatively impoverished stimuli (e.g., Fox et al., 2000), whereas more complex stimuli tend to provide weaker evidence (e.g., Nothdurft, 1993; Öhman et al., 2001). Given that real faces are even more complex, this may indicate that with real faces, no compelling threat advantage exists. Second, nearly all studies appear to rely on a facial feature—the frowning mouth—which is not usually considered a feature of an angry face (e.g., Ekman & Friesen, 1976; see also below). Thus, to evaluate any evolutionary inspired hypothesis, it is of primary importance to know which of the facial stimulus pairs used in prior research really represent real faces, and is thus relevant for the evolutionary hypothesis.

The external validity issue—which of the different faces used as stimuli best represents real faces—becomes even more problematic if it is recalled that the threatening and friendly faces are, by necessity, perceptually different (see also, Eastwood et al., 2001). Because of this fact, it is always debatable whether different search slopes for positive or negative faces are because of differences in facial expression or to more general differences in perceptual features. Of course, vision research has revealed quite a number of search asymmetries that have little to do with facial expressions or emotional stimuli in general (cf. Wolfe, 2001).

Researchers have often considered the "perceptual difference" alternative hypothesis in their research, for example by scrambling or inverting the faces in control experiments. In scrambled-face control experiments, the individual components of the faces are arranged in a way that does not give rise to the impression of a face (e.g., Nothdurft, 1993), thus controlling for the possibility that some of the components, for example, the orientation of the line representing the mouth (an upward or a downward pointing curve), alone causes the asymmetry. If the asymmetry is present also with the scrambled faces, this is prima facie evidence that isolated components support the asymmetry. However, the elimination of the asymmetry is more ambiguous because the scrambling does not only destroy the impression of a face, but also additional perceptual supraelement cues (cf. Pomerantz et al., 2003).

The logic of inverting the face relies on the finding that face processing is disturbed if the face is presented upside down (e.g., Thompson, 1980). Thus, this procedure tries to control for the possibility that the configuration of stimulus features is sufficient for producing the asymmetry even if face-processing is handicapped. This control procedure is also problematic, however, because orientation asymmetries are known that have nothing to do with faces or emotional stimuli (Enns & Rensink, 1990; Kleffner & Ramachandran, 1992; Wolfe, 2001). For example, Enns and Rensink (1990) found that objects that appeared to be lit from below were more efficiently detected among objects lit from above than vice versa. Obviously, object inversion has its own effect on visual search performance, and for this reason, an orientation asymmetry is not easy to interpret.

To summarize, (a) existing studies have not unequivocally revealed efficient search for angry faces, (b) these studies have mostly used schematic faces, whose ecological relevance (and the relevance for the evolutionary hypothesis) is questionable, and (c) suffer from the unavoidable confound of emotional and perceptual differences.

We suggest two solutions to these problems, one theoretical and one empirical. Theoretically, we would like to recall that the hypothesis that the visual system has evolved special capabilities to preattentively detect facial threat covers only one part of a broader evolutionary grounded hypothesis. In essence, this preattentive threat-detector hypothesis assumes that the visual system has changed during human evolution to adapt to ecologically important facial stimuli. The complementary hypothesis is that facial expressions of emotion evolved in ways that exploit the perceptual capabilities of the extant perceptual system. This sensory-bias hypothesis is consistent with the insight that there is always a coevolution between sender and perceiver mechanisms in the evolution of communicative signals (e.g., Fridlund, 1994; Krebs & Dawkins, 1984). Signals should evolve to be conspicuous and easily detected by the general perceptual capabilities of the intended receiver. If the sensory-bias hypothesis is considered along with the threat-detector hypothesis, the possibility that perceptual factors contribute to a threat advantage do not inevitably conflict with the broader idea that facial threat may be easily detected. Moreover, with the sensory-bias hypothesis, we do not have to predict efficient search for angry-target faces because we do not have to assume a specialized, preattentive threat-detector (recall that the demonstration of efficient search is a necessary condition to establish a basic or attention-guiding feature). Rather, the demonstration of a relative advantage in search would be sufficient.

Empirically, the most straightforward way out of this dilemma appears to be to restore the original strategy by Hansen and Hansen (1988) and use facial stimuli that are derived from real facial expressions. Because the anger-superiority hypothesis is concerned with real facial expressions, ecological validity is less of a concern with photographs of faces than with schematic faces. Results obtained with stimuli derived from real facial expressions can therefore be better related to the general threat-advantage hypothesis than results obtained with schematic stimuli. Second, according to our analysis, perceptual differences are problematic mainly for schematic stimuli. Because these stimuli are artificial, it is always a concern whether differences in effects (e.g., in search slopes) are genuinely connected to the facial expression or are because of the geometrical forms that are used to portray these expressions. This problem is attenuated if the experimental stimuli are directly derived from real facial expressions.

Moreover, as we have substantiated before, whether differences in search slopes are ultimately because of differences in components of the facial expressions, the facial expressions as configurations or "wholes," or to the emotions evoked, are not a primary problem for the general threat-advantage hypothesis, but rather a secondary question. Even a purely perceptually founded search advantage for angry faces can be related to the threat-advantage hypothesis—given that it is established with stimuli that come as close as possible to real faces.

The use of schematic facial stimuli has been mostly motivated by noting the confound in the original Hansen and Hansen (1988) study revealed by follow-up studies (Hampton, Purcell, Bersine, Hansen, & Hansen, 1989; Purcell et al., 1996). We think that this response to the confound-problem is an overreaction. With respect to the anger-superiority hypothesis, a confound is a stimulus feature that discriminates between stimuli but is unrelated to the facial expression displayed. However, with available computer applications it is possible to construct stimulus pairs that differ only in their facial expression, but not in other aspects. One way would be to use faces that are totally digitally constructed. Another way (the procedure used in our present experiments) is to use digital image manipulation to eliminate all (or at least, most) differences between faces that are unrelated to the facial expression displayed.

With this procedure, it is even possible to examine a possible anger-superiority effect in a more meaningful way than before. For example, researchers in the past have been keen to keep the differences between the tested stimuli as minimal as possible, and many researchers have used an upwardly versus downwardly pointing curve as the sole feature that differentiates the happy from the angry face (e.g., Eastwood et al., 2001; White, 1995). However, when consulting the literature on facial expressions (e.g., Ekman & Friesen, 1976), it is not quite obvious that the drooped mouth corners are regarded typical anger expressions. Actually, Ekman and Friesen (1976) distinguish two variants for the lower face for anger—one with the lips pressed together, and one with an open mouth, sometimes with visible teeth. With several variants of the anger face, which additionally can vary in intensity and idiosyncratic factors, it is of interest to note which of the displays (if any) shows an advantage compared to a smiling face. Moreover, if a search asymmetry is obtained, it can be furthermore asked whether single components of the face, for example, the mouth, the eye-region, and so forth, are sufficient or necessary for the obtained effect.

A recent study by Williams et al. (2005) exemplifies the approach of testing realistic facial stimuli with special care to avoid perceptual confounds. In their Experiment 3, they tested photos of happy, surprised, fearful, and angry faces among emotionally neutral faces and found more efficient search with angry faces compared to the other faces. This threat-advantage supports the hypothesis that some aspect of the angry face can be used to guide attention to its position. However, search was inefficient even for the angry target face, and the results therefore fail to support the preattentive-detector hypothesis.

