

Visual search for singleton feature targets within and across feature dimensions

HERMANN J. MÜLLER

Birkbeck College, University of London, London, England

and

DIETER HELLER and JOHANNES ZIEGLER

Rheinisch-Westfälische Technische Hochschule, Aachen, Germany

Three experiments investigated visual search for singleton feature targets. The critical dimension on which the target differed from the nontargets was either known in advance or unknown—that is, the critical difference varied either *within* a dimension or *across* dimensions. Previous work (Treisman, 1988) had shown that, while the search reaction time (RT) functions were flat in both conditions, there was an intercept cost for the cross-dimension condition. Experiment 1 examined whether this cost would disappear when responses could be based on the detection of *any* (target–nontarget) difference in the display (by requiring a “heterogeneity/homogeneity” decision). The cost remained. This argues that pop-out requires (or involves) knowledge of the particular dimension in which an odd-one-out target differs from the nontargets; furthermore, that knowledge is acquired through the elimination of dimensions not containing a target. In Experiment 2, the subjects had to eliminate (or ignore) one potential source of difference in order to give a positive response (displays could contain a “noncritical” difference requiring a negative response). The result was a comparatively large cost in the within-dimension (positive) condition. This can be taken to indicate that pop-out as such does not make available information as to the particular feature value in which the target differs from the nontargets. Experiment 3 examined whether search priorities can be biased in accordance with advance knowledge of the likely source of difference. The subjects were found to have a high degree of top-down control over what particular dimension to assign priority of checking to. The implication of the results for models of visual search and selection are discussed.

Much recent research in visual perception has attempted to understand how humans detect a (usually known) target stimulus in a background of nontargets (distractors; see, e.g., Cave & Wolfe, 1990; Treisman & Gelade, 1980; Treisman & Sato, 1990; Wolfe, 1994). This work has used the analysis of reaction times (RTs) as a function of the number of items in the display (the display size) as its major tool. Two different patterns of RT/display size functions are typically observed. Which pattern is found depends on the nature of the target and nontarget stimuli. When the search involves a target that differs from nontargets in a single salient feature (e.g., a vertical line amongst horizontal lines), the search RT functions are flat, supporting the notion that targets are detected by a spatially parallel process. In contrast, when the target is defined by a conjunction of features, each of

which is separately present among the nontarget items (e.g., a red X among red Os and green Xs), the search RT functions are linearly increasing and the slope ratios between target absent and present responses are approximately 2:1. This pattern has been taken as indicative of a spatially serial search process, where search is exhaustive on absent trials and self-terminates upon detecting a target on present trials. However, there is no simple dichotomy between parallel and serial search RT functions. Rather, the functions form a continuum ranging from flat to linear with various search rates (slopes). To account for this family of functions, two-stage models have been proposed in which a rapid, but error-prone, spatially parallel analysis stage is followed by an accurate serial decision stage, where the parallel stage serves to compute the selection priorities for the subsequent stage by indexing the locations of likely targets (i.e., items that are dissimilar to the nontargets and/or are similar to the target; see, e.g., Cave & Wolfe, 1990; Treisman & Sato, 1990).

All accounts of visual search make explicit or implicit assumptions about the architecture of the underlying processing system. A review of the processing architectures proposed by the above two-stage models shows that they have difficulty accounting for pop-out in visual search for unknown feature targets (in particular, when

This research was supported by Science and Engineering Research Council Grant GR/H/54966 to H. Müller. It benefitted from remarks made by A. M. Jacobs and D. Reiffenrath (1993). The authors would like to thank A. Kramer, J. Theeuwes, J. M. Wolfe, and one anonymous reviewer for their extremely helpful comments on an earlier version of the paper. Correspondence concerning this article should be addressed to H. J. Müller, Department of Psychology, Birkbeck College (University of London), Malet Street, London WC1E 7HX, UK (e-mail: ubjta52@cu.bb.k.ac.uk).

the dimension on which the target differs from the nontargets cannot be predicted). Three experiments are presented addressing this difficulty, and modifications to the above accounts are suggested to accommodate the experimental findings.

Processing Architectures

One architectural assumption shared by virtually all accounts of visual search is that, at the lowest level of the system, there are dimension-specific input modules, such as orientation, brightness, color, size (spatial frequency), and motion. Each input module consists of a set of analyzer units at each spatial location that are tuned to particular values that the input can take on. These feature analyzers are spatiotopically organized, forming a feature map. In brief, input dimensions consist of sets of spatiotopically organized feature maps. (With some dimensions, in particular color, this is an acknowledged oversimplification.)

Treisman's influential feature integration theory (FIT; see, e.g., Treisman & Gelade, 1980) is concerned with the way in which the features of an object, that are represented in separate maps, are integrated (i.e., bound together) into a single object representation. Feature integration is achieved by focal attention, which operates on a master map of locations. The master map receives input from all the feature maps in the various modules. Its units code whether or not a stimulus is present at a location, rather than what that particular stimulus is (i.e., the location map units respond in all-or-none, rather than graded, fashion). Directing focal attention to a particular location effectively enables the gating of whatever features (in the various input modules) are active at that location into a temporary object representation system. These object representations are then matched with stored descriptions in the recognition network.

Revised versions of FIT (e.g., Treisman & Sato, 1990) also incorporate a mechanism of top-down control, which serves to exclude master map locations from attentional scanning that do not represent objects possessing target features (e.g., if the target is not green, locations with green objects should be excluded from the search). Master map locations are excluded from search by means of inhibition from the respective feature map(s). If inhibited below threshold, location units no longer signal the presence of a stimulus and will not be attended. Note that, according to revised FIT, the location map units no longer code location in an all-or-none manner, but rather in a graded fashion. In other words, the map of locations serves as a kind of saliency map that guides the allocation of focal attention—in particular, the higher the activation of a location unit, the more likely that location is to receive focal attention (see Treisman, 1988, p. 226).

Furthermore, according to FIT, if a target object differs from the nontarget objects in a single feature (a single-feature target), “distractors [nontargets] can be rejected in parallel.” As a result, the target pops out and “will then normally ‘call attention’ to its location”

(Treisman, 1988, p. 205). Nontargets can be rejected in parallel using inhibitory links from the nontarget feature map to the master map of locations. The only location to remain active (salient) would then be that corresponding to the target, which would thus receive focal attention automatically (without requiring serial scanning).

The revised FIT account copes well with feature (as well as conjunction) search when the target is known in advance. In this case, inhibition from the nontarget feature map(s) suppresses (at least to some extent) activity of all master map location units that represent nontargets. As a result, the target gains in salience and pops out. However, it is not clear how FIT would explain pop-out of an unknown target.

One possible answer is provided by the guided search (GS) model of Wolfe and his colleagues (Cave & Wolfe, 1990; Wolfe, 1994). GS assumes that each dimension-specific module (consisting of maps of analyzer units) computes its own saliency map. The saliency maps of all modules pass their activity to a master map of activations, which integrates (sums), for each location, the saliency signals from the separate modules. The (serial) allocation of focal attention is guided by the activity of units in the master map of activations (which is essentially equivalent the master map of locations in revised FIT).¹

The most interesting aspect of GS is that it provides a means of computing a dimension-specific saliency map in a bottom-up manner. This is done by similarity comparisons between each active unit in a particular feature map with all other active feature units at different locations within the module.² The overall difference (between a feature at one location relative to all other activated features) computed in these comparisons determines the activity of units in the module-specific saliency (location) map. In other words, the dimension-specific saliency map computes a signal (i.e., a distributed representation), which indicates loci of difference (or the absence of such loci) in the input to the module. A dimension-specific saliency unit thus only knows that there is a difference at one location relative to the others, but not *what* that difference is (i.e., on what particular feature value that difference exists). The difference signals are then integrated by units in the master map of activations, at which level they may compete for focal attention. Using this bottom-up mechanism, the presence of a feature target can be detected even if the identity of that target is not known in advance.

Note that GS also incorporates a top-down mechanism that comes into play when target identity is known (this is similar to revised FIT). In particular, top-down connections excite units coding features possessed by the target (the Cave & Wolfe, 1990, version of GS is agnostic as to whether the top-down modulation inhibits nontarget features or excites target features; in contrast, GS version 2.0 [Wolfe, 1994] expresses a theoretical preference for top-down excitation). The essential difference to FIT is that top-down excitation affects (in the first instance) the saliency signal computed by an input dimension.

In sum, GS can explain how a dimension-specific saliency signal is derived (in bottom-up manner) that may form the basis for pop-out even when the target identity is unknown. GS assumes that focal attention operates on the overall saliency map (which is computed by summing the dimension-specific saliency signals). In this respect, GS may present an advance on FIT. Note, though, that Treisman (1988) admits the possibility of “mutual inhibition between replicated features [in a non-target feature map] which might be fed down to the master-map locations that contain [those features]” (p. 226). This would serve a similar function as the bottom-up mechanism in GS, without, however, assuming the existence of dimension-specific saliency maps. One interesting implication of GS and extended FIT is that, in the case of an unknown single-feature target, the overall activation/location map would only know *that* (and *where*) there is an odd-one-out item, but not in *what* dimension (and what particular value in that dimension) the odd-one-out item differs from the others.

Visual Search for Unknown Feature Targets

Though clearly theoretically important, the question of how unknown feature targets are detected has received little attention. Treisman addressed this issue in a set of experiments reported in her Bartlett memorial lecture (Treisman, 1988, pp. 207–210). In her experiments, the subjects had to search for a target defined by a single-feature difference relative to the nontargets. The task differed from the typical feature search in that the subjects were unable to predict in what particular way the target would differ from the nontarget items. The critical difference could be of two types: within- and cross-dimensional. In the within-dimension condition, the nontarget items were all short black vertical lines, and the target could be a left-oriented diagonal line, a right-oriented diagonal line, or a horizontal line—that is, the critical difference was one within the orientation dimension. In the cross-dimension condition, the nontarget items were again short black vertical lines, and the target could be a long black vertical line, a short red vertical line, or a short black horizontal line—that is, the difference was one in the size, color, or orientation dimension. Treisman found that the search RT functions were flat (parallel) with both types of critical difference (within- and cross-dimensional). However, the *y*-intercepts of the functions were increased when the target was defined by a critical difference that could vary across functionally independent dimensions.

According to Treisman (1988), this RT cost can be explained as follows: “the ‘odd one out’ [item] pops out *within* a single, pre-specified dimensional module, but each different module may need to be separately checked to determine which of them contains it” (p. 207). Restated, detection of a (target–nontarget) difference in an unknown dimension requires the separate search of each dimension, where search is parallel within a dimension but serial across dimensions.

Note that theories that assume an overall master map of activations/locations to guide focal attention (such as GS and FIT) seem to have difficulty with this finding. According to these theories, interitem differences within dimensions are analyzed in parallel (bottom-up mechanisms) and passed, as differential activation values, to units in the master map, which sums these signals.³ Accordingly, pop-out (or, at least, focal attentional access to the odd-one-out item) should be based on the emergence of a high-activation unit at the level of the master map. Assuming that processing within dimensions is equally fast, the appearance of a high-activation unit should not depend on the particular dimension on which a critical featural (target–nontarget) difference exists. Yet Treisman’s data seem to show that master map (saliency) representations cannot form the basis of pop-out.

Note that the above difficulty only arises if processing within the various dimensions (size, color, orientation) is indeed equally efficient. This is not entirely clear from Treisman (1988), who reported search RT data averaged across the three different dimensions (though it appears that the within-dimension data were also averaged across data from searches within different dimensions). Any increased processing time within, say, the size or color dimensions relative to the orientation dimension would increase the average RT in the cross-dimension task (size, color, orientation), relative to the within-dimension task (orientation only). The experiments to be reported addressed this problem by making certain that processing for a target–nontarget difference was indeed equally efficient within all dimensions (see pilot experiment in the Method section of Experiment 1).

Assuming that processing efficiency did not differ between dimensions, there is a difficulty for GS and FIT. One possible way out would be to argue that the subjects in Treisman’s (1988) experiments attempted not only to detect an odd-one-out item (“that” response) but also to discern on what dimension the critical difference existed (“what” response). This would require additional processing if that difference could exist on various dimensions, such as serial checking of the relevant dimensions.

In summary, the problem for GS and FIT is that, it seems, the master map cannot form the basis of pop-out. An additional problem for FIT is that, as it stands, it has no explicit mechanism for dimension-specific pop-out (such as dimension-specific saliency maps that may be checked in serial order).

The present experiments were designed to further investigate visual search for unknown feature targets within and across dimensions. Experiment 1 examined whether master map representations can form the basis for pop-out when the subjects can respond positively upon detecting any difference among the display items (i.e., by requiring a homogeneity/heterogeneity decision rather than a target-absent/target-present decision). The answer was negative. Experiments 2 and 3 went on to examine in more detail the processes involved in within- and cross-dimension feature search—in particular,

whether, in the within-dimension case, pop-out is equivalent to knowledge of the particular feature value in which the target differs from the nontargets, and to what extent, in the cross-dimension case, the checking of the relevant dimensions can be top-down controlled.

EXPERIMENT 1

Can the distribution of saliency values over the master map of locations/activations (a representation assumed to exist by both GS and FIT) be used to produce pop-out, without the need to check dimension-specific saliency representations? Both GS and FIT predict that this should be possible in principle. In particular, if pop-out can arise from the overall saliency map, the (y -intercept) difference between within- and cross-dimension feature search should be abolished (or, at least, substantially reduced) if the subjects do not attempt to discern the source dimension of the target–nontarget difference. This prediction was tested in Experiment 1. The subjects were no longer asked to detect the presence/absence of a long, red, or horizontal (odd-one-out) target (as in Treisman's experiments), but rather to simply make a heterogeneity/homogeneity decision. That is, a positive response could be given upon detection of any heterogeneity (due to any odd-one-out target) in the display, regardless of the source of the heterogeneity.

Method

Subjects. Twelve subjects participated in Experiment 1 (unpaid volunteers from the subject panel of the RWTH Aachen Institute of Psychology). Five subjects were female, and 7 were male. Their ages ranged from 21 to 27 years. They all had normal or corrected-to-normal vision (as assessed by Titmus Vision Tester).

Apparatus. The subjects viewed the display from a distance of 50 cm, with their heads supported by a chin/forehead rest. The stimuli were presented within the central $6.9^\circ \times 6.9^\circ$ area (measured in degrees of visual angle) of a monochrome Atari SM 124 monitor (70-Hz refresh rate), which was controlled by an Atari 1024 ST computer. The subjects responded by pressing one of two designated (positive/negative) response keys on the Atari keyboard in front of them. The experiment was conducted under near-photopic ambient lighting conditions. The screen background was white. The stimulus/screen-background contrast ratio was 1:2.2 for gray stimuli and 1:31.8 for black stimuli. Stimulus size was 0.69° for small stimuli and 1.15° for large stimuli.

Stimuli. Displays consisted of 2, 4, 8, 16, or 24 items (the display size). On half the trials, the display contained one (of three possible) targets; on the other half, the display contained only nontarget items (target absent/target present). In all conditions, the nontarget items were small gray vertical lines. The target–nontarget difference could be one of two types: within- or cross-dimensional (see Figure 1). In the within-dimension condition, the target was a left-oriented diagonal, a right-oriented diagonal, or a horizontal line (small and grey in all cases). In the cross-dimension condition, the target could be a large gray vertical line, a small black vertical line, or a small gray right-oriented diagonal line. There was also a control condition in which only small gray right-oriented targets were presented (on positive trials) among small gray vertical nontargets—that is, the target was known in advance.

A pilot experiment examined whether search was equally efficient for each of these targets. Nontarget displays consisted of varying numbers of small gray vertical lines, and the targets were

(1) a small gray left-oriented diagonal, (2) a right-oriented diagonal, (3) a horizontal line, (4) a large gray vertical line, or (5) a small black vertical line. Note that each of the five targets was presented within separate blocks of trials (so target identity was perfectly predictable). Seven unpaid volunteers participated in the pilot experiment, searching for each of the five targets (in an order that was counterbalanced as well as possible). A repeated measures analysis of variance (ANOVA) of the search RTs failed to find any significant effects of target [$F(4,24) = 0.79$, n.s.] or display size [$F(4,24) = .28$, n.s.]. The search RT functions were flat (in all conditions) and did not differ (in terms of their y -intercepts) between the various targets. The average positive RTs were 492, 496, 502, 499, and 487 msec, respectively, for Targets 1–5 (see above). Averaged across the (three) within- and (three) cross-dimension conditions of Experiment 1, the RTs were 497 and 493 msec for within- and cross-dimension targets, respectively. Thus, the within- and cross-dimension conditions can be regarded as equal in terms of (single-target) search difficulty.

Procedure. All (12) subjects participated in the within-dimension, cross-dimension, and control conditions. These three types of condition were blocked and presented in counterbalanced order. Each experimental condition consisted of 300 trials (each with 150 target-present and target-absent trials). On target-present trials, each of the three targets in the within- and cross-dimension conditions was presented equally often (i.e., 50 trials per target). All types of trial were presented in randomized order. A central fixation cross was presented at the beginning of each trial, which was followed, after 500 msec, by the stimulus display. The display stayed on the screen until the subject responded. Half the subjects responded with the preferred hand and half with the nonpreferred hand.

The subjects were told to respond positively upon detection of any display heterogeneity and negatively otherwise. They were encouraged to keep their error rates below 5%. Erroneous responses generated immediate feedback to the subjects by a brief computer-generated beep; the subjects received further accuracy feedback during breaks (after each 40 trials), when their error percentage was displayed on the computer monitor.

Results

Of the positive responses, only those to right-oriented targets were analyzed, because this was the only target that was the same in the within-dimension, cross-dimension, and control conditions. (A preliminary analysis of the positive RTs to the three cross-dimension target alternatives failed to reveal any significant effect due to target dimension.) Figure 2A presents the correct (group) mean RTs and Figure 2B presents the error rates, as a function of display size, separately for positive and negative responses. The left-hand panel shows the functions for the control condition (target known); the middle and right-hand panels show the functions for the within- and cross-dimension conditions, respectively. The findings can be summarized as follows: The search RT functions were flat in all conditions. The RTs showed little difference between the control and within-dimension conditions, but they were significantly raised in the cross-dimension condition (see RT analysis below). The error rates were low overall (<2%) and tended to be higher with longer RTs, arguing against speed–accuracy tradeoffs (see error analysis below).

RT analysis. The RT data were examined using a repeated measures ANOVA with main terms for condition (control, within dimension, across dimension), response

(negative, positive), and display size. There was a highly significant main effect of condition [$F(2,22) = 24.95$, $p < .001$]. The main effects of response [$F(1,11) = 2.61$] and display size [$F(4,44) = 0.96$] were not significant, nor were there any significant interactions.

The search RT functions were flat (there was no display size effect), indicating that search was parallel in all cases. There was some (nonsignificant) tendency for negative responses to be slower than positive responses (in particular, in the within-dimension condition), as is typically found in feature search.

The search RT functions for the (target-known) control condition were very similar to those in the within-dimension condition (in which only the target dimension was known). A separate ANOVA comparing the search RTs in the two conditions revealed no significant ef-

fects. [Note, however, that there was a marginally significant control/within-dimension \times negative/positive response interaction, $F(1,11) = 4.02$, $.075 > p > .05$, due to a tendency for positive RTs to be faster in the within-dimension condition.]

In contrast, the cross-dimension condition (in which the dimension on which the target differed from the nontargets could not be predicted) exhibited a cost of some 55–60 msec, on average, relative to the control and within-dimension conditions. This cost accounts for the significant effect of condition.

Error analysis. An ANOVA of the arcsine-transformed error rates revealed significant main effects of control/within/cross condition [$F(2,22) = 8.05$, $p < .005$] and negative/positive trial [i.e., false alarms/misses; $F(1,11) = 16.5$, $p < .005$]. Furthermore, there




Figure 1. Pattern of activity across feature analyzers in different dimensions (size, color, orientation) produced by a small black horizontal line (the target) amongst small black vertical lines. A critical difference singling out the target from the nontargets exists only in the orientation dimension. In the within-dimension condition, the target was, unpredictably, a left-tilted, right-tilted, or horizontal line. In the cross-dimension condition, the target was, unpredictably, a difference in terms of stimulus orientation, color, or size.

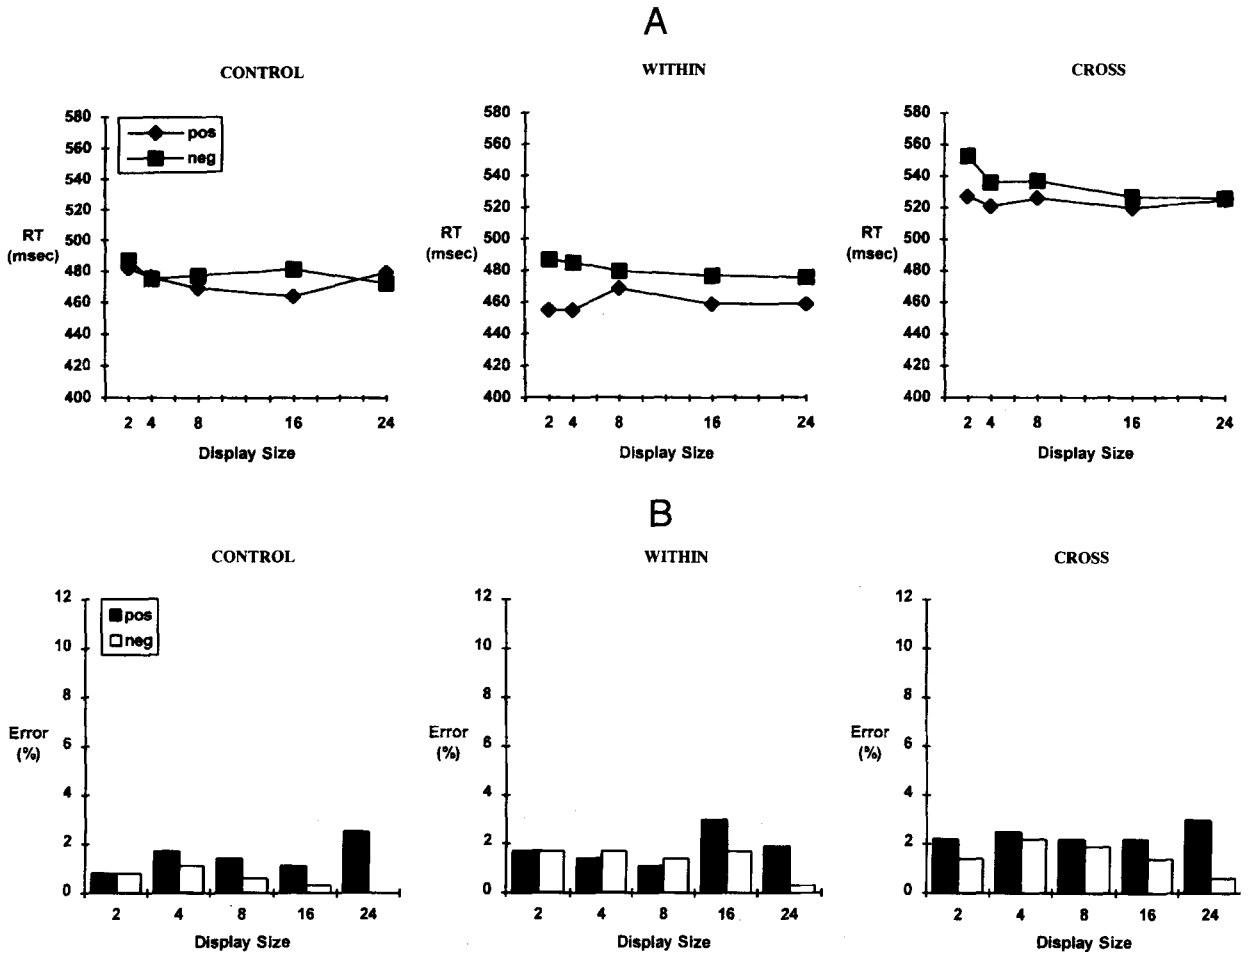


Figure 2. Positive and negative search RT/display size functions (Panel A) and error rates (Panel B) in the control, within-dimension, and cross-dimension conditions of Experiment 1.

was a significant negative/positive trial \times display size interaction [$F(4,44) = 3.56, p < .025$]. The error rates were slightly lower in the control condition [.010, on average; control/within, $F(1,11) = 7.73, p < .025$; control/cross, $F(1,11) = 11.10, p < .01$] and did not differ between the within and cross conditions [.016 and .020, respectively; $F(1,11) = 2.49, n.s.$]. Misses (positive trials, .019) were made more frequently than were false alarms (negative trials, .011), and they increased/decreased with increasing display size.

Discussion

The search RT functions for the (target-known) control condition were not different from those in the within-dimension condition. This is in contrast to Treisman (1988), who found a small cost (of around 20 msec) when the target was not known in advance. Recall that, unlike Treisman's, the present experiment required the subjects to make only a heterogeneous/homogeneous decision—that is, they could respond positively upon detecting “that” there was a (target–nontarget) difference, without any need to know “what” the difference

was. The absence of a difference between the control and within-dimension conditions thus indicates that, within a dimension, a successful “that” decision (i.e., responding on the basis that there is a critical target–nontarget difference) is feasible. Presumably, within-dimension pop-out is based on a single signal that integrates the output of dimension-specific feature analyzers, as is assumed explicitly in GS (see also Treisman, 1988, p. 207, who stated that “the ‘odd one out’ pops out *within* a dimension”).

However, this does not apply to the cross-dimension condition, which showed a significant cost (of some 55–60 msec) when the dimension on which the target differed from the nontargets could not be predicted. This cost is similar to, though less marked than, the 100-msec cost found by Treisman (1988). Thus, it appears that, across dimensions, a heterogeneity/homogeneity decision can be made only if information as to the source (dimension) of the critical difference is available. This finding is difficult to accommodate by models that assume that decisions can be based on the detection of any differential activity in a master map of activation/

location units that simply sum the parallel output from dimension-specific saliency units (or feature analyzers).

One radical response to this difficulty might be to suggest that the master map plays no role at all in singleton feature search. Rather, feature search relies solely on dimension-specific mechanisms that, in the cross-dimension case, “need to be separately checked to determine which of [the dimensions] contains [the target]” (Treisman, 1988, p. 207). The alternative is that the master map does play a role, using some procedure to obtain information as to the dimension of critical difference.

For instance, Procedure 1 might be to permit only one of the three relevant dimensions at a time to pass activation to the master map. If a critical activation difference is detected at the level of the master map, the difference can be attributed to the searched dimension. If not, that dimension is eliminated, and checking continues until a definitive decision can be made. In brief, the procedure would involve top-down modulation, and serial self-terminating checking of input dimensions (similar to Treisman’s, 1988, separate checking account). An alternative, Procedure 2, would involve the decomposition of the master map saliency signal by serial inhibition of dimensions (preventing them from activating the master map). If the master map signal is eliminated by the inhibition of one of the three dimensions, the difference signal must originate from the inhibited dimension. If it persists, the difference signal must originate from one of the two remaining dimensions. So only one further dimension would need to be eliminated to discern the source dimension of the master map saliency signal. Both procedures assume that the master map saliency signal can be used for response only when information as to the source (dimension) of the critical difference becomes available. (This information is provided in advance in the within-dimension condition dimension.)

These possibilities were examined by analyzing the standard deviations associated with each subject’s mean RTs. The RT variability was significantly increased in the cross-dimension condition (114 msec), relative to the control and within-dimension conditions [95 and 91 msec, respectively; main effect of control/within/cross, $F(2,22) = 5.58, p < .025$; control/cross, $F(1,11) = 8.15, p < .025$; within/cross, $F(1,11) = 6.06, p < .05$]. Furthermore, negative RTs (106 msec) exhibited greater variability than did positive RTs [93 msec; $F(1,11) = 8.05, p < .025$]. The increased variability in the cross-dimension condition, relative to the control and within-dimension conditions, was more marked for negative RTs than for positive RTs [control/within/cross condition \times negative/positive response interaction, $F(2,22) = 4.11, p < .05$]. This pattern of effects appears to be consistent with serial checking of dimensions.

To further examine the checking procedure used, the increase in variability in the cross-dimension (positive and negative) condition was compared with that expected on the assumption that the checking of the three relevant dimensions is serial (in random order) and self-terminating (see Schneider & Shiffrin, 1977, for the deri-

vation of variability predictions of serial self-terminating search models). The mean time required to check a dimension was assumed to be half the difference between the mean positive RTs in the cross- and within-dimension conditions (since two checks, on average, are necessary to find the target dimension); furthermore, the variability of the time per check was assumed to equal to that of the mean RT in the within-dimension condition. The increase in variability in the cross-dimension condition was found to be only about half that expected on the assumption of serial self-terminating checking of the three relevant dimensions (Procedure 1 above). The increase fits best with the assumption that only one or two of the three relevant dimensions need to be checked, consistent with Procedure 2. Other ways in which the critical dimension might be discerned in the cross-dimension condition (including parallel checking) are considered in the General Discussion. For the time being, it is sufficient to assume that this involves some form of elimination process.

This assumption presupposes that the RT cost in the cross-dimension condition, relative to the within-dimension condition, arises at an early perceptual stage of processing (see also Treisman, 1988) rather than a later response selection stage. The number of target stimuli mapped to the positive response was the same in both conditions, equating the demands placed on response selection (this also applies to Experiments 2 and 3 below). A response selection account of the cross-dimension cost is therefore unlikely to be true. Of course, it cannot be definitively ruled out that, for some reason, mapping of three cross-dimension stimuli to response may be more difficult than mapping three within-dimension stimuli. However, the error data of Experiment 1 showed no evidence of this being the case.

EXPERIMENT 2

The above elimination account predicts that, when there is a possibility of displays containing a single irrelevant heterogeneity that needs to be eliminated (i.e., a single interitem difference that does not signal the presence of a target), feature search across dimensions should be less affected than feature search within a dimension.

In the within-(orientation-)dimension case, an irrelevant difference could be, for example, a horizontal line among vertical lines, where the horizontal line is not to receive a positive response (as in Experiment 1) but rather a negative response in the same way as a homogeneous display of vertical lines; positive responses are to be made only to left- and right-oriented line targets. In the cross-dimension case, an irrelevant difference could be one in terms of size, such as a large vertical line among small vertical lines, where the large vertical line is to be responded to negatively in the same way as a uniform display of vertical lines; positive responses are reserved for orientation (left-oriented lines) and color (black line) targets.

Assuming that some form of elimination of dimensions is the normal procedure in cross-dimension feature search, then redefining one dimension of variation to be irrelevant (requiring a negative response) should have little effect on the search, because that dimension may either be ignored or eliminated with priority. In fact, there may even be a facilitatory effect on the search, which may, however, be outweighed by the generally increased difficulty of the task. The effect on performance should be more pronounced for within-dimension search—in particular, if one assumes that, normally, within-dimension pop-out is based on a dimension-specific saliency signal. The reason is that successful task performance requires not only the detection of any within-dimension heterogeneity but also the time-consuming determination of the featural source of the heterogeneity.

Dimension-specific saliency maps are explicitly assumed by GS, but not FIT. Recall that, in FIT, a dimension consists of a set of maps each tuned to a particular feature value. There may also be pooled (map) activity units, which sum the activation within each feature map (Treisman & Gormican, 1988). If active, such a unit would signal the presence of a stimulus (or stimuli) of, say, a particular orientation. Furthermore, the master map of locations receives direct input from each feature map within a dimension. The master map of locations would, thus, know *that* (and *where*) there is a difference in activation (signaling the presence of a target), but not *what* constitutes the difference. To find out *what* the difference is may require time-consuming checking of the pooled map activity units. Thus, according to both GS and FIT, the requirement to eliminate one particular heterogeneity in, say, orientation should have a comparatively marked effect on within-dimension search because of the need to determine the source of the heterogeneity (e.g., by decomposing an integrated saliency signal either at master map or dimension level). This prediction was tested in Experiment 2.

Note that a positive finding would be new. Treisman (1988) observed that heterogeneity within a dimension was more harmful than heterogeneity across dimensions. However, unlike in the present Experiment 2, heterogeneity in her experiment meant that there was not just a single item that differed from all the others (which were themselves the same) but rather that the nontargets were themselves heterogeneous. (See also Duncan & Humphreys, 1989, 1992, and Humphreys & Müller, 1993.)

Method

Subjects. Twelve unpaid volunteers participated in Experiment 2 (none of them had taken part in Experiment 1). Six subjects were female, and 6 were male. Their ages ranged from 23 to 29 years. They all had normal or corrected-to-normal vision.

Stimuli. The stimuli were the same as those in Experiment 1. However, there was a difference in the response required by some of the stimuli. In particular, in the cross-dimension condition, a positive response was required only when the display contained either a small black vertical line or a small gray right-oriented diag-

onal line; but a negative response was required when there was a large gray vertical line, as well as when there were only small gray vertical lines. In the within-dimension condition, a positive response was required when the display contained either a left-oriented or a right-oriented diagonal line; a negative response was required, when there was a horizontal line, as well as when there were only vertical lines (all stimuli were small and gray). The control (target-known) condition was the same as that in Experiment 1.

Procedure. The subjects were told to give a positive response only to relevant targets. They were not explicitly instructed to ignore irrelevant targets. In all other respects, the procedure in Experiment 2 was the same as that in Experiment 1.

Results

Again, only positive responses to right-oriented targets were analyzed (because this was the only target that was the same in the within-dimension, cross-dimension, and control conditions). A preliminary analysis of the positive RTs to the two cross-dimension target alternatives failed to reveal any significant effect due to target dimension. The negative responses analyzed were those to homogeneous displays with vertical lines only. Figures 3A and 3B present the correct group mean RTs and the error rates, respectively, as a function of display size, separately for positive and negative responses. The left-hand panel shows the function for the control condition (target known); the middle and right-hand panels show the functions for the within- and cross-dimension conditions, respectively. In brief, the results were as follows: Positive RTs were slower than negative RTs in the within- and cross-dimension conditions, which is unusual for visual search experiments. Positive RTs were slower in the within-dimension condition than those in the control condition (in contrast with the results of Experiment 1). Positive RTs were also affected in the cross-dimension condition, but less markedly so than were those in the within-dimension condition (see RT analysis below). The error rates were low overall (<2%) and tended to be higher with longer RTs (arguing against speed-accuracy tradeoffs) (see error analysis below).

RT analysis. The data were examined using a repeated measures ANOVA with main terms for condition (control, within dimension, across dimension), response (negative, positive), and display size. There were significant main effects of condition [$F(2,22) = 7.90, p < .005$] and response [$F(1,11) = 6.81, p < .025$]. The main effects of display size were not significant [$F(4,44) = 0.71$]. Furthermore, there was a significant control/within/cross condition \times negative/positive response interaction [$F(2,22) = 7.96, p < .005$].

The search RT functions were flat (there was no display size effect), indicating that search was parallel in all cases, as in Experiment 1. However, unlike the results of Experiment 1, in the within- and cross-dimension conditions, positive RTs were slower than negative RTs (whereas the control condition showed no difference, as in Experiment 1) (control/within/cross condition \times negative/positive response interaction). This reversal tended to be particularly marked in the within-dimension condition [an ANOVA excluding the control condition

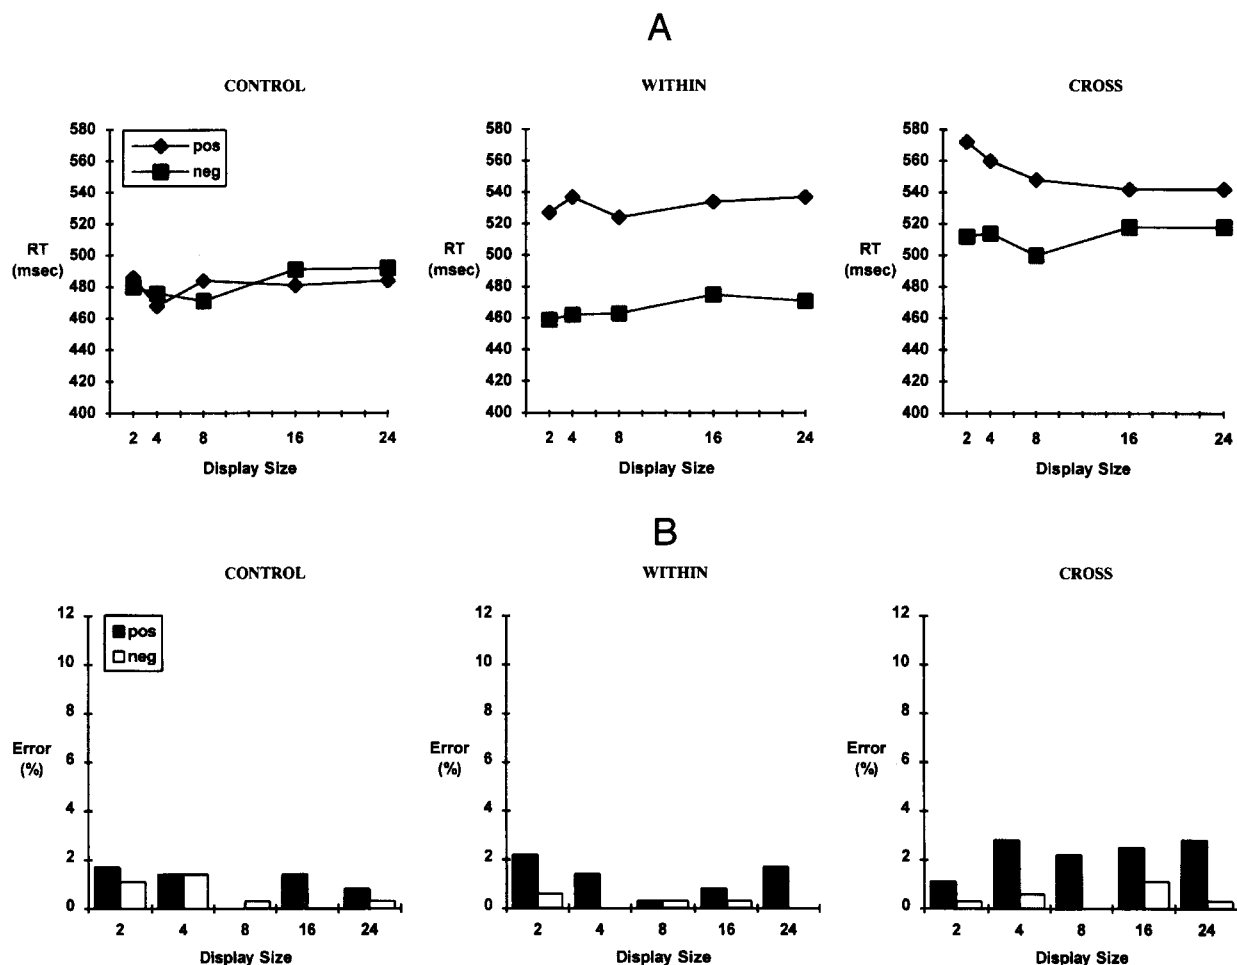


Figure 3. Positive and negative search RT/display size functions (Panel A) and error rates (Panel B) in the control, within-dimension, and cross-dimension conditions of Experiment 2.

revealed a marginally significant within/cross condition \times negative/positive response interaction, $F(1,11) = 4.17, .075 > p > .05$.

Also unlike the results of Experiment 1, positive RTs were slower in the within-dimension condition than those in the (target-known) control condition, whereas there was little difference between the negative RTs [an ANOVA comparing the two conditions revealed a significant condition \times response interaction, $F(1,11) = 11.16, p < .01$].

Moreover, there was a cost of some 35 msec, on average, when the dimension on which the target differed from the nontargets could not be predicted (i.e., for the cross-dimension condition relative to the within-dimension condition). An ANOVA that excluded the control condition revealed this cost to be significant [main effect of within/cross condition, $F(1,11) = 6.65, p < .025$]. Note that the cost tended to be more marked with negative responses (48-msec difference) than with positive responses (21-msec difference) [within/cross

condition \times negative/positive response interaction, $F(1,11) = 4.17, .075 > p > .05$].

The effect of the changed task (elimination requirement) in Experiment 2 was further examined by comparing the within- and cross-dimension conditions in Experiments 1 and 2 (separate ANOVAs for positive and negative RTs, with experiment [Experiments 1 and 2] as between-subject variable). The condition to be most affected in comparison with Experiment 1 was the within dimension positive response, which showed an RT increase of 73 msec. The cross-dimension positive response was much less affected, showing an increase of only 29 msec. The ANOVA of the positive RTs revealed a significant experiment \times condition (within dimension, across dimension) interaction [$F(1,22) = 5.55, p < .05$]. In contrast, the negative responses did not differ significantly between Experiments 1 and 2. The ANOVA of the negative RTs revealed no significant main effect of experiment [$F(1,22) = 0.51$] and no significant experiment \times condition (within dimension, across dimension)

interaction [$F(1,22) = 0.13$]. Since the negative responses were not affected by the changed task demands (with respect to Experiment 1), it is preferable to think of the reversed order of positive RTs (slower) and negative RTs (faster) in terms of *slow* positive, rather than *fast* negative, responses.

The standard deviations of each subject's correct mean RTs were also examined using an ANOVA. As in Experiment 1, the variability was significantly greater for the cross-dimension condition (111 msec) than for the control and within-dimension conditions (91 and 87 msec, respectively) [main effect of control/within/cross, $F(2,22) = 6.74, p < .01$; control/cross, $F(1,11) = 5.58, p < .05$; within/cross, $F(1,11) = 15.67, p < .005$]. Note that the standard deviations were not increased for the within-dimension condition relative to those for the control condition [$F(1,11) = .04, n.s.$].

Error analysis. An ANOVA of the arcsine-transformed error rates revealed significant main effects of control/within/cross condition [$F(2,22) = 4.05, p < .05$] and negative/positive trial [i.e., false alarms/misses; $F(1,11) = 8.07, p < .025$]. Furthermore, there was a significant control/within/cross condition \times negative/positive trial [$F(2,22) = 3.57, p < .05$]. The error rates did not differ between the control and within-dimension conditions [.008 and .007, respectively; $F(1,11) = 0.33, n.s.$], but they were slightly raised in the cross-dimension condition [.014; control/cross, $F(1,11) = 3.94, p < .075$; within/cross, $F(1,11) = 7.02, p < .025$]. Misses (positive trials, .015) were more frequent than were false alarms (negative trials, .004), particularly in the cross-dimension condition (false alarms/misses, .004/.023).

Discussion

The disproportionately large RT effect on the within-dimension (positive) condition is consistent with the expectation. Normally, within-dimension pop-out may be based on a dimension-specific saliency signal (communicated via the master map). This signal only indicates that there is a difference, but not what the difference is. If there is no difference, no further processing is required, so negative decisions are not affected. But if there is a difference, and one specific source of difference must be ruled out, further processing is necessary, producing costs in the time required to respond positively (see Sagi & Julesz, 1985a, 1985b). (Interestingly, the variability of each subject's RTs was not increased in the within-dimension condition relative to that in the control condition, suggesting, perhaps, that any further processing required in the within-dimension condition is not serial in nature.)

In contrast, with cross-dimension search, decision making always depends on information as to the dimension of the critical difference (i.e., it always involves some form of elimination), so that the explicit elimination requirement in Experiment 2 produced comparatively little added cost. (The increased variability of each subject's RTs in the cross-dimension condition relative to that in the control and within-dimension conditions

supports the idea that the former condition involves some form of serial elimination of input dimensions.)

Nevertheless, slow positive responses (i.e., positive responses being slower than negative responses) were also found in the cross-dimension condition of Experiment 2. This suggests that the subjects did not simply ignore the irrelevant dimension (size) and respond negatively (by default) if there was no difference on the other (form and color) dimensions. Rather, it appears that the subjects checked the irrelevant dimension. Not ignoring the irrelevant dimension (or even giving it a high priority) can produce slow positive decision, because the target would be found only after elimination of the irrelevant dimension.

There are two possible reasons why the subjects did not ignore the irrelevant dimension. (1) They may have chosen not to ignore the irrelevant dimension (as a matter of strategy). Detection of a difference in the irrelevant dimension immediately signaled that a negative decision was required, so checking the irrelevant dimension could lead to fast negative responses (the other dimensions needed only be searched if the irrelevant dimension was found not to signal the presence of a difference). (2) Alternatively, the subjects may have been unable to ignore the irrelevant dimension, for example, because any dimension that is a potential source of difference cannot simply be disregarded. In other words, there is a limit on subjects' ability to top-down control the checking of dimensions.

EXPERIMENT 3

Experiment 3 examined to what extent the subjects can top-down control the checking of relevant dimensions (cross-dimension condition). If they do have control, then it should be possible to bias the search in such a way that one of the possible alternatives is assigned priority of checking—for instance, by making one particular target much more likely than the alternative targets. Under conditions similar to those in Experiment 1, this manipulation should have relatively little effect in the within-dimension condition, since the response can be based on the detection of any difference (without the need to identify the exact source of the difference). In contrast, in the cross-dimension condition, this should have a marked effect on the likely target (dimension), which can always be checked first. In fact, performance for that target should approach the level achieved in the within-dimension condition. However, unlikely targets should be subject to costs, due to being assigned a low checking priority. This prediction was tested in Experiment 3.

Method

Subjects. Eight unpaid volunteers participated in Experiment 3 (some of them had taken part in the previous experiments). Four subjects were female, and 4 were male. Their ages ranged from 21 to 30 years. They all had normal or corrected-to-normal vision.

Stimuli. The stimuli were the same as those in the previous experiments. However, only display sizes larger than two items were

presented. Furthermore, in both the within- and the cross-dimension condition, the right-oriented diagonal line was more likely to be a target ($p = .8$) than were the respective alternatives ($ps = .1$). That is, in the within-dimension condition, 80% of the targets were right-oriented diagonal lines, 10% were left-oriented diagonal lines, and 10% were horizontal lines. Similarly, in the cross-dimension condition, 80% of the targets were small gray right-oriented diagonal lines, 10% were large gray vertical lines, and 10% were small black vertical lines. There was no control condition.

Procedure. The instruction was the same as that in Experiment 1 (homogeneity/heterogeneity decision). The subjects were fully informed of the probabilities with which the different targets were presented. All other procedural details were the same as those in Experiments 1 and 2.

Results

Figures 4A and 4B present the correct group mean RTs and the error rates, respectively, as a function of display size, separately for 80% and (2 × 10% =) 20% positive responses and for the negative responses. (The 20% positive RTs refer to the averaged RTs to the two unlikely [10%] targets. Preliminary analyses of the RTs to the two unlikely targets in the within- and cross-dimension conditions failed to find any significant differences due to unlikely target [within, $F(1,7) = 0.00$,

n.s.; cross, $F(1,7) = 2.10$, n.s.] and so were pooled [20% positive responses].) The left-hand and right-hand panels show the functions for the within- and cross-dimension conditions, respectively. The main findings were as follows: In the within-dimension condition, there were no benefits in RT to the likely target. In contrast, there were marked benefits in the cross-dimension condition; indeed, RTs to likely targets were not significantly slower than those in the within-dimension condition (see RT analysis below). The error rates were low overall and tended to be higher with longer RTs (arguing against speed-accuracy tradeoffs) (see error analysis below).

RT analysis. The data were examined using a repeated measures ANOVA with main terms for within/cross-dimension condition, negative/80% positive/20% positive/negative response and display size. There were significant main effects of within/cross-dimension condition [$F(1,7) = 14.78$, $p < .01$] and negative/80% positive/20% positive response [$F(2,14) = 20.89$, $p < .001$]. The main effect of display size was not significant [$F(3,21) = 1.15$]. Furthermore, within/cross-dimension condition interacted significantly with negative/80%

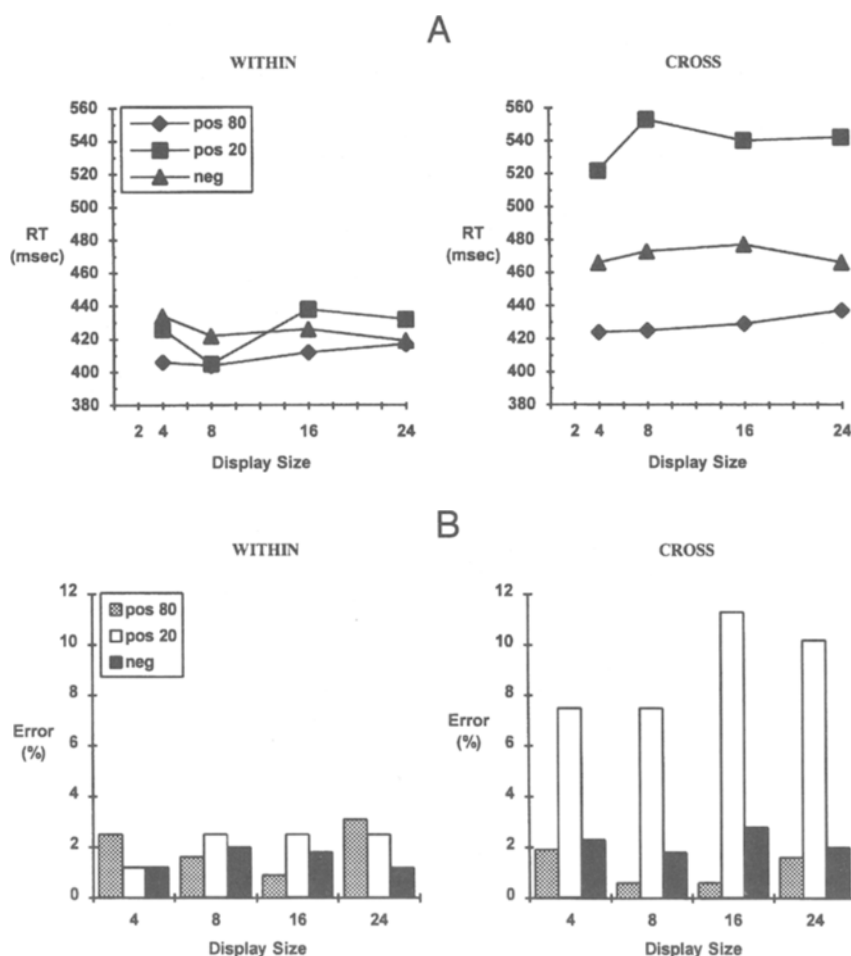


Figure 4. Positive and negative search RT/display size functions (Panel A) and error rates (Panel B) in the within- and cross-dimension conditions in Experiment 3.

positive/20% positive response [$F(2,14) = 19.75, p < .001$] and display size [$F(3,21) = 3.67, p < .05$].

The search RT functions were generally flat (there was no main effect of display size), indicating that search was parallel, as observed previously.

In the within-dimension condition, there was some nonsignificant advantage (15 msec, on average) for the likely (right-oriented diagonal line) target relative to the unlikely targets and the absent conditions. [A separate ANOVA of the within-dimension RTs failed to reveal a main effect of negative/80% positive/20% positive response, $F(2,14) = 1.31$]. In the cross-dimension condition, there was a large advantage for the likely (right-oriented line) target relative to both the negative (41 msec) and the unlikely target conditions (111 msec). The latter showed a disadvantage (70 msec) relative to the negative condition. [A separate ANOVA of the cross-dimension RTs revealed a significant main effect of negative/80% positive/20% positive response, $F(2,14) = 32.52, p < .001$]. Most interestingly, RTs to the likely (right-oriented diagonal line) targets showed little difference (only 19 msec) between the within- and cross-dimension conditions. [A separate ANOVA of the 80% positive response RTs failed to reveal a significant main effect of within/cross-dimension condition, $F(1,7) = 2.39, n.s.$]. This pattern of effects is exactly as expected.

Note that the overall ANOVA revealed a negative/80% positive/20% positive response \times display size interaction. This was due to the within-dimension condition, in particular fast positive RTs to unlikely targets with eight-item displays. [A separate ANOVA of the within-dimension condition showed a main effect of display size, $F(3,21) = 3.80, p < .025$, and a negative/80% positive/20% positive response \times display size interaction, $F(6,42) = 2.73, p < .025$. A separate ANOVA of the cross-dimension condition failed to reveal any significant display size effects.]

The standard deviations of each subject's correct mean RTs were also examined using an ANOVA. As in Experiments 1 and 2, the variability was significantly greater for the cross-dimension condition (83 msec) than for the within-dimension condition (67 msec) [main effect of within/cross, $F(1,7) = 6.74, p < .01$], consistent with the idea that the former condition involves serial elimination of input dimensions. Note, however, that the standard deviations of the positive RTs to likely targets were not increased for the cross-dimension condition relative to those for the within-dimension condition (74 vs. 69 msec) [$F(1,7) = .04, n.s.$], consistent with the idea that the likely dimension is given priority of checking (cross-dimension condition).

Error analysis. An ANOVA of the arcsine-transformed error rates revealed significant main effects of within/cross-dimension condition [$F(1,7) = 16.77, p < .005$] and negative/80% positive/20% positive trial [$F(2,14) = 13.31, p < .001$]. Furthermore, there was a significant within/cross condition \times negative/80% positive/20% positive trial interaction [$F(2,14) = 9.10, p < .005$]. These effects were due to the unlikely targets

(20% positive trials) in the cross-dimension condition, which were missed on a disproportionately large percentage of trials (9.1%). [Separate ANOVAs of the error rates in the within- and cross-dimension conditions revealed a significant trial effect only in the latter condition; within, $F(2,14) = 0.34$; cross, $F(2,14) = 13.40, p < .001$.]

Discussion

The results of Experiment 3 were as expected, providing clear evidence of priming of a likely dimension (in the cross-dimension condition), but no reliable evidence of priming of a featural value on a dimension (in the within-dimension condition). This pattern is consistent with the idea that, in both conditions, the response can be based on the detection of any difference (regardless of the exact featural value of that difference) in a primed (cross-dimension condition) or prespecified dimension (within-dimension condition). It is not clear from the within-dimension data of Experiment 3 whether feature priming is not possible⁴ or just not advantageous when a response can be made upon detection of any dimension-specific difference.

GENERAL DISCUSSION

Taken together, the results of the three experiments are consistent with the idea that, in the within-dimension condition, the response can be based on the detection of any difference—that is, on a module's integrated saliency signal that does not carry information as to exactly how the target differs from the nontargets. In the cross-dimension condition, however, a response cannot normally be given without information as to the source (dimension) of the target–nontarget difference; furthermore, this source must be identified in some form of elimination process. This applies even when the task requires only detection of any difference, rather than identification of what constitutes that difference. In contrast, in the within-dimension condition, the source dimension of the critical difference is known in advance, so no elimination process is required.

In the cross-dimension condition, if one dimension is given search priority and there is a critical (target–nontarget) difference in that dimension, positive responses can be made almost as rapidly as in the within-dimension condition. However, this is not true of negative responses. If one dimension is checked and found not to exhibit a difference (thus favoring a negative response), it cannot be ruled out (without further checking) that there may be a difference within the alternative dimensions. In contrast, in the within-dimension condition, negative responses can be based on the absence of a difference in the relevant dimension only. As a result, negative responses are faster in the within-dimension condition than in the cross-dimension condition.

In some situations (as in Experiment 2), positive responses require the elimination (or ignoring) of one potential source of difference (i.e., of a difference that de-

mands a negative response). In the within-dimension condition, the integrated saliency signal of the relevant dimension cannot then be used as the basis for response (because that signal only indicates that there is a difference). Rather, some process of elimination must take place within the dimension (to rule out the task-irrelevant source of difference), producing costs in the time required to respond positively. The negative RTs are not affected (no elimination is required when the dimension-specific saliency signal does not indicate the presence of a difference).

The present results have implications for the underlying processing architecture, which will be discussed in the next section. A more complete understanding of the findings requires consideration of two issues (to be discussed in the subsequent sections)—namely, whether the elimination of dimensions is serial or parallel and to which extent feature analysis is top-down “penetrable.”

Implications for the Processing Architecture

How well can the present findings be accommodated within the frameworks provided by GS and FIT? Recall that, in GS, modules (dimensions) are organized into sets of feature maps. However, feature map units do not communicate directly with the master map of activations, but through a domain-specific saliency map that effectively represents differences between neighboring stimuli (i.e., feature analyzers). The domain-specific saliency signals are passed to units in the master map of activations, which integrate (sum) the parallel outputs from the different dimensions. This can result in enhanced differential activation over the master map units (e.g., if an odd-one-out target differs from the other items in terms of both orientation and color, the activation of the corresponding master map unit would be greater than when the target differs from the other items in only orientation or color). The distribution of activation across the master map guides higher order processes, such as focal attentional scanning of particular stimulus locations. FIT also assumes the existence of a master map of locations/activations that receives parallel input from all dimensions. But it does not assume the existence of domain-specific saliency maps. However, it appears difficult to account for the present results without appealing to the existence of domain-specific saliency maps. The consequence for FIT is that it would need to be extended to include such maps.

Furthermore, the present data affirm the need for a master map (saliency signal). In FIT and GS, the function of the master map is to allow, respectively, the cross-referencing of or simultaneous access to the information available at a particular location in the various dimension-specific analyzer modules (i.e., it plays a role primarily in conjunction search). The present findings argue that the master map is also important for singleton feature search. In particular, a master map saliency is necessary, though possibly not sufficient, for pop-out and detection to occur (see below for a more detailed discussion of this point). Detection minimally requires (or involves) knowl-

edge on what dimension a critical target–nontarget difference exists. The master map representation plays a role in retrieving this information.

However, the knowledge of the source dimension of the critical difference does not yet specify on what particular value (i.e., feature map) on that dimension the difference exists. For that information to become available, “checking” is required of the different feature maps for the specific map responsible for the difference to be discerned. The mechanisms subserving this process are as yet poorly understood. Presumably, the checking of feature maps operates in parallel (see Discussion section of Experiment 2) at the location of the odd-one-out item (see Johnston & Pashler, 1990; Monheit & Johnston, 1994). It may also be aided by pooled (map) activity units (Treisman & Gormican, 1988). If active, such a unit would signal the presence of a stimulus (or stimuli) of, say, a particular orientation. So checking could be restricted to those maps whose map activity units are active.

Elimination of Dimensions: Serial or Parallel?

If one admits the existence of domain-specific (local) saliency maps as well as a master map, one aspect of the present findings becomes puzzling. If the master map receives, in parallel, copies of the local saliency maps, why should it not be possible, in the cross-dimension condition, to simply respond (positively or negatively) on the basis of the master map representation (see Experiment 1)? In other words, why should there be a need to eliminate those (of the relevant) dimensions that do not contribute to any differential activation over the master map? The answer depends on the nature, serial or parallel, of the elimination process.

So far, only serial elimination of dimensions (e.g., involving the decomposition of the master map saliency signal) has been considered. The mean RTs in the cross-dimension condition were raised and the RT variances increased, but this does not necessarily imply seriality. It is possible that relevant dimensions are checked in parallel, and checking multiple dimensions takes longer than does checking just one. (Even in Experiment 3, dimensions could be checked in parallel, with a greater weight assigned to the more likely dimension.)

One promising parallel account can be derived from GS Model 2.0 (Wolfe, 1994), which assumes that the overall activation map computes the weighted sum of dimension-specific activations (see also Koch & Ullman, 1985). In other words, there is parallel weighting that modulates how strongly a particular dimension determines the overall activation. The weight setting can be top-down controlled such that only relevant dimensions are sampled and greater weight is assigned to likely target dimensions. To accommodate the present findings, this model could make the additional assumption that the weight setting also depends on the “history” of the previous targets (Wolfe, personal communication, January 24, 1994). In particular, if the target on Trial N is in the same dimension as that on Trial N-1,

there is an RT benefit due to the correct dimension still having the largest weight. The overall slower RTs in the cross-dimension condition (Experiment 1) could then be attributed to changes in dimensional weights occurring on consecutive trials with targets in different dimensions. No such cost arises in the within-dimension condition in which the weights are set for orientation and do not change. One possible reason why there should be such intertrial dependencies in the cross-dimension condition is that singleton detection requires a weight shift to the target dimension (see below).

Reexamination of the cross-dimension data of Experiment 1 indeed showed the intertrial dependencies predicted by Wolfe. The data of only 11 of the 12 subjects in the experiment could be reexamined (the floppy disk containing the data of the 12th subject was corrupted). Furthermore, consecutive cross-dimension trials with an orientation target were relatively rare (indeed, for 2 subjects, there was only one such case). Intertrial dependencies were therefore analyzed by examining all available data (including the color and size targets). When the target on Trial N was in the same dimension as that on Trial N-1, the RTs were faster than on the preceding trial (484 msec, on average, vs. 520 msec; 10 of the 11 subjects exhibited the difference). Conversely, when the target on Trial N was in a different dimension, the RTs were slower (541 msec vs. 516 msec; all 11 subjects showed the difference). A similar pattern was observed for the cross-dimension data of Experiment 2 (see Note 4).

Serial accounts. Note that this pattern of intertrial dependency could also be explained by a serial elimination account, according to which that dimension is checked first on Trial N in which a target–nontarget difference was detected on Trial N-1. Although a serial account is possible, would it be plausible? It would be if a master map (saliency) representation did not exist and responses were based directly on dimension-specific saliency mechanisms. However, the best-fitting serial account of the present data would not be feasible without the existence of a master map. The variance data indicated that only one or two dimensions were serially checked on a trial, even though there were three relevant dimensions. To reliably identify the target dimension using only one or two checks would require an integrated master map saliency signal that is decomposed by the serial elimination (inhibition) of dimensions not contributing to master map signal (see Discussion section of Experiment 1).

This raises an awkward question: Why should a master map saliency signal be necessary, but not sufficient, for response? A possible (though somewhat contrived) answer would be that the master map signal serves simply to guide the allocation of focal attention to locations at which a difference was detected, but cannot, as such, be used for response. Responses depend on “perceptual content.” However, an attention shift does not directly make available perceptual information, because the master map units cannot distinguish between the dimension-specific (content) signals that feed into them. Perceptual

information as to the dimension(s) from which the saliency signal originated becomes available only through the serial elimination of noncontributory dimensions. Elimination of dimensions is also required when there is no differential master map saliency signal (i.e., negative trials), because the absence of a master map saliency signal cannot, as such, be used for response. Rather, a negative response requires information that specifies on what dimension(s) a target is absent, and this involves elimination of all relevant dimensions.

Parallel accounts. A more convincing explanation may be provided by a revised parallel-weighting account (along the lines of Wolfe, 1994). The question of why a master map saliency signal should be necessary, but not sufficient, for response would not arise. Pop-out would be based on the master map signal; for that signal to reach sufficient strength, a weight shift to the target dimension would have to occur.

As it stands, Wolfe’s (1994) notion of parallel weighting does not specify how a target that differs from that on the preceding trial may actually be detected. On such trials with a different target, the dimension that contained the target on Trial N-1 would start off with the largest weight; but since it does not contain a target on Trial N, it would contribute nothing to the master map saliency signal. The dimension containing the target would start off with a small weight, producing only a weak master map signal. Assuming that the master map signal must exceed some threshold value to permit a response (a revision to GS 2.0, accepted by Wolfe [personal communication, January 24, 1994]), there must be a weight shift from the previous to the current target dimension, giving rise to a more salient master map signal. This shift might then persist into the subsequent trial.

Wolfe’s account does not specify how the weight redistribution might actually be accomplished. It could be possible that (1) the weights assigned to the relevant dimensions are linked, (2) if the weight for one dimension is at maximum without producing an above-threshold master map signal, that weight is reduced, (3) the reduction in weight is redistributed to other relevant dimensions that have not reached maximum, and (4) that process terminates as soon as the master map signal reaches threshold or the weights of all relevant dimensions have reached maximum. This is essentially Wolfe’s parallel-weighting account coupled with a process of serial elimination. An alternative is that active dimensions (i.e., dimensions exhibiting a dimension-specific saliency signal) attract weight away from noncontributory dimensions. Restated, relevant dimensions compete for weight, with the target dimension gradually accumulating sufficient weight to permit a response. This might involve rapid repeated sampling of the relevant dimensions as to the strength of their difference signal and competitive (“winner-take-all”) updating of weight units modulating the transmission of signals from dimension-specific saliency maps to the master map (analogous to the updating of template units in Humphreys and Müller’s, 1993, SERR model of visual search). If rapid

enough, repeated serial sampling of dimensions may not be distinguishable from parallel sampling.

In summary, parallel-weighting accounts would require some process operating within the time scale of a trial that would allow the target dimension to be discerned in the cross-dimension condition. Any such procedure would involve a weight shift to the target dimension. This shift might be accomplished in a purely bottom-up manner—for example, involving competition between weight units. Such units might also be biased in a top-down manner, to restrict the competition to relevant dimensions (see below) or, more generally, to reflect the probabilities with which dimensions contain a target.

The present data do not allow a decision to be made between the serial-checking and parallel-weighting accounts developed above. Further empirical work is necessary to do so. However, on the grounds of plausibility, some version of parallel weighting seems to be preferable. One problem for the serial-checking account arises from the fact that it can only awkwardly answer the question why a master map signal should be necessary, but not sufficient, for response. The parallel-weighting account does not encounter this difficulty. A further problem for the (strictly) serial-checking account would arise in visual search for singleton (unknown) conjunction targets: There would be no straightforward way of ascertaining on what dimensions the target may be defined, so checking could not be restricted to relevant dimensions only. A parallel-weighting account in which no dimension is assigned zero weight (see below) would avoid this difficulty.

Top-Down Penetrability of the Feature Analysis Stage

One further important issue concerns whether or not the search of the various dimensions is top-down penetrable. There seem to be at least three studies relevant to this question: Theeuwes (1991, 1992) and Pashler (1988). Theeuwes observed that, when subjects searched for, say, an odd-one-out form target, an odd-one-out color nontarget could slow RTs to the target, even though the color dimension was known to be irrelevant. Whether or not the nontarget would interfere with target detection depended on the relative discriminability of the target and nontarget within their dimensions. Theeuwes proposed that difference signals are computed (independently and in parallel) within each dimension. If there are two (or more) difference signals, there will be a “horse race” between the signals to capture focal attention (where the competition occurs at the level of the master map). What signal is computed first depends, in part, on the particular dimension and the discriminability of an odd-one-out item within that dimension (see also Quinlan & Humphreys, 1987). Theeuwes (1991) concluded that (1) “attention is unintentionally switched based on the order of the availability of the local difference in features,” and (2) “features that are relevant to the task do not receive a higher priority” (p. 137). That is, “no *intentional* selection is possible at the second

stage of focal attention” (p. 137), because “the preattentive stage is top-down impenetrable” (Theeuwes, 1992, p. 604).

Theeuwes’s finding that an irrelevant (nontarget) saliency signal could not be ignored in his experiments does not rule out that dimensions can be ignored when they do not give rise to a saliency signal. In this case, the weights for irrelevant dimensions would be set to near zero (such dimensions would be strongly inhibited), effectively restricting sampling to the relevant dimensions. It may not be possible to completely inhibit dimensions on which stimuli are defined; rather, all dimensions on which stimuli are defined may be sampled, but with greater weight assigned to the response relevant dimensions. This would leave the possibility for an irrelevant (to-be-ignored) signal from a nontarget dimension to affect response. Although that dimension would not normally receive much sampling, a strong saliency signal on that dimension may compete for, and attract, weight in a bottom-up manner. On this account, an irrelevant saliency signal would not automatically capture focal attention, but would do so only if it is relatively more salient than the signal produced by a relevant dimension. Note that Theeuwes’s (1991, 1992) data are not incompatible with such an account. In contrast, Theeuwes’s own account would have difficulty explaining why search for an odd-one-out target (in the absence of an odd-one-out nontarget) is slower when the critical target dimension is unknown than when it is known (see present experiments; Treisman, 1988).

Another study that may pose a difficulty for the present account is Pashler (1988). Pashler’s subjects had to localize (left/right decision) a singleton target that was unique in one dimension (e.g., form), while items “varied randomly” in the other, nontarget, dimension (e.g., color; i.e., half the items were green and half were red). In Pashler’s Experiment 5, the relevant dimension was either specified or not at the beginning of a trial. Pashler found that foreknowledge of the relevant dimension produced only a minimal (though significant) RT advantage (8.8 msec). Pashler (1988) concluded that “The observer can voluntarily attend to . . . pooled representations [i.e., pooled across dimensions], but cannot effectively either (1) suppress the outputs of the specified irrelevant dimensions, or (2) attend to only those outputs representing comparisons within a particular dimension” (p. 317).

However, Pashler’s conclusion faces a similar difficulty as that of Theeuwes. It is generally assumed that dimension-specific saliency signals are computed on the basis of local interitem differences (i.e., differences are weighted according to the distance between compared items, with near-neighbor differences being assigned a greater weight; e.g., see Wolfe, 1994). Pashler’s items varied randomly on the irrelevant dimension. Thus, if color was irrelevant, it could happen that there were fewer red items than green items in one area of the field (or vice versa). As a result, local saliency would have been computed for that dimension, which may have in-

terfered with the sampling of the relevant dimension. That is, Pashler's findings do not challenge the present proposal that sampling can be effectively restricted to one dimension when there is no interference from other dimensions. (One further problem with Pashler's study may be that it did not use a standard visual search task, complicating comparisons with search experiments.)

The present proposals challenge strong claims that the parallel stage is top-down impenetrable. Support for this conclusion comes from other recent work by Bacon and Egeth (1994; see also Yantis, 1993), whose findings suggest that salient irrelevant singletons only capture attention when the target of search is a featural singleton.

Conclusions

The present experiments examined search for an odd-one-out feature target under conditions in which the critical (target) dimension was either known or unknown. The results indicated that pop-out in these tasks cannot be based simply on a master map saliency signal that does not specify the dimension(s) from which saliency signals originate. Rather, for pop-out and detection to occur, minimally, the source dimension of the saliency signal must be known (this also applies to negative, no-pop-out, decisions). That is, pop-out is ultimately based on the output of dimension-specific saliency maps. It appears that these maps are sampled in parallel by the master map. Sampling is initially biased (i.e., weighted) toward the dimension that contained a target on the previous trial. If the current target is contained in a different dimension, there must be a weight shift toward that dimension, specifying the origin of the saliency signal. Normally, sampling may be restricted to the relevant dimensions, on which the target is likely to differ from the nontargets (top-down control over sampling). However, a difference signal from an irrelevant dimension may also attract sampling if it is sufficiently salient (bottom-up driven sampling). For a pop-out response to be given, the dimension of an odd-one-out target must be known, but this is not equivalent to knowing the particular feature value on that dimension in which the target differs from the nontargets. To acquire this information, further processing of the critical dimension is required. Further work is necessary to investigate the processes involved in that (postpop-out) stage.

REFERENCES

- BACON, W. F., & EGETH, H. E. (1994). Overriding stimulus-driven attentional capture. *Perception & Psychophysics*, *55*, 485-496.
- CAVE, K. R., & WOLFE, J. M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology*, *22*, 225-271.
- DUNCAN, J., & HUMPHREYS, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*, 433-458.
- DUNCAN, J., & HUMPHREYS, G. W. (1992). Beyond the search surface: Visual search and attentional engagement. *Journal of Experimental Psychology: Human Perception & Performance*, *18*, 578-588.
- HUMPHREYS, G. W., & MÜLLER, H. J. (1993). SEarch via Recursive Rejection (SERR): A connectionist model of visual search. *Cognitive Psychology*, *25*, 43-110.
- JOHNSTON, J. C., & PASHLER, H. (1990). Close binding of identity and location in visual feature perception. *Journal of Experimental Psychology: Human Perception & Performance*, *16*, 843-856.
- KOCH, C., & ULLMAN, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, *4*, 219-227.
- MALJKOVIC, V., & NAKAYAMA, K. (1994). Priming of popout: I. Role of features. *Memory & Cognition*, *22*, 657-672.
- MONHEIT, M. A., & JOHNSTON, J. C. (1994). Spatial attention to arrays of multidimensional objects. *Journal of Experimental Psychology: Human Perception & Performance*, *20*, 691-708.
- PASHLER, H. (1988). Cross-dimensional interaction and texture segregation. *Perception & Psychophysics*, *43*, 307-318.
- QUINLAN, P. T., & HUMPHREYS, G. W. (1987). Visual search for targets defined by combinations of color, shape, and size: An examination of the task constraints on feature and conjunction searches. *Perception & Psychophysics*, *41*, 455-472.
- REIFFENRATH, D. (1993). *Zur Merkmalsanalyse bei visuellen Suchaufgaben—ein simulationsgestützter Ansatz [Feature analysis in visual search: A simulation-based approach]*. Doctoral dissertation, Rheinisch-Westfälische Technische Hochschule Aachen. Aachen: Shaker.
- SAGI, D., & JULESZ, B. (1985a). Detection versus discrimination of visual orientation. *Perception*, *14*, 619-628.
- SAGI, D., & JULESZ, B. (1985b). "Where" and "what" in vision. *Science*, *228*, 1217-1219.
- SCHNEIDER, W., & SHIFFRIN, R. M. (1977). Controlled and automatic human information processing: I: Detection, search, and attention. *Psychological Review*, *84*, 1-66.
- THEEUWES, J. (1991). Cross-dimensional perceptual selectivity. *Perception & Psychophysics*, *50*, 184-193.
- THEEUWES, J. (1992). Perceptual selectivity for color and form. *Perception & Psychophysics*, *51*, 599-606.
- TREISMAN, A. (1988). Features and objects: The fourteenth Bartlett memorial lectures. *Quarterly Journal of Experimental Psychology*, *40A*, 201-237.
- TREISMAN, A., & GELADE, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*, 97-136.
- TREISMAN, A., & GORMICAN, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, *95*, 15-48.
- TREISMAN, A., & SATO, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception & Performance*, *16*, 459-478.
- WOLFE, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, *1*, 202-238.
- YANTIS, S. (1993). Stimulus-driven attentional capture. *Current Directions in Psychological Science*, *2*, 156-161.

NOTES

1. The distinction between *local* dimension-specific saliency maps and a *global* integrated map was anticipated by Koch and Ullman (1985). They suggested that "the different feature maps code for the conspicuity *within* a dimension"—for instance in terms of the "firing frequency" of individual analyzer units. In addition, they assumed "the existence of another ... *saliency map*, which combines the information of the individual maps into one global measure of conspicuity" (p. 221). This representation drives attention by means of a winner-take-all (WTA) mechanism. "Once the most conspicuous location ... has been detected and examined, its corresponding output [saliency signal] decays and the WTA mechanism shifts to the next most salient location" (p. 222). Koch and Ullman speculated that "it is possible ... that the relative weight of the different properties [dimensions] contributing to this [global saliency] representation can be modulated by the activity of some higher cortical centers" (p. 221).
2. Feature maps were not specified in the original architecture of GS, which was dimension-based rather than feature-based (Cave & Wolfe, 1990). However, a more articulated version of GS might assume that input modules consist of separate feature maps plus a saliency map (e.g., see Wolfe, 1994). Similarity comparisons would then be computed between each active feature unit and all other active

units in both the same and the different feature maps within the module. Note further that in Guided Search Version 2.0 (Wolfe, 1994), the similarity comparisons between any two feature units are scaled according to their distance (with their weight decreasing with increasing distance).

3. GS Version 2.0 (Wolfe, 1994) assumes that master map units compute a weighted sum of the dimension-specific saliency signals. Weights are set independently for dimensions (i.e., the total weight is not limited and weights are not linked). Dimensions have a normal weight that may be reduced. For instance, if nontargets share a feature, say, redness, with the target, the bottom-up activation for color is reduced proportionally to the ratio of nontarget types. However, bottom-up information is never reduced to zero (to account for the finding that

irrelevant singletons can attract attention). Accordingly, in cross-dimension singleton feature search, possible (and equally likely) target dimensions should be assigned equal weight, and all other dimensions should be assigned a weight close to zero.

4. Interestingly, no such pattern was observed in the within-(orientation-)dimension data of Experiments 1 and 2. This is in contrast to recent findings by Maljkovic and Nakayama (1994), who found systematic intertrial dependencies with color targets. There may well be a basic architectural difference between the color and orientation feature mechanisms.

(Manuscript received November 22, 1993;
revision accepted for publication July 26, 1994.)