

Searching for unknown feature targets on more than one dimension: Investigating a “dimension-weighting” account

ANDREW FOUND and HERMANN J. MÜLLER
Birkbeck College, University of London, London, England

Search for odd-one-out feature targets takes longer when the target can be present in one of several dimensions as opposed to only one dimension (Müller, Heller, & Ziegler, 1995; Treisman, 1988). Müller et al. attributed this cost to the need to discern the target dimension. They proposed a *dimension-weighting* account, in which master map units compute, in parallel, the weighted sum of dimension-specific saliency signals. If the target dimension is known in advance, signals from that dimension are amplified. But if the target dimension is unknown, it is determined in a process that shifts weight from the nontarget to the target dimension. The weight pattern thus generated persists across trials, producing intertrial facilitation for a target (trial $n+1$) dimensionally identical to the preceding target (trial n). In the present study, we employed a set of new tasks in order to reexamine and extend this account. Targets were defined along two possible dimensions (color or orientation) and could take on one of two feature values (e.g., red or blue). Experiments 1 and 2 required absent/present and color/orientation discrimination of a single target, respectively. They showed that (1) both tasks involve *weight shifting*, though (explicitly) discerning the dimension of a target requires some process additional to simply detecting its presence; and (2) the intertrial facilitation is indeed (largely) dimension specific rather than feature specific in nature. In Experiment 3, the task was to count the number of targets in a display (either three or four), which could be either dimensionally the same (all color or all orientation) or mixed (some color and some orientation). As predicted by the dimension-weighting account, enumerating four targets all defined within the same dimension was faster than counting three such targets or mixed targets defined in two dimensions.

It is well established that targets which differ from distractors in a single salient attribute can be rapidly discerned irrespective of the number of items in the display (the set size). Typical examples are a horizontal line among vertical lines or a “big” bar among “small” bars. Subjectively, the *odd-one-out* item (the target) appears to *pop out* of the display (the horizontal line and big bar in the preceding examples; it seems impossible for its presence to escape the observer’s attention. This *pop-out* phenomenon, together with the flat search reaction time (RT)/set size functions, has been taken as support for the idea that certain visual attributes (i.e., features) are registered and compared in parallel across the entire visual field. For example, see Treisman and Gelade’s (1980) influential feature integration theory (FIT).

However, it is not clearly specified in FIT precisely how the activity of a unique feature leads to the detection of a target and the elicitation of a “present” response. Treis-

man (1988) examined the search RTs for a single blue feature target in experiments in which, unpredictably on a trial, there could be two alternative (single) feature targets as well. In one experiment—the within-dimension experiment—the possible targets were blue, green, or red (i.e., the featural differences were defined along the same—color—dimension¹). In the other experiment—the across-dimension experiment—the targets could be blue (to allow comparison with the previous experiment), horizontal, or large in comparison with the distractors (i.e., the featural differences were defined along separate—color, orientation, or size—dimensions). In both these experiments, the search was spatially parallel (no set size effects). But the RTs to the blue target were 90 msec longer in the across-dimension experiment than in the within-dimension experiment. Treisman suggested that targets popped out within a single, prespecified dimension, but that the cost in the across-dimension condition reflected the (serial) checking of separate dimensions to determine which of them contained the target.

Additional evidence for a dimension-checking account comes from other experiments by Treisman (1988) in which the heterogeneity of the distractors was varied. Treisman found that, when the distractors were heterogeneous in dimensions not relevant to the search for the target (e.g., a blue target bar among green distractor bars of different sizes and orientations), search remained parallel. Conversely,

This research was supported by Science and Engineering Research Council Grant GR/H/54966 and a Royal Society research grant to H. J. M. The authors would like to thank C. Folk, A. Kramer, J. Wolfe, and one anonymous reviewer for their helpful comments on earlier versions of this paper. Correspondence concerning this article should be addressed to A. Found or H. J. Müller, Department of Psychology, Birkbeck College, University of London, Malet Street, London WC1E 7HX, England (e-mail: ubjta52@uk.ac.bbk.cu).

when the distractors were heterogeneous within the same dimension on which the target was specified (e.g., a blue target among distractors colored green, red, and white), search was slow and set size dependent. Treisman suggested that heterogeneity in irrelevant dimensions would not affect search because subjects simply check for activity signaling a contrasting item in the relevant, target-defining, dimension. But heterogeneity within the target-defining dimension would slow search because it becomes necessary to locate the feature map, not just the dimension, containing the unique target item.

Treisman has since revised her account of the feature integration process to incorporate top-down influences of subjects' foreknowledge of the target identity (Treisman & Sato, 1990). The addition of top-down processes to revised FIT allows it to account for feature searches when the target is known in advance, but it is not clear how it could account for pop-out of unknown targets.