The present experiments are intended to be a further step in the indicated direction. We used faces from the Ekman and Friesen (1976) series, so that the results could be easily connected to the relevant literature. This is an advantage with respect to the preceding studies (including Williams et al., 2005), where it is more difficult to determine to what stimuli the obtained results could be generalized. A second difference to the Williams et al., (2005) study is that we used the classical search-asymmetry design instead of the constant distractor paradigm (cf. Eastwood et al., 2001). Although the constant distractor paradigm also has its virtues (see General Discussion), the search for angry face targets among happy face distractors has the advantage of contrasting stimuli that clearly lack indications of threat (happy faces) with ones that clearly show indications of threat (angry faces). Neutral faces are often perceived as mildly hostile (cf. Öhman et al., 2001), possibly explaining the nonefficient searches revealed by Williams et al. (2005).

The present study tested the angry face variant with an open mouth because we felt that this variant is more distinct from the happy face and should thus have a better chance to pop-out in visual search. Experiment 1 tests an expression that shows anger in the eye-region as well as in the mouth- region, but with the mouth only slightly opened. Experiment 2, which presents pictures from a different model, tests expressions with an even wider opened mouth. Experiment 3a eliminates some of the facial features indicating anger, in particular the eye region, in order to test a more reduced stimulus; Experiment 3b in turn tests faces in which everything but the eye region was removed. Experiment 4 was an additional control experiment, aimed to test whether the relative search asymmetry is due to the different emotions conveyed, or to perceptual differences between the faces. More precisely, this experiment employed a new variant of the inverted-face control condition which is argued to circumvent some problems with the standard inverted-face control condition.

Experiment 1

Method

Experiment 1 used two expressions (happy and angry) provided by the male Model J.J. (Ekman & Friesen, 1976). The happy face shows a Duchenne smile, that is, a smile accompanied by laugh wrinkles. The angry face has narrowed brows, a staring gaze, and slightly opened mouth with bared teeth. The stimuli were digitized as gray-scale images and subsequently modified using Paint Shop Pro. The number of different gray shades were reduced to 16 before further modifications. Next, the pictures were cropped, that is all pixels outside the face-outline were changed to white. In the next step, most pixels inside the face, except those belonging to the eyes, nose, mouth, and a shadow on the right side that resulted from the cheek, were also turned white (see Figure 2). In the happy face, the wrinkles near the eye were retained, as well as a furrow between the nose and the corner of the mouth on the left side of the face. To further eliminate accidental differences between the faces, the facial features (eyes with brows, nose, and mouth) of the angry face were cut out and pasted into a copy of the face outline of the happy face. As a result, the two faces differed perceptually, but the differences were restricted to the elements that distinguished the two facial expressions. Note that the nose is also perceptually different in the two faces. This is because the somewhat widened nostrils were apparently the result of the *zygomaticus major* activity, and thus not accidental.

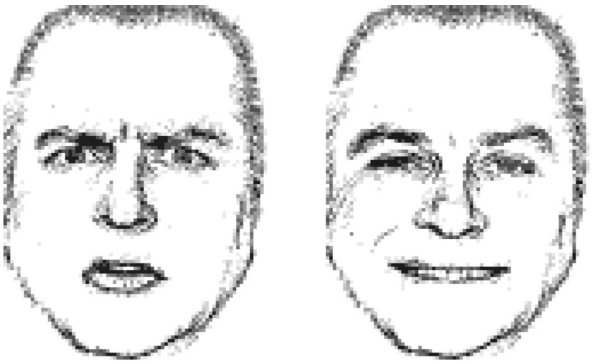

Figure 2. Stimulus pair used in Experiment 1.

Participants. Sixteen students were tested in individual 30-min sessions. Four additional participants were tested but discarded because they made more than 12% errors in at least one of the experimental conditions[1] (Here, and in the following experiments, an inclusion of all participants does not alter the pattern of results in theoretically significant ways). The mean age of the remaining 13 women and 3 men was 23 years ($SD = 6.0$). They participated voluntarily and received either €3 or course credits.

Design. The design was modeled after the classical visual search experiments by Treisman and colleagues (e.g., Treisman & Souther, 1985; Treisman & Gormican, 1988). Participants completed two blocks of trials. In each block, they were presented with display sizes of 1, 6, and 12 facial stimuli. Blocks differed with respect to the identity of the target and the identity of the distractors. In addition to set size, trials differed depending on whether a target was presented (target-present trials) or not (target-absent trials or blank trials). Each of the 12 conditions that resulted from the orthogonal combination of set size (1 vs. 6 vs. 12), target identity (angry vs. happy), and target presence (target present vs. target absent) was repeated 25 times. Dependent variables were RT and errors percentages. The order of blocks was balanced, as was the stimulus-response mapping (i.e., half of the participants responded with the left response key when the target was present and with the right response key when the target was absent, while for the other half of the participants, this mapping was reversed). Set sizes and target presence varied randomly within blocks.

Apparatus. The experiments were controlled by a personal computer equipped with a 80486 CPU, connected to a 15″ color monitor run with a resolution of 1024 × 768 pixels for stimulus presentations and a keyboard used to collect the manual responses.

Stimuli. The faces measured 2.2 cm (width) × 3.1 cm (height). In each trial, either 1, 6, or 12 facial stimuli were presented inside a monitor area of 9 cm × 9 cm (viewing distance was 120 cm). Individual faces were presented in a (invisible) 4 (horizontal) × 3 (vertical) matrix. The mean distance between the centers of the faces was 3.0 cm in the vertical axis and 3.9 cm in the horizontal axis. Average positions were altered by random displacement, separately computed for each position in each given trial. In particular, the average position of a stimulus was the center of a 3 × 3 grid, and the actual position of the stimulus was randomly chosen from the resulting 9 positions. The distance of adjacent positions in the 3 × 3 grid was 3 mm. The effect of this manipulation was a moderately irregular arrangement of the stimuli, intended to eliminate possible suprastimulus cues to the target's position (Duncan & Humphreys, 1989). In none of the displays, however, did the jitter lead to a spatial overlap between adjacent stimuli. Background color was white.

Procedure. Participants were fully informed about the structure of the experiment by written instructions. On request, the experimenter provided additional information. The experiment was divided into two main blocks, each of which was preceded by a screen that informed about the identity of the target in the following trials. For example, they were told that in the following block, the target was the happy face, that is, they should search for the happy face and indicate with the correct response key whether the happy face was present or absent. After this instruction screen, participants worked on 20 practice trials, which were followed by 150 experimental trials. The second block had exactly the same structure.

Each trial began with the 1,000 ms fixation-cross presentation, immediately followed by the faces display. The face display was on until a response was made, but a trial was aborted if no response was made within 6 seconds. If participants pressed the wrong key, a 100-ms tone was given as error feedback. The intertrial interval was 1,100 ms.

Results

Data treatment. For the analysis of RTs, RTs <200 ms and >3,000 ms, and errors, were excluded. Mean reaction times for each of the 12 experimental conditions were calculated. Grand means are reported in Figure 3. Because the predictions for preattentive processing concerns the slope of the RT - set size function, separate linear regressions with RT as the dependent variable and set size as the independent variable were computed for each participant, to obtain individual estimates of the two parameters b (slope) and a (intercept) of the linear regression equation $y = bx + a$. Further analysis was done using the regression parameters.

For the analysis of the errors, error scores were computed as the proportion of false responses. Analogous to the RT-analysis, the statistical tests were performed on the slope and intercept parameters obtained by linear regression.

Two critical results are predicted on the view that threat detection proceeds preattentively. First, the target present slope for angry targets should be near zero, and second, the target present slope for angry targets should be shallower than the target present slope for happy targets. The first criterion is conventionally tested assuming that a slope of 10 ms/display item instantiates a "flat" slope. The second prediction of a shallower slope with angry than with happy targets was tested with a directed t test to achieve maximum statistical power. ANOVAs were also conducted to provide an exhaustive analysis of the entire results pattern.

Slopes. Table 1 shows the results of the ANOVAs, Table 2 shows the mean slopes and intercepts. The ANOVA with the variables target presence (target present vs. absent) and target identity (angry vs. happy target) of the slopes for RTs revealed a significant main effect for target presence only, revealing shallower slopes for target present than for target absent trials (35 vs. 75 ms). In a similar vein, the ANOVA of the slopes for errors revealed that errors were less dependent on set size in target absent than in target present trials.

The predicted superiority for the angry target faces was revealed by a one-tailed t test for the difference between the angry versus happy target present trials, $t(15) = 1.78$, $p < .05$. However, the slope in the angry target present condition was considerably greater than 10 ms, thus indicating inefficient search (cf. Wolfe, 1998).

---

[1] Initially, we chose this rather stringent exclusion criterion because we feared that speed–accuracy trade-offs that are specific for certain conditions would level out differences in RTs between these conditions. Later it turned out that the exclusion did barely change the patterning of the RTs. However, because the experimental design assumes that method factors (order of conditions and response mapping) are balanced across participants, we present the data as we originally collected them.

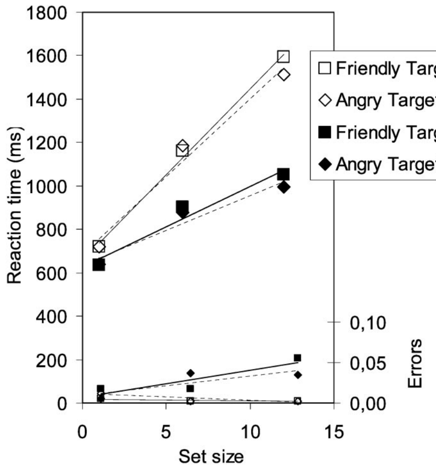

Figure 3. Mean correct RTs and error rates for each of the 12 conditions in Experiment 1. The figure also displays the linear trends obtained by linear regression analysis.

Intercepts. The ANOVA of the intercepts for RTs revealed a significant main effect for target presence, revealing faster RT in the target present than in the target absent condition (631 vs. 673 ms). No further effects emerged. The ANOVA of the error intercepts revealed no significant results.

Discussion

Experiment 1 revealed a weak but significant anger-superiority effect that is consistent with the general threat-advantage hypothesis. However, there was no evidence for a preattentive-threat detector: All slopes were relatively steep, and there was thus little evidence for preattentive processing of angry faces. Also, the target-presence to target-absence ratio of the RTs was almost perfectly 2:1 with both the angry and the happy targets, which is consistent with serial self-terminating search. That is, both the overall size of the slopes and the present to absent ratio was completely consistent with a serial search. To summarize, both target faces were apparently detected during a serial search, although the angry face was detected a little more efficiently than the happy face.

Table 1

Summary of the F-Values From the ANOVAs on the Search Slopes and the Intercepts, For the RT Data and the Error Data, Respectively, in Experiment 1

	Slope		Intercept	
	RTs	Errors	RTs	Errors
Target presence (1)	68.12	12.10	14.49	0.07
Target identity (2)	1.81	3.00	1.34	0.33
(1) × (2)	0.02	3.00	0.43	0.69

Note. For all *F*s, nominator *df* were 1 and denominator *df* were 15. Underlined values exceed the critical $F = 4.54$, $p < .05$.

Table 2

Summary of the Search Slopes and the Intercepts for the RT Data and the Error Data in Experiment 1

	Slope		Intercept	
	RTs	Errors	RTs	Errors
Friendly target present	37.3	0.004	626	0.007
Angry target present	31.8	0.002	637	0.008
Friendly target absent	77.9	0.000	659	0.013
Angry target absent	71.4	0.000	686	0.005

Experiment 2

Experiment 2 was intended to verify the results of Experiment 1 with a different model (see Figure 4). The new stimuli differed in several respects from the stimuli used in Experiment 1: (a) the model was a woman; (b) both mouth and eye region appeared subjectively more expressive to us; in particular, the shape of the mouths was more different for the two expressions, with a widely opened mouth in the anger display; and (c) we allowed for different shapes of the face because the somewhat elongated angry face was because of the lowered jaw.

Method

Participants. Twenty-three students were tested in individual 30-min sessions. Of these, three were discarded because they made more than 12% errors in at least one of the experimental conditions. The mean age of the remaining 16 women and 4 men was 25 years (*SD* = 9.3). They participated voluntarily and received either €3 or course credits. Part of the participants had also served in Experiment 1 (these were, however, also naïve as to the aim of the study). Practice effects in discrimination from prior participation should be negligible because these show very little transfer from one stimulus to another (cf. Wolfe, 1998).

Stimuli. The two expressions (happy and angry) were provided by the female Model F.M. (Ekman & Friesen, 1976). The picture manipulation procedure was very similar to that in Experiment 1, with the important exception being that the shape of the face outline was not the same for the two expressions (see Figure 4). In particular, the angry face was longer than the happy face (3.3 vs. 3.0 cm; both faces were 2.2 cm in width). Two

Figure 4. Stimulus pair used in Experiment 2.

further differences involved the background color, which was black in the present experiment, and the number of different gray shades (15) used.

Design, apparatus, and procedure. These were the same as for Experiment 1.

Results

Data treatment. Mean correct RTs above 200 ms and below 3,000 ms were analyzed as in Experiment 1.

Slopes. The grand means are presented in Figure 5. Table 3 shows the results of the ANOVAs, Table 4 shows the mean slopes and intercepts. The ANOVA of the slopes for RTs revealed significant main effects for target presence, revealing shallower slopes for target present than for target absent trials (14 vs. 32 ms), for target identity, revealing shallower slopes for angry than for happy targets (16 vs. 29 ms), and a significant two-way interaction, indicating that the crowd effect was stronger in the target absent trials (see Table 4). As predicted on the threat advantage hypothesis, the target present slope was significantly shallower for the angry than the happy face, $t(19) = 3.0$, $p < .05$. The ANOVA for the slopes of the errors revealed no significant results.

Intercepts. The ANOVA of the intercepts for RTs revealed significant main effects for target presence, with a lower intercept for target present than for target absent trials (570 vs. 602 ms), for target identity, revealing a higher intercept for happy than for angry targets (606 vs. 567 ms), and a significant two-way interaction (see Table 4). The ANOVA for the intercepts of the errors revealed no significant results.

Discussion

The threat superiority hypothesis was supported by Experiment 2. The target present slopes for the angry faces were significantly lower than for the happy faces, revealing a more efficient search for angry than for happy faces as targets. This time, the advantage was not only significant, but also numerically sizable.

The target present slope for the angry face was around 10 ms, revealing fairly efficient search for this target (Wolfe, 1998). This indicates that the angry face can be detected by searching for the presence of a preattentively available feature. For highly salient features, search is often also efficient in the target absent trials, whereas in the present case, the present-absent ratio was almost perfectly 1:2, indicating a serial search. However, given that the target-present slope tends to be 0, the present-absent ratio of 1:2 does not inevitably imply serial search because longer search times for the target absent trials may indicate that the observers were uncertain whether the target was really absent (Treisman & Souther, 1985).

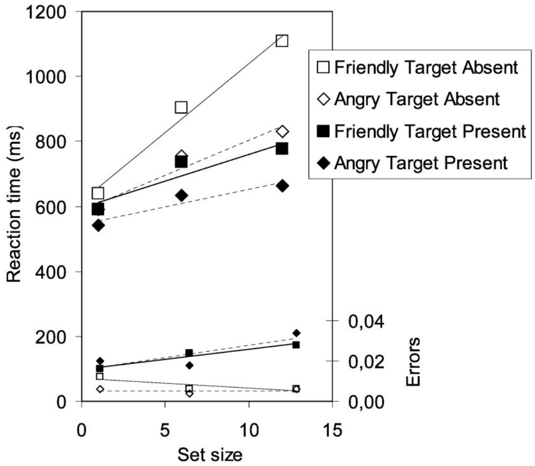

Figure 5. Mean correct RTs and error rates for Experiment 2 (see also the caption of Figure 3).

Table 3

Summary of the F Values From the ANOVAs on the Search Slopes and the Intercepts, for the RT Data and the Error Data, Respectively, in Experiment 2

	Slope		Intercept	
	RTs	Errors	RTs	Errors
Target presence (1)	60.81	4.13	20.50	2.21
Target identity (2)	36.72	0.14	11.81	0.32
(1) × (2)	22.72	0.14	5.03	0.16

Note. For all *F*s, nominator *df* were 1 and denominator *df* were 19. Underlined values exceed the critical $F = 4.38$, $p < .05$.

The slope for the target-present trials with happy faces (17 ms per item) was also not too steep, indicating an easy serial search. This result is relatively ambiguous with respect to the mode of processing. One possibility is that the two faces also differ on a second preattentively available feature, which is higher in the happy than in the angry face, but with only a small difference. Because of the relatively low signal-to-noise ratio, however, the presence of the happy face cannot be detected immediately. Rather, the display must be serially scanned using focused attention. However, although in difficult searches where each location must be checked separately for presence or absence of the target, in an easy search larger chunks of stimuli are attended to at once; this reduces the influence of the distractors and enhances the signal-to-noise ratio. Alternatively, finding the happy face target among angry face distractors may require the scanning of each face in turn, with the rejection of the angry face distractors being relatively easy.

The facial stimuli tested in Experiment 2 revealed a threat advantage. As faces are relatively complex stimuli, the question arises whether the threat advantage is because of the whole face stimulus or to more specific regions of the compared stimuli. The facial stimuli used in the present experiment differed on a number of characteristics, at least (a) the eyes region, (b) the mouth region,

Table 4

Summary of the Search Slopes and the Intercepts for the RT Data and the Error Data in Experiment 2

	Slope		Intercept	
	RTs	Errors	RTs	Errors
Friendly target present	16.5	0.001	596	0.016
Angry target present	10.8	0.002	545	0.016
Friendly target absent	42.4	0.000	616	0.006
Angry target absent	21.5	0.000	589	0.012

and (c) the overall shape of the face or its vertical extension. Based on the impression that the mouth region is an especially salient discriminating feature, Experiment 3a presented the mouth shapes from the stimuli in Experiment 2 in an otherwise empty face area derived from the happy face stimulus. Experiment 3b tests whether the eyes region supports attentional guidance.

Experiment 3a

Method

Participants. Eighteen students were tested in individual 30-min sessions. Of these, two were discarded because they made more than 12% errors in at least one of the experimental conditions. The mean age of the remaining 14 women and 2 men was 24 years (*SD* = 5.6). They participated voluntarily and received either €3 or course credits.

Stimuli. The stimuli were derived from that of Experiment 2 by eliminating all facial features in a copy of the happy face and pasting the mouth region from the happy face or from the angry face inside this copy (see Figure 6).

Design, apparatus, and procedure. These were the same as in Experiment 2. Although the stimuli were no longer happy and angry faces (because only the mouth was presented), it was decided to hold constant the instructions to be better able to compare results between experiments.

Results

Data treatment. Mean RTs were computed as before on the basis of correct responses below 200 ms and above 3,000 ms. The grand means are presented in Figure 7.

Slopes. Table 5 shows the results of the ANOVAs, Table 6 shows the mean slopes and intercepts. The ANOVA of the slopes for RTs revealed significant main effects for target presence, revealing shallower slopes for target present than for target absent trials (6 vs. 10 ms), for target identity, revealing shallower slopes for angry than for happy targets (4 vs. 12 ms), and a significant two-way interaction (see Table 6). As predicted, the target present slope was significantly smaller for the angry than for the happy target, $t(15) = 3.1$, $p < .05$. The analysis of the error rates revealed a significant main effect for target presence only, revealing less dependence of errors from set size with in target absent than present displays.

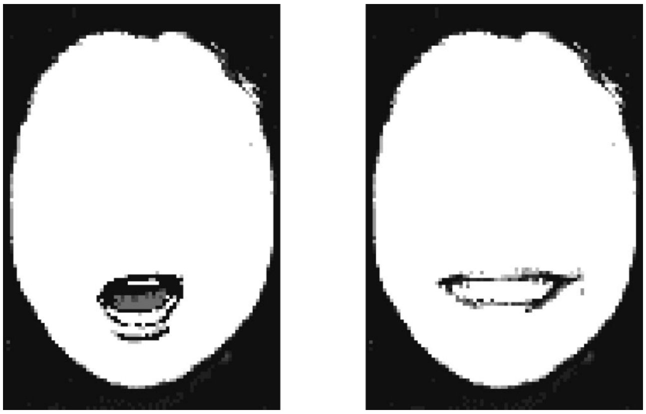

Figure 6. Stimulus pair used in Experiment 3a.

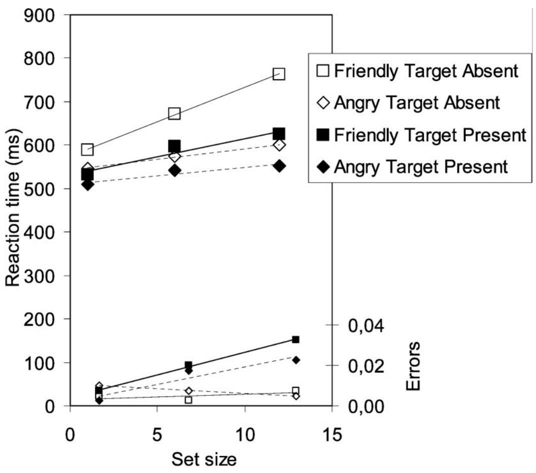

Figure 7. Mean correct RTs and error rates for Experiment 3 (see also the caption of Figure 3).

Intercepts. The ANOVA of the intercepts for RTs revealed significant main effects for target presence, due to a higher intercept RT for target absent versus present trials (520 ms vs. 558 ms), and target identity, revealing shorter RTs with angry than with happy target faces (526 vs. 553). The ANOVA of the intercept for errors rendered no significant differences.

Discussion

Experiment 3 reveals that the mouth region of the whole face stimuli used in Experiment 2 is sufficient to foster the anger-superiority effect. In fact, search rates were even faster in Experiment 3 compared to Experiment 2. This amounts to a face-inferiority effect being that single salient features are processed less efficiently in a whole face stimulus than in relative isolation (see Suzuki & Cavanagh, 1995), possibly because in the whole face stimulus, the entire stimulus configuration dominates the more basic features. If this interpretation is accepted, it would

Table 5

Summary of the F Values From the ANOVAs on the Search Slopes and the Intercepts, for the RT Data and the Error Data, respectively, in Experiment 3

	Slope		Intercept	
	RTs	Errors	RTs	Errors
Experiment 3a				
Target presence (1)	8.54	5.87	10.87	0.52
Target identity (2)	16.75	1.00	5.56	0.03
(1) × (2)	8.25	0.00	1.07	1.25
Experiment 3b				
Target presence (1)	34.84	11.67	2.00	0.02
Target identity (2)	0.00	4.20	0.02	0.20
(1) × (2)	0.70	1.40	12.11	0.39

Note. For all *F*s in Experiment 3a, nominator *df* were 1, and denominator *df* were 15; underlined values exceed the critical $F = 4.54$, $p < .05$. For all *F*s in Experiment 3b, nominator *df* were 1, and denominator *df* were 7; underlined values exceed the critical $F = 5.59$, $p < .05$.

Table 6
Summary of the Search Slopes and the Intercepts for the RT Data and the Error Data in Experiment 3a and 3b

	Slope		Intercept	
	RTs	Errors	RTs	Errors
Experiment 3a				
Friendly target present	8.4	0.002	531	0.006
Angry target present	3.9	0.001	509	0.001
Friendly target absent	15.8	0.000	574	0.004
Angry target absent	4.9	−0.001	542	0.011
Experiment 3b				
Friendly target present	52.3	0.003	701	0.023
Angry target present	54.8	0.009	663	0.023
Friendly target absent	106.1	−0.001	703	0.015
Angry target absent	103.8	0.000	732	0.026

mean that the reported effect is not tightly connected to face perception but is mainly supported by more simple perceptual features. However, other interpretations are possible. For example, one may argue that with three major differences between the faces in Experiment 2, participants used the relatively inferior strategy of relying on a less optimal feature contrast between the faces, or of checking more than one of the correlated features in at least some of the faces. From a more general perspective, however, the differences between Experiment 2 and 3 indicate that the specifics of search performance can be altered by using different stimulus configurations. This latter statement is in agreement with the heterogeneity of target present slopes in the other published studies on the threat-superiority effect (e.g., Eastwood et al., 2001; Fox et al., 2000; Hansen & Hansen, 1988; Nothdurft, 1993; White, 1995; Öhman et al., 2001), as well with a recent study by Tipples, Atkinson, and Young (2003), revealing that RTs are strongly influenced by the specific composition of component features of schematic faces.

Experiment 3b

Experiment 3b was a follow-up study to Experiment 3a, testing whether angry versus happy eyes and brows also show up with a search asymmetry.

Method

Participants. Four women and 4 men were tested in individual 30-min sessions. Because it became immediately clear that search for angry versus happy eyes is very inefficient, we did not exclude participants with more than 12% errors in at least one condition (otherwise 75% of the participants had to be excluded). The mean age was 26 years (*SD* = 4.8), and participation was voluntary and paid.

Stimuli. The stimuli were derived from that of Experiment 2 by inserting the eyes region of the happy and the angry faces, respectively, at corresponding positions within a white oval shape (same for both faces).

Design, apparatus, and procedure. These were the same as before.

Results

Data treatment. Mean correct RTs, computed from the correct responses below 200 ms and 3,000 ms, were analyzed as in the preceding experiments.

Slopes. Table 5 shows the results of the ANOVAs, Table 6 shows the mean slopes and intercepts. The ANOVA of the slopes for RTs revealed significant main effects for target presence only, revealing shallower slopes for target present than for target absent trials (54 vs. 105 ms). Planned *t* tests did not reveal any differences depending on the target, $t < 1$. The analysis of the error rates revealed a significant main effect for target presence only, revealing less dependence of errors from set size with in target absent than present displays.

Intercepts. The ANOVA of the intercepts for RTs revealed a significant interaction between target presence and target identity, revealing no effect for target presence with the happy target, but a lower intercept for target present trials than for target absent trials with the angry target (see Table 6). The ANOVA of the intercept for errors rendered no significant differences.

Discussion

The main result of Experiment 3b is that search for angry versus happy eyes and brows is very inefficient. Obviously, to discriminate these two stimuli, an effortful search has to be conducted, and efficient pop-out search cannot be used. One implication of this result is that the eyes (alone) are not responsible for the search performance with the whole-face stimuli in Experiment 2, where search efficiency was much better.

Experiments 4a and 4b

In Experiments 1–3 we have seen that photos of expressions of anger are searched for more efficiently than photos of expressions of happiness. This is in accordance with the hypothesis that facial threat is salient, promoting its fast detection. Of course, as such, the result is compatible with at least two accounts. One possibility is that the sensory system evolved some special capabilities to detect facial threat (the threat-detector hypothesis). On the other hand, it might be that evolved facial signals have adapted to the human sensory apparatus, exploiting its tendencies to respond easily to perceptually salient stimuli (the sensory-bias hypothesis). A comparison between Experiments 2 and 3 tends to support the sensory-bias hypothesis more than the threat-detector hypothesis because a relative search asymmetry was found to be more pronounced with stimuli less similar to facial expressions (Experiment 3) than with stimuli more similar to facial expressions (Experiment 2). However, this result alone would probably not fully convince a proponent of the threat-detector hypothesis.

Previous experiments have often used face-inversion as a manipulation to distinguish between the effects of the natural face (upright face) and its individual features (inverted face). It is usually reasoned that if the search asymmetry disappears with the inverted faces, it has been shown that the search asymmetry is an effect of the face as an integrated stimulus (i.e., a Gestalt)—which is altered by inversion—rather than by the perceptual effects of its isolated components—which are preserved by inversion.

It is interesting that although the idea is intuitively appealing, its rationale turns out to be somewhat more complicated on a closer look. Clearly, the upright face is presented in its ecologically familiar orientation, while an inverted face is rarely seen. Furthermore, there are tasks where stimulus inversion slows performance dramatically (e.g., when the identity of a face has to be determined, cf. Yin, 1969; Valentine, 1988). However, when it comes to the perception of emotion conveyed by the face, things are not that straightforward. This is because, in many instances, it is very easy to determine the emotion conveyed by an inverted face, as can be easily demonstrated (e.g., look at Figure 2 when holding the page upside-down; see also, McKelvie, 1995). For this reason, one might even predict that the search advantage for angry faces, if in fact mediated by the emotion conveyed, should not be changed by inverting the face.

Similar holds true for faces that are thatcherized. In thatcherized faces, only parts of the face, typically the eyes and the mouth, are inverted (Thompson, 1980). If these faces (as was 1st demonstrated with a picture of the smiling Margret Thatcher, then British prime minister, thus the name) are presented upright, the resulting impression is grotesque. However, if these faces are presented inverted, the grotesque impression is strongly reduced: The thatcherized Margret Thatcher appears quite normal and smiling. What is important for the present discussion is that the emotion conveyed is preserved by inverting a thatcherized face, but it is the specific unusual conjunction of the components of the face that is concealed by the inversion. Thus, although it is certainly correct to call the effect of inversion a distortion of normal holistic face processing, it is incorrect to conclude that face processing, including the processing of facial affect, is entirely abolished.

Because of these considerations, we decided not to use the inversion of the faces presented in Experiments 1–3 as a further means to elucidate the determinants of the relative search asymmetry. Rather, we used upright and inverted thatcherized variants of the stimuli used in Experiment 2 to test the effects of facial components versus whole face configurations. The reasoning was as follows.

In Experiment 4a, we tested inverted thatcherized happy and angry faces. In short, these stimuli were composed by creating a symmetrical oval shape approximately of the same size as the smiling face of the Experiment 3 stimuli, inserting a nose approximately at its normal position, and pasting in the inverted eyes and the mouth regions of the happy and the angry face from Experiment 2, respectively, at corresponding spatial positions. The mouth region of the happy face included the labial furrows that we regarded to be an integral part of a natural smile. Finally, these stimuli were inverted. As can be seen in Figure 8, these stimuli look like inverted, moderately odd pictures of facial expressions of happiness and anger. Because the emotional expressions are not unduly changed, we reasoned that the threat-detector hypothesis predicts a similar asymmetry as in Experiment 2. Analogously, the sensory-bias hypothesis also predicts a similar asymmetry, because the perceptual characteristics of the components of the face are not unchanged (except for their orientation). This was the control experiment.

The critical experiment was Experiment 4b. Here we presented the same stimuli as in Experiment 4, but this time in upright orientation. As can be seen in Figure 8, the impression is quite strange, with the meaning of the facial expression changed dramatically. The change is most severe with the happy expression, though the angry expression is altered as well. In fact, both facial configurations now appear negative, or even threatening. According to the threat-detection hypothesis, no asymmetry is predicted because the difference in these faces in their indication of threat (or nonthreat) is minimal. In contrast, the sensory-bias hypothesis still predicts the asymmetry with the thatcherized face based on the angry face being searched for more efficiently than the thatcherized face based on the happy face.

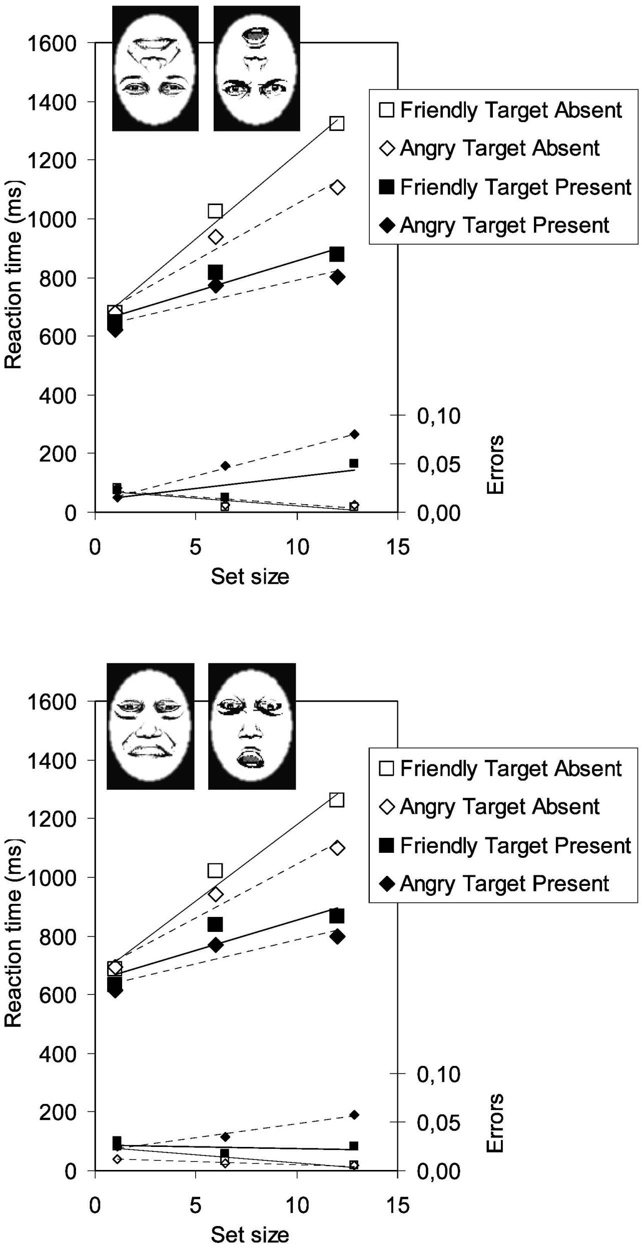

Figure 8. Mean correct RTs and error rates for Experiment 4 (see also the caption of Figure 3).

Method

Participants. Sixteen participants were tested in individual 50-min sessions. The mean age of the 8 women and 8 men was 26 years (*SD* = 6.9). They participated voluntarily and received either €5 or course credits.

Stimuli. The stimuli were derived from that of Experiment 2 as explained before (see introduction). Experiment 4a presented the thatcherized happy and angry faces inverted, whereas Experiment 4b presented them upright (see the miniatures in Figure 8a and b).

Design, apparatus, and procedure. The apparatus was the same as before. Experiments 4a and 4b, respectively, were procedurally identical to Experiments 2 and 3. Half of the participants completed Experiment 4a first and Experiment 4b second, while this order was reversed for the other half of the participants. Stimulus-response mapping, and the order of conditions (search for angry target face first and for friendly target face second or vice versa) were balanced over participants.

Results

Data treatment. This was the same as before, with one exception: To prevent selective dropout in one of the two experiments, and because the overall error rates were low (no participant had more than 7% errors on average), we did not exclude two participants who had each 1 condition with error rates exceeding 12%. The grand means are presented in Figure 8.

Slopes. A first analysis was done analogously to the preceding experiments to facilitate interexperiment comparisons. Table 7 shows the results of the ANOVAs; Table 8 shows the mean slopes and intercepts.

With the inverted faces (Experiment 4a), the ANOVA of the slopes for the RTs revealed significant main effects for target presence, revealing shallower slopes for target present than for target absent trials (18 vs. 48 ms), for target identity, revealing shallower slopes for angry than for happy targets (27 vs. 39 ms), and a significant two-way interaction (see Table 8). The target present slope was significantly smaller for the angry than for the happy target, $t(15) = 2.4$, $p < .05$, as was the target absent slope, $t(15) = 3.4$, $p < .01$. The analysis of the error rates revealed a significant main effect for target presence only, revealing less dependence of errors from set size with target absent than present displays (0.4% vs. −0.1% errors per item).

With the upright face (Experiment 4b), the ANOVA revealed the same pattern of results: Slopes were shallower with targets present rather than absent (18 vs. 44 ms); slopes were shallower with an angry rather than a happy target (26 vs. 36 ms), and there was a significant two-way interaction (see Table 8). This time, the target present slope was not significantly smaller for the angry than for the happy target, $t(15) = 1.3$, whereas the target absent slope was, $t(15) = 3.4$, $p < .01$.

Table 7
Summary of the F Values From the ANOVAs on the Search Slopes and the Intercepts, for the RT Data and the Error Data, Respectively, in Experiment 4a and 4b

	Slope		Intercept	
	RTs	Errors	RTs	Errors
Experiment 4a				
Target presence (1)	27.5	20.9	0.8	4.0
Target identity (2)	12.3	2.5	0.0	0.1
(1) × (2)	8.8	2.5	5.2	0.1
Experiment 4b				
Target presence (1)	33.1	6.0	5.9	0.2
Target identity (2)	6.3	0.0	0.0	0.0
(1) × (2)	6.1	2.9	5.7	0.1

Note. For all *F*s, nominator *df* were 1 and denominator df were 15. Underlined values exceed the critical $F = 4.54$, $p < .05$.

Table 8
Summary of the Search Slopes and the Intercepts for the RT Data and the Error Data in Experiment 4a and 4b

	Slope		Intercept	
	RTs	Errors	RTs	Errors
Experiment 4a				
Friendly target present	20.7	0.003	648.7	0.013
Angry target present	15.9	0.006	630.8	0.009
Friendly target absent	58.1	−0.001	641.3	0.024
Angry target absent	38.5	−0.001	665.8	0.023
Experiment 4b				
Friendly target present	20.7	0.001	648.1	0.028
Angry target present	16.3	0.003	624.5	0.019
Friendly target absent	52.1	0.000	659.7	0.014
Angry target absent	36.6	−0.002	679.3	0.025

In order to test whether the orientation of the thatcherized happy and angry faces had a significant effect on the RT-slopes, a conjoint analysis was performed on the data of Experiments 4a and 4b, with stimulus orientation as a third factor in the ANOVA. As before, the main effects for target presence, $F(1, 15) = 33.3$, $p < .001$ and for target identity, $F(1, 15) = 11.6$, $p < .01$, as well as their interaction, $F(1, 15) = 13.7$, $p < .01$, were significant. Importantly, neither the main effect for stimulus orientation, $F < 1$, nor any interaction involving the factor stimulus orientation (with target presence: $F = 1.4$; with target identity: $F < 1$; with target presence and target identity: $F < 1$) was significant.

Intercepts. The ANOVAs of the intercepts for RTs revealed a significant main effect for target presence with upright faces (Experiment 4b) because of a higher intercept RT for target absent versus present trials (636 ms vs. 669 ms). In Experiment 4a and 4b, there was a significant 2-way interaction, revealing that this effect was somewhat more pronounced for the angry faces (see Table 8). The ANOVA of the intercept for errors rendered no significant differences.

Discussion

Experiment 4 supported and qualified the anger-superiority hypothesis. The overall pattern of results with both inverted and upright thatcherized faces was similar as in Experiment 2 with upright and nonthatcherized faces. Importantly, there was no striking difference between the results pattern obtained with inverted versus upright thatcherized stimuli. We interpret this result as evidence that the emotion perceived in the pictures is not the most important variable for the relative search asymmetry in the present experiments: The difference in happiness or threat is preserved in the inverted thatcherized faces but is distorted in the upright thatcherized faces. In contrast, perceptual differences between the two pictures are mostly unaffected by the stimulus inversion, which is consistent with the assumption that perceptual differences in the components of the facial display can be used in visual search. Thus, the present experiment lent more support to the sensory-bias hypothesis than to the threat-detector account of the relative search asymmetry.

General Discussion

The threat-advantage hypothesis, which is based on ecological or evolutionary considerations, states that angry faces are processed especially efficiently. Presuming that this prediction is correct, the threat advantage can be accounted for by at least two specific hypotheses. First, the threat-detector hypothesis assumes that selection pressures changed the human visual system, which is nowadays endued with a specialized facial-threat detector. Second, the sensory bias hypothesis assumes that selection pressures molded facial signals so as to exploit extant capabilities of the human visual system, and that the most important signals (such as facial threat), became perceptually most salient. Four experiments tested these hypotheses with photographic reproductions of facial expressions of anger and happiness with special care for the elimination of perceptual artifacts that were unrelated to the facial expressions.

The threat-advantage hypothesis was supported in all four experiments, testifying a relative search asymmetry, with more efficient detection of angry among happy faces than vice versa. The effect was more pronounced in Experiment 2 than in Experiment 1. This difference may be partially because the stimulus pair used in Experiment 1 was of lower contrast and more detail than the Experiment 2 stimulus pair. Both factors would slow down search because the processing of details consumes time and the lower contrast would complicate the segregation of the stimulus into relevant discriminative regions. Arguably, contrast should be even greater in vivo faces than in the gray-scale images used in the present experiments because color perception adds contrast to luminance differences. Another contributing factor is indicated by Experiment 3, which showed that a salient difference in the mouth region alone can foster a threat advantage. Tentatively assuming that this region is, in fact, the most important difference between the faces (as far as search efficiency is concerned), this would explain the negligible advantage for angry faces in Experiment 1, where the difference in the mouth region was less obvious.

Experiments 3 and 4 were conducted to elucidate the obtained effects in more detail. Experiment 3 was primarily concerned with the question whether the eye regions or the mouth regions of the Experiment 2 stimulus pairs alone would foster the threat advantage. The results showed that the mouth region supports quick discrimination of the two faces, with a search advantage for the angry mouth relative to the happy mouth. In contrast, the eyes region does not support attentional guidance because search performance was very inefficient. This is an unanticipated, but theoretically significant result, insofar as several researchers, basing their conclusions on research using schematic faces or other artistic stimuli, have proposed that the eyes region is the most important indicator for facial threat (e.g., Aronoff, Barclay, & Stevenson, 1988; Lundquist, Esteves, and Öhman, 1999, 2004; and Öhman et al., 2001; Tipples, Atkinson, & Young, 2003). This discrepancy is a further indication that it is uncertain what can be learned about the perception of real faces by the use of schematic facial stimuli.

Experiment 3 also reveals that a reduced stimulus can be searched for more efficiently than a whole-face stimulus, indicating that the presence of other facial features may hamper performance. This indicates that the reported effect may not be due to face perception proper because otherwise we would have expected the results to be better with stimuli that are more similar to a complete face, relative to stimuli that are less similar.

Experiment 4 tested whether the results are better explained by differences in valence or threat (threat-detector hypothesis) or by perceptual differences (sensory-bias hypothesis). Using thatcherized faces, we found the pattern of results virtually uninfluenced by face orientation. Because inverting the thatcherized face has little impact on lower level perceptual features but a larger impact on the facial expression, this result is more favorable to the sensory bias-hypothesis than to the threat-detector hypothesis.

Consistent with the literature, we found no strong evidence for a preattentive processing of threat, as would have been indicated by a search asymmetry favoring angry faces as targets among happy face distractors. In fact, only Experiment 3a revealed nearly flat slopes, but simultaneously, the search asymmetry was scarcely pronounced (i.e., search for the happy face was relatively efficient as well). We have argued that the search-asymmetry design is well suited to test the preattentive-threat-detector hypothesis because with happy and angry faces, stimuli are contrasted that constitute extreme values on the dimension of facial threat. In light of these near optimal conditions, one might have expected a clear-cut search asymmetry with flat slopes for target present and target absent trials, and serial search for happy faces. In contrast, the entire data patterns in Experiments 1, 2, and 4, resemble a serial-search pattern better than a parallel-search pattern, with target-present slopes being approximately half the value of the target-absent slopes. Only in Experiment 3a does the data pattern indicate the preattentive detection of the "angry" face, with flat target-present and flat-target absent slopes.

With nearly all search slopes indicating inefficient rather than efficient search, we want to consider the possibility that angry faces are not searched for in parallel, but in serial. Most researchers interested in the attentional effects of emotional faces have implicitly or explicitly assumed a working hypothesis where faces and facial expressions act like visual primitives like colors or line orientation, and may be explained by models like FIT (e.g., Treisman & Gelade, 1980) or Guided Search 2.0 (Wolfe, 1994). This, however, is not very probable. Visual primitives are at a different level of organization from facial expressions. Facial expressions are composed of complex shapes, and even shape is not a visual primitive (though some very simple configurations might be, cf. Wolfe & Horowitz, 2004), but is composed of visual primitives (shape is not even a conjunction of features—all at the same spatial position—but a spatial concatenation of features, such as lines of different orientation). Models like FIT or Guided Search 2.0 explain performance as being based on the processing of visual primitives; they do not allow precise predictions about serial search for complex forms. Thus, future research might want to consider more adequate theories to elucidate serial search performance with angry and happy faces, such as the Attentional Engagement Theory proposed by Duncan and Humphreys (1989).

We have used the search asymmetry design in the present experiments, although recently, some researchers have claimed that the constant distractor paradigm is better suited to test attentional guidance (e.g., Eastwood et al., 2001). These authors note that in the search asymmetry design, with changing roles of two stimuli as target or distractors, the effects of guidance by the target and the speed of the serial scanning through the crowd of distractors are confounded. In fact many experiments (e.g., Fox et al., 2000), including the present ones, reveal a pronounced search asymmetry in the target absent trials, that is, crowds consisting of only happy faces are scanned faster than crowds consisting of only angry faces. This had led some researchers to propose that the asymmetry arises because it takes longer to search through an angry crowd, probably because attention dwelt longer on the angry than on the friendly stimuli (see also, Fox et al., 2000). In contrast, Eastwood et al. (2001) argue that with a constant crowd of neutral faces, all effects can be unambiguously attributed to the target (because crowd effects are held constant). Although this reasoning is formally correct, it faces several problems. First, what is the adequate neutral stimulus? Consider, for example, a search for an F versus an O among Es. Presumably, there would be a search asymmetry favoring the O, but what would this result indicate except that the O is more dissimilar to the E than the *F* (see also Duncan & Humphreys, 1989). Second, insofar as the target (happy or angry) changes randomly between trials (as in Eastwood et al., 2001), whether the 2 targets are searched for with equal priority it is left uncontrolled. In a blocked design, as exemplified in our experiments, only 1 single target is searched for in every trial of a block, eliminating the priority-problem. Third, the rejection of distractors during serial search is dependent on their physical appearance as well as on the identity of the targets (cf. Duncan & Humphreys, 1989). To illustrate, Williams et al. (2005, Figure 6), found different search rates among identical distractors, depending on the identity of the searched-for target (or more precisely, the target template). Moreover, if 1 stimulus (e.g., the happy face) is perceptually more similar to the neutral distractor than to the other stimulus (e.g., the angry face), it may more often be erroneously rejected as a distractor, giving a disadvantage in search efficiency.

To conclude, we have diagnosed several conceptual and methodological problems with the use of schematic facial stimuli that have been circumvented in the present study by using stimuli derived from photographs. We have also pointed out that the general threat-advantage hypothesis (that it is of evolutionary advantage to detect threatening faces easily) does not enforce the acceptance of the preattentive-threat-detector hypothesis (that facial threat is a preattentively available feature similar to the known basic perceptual features). The general idea is also compatible with a sensory-bias hypothesis (that facial expressions have adapted to the extant perceptual mechanisms of the intended observer). In our view, the evidence—including the present Experiments 3 and 4—favors the sensory bias hypothesis, which is also easier to bring in line with present theories of visual search. Finally, consistent with most of the previous research, facial expressions of anger were not searched for efficiently; thus, we obtained no strong evidence for a preattentive detection of threat.

References

Aronoff, J., Barclay, A. M., & Stevenson, L. A. (1988). The recognition of threatening face stimuli. *Journal of Personality & Social Psychology, 4,* 647–655.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96,* 433–458.

Eastwood, J. D., Smilek, D., & Merikle, P. M. (2001). Differential attentional guidance by unattended faces expressing positive and negative emotion. *Perception & Psychophysics, 63,* 1004–1013.

Ekman, P. (1972). Universals and cultural differences in facial expressions of emotion. In J. Cole (Ed.), *Nebraska Symposium on Motivation* (Vol. 19, pp. 207–283). Lincoln: University of Nebraska Press.

Ekman, P., & Friesen, W. V. (1976). *Pictures of facial affect.* Palo Alto, CA: Consulting Psychologists Press.

Enns, J. T., & Rensink, R. A. (1990). Influence of scene-based properties on visual search. *Science, 247,* 721–723.

Fenske, M. J., & Eastwood, J. D. (2003). Modulation of focused attention by faces expressing emotion: Evidence from flanker tasks. *Emotion, 3,* 327–343.

Fox, E., Lester, V., Russo, R., Bowles, R. J., Pichler, A., & Dutton, K. (2000). Facial expressions of emotion: Are angry faces detected more efficiently? *Cognition and Emotion, 14,* 61–92.

Fox, E., Russo, R., & Dutton, K. (2002). Attentional bias for threat: Evidence for delayed disengagement from emotional faces. *Cognition and Emotion, 16,* 355–379.

Fridlund, A. J. (1994). *Human facial expression: An evolutionary view*. San Diego, CA: Academic Press.

Hampton, C., Purcell, D. G., Bersine, L., Hansen, C. H., & Hansen, R. D. (1989). Probing "pop-out": Another look at the face-in-the-crowd effect. *Bulletin of the Psychonomic Society, 27,* 563–566.

Hansen, C. H., & Hansen, R. D. (1988). Finding the face in the crowd: An anger superiority effect. *Journal of Personality and Social Psychology, 54,* 917–924.

Horstmann, G. (2003). What do facial expressions convey: Feeling states, behavioral intentions, or action requests? *Emotion, 3,* 150–166.

Horstmann, G. (submitted). What do we learn from visual search experiments with emotional faces?

Horstmann, G., Borgstedt, K., & Heumann, M. (in press). Flanker effects with faces may depend on perceptual rather than emotional differences. *Emotion.*

Kleffner, D. A., & Ramachandran, V. S. (1992). On the perception of shape from shading. *Perception & Psychophysics, 52,* 18–36.

Krebs, J. R. & Dawkins, R. (1984). Animal signals: Mind reading and manipulation. In J. R. Krebs & N. B. Davies (Eds.), *Behavioural Ecology: An Evolutionary Approach* (2nd ed., pp. 380–402). Oxford, U.K.: Blackwell Scientific Publications.

LeDoux, J. (1998). *The emotional brain.* London: Phoenix.

Lundqvist, D., Esteves, F., & Öhman, A. (1999). The face of wrath: Critical features for conveying facial threat. *Cognition & Emotion, 13,* 691–711.

Lundqvist, D., Esteves, F., & Öhman, A. (2004). The face of wrath: The role of features and configurations in conveying facial threat. *Cognition & Emotion, 18,* 161–182.

McKelvie, S. J. (1995). Emotional expression in upside-down faces: Evidence for configurational and componential processing. *British Journal of Personality & Social Psychology, 4,* 647–655.

Mogg, K., & Bradley, B. P. (1999). Selective attention and anxiety: A cognitive-motivational persepective. In T. Dalgleish & M. Power (Eds.), *Handbook of cognition and emotion* (pp. 145–170). New York: Wiley.

Nothdurft, H. C. (1991). Texture segmentation and pop-out from orientation contrast. *Vision Research, 31,* 1073–1078.

Nothdurft, H. C. (1993). Faces and facial expressions do not pop out. *Perception, 22,* 1287–1298.

Öhman, A., Lundqvist, D., & Esteves, F. (2001). The face in the crowd revisited: A threat advantage with schematic stimuli. *Journal of Personality and Social Psychology, 80,* 381–396.

Pomerantz, J. R., Agrawal, A., Jewell, S. W., Jeong, M., Khan, H., & Lozano, S. C. (2003). Contour grouping inside and outside of facial contexts. *Acta Psychologica, 114,* 245–271.

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.

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.

Thompson, P. (1980). Margaret Thatcher: A new illusion. *Perception, 9,* 483–484.

Tipples, J., Atkinson, A. P., & Young, A. W. (2003). The eyebrow frown: A salient social signal. *Emotion, 2,* 288–296.

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.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12,* 97–136.

Valentine, T. (1988). Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology, 79,* 471–491.

White, M. (1995). Preattentive analysis of facial expressions of emotion. *Cognition and Emotion, 9,* 439–460.

Williams, M. A., Moss, S. A., Bradshaw, J. L., & Mattingley, J. B. (2005). Look at me, I'm smiling: Searching for threatening and non-threatening facial expressions. *Visual Cognition, 12,* 29–50.

Wolfe, J. M. (1994). Guided Search 2.0. A revised model of visual search. *Psychonomic Bulletin & Review, 1,* 202–238.

Wolfe, J. M. (1998). Visual search. In Pashler H. (Ed.), *Attention* (pp. 13–73). London, UK: UCL Press.

Wolfe, J. M. (2001). Asymmetries in visual search: An introduction. *Perception and Psychophysics, 63,* 381–389.

Wolfe, J. M., & Horowitz, T. S. (2005). What attributes guide the deployment of visual attention and how do they do it? *Nature Review Neuroscience, 5,* 495–501.

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

Received February 17, 2005
Revision received October 14, 2005
Accepted October 17, 2005 ■