One model of visual search which can account for pop-out of unknown targets is the Guided Search (GS) model of Wolfe and colleagues (e.g., Cave & Wolfe, 1990; Wolfe, 1994). Similar to FIT, GS assumes that the visual field is initially represented, in parallel, as a set of basic stimulus attributes in separate "modules" (such as color, orientation, etc.). Each dimension-specific module computes a saliency signal for each item, indicating how *dissimilar* it is in comparison with the other items represented in the module (the more dissimilar an item is, the greater its saliency²). Dimension-specific saliency signals are computed in parallel for all stimulus locations, and these signals are then summed across dimensions by units in an "overall map of activations/locations." The activity on this master map is used to guide focal attention, the most active locations being sampled with priority. Focal attention "gates" the passage of visual information to the higher stages of processing (i.e., the object recognition and response systems). GS also incorporates *top-down* control of attention, which can be used to make search more efficient when the target identity is known in advance. Top-down control is dimension specific and involves parallel comparisons of the feature values at each location with the *desired* feature value defining the target (Wolfe, 1994). The better the match between the desired and actual feature values, the larger the top-down enhancement.

Top-down knowledge is a useful means of guiding search, in particular for targets defined by a unique conjunction of features. However, it is not necessary for efficient (parallel) search for unspecified (singleton) feature targets. Any odd-one-out feature will generate a strong dimension-specific salience signal by virtue of its contrast (dissimilarity) to the other features represented in the same dimension. Given some variability due to noise, the target's salience on the overall map of activations will always be larger than those of distractor items, and so attention will be deployed first to its location. Attention is only deployed to locations on the master map whose activation exceeds a certain threshold. On target-absent trials, there will not be any (above-threshold) peaks of activation, and so a default absent response will be given.

According to GS, no item in the display is responded to until it has been sampled by focal attention and passed on to the later stages of processing—in particular, object recognition (see Wolfe, 1994). Thus, knowing that a target is present would entail knowing its location and identity. Similarly, in Treisman's (1988; Treisman & Sato, 1990) scheme, (at least) conscious perceptions and actions depend on temporary object representations (specifying time, location, spatial and temporal relations, identity, etc. of an object) which are matched to stored descriptions held in long-term visual memory. However, earlier work by Treisman and Gelade (1980) suggests that the detection of a unique activity on a feature map can "directly" elicit a response (subjects were reported to be able to indicate the presence of a feature target without having localized it; but Johnston & Pashler, 1990, pointed out possible artifacts in this result).

Müller, Heller, and Ziegler (1995) recently reinvestigated search for feature targets within and across stimulus dimensions. In their first experiment, they compared search for three possible targets all defined within the dimension of orientation (a left-tilted, horizontal, or right-tilted small gray bar) against search for three possible targets defined across three different dimensions (a right-tilted small gray bar, a black vertical bar, or a large gray vertical bar).³ The distractors in both cases were the same: small gray vertical bars. The subjects were instructed to respond to any "heterogeneity" (i.e., difference) in the display, which invariably signaled the presence of a target, without attempting to further analyze the source of the difference. Müller et al. found that search was parallel in both conditions. However, the mean RT for the right-tilted target was slower by 60 msec in the across-dimension condition than in the within-dimension condition. This cost confirmed the earlier finding obtained by Treisman (1988), though it was 30 msec less than her cost (see above). Furthermore, Müller et al. found no reliable cost for RTs in the within-dimension condition relative to a control condition in which not only the dimension, but also the featural identity, of the target was fixed.

Müller et al. (1995) initially reasoned that the cost in the across-dimension condition was due to the subjects', for some reason, having to determine in which dimension the target was present (color, orientation, or size), even though they were not explicitly required to do so. In the within-dimension condition, the target-defining dimension was known in advance because all targets were prespecified to be orientation differences. In the across-dimension condition, the target-defining dimension was not known and so had to be determined, producing an RT cost. Müller et al. concluded that, when subjects were required to respond to a target-distractor difference, they (minimally) required knowledge of the dimension the target resided in. However, responding positively did not require precise knowledge of the featural value of the target (because the within-dimension condition produced no cost relative to the control condition; see above).

In a second experiment, Müller et al. (1995) repeated the within- and across-dimension searches described above,

but with changed response requirements. In the first experiment, subjects could respond “present” once they had detected “that” there was a different item in the display; they did not have to determine the dimensional and/or featural identity of the different item (although, it seemed, they required dimensional knowledge in order to respond). In the second experiment, the task forced the subjects to identify the target before a response could be made. In the within-dimension condition, subjects responded “present” to a bar tilted left or right, but a horizontal bar required an “absent” response. Similarly, in the across-dimension condition, subjects responded “present” to black and tilted bars, but a large bar required an “absent” response. Again, Müller et al. compared the RTs for the (common) right-tilted target in the two conditions. If a response indeed requires knowledge of the target-defining dimension (as suggested by Experiment 1), then the across-dimension condition of Experiment 2 should be only slightly affected by the explicit requirement to know which *dimension* the target was in (the dimension needed to be known in order to reject the size “target”). However, there should be a large cost in the within-dimension condition, because the *featural* identity of the target had to be determined in an extra time-consuming process (in order to reject the horizontal “target”). This was indeed what Müller et al. found. In the across-dimension conditions, the RTs to right-tilted targets (on present trials) increased by only about 20 msec relative to those in Experiment 1. However, in the within-dimension condition, they increased by about 70 msec. Interestingly, in both the across- and within-dimension conditions, *absent* RTs (to homogeneous displays not containing an odd-one-out item) were faster than *present* RTs.

Two further aspects of the across-dimension results of Müller et al. (1995) are noteworthy. First, there was an RT advantage if the target on trial n was in the same dimension as that on trial $n-1$, relative to when it was in a different dimension (no such advantage was evident for identical consecutive targets in the within condition). Second, the variability of each subject’s RTs was increased relative to the within-dimension and control conditions. However, the increase was less than that expected if the checking and elimination of dimensions were serial and self-terminating (as implied by Treisman, 1988), suggesting some parallelism in across-dimension search.

To explain their findings, Müller et al. (1995) proposed a *dimension-weighting* account, according to which master map units compute the *weighted* sum of dimension-specific saliency signals in parallel. If the target dimension is known in advance (or is very “likely”; see Experiment 3 of Müller et al.⁴), that dimension is assigned a large weight and thus strongly determines the master map saliency signal which Müller et al. proposed can form the basis of a response. However, if the target-defining dimension is unknown and if it starts off with a small weight, the saliency of the target at the level of the master map will be below the threshold required for response. Thus, for the master map saliency signal to exceed the threshold, there must be a redistribution of weight to the target-defining

dimension, thereby providing knowledge about the dimensional identity of the target. This explains the cost in the across-dimension condition. The weight shift persists into the next trial, producing an RT advantage when the next trial contains a target defined in the same dimension, but a disadvantage when it is defined in a different dimension (since weight must be allocated to the new dimension).

In the present study, we reinvestigated the dimensional weighting account proposed by Müller et al. (1995), using a set of convergent procedures.

The first aim was to reexamine whether target detection (minimally) “requires” knowledge of the target’s dimensional identity. A “strong” interpretation of the data of Müller et al. (1995) would equate target detection with knowledge of its dimensional identity. Acquisition of this information requires the elimination of dimensions not containing a target, by means of a dimensional weight shifting process. A weaker interpretation would be that the across-condition involves weight shifting (allocating “attention” to the target dimension), but that explicitly determining the target dimension involves further time-consuming processing.⁵

In Experiments 1 and 2, we reinvestigated this issue by using two single-target tasks that directly addressed the detection and dimension identification stages. In Experiment 1, 50% of the trials contained only distractors (white vertical bars); the other 50% contained one of four possible targets: either a left- or a right-tilted white bar (orientation target), or a red or a blue vertical bar (color target) (with each alternative being equally likely). Subjects were required to respond simply to the presence or the absence of a target (present/absent discrimination task). In Experiment 2, the display always contained one target (one of the four alternatives mentioned above), and the subjects had to respond according to its dimension: color or orientation (color/orientation discrimination task). According to the “strong” interpretation of Müller et al., there should be no difference between the (positive) RTs in Experiment 1 and the RTs in Experiment 2.

The second aim in Experiments 1 and 2 was to determine whether the intertrial dependencies reported by Müller et al. (1995) were indeed due to the dimensional identity of consecutive targets rather than their featural identity. Müller et al. favored an account in terms of dimensional identity for two reasons: Only their across-dimension conditions, but not their within-dimension conditions, showed significant intertrial facilitation, and there was no significant advantage in the foreknowledge of the precise featural identity of the target. However, Malkjovic and Nakayama (1994) recently reported feature-specific intertrial effects with color targets. Such color-specific effects could account for the pattern of intertrial dependencies reported by Müller et al. Their failure to find intertrial facilitation in within-dimension search might be due to the use of only orientation, but not color, targets in these conditions; and their finding of intertrial facilitation in across-dimension search might be more or less entirely caused by consecutive color targets (Müller et al. did not analyze intertrial effects separately for each dimension, owing to in-

sufficient numbers of trials). Since consecutive color targets were always featurally identical, it is possible that what Müller et al. took to be a general dimension-specific effect is really a feature-specific effect confined to the color dimension.

Since the dimensional specificity of the intertrial effects is crucial to the dimensional weighting account of Müller et al. (1995), Experiments 1 and 2 were designed to allow a stronger test of this assertion. In Experiments 1 and 2, targets were defined along one of two possible dimensions (color or orientation), and each dimension could contain one of two possible featural differences (red or blue; left or right tilted). This allowed for a direct comparison of intertrial effects when consecutive targets were dimensionally and featurally identical (e.g., red followed by red) with intertrial effects when consecutive targets were dimensionally identical but featurally different (e.g., red followed by blue). If the intertrial effects are dimension specific, as asserted by the dimensional weighting account, they should not differ between these conditions.

Experiment 3 extended the predictions of the dimensional weighting account to a subitization task. It is well established that subjects can enumerate (count) up to 4 items in a display rapidly and with little cost associated with increasing numbers of items (Sagi & Julesz, 1985a; Trick & Pylyshyn, 1993). Furthermore, Trick and Pylyshyn (1993) demonstrated that subjects could rapidly enumerate target items among distractor items, provided that the targets were featurally distinct from the distractors. The displays in Experiment 3 consisted of a fixed number of 25 items. Among these were either 3 or 4 "target" items, which were distinct from the remaining 22 or 21 identical background (distractor) items. Subjects were required to determine whether displays contained 3 or 4 targets. On half the trials, all targets were specified within the same dimension (all color or all orientation); on the other half, targets were specified on two dimensions (color-orientation mixture). The dimension-weighting account predicts, somewhat counterintuitively, that, with dimensionally identical targets, the enumeration of 4 such items should be faster than the enumeration of 3 items. However, no such difference should be apparent when targets are a mixture of color and orientation items. The advantage for "4" responses to dimensionally identical targets is expected, because all response-relevant information is available in one dimension. In contrast, "3" responses and mixture trial responses incur a cost, owing to the necessity of checking both dimensions.

EXPERIMENTS 1 AND 2

In Experiments 1 and 2, subjects' responses to the presence of a target were compared with their responses to the target's dimensional identity. In both experiments, identical targets were presented in identical displays in two-alternative forced-choice RT tasks, but the response requirements differed: In Experiment 1, a positive response could be based on simply detecting the presence of a tar-

get (absent/present discrimination task). In Experiment 2, each display contained one of these four target items. Subjects were required to respond "color" or "orientation," depending on the dimensional identity of this target item (color/orientation discrimination task).

Method

Subjects. The same 8 subjects (including one of the experimenters) took part in both experiments. All subjects were right-handed. Some of the subjects had prior experience with visual search tasks but (with the exception of the experimenter) were naive as to the purpose of the experiments. The subjects' ages ranged from 19 to 30 years, and they all had normal or corrected-to-normal vision. All subjects (except the experimenter) were paid £4 per hour. The subjects were fully briefed about the task they were to perform.

Apparatus. Stimuli were presented on a high-resolution color CRT monitor driven by a Silicon Graphics UNIX workstation. The laboratory was dimly illuminated to prevent reflections on the CRT. The brightness of the CRT was adjusted to a comfortable level and remained constant throughout the experiments. Responses were recorded from a four-button response box (only the left- and right-most buttons were used). The subjects viewed the displays from a distance of about 70 cm, with their heads supported by a chinrest.

Stimuli. The size of each display element (bar) was $0.45^\circ \times 0.06^\circ$ of visual angle. All displays were presented on a black background. The luminance of the display items was approximately 5 cd/m^2 (i.e., the luminance was equated between the different items), and that of the background brightness, 0.5 cd/m^2 . The elements were arranged in a grid-like pattern, with some spatial "jittering" to prevent the formation of collinear groups in the display. Distractor items were white vertical bars. Target items could be a white bar tilted left (-45° from vertical), a white bar tilted right ($+45^\circ$ from vertical), a red vertical bar, or a blue vertical bar and were randomly assigned to any location in the array of items. Displays contained 4, 9, or 16 items. Elements were arranged in a grid of 2×2 , 3×3 , or 4×4 items. The corresponding displays subtended areas of approximately $1.65^\circ \times 1.65^\circ$, $3.30^\circ \times 3.30^\circ$, and $4.90^\circ \times 4.90^\circ$.

Procedure. Trials were presented randomly in blocks of 50. At the start of each block, a large red cross was displayed in the center of the screen. The subjects initiated a block by pressing the space bar of a keyboard placed in front of them. Each trial began with a small white fixation cross presented in the center of the screen, around which the display elements would appear later. The subjects were instructed to fixate the cross, which remained on for 1,000 msec. The cross was followed by a short blank interval (of 400 msec), which was itself followed by the display. At the end of each trial, the subjects received visual feedback about their responses: a white plus sign for a correct response, or a white minus sign for an incorrect response. After this feedback, the screen went blank for 1,000 msec. Then the next trial began (unless the end of a block had been reached).

In Experiment 1, 50% of the trial displays contained one target item; the other 50% contained only distractors. The four possible targets (see above) were equally likely on *present* trials. Subjects were instructed to respond "absent" or "present" as fast as possible, while keeping their error rates low (absent/present discrimination). One experimental run consisted of 480 trials (presented in blocks of 50 trials, with the last block containing only 30 trials). Each subject performed two runs, giving a total of 960 trials. Of these, half were *absent* trials (160 at each of the three set sizes), and each target appeared in a total of 120 trials (40 at each of the three set sizes). Half of the subjects responded present using the left button of the response box; the other half used the right button.

In Experiment 2, a target was always present in the display, with the four possible targets (see above) being equally likely. The subjects had to decide, as rapidly as possible, whether the target was a an odd-one-out color item or an odd-one-out orientation item (color/

orientation discrimination). Half the subjects responded "color" by using the left response button; the other half used the right button. One run of Experiment 2 consisted of 480 trials. Each target appeared in 120 trials (40 at each of the three set sizes). Subjects again performed two runs, making a total of 960 trials.

The order of the Experiments 1 and 2 was counterbalanced across subjects. To prevent any carryover effects (the experiment performed second required subjects to respond to some of the targets with a different hand from that used in the previous experiment), the experiments were conducted in two sessions, which were separated by a period of several days. Subjects were familiarized with the task in (at least) 100 practice trials given prior to commencing each experiment. Three subjects were given extra practice (another 50–100 trials) before their second experiment until they felt comfortable with the task.

Results

The RT and response error data for individual trials were grouped according to trial type (left-tilted, right-tilted, red, or blue target in both experiments, and target absent in Experiment 1) and set size (4, 9, or 16 items). RTs slower than 1,200 msec were excluded from the analysis (less than 2% of the data were eliminated in this way). Table 1 shows the (correct) RTs averaged across the 8 subjects (and associated standard deviations) for each target condition in Experiments 1 and 2. Figure 1 presents the mean RTs for targets averaged according to their dimensional identity (color or orientation). The bars in Figure 1 depict the error rates (error rates for absent trials in Experiment 1 are not shown).

RT analysis. Two separate two-way analyses of variance (ANOVAs) were performed for Experiments 1 and 2, with set size (4, 9, or 16 items) and trial type (Experiment 1, target absent plus 4 targets; Experiment 2, 4 targets) as factors. In both experiments, RTs increased with increasing set size [Experiment 1, $F(2,14) = 6.65$, $MS_e = 240.10$, $p < .01$; Experiment 2, $F(3,21) = 4.77$, $MS_e = 371.86$, $p < .05$]. But the slopes of the RT/set size functions (estimated in linear regression analyses) were all less than 2 msec/item. These values are clearly within the range taken to be indicative of parallel search. There were also significant effects of trial type on RT [Experiment 1, $F(4,28) = 6.33$, $MS_e = 779.34$, $p < .01$; Experiment 2, $F(3,21) = 4.18$, $MS_e = 970.53$, $p < .02$]. The trial type \times set size interactions were not significant in either experiment [$F(8,56) = 1.03$ and $F(6,42) = 1.17$, respectively].

One potential source of the significant trial type effects is the dimension of the targets (color or orientation). This was examined by grouping the RTs to the four targets ac-

ording to dimension (i.e., for each subject, RTs were averaged across left- and right-tilted targets and across red and blue targets) and subjecting them to two further ANOVAs (one for each experiment) with set size (4, 9, or 16) and dimension (color or orientation) as factors.

In Experiment 1, the exclusion of *absent* trials (and the grouping by dimension) rendered the set size effect insignificant [$F(2,14) = 2.57$, $MS_e = 122.54$, $p > .10$]. The set size effect remained significant in Experiment 2 [$F(2,14) = 5.11$, $MS_e = 149.46$, $p < .025$] (but the slope was shallow, indicating parallel search). Both experiments showed significant effects of target dimension (color or orientation) [Experiment 1, $F(1,7) = 6.05$, $MS_e = 480.45$, $p < .05$; Experiment 2, $F(1,7) = 5.71$, $MS_e = 1284.80$, $p < .05$].

The significant dimension effects were due to RTs' being faster when displays contained a color target in comparison with an orientation target (regardless of whether the task required an absent/present or a color/orientation discrimination).

Error analysis. The errors rates were also subjected to set size \times target dimension ANOVAs, to examine whether the dimension effects on the RTs were caused by speed-accuracy tradeoffs.

In Experiment 1, more errors (misses) tended to be made with orientation than with color targets [$F(1,7) = 4.59$, $MS_e = 11.42$, $.075 > p > .05$]. This suggests that the RT difference between orientation and color targets (longer RT for orientation targets) underestimates the true difference (compensating for the increased error rate with orientation targets would increase the orientation RTs and, thus, the difference to the color RTs).

In Experiment 2, there was no difference in errors between orientation or color targets [$F(1,7) = 0.03$]. The set size effect approached significance [$F(2,14) = 3.26$, $MS_e = 1.72$, $.075 > p > .05$], but this was due to an overall decrease in errors as set size increased. So the small set size effect on the RTs may be due to subjects' responding more accurately with larger set sizes.

RT comparison between Experiments 1 and 2. One aim of Experiments 1 and 2 was to determine whether a cost was associated with responding to the dimensional identity of a target in comparison with responding to its presence. This was examined by subjecting the (target trial) RTs of Experiments 1 and 2 to a three-way ANOVA with experiment, target dimension, and set size as factors. This ANOVA revealed significant effects of experiment

Table 1
Group Mean Reaction Times and Associated Standard Deviations (Both in Milliseconds)
as a Function of Set Size and Trial Type in Experiments 1 and 2

	Set Size											
	Experiment 1 (Absent/Present)						Experiment 2 (Color/Orientation)					
	4		9		16		4		9		16	
	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>	<i>M</i>	<i>SD</i>
Left	412.3	60.0	428.9	79.0	433.4	71.5	469.6	57.8	461.8	52.1	476.6	64.3
Right	420.9	68.3	415.0	71.8	432.6	69.0	461.6	64.3	463.4	59.0	479.9	55.3
Red	407.9	60.3	403.1	61.9	408.9	66.9	437.9	51.2	457.9	50.5	456.5	48.7
Blue	403.5	52.1	408.3	57.2	413.5	55.7	433.3	53.3	447.4	42.5	448.9	56.9
Absent	433.4	76.2	441.4	75.0	451.1	74.4						

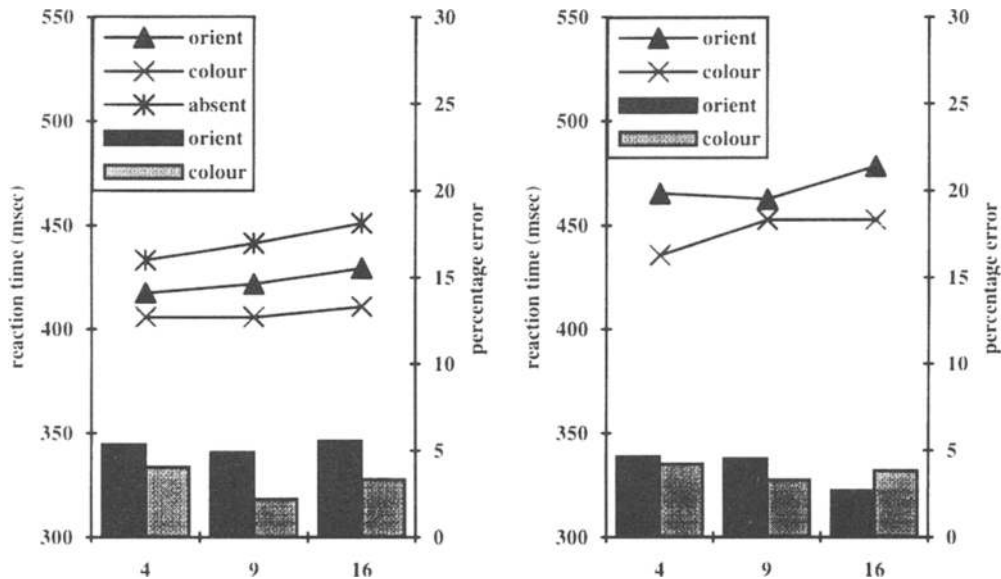


Figure 1. Mean reaction times (RTs), and error percentages, to color and orientation targets as a function of set size in Experiment 1 (absent/present task) (left graph) and Experiment 2 (color/orientation task) (right graph). The left-hand graph also shows the target-absent RTs.

$[F(1,7) = 6.13, MS_e = 7,091.39, p < .05]$, target dimension $[F(1,7) = 14.56, MS_e = 668.95, p < .01]$, and set size $[F(2,14) = 11.08, MS_e = 90.77, p < .01]$. None of the interactions were significant.

RTs to identical targets in identical displays were slower overall in Experiment 2 than in Experiment 1 (see Figure 1). However, experiment did not affect the RT advantage for color over orientation targets. Note that the effect of experiment on RTs was not due to a speed-accuracy tradeoff. An equivalent ANOVA of the error rates failed to reveal any difference between the two experiments $[F(1,7) = 0.51]$ (the only significant result uncovered by the error ANOVA was a set size effect $[F(2,14) = 5.50, MS_e = 1.26, p < .05]$, with errors decreasing overall as set size increased).

The preceding comparison of Experiments 1 and 2 excluded the absent responses in Experiment 1. To examine whether RTs differ overall between the two experiments, all correct responses (including absent responses in Experiment 1) for each of the three set sizes were averaged and compared in a two-way ANOVA with experiment and set size as factors. This ANOVA failed to reveal a significant effect of experiment $[F(1,7) = 2.46, n.s.]$. The set size effect was still evident $[F(2,14) = 14.64, MS_e = 65.67, p < .01]$.

Intertrial facilitation. For both Experiments 1 and 2, the effect on RT to a target on a given trial n was analyzed according to the identity of any target present on the previous trial $n-1$ (trials on which an incorrect response was given on the preceding trial were excluded from the analysis). The responses on consecutive target present trials were categorized as follows:

Different dimension. RTs to targets (on trial n) with the preceding trial ($n-1$) containing a target defined in a dif-

ferent dimension (e.g., a left-tilted target preceded by a red target).

Same dimension, different feature. RTs to targets (on trial n) with the preceding trial ($n-1$) containing a target specified on the same dimension, but with a different feature value on that dimension (e.g., a left-tilted target preceded by a right-tilted target).

Same dimension, same feature. RTs to targets (on trial n) with the preceding trial ($n-1$) containing a target with the same feature value on the same dimension (e.g., a left-tilted target preceded by a left-tilted target).

Table 2 shows the group mean RTs for each of these categories, separately for Experiments 1 and 2. RTs were slower in the *different-dimension* category than in any of the *same dimension*; that is, there was intertrial facilitation (ITF). In both experiments, ITF tended to be greater for same-dimension targets preceded by a featurally identical target. Nevertheless, approximately three quarters of the ITF was attributable to the dimensional identity of con-

Table 2
Intertrial Dependencies in Experiments 1 and 2

	Experiment 1		Experiment 2	
	RT	ITF	RT	ITF
Different dimension	428.3		473.4	
Same dimension/ different feature	392.9	35.4	446.6	26.8
Same dimension/ same feature	382.6	45.7	433.0	40.3

Note—The table shows the group mean reaction times (RTs, in milliseconds) to targets on trial n dependent on the identity of the target on trial $n-1$ (different dimension, same dimension/different feature, same dimension/same feature). The intertrial facilitation (ITF) for same-dimension relative to different-dimension targets is shown in bold.

secutive targets. A two-way ANOVA, with the factors experiment and category, revealed the category main effect to be significant [$F(2,14) = 31.79$, $MS_e = 249.09$, $p < .001$]. There was also a significant main effect of experiment [$F(1,7) = 8.48$, $MS_e = 3,542.69$, $p < .025$]. Comparisons between pairs of categories in each experiment (using the Tukey HSD test) showed that in both experiments, the two same-dimension RTs were significantly faster than the different-dimension RTs.

Table 3 shows the intertrial effects analyzed separately for the color and orientation dimensions in Experiments 1 and 2. Category \times dimension ANOVAs performed separately for each experiment revealed significant ITF (i.e., main effects of category) in both cases [Experiment 1, $F(2,14) = 14.26$, $MS_e = 548.06$, $p < .001$; Experiment 2, $F(2,14) = 26.12$, $MS_e = 243.29$, $p < .001$]. In addition, Experiment 2 showed a marginally significant category \times dimension interaction [$F(2,14) = 3.59$, $MS_e = 255.89$, $.075 > p > .075$]. There was also a significant main effect of dimension [$F(1,7) = 7.18$, $MS_e = 1,191.88$, $p < .05$]—that is, overall faster RTs to color targets. In Experiment 1, there was no significant extra ITF for same-dimension, same-feature targets relative to same-dimension, different-feature targets, for either the orientation or the color dimension. In Experiment 2, there was some evidence of extra facilitation for the color dimension, but not for the orientation dimension (marginal category \times dimension interaction).

Summary. In both Experiment 1 (present/absent discrimination) and Experiment 2 (color/orientation discrimination), the RT data were consistent with spatially parallel processing of the display, with color targets being responded to faster than orientation targets. Comparison of RTs between Experiments 1 and 2 showed a cost for responding to the dimensional identity of the feature targets relative to responding to their presence. This cost was significant when only the present RTs of Experiment 1 were considered, but not when the absent RTs were included in the comparison.

An analysis of the effect on RT to a given target (trial n) of the identity of a target on the preceding trial ($n - 1$) revealed significant ITF—faster RTs when the current target was dimensionally identical to the preceding target than when it was different. This was the case for both experiments and dimensions. In Experiment 2, the color di-

mension exhibited a tendency toward extra ITF when the current (color) target was dimensionally and featurally identical to the preceding (color) target relative to when it was dimensionally identical, but featurally different.

Discussion

Does the detection of a singleton feature target “require” knowledge of its dimension? According to the dimension-weighting account, responses in both experiments are based on a master map representation of summed dimension-specific saliency signals (such as the “overall map of activations” in GS; e.g., Wolfe, 1994). Peaks of activation on this map can be used to determine the presence and location of salient feature differences (presumably up to four differences simultaneously; see Sagi & Julesz, 1985a, 1985b; Trick & Pylyshyn, 1993).

Müller et al. (1995) reported evidence suggesting that, when an unknown feature target can be present on one of several (unpredictable) dimensions, responding to the presence of a target requires (implicit) knowledge of the source dimension of the critical (target-distractor) difference. Experiments 1 and 2 above showed a difference between responses to the presence and dimension of a target that was not predicted by a “strong” interpretation of the data of Müller et al. However, when all responses (including the absent responses of Experiment 1) were taken into account, this difference was *not* significant. Whether the absent RTs in Experiment 1 should be taken into consideration in the comparison between the two experiments depends on how “absent” responses are produced. “Absent” responses may simply be a default (negative) response elicited after failing to find an odd-one-out target (within a set time limit) (see, e.g., Kaptein, Theeuwes, & van der Heijden, in press; Wolfe, 1994). Alternatively, they may be based on some (positive) signal indicating a homogeneous field of items (see, e.g., Duncan & Humphreys, 1989, p. 449). In the first view, “absent” responses are not very interesting theoretically, so they can be ignored in the comparison between Experiments 1 and 2. In the second view, the “absent” responses should be taken into account because the task would truly involve present/absent discrimination (i.e., equivalent decision making and response selection demands to color/orientation discrimination). The fact that “absent” responses may be faster than “present”

Table 3
Intertrial Dependencies in Experiments 1 and 2, Analyzed Separately
for the Color and Orientation Dimensions in Each Experiment

	Experiment 1				Experiment 2			
	Color		Orientation		Color		Orientation	
	RT	ITF	RT	ITF	RT	ITF	RT	ITF
Different dimension	420.5		431.3		464.8		479.6	
Same dimension/ different feature	384.5	36.0	402.5	28.8	436.4	28.4	458.0	21.6
Same dimension same feature	374.4	46.1	392.5	38.8	411.0	53.8	454.8	24.8

Note—The table shows the group mean reaction times (RTs) and intertrial facilitation (ITF) (in milliseconds) analogously to Table 2.

responses (see Experiment 2 of Müller et al., 1995) lends some support to the second view. However, even when the absent responses are included in the comparison, a (non-significant) RT difference (of some 35 msec) between Experiments 1 and 2 is still evident. This cost in responding to the dimensional identity of a feature target relative to responding to its presence should not be too readily ignored. Recall that Müller et al. observed a similar cost (of some 20 msec) when responding present in across-dimension feature search required explicit knowledge of the target dimension (because a feature difference in one dimension required a negative response; see their Experiment 2).

Müller et al. (1995) attributed the cost in responding to feature targets defined on variable dimensions (across-dimension condition) to the need to shift dimensional weight to the target dimension. If the target-defining dimension is known in advance, this dimension receives a strong weighting. Consequently, saliency signals from this dimension are amplified relative to signals from other dimensions, resulting in fast RTs. But when a target can be present on one of several dimensions, the target dimension cannot be weighted in advance. To produce pop-out of the target, weight needs to be shifted between dimensions until signals from the target dimension are strong enough to produce an above-threshold peak of activation on the master map. In the across-dimension condition of Müller et al., this process of shifting weight to the target dimension may have resulted in (implicit) knowledge of the target's dimensional identity. However, the present experiments (together with Experiment 2 of Müller et al.) suggest that there may well be a cost associated with explicitly determining the target dimension.⁶

Are the intertrial effects dimension specific? Although responses to the presence (Experiment 1) and dimension (Experiment 2) of a target may be different in some way, the results of the intertrial analysis strongly suggests a common underlying process between the two tasks. In both experiments, RTs were faster when the current target was dimensionally identical to a preceding target than when it was different. This was the case for both color and orientation targets. In Experiment 2, this dimension-specific ITF could be viewed as an artifact of the color/orientation response, since subjects used one hand to respond to color targets and the other hand to respond to orientation targets. However, in Experiment 1, all responses to targets were made with the same hand, ruling out a response-based account. By implication, the similarity of the intertrial effects in the two experiments argues against a response-mapping account of the ITF in Experiment 2. Thus, the ITF observed was largely dimension specific in origin.

According to the dimensional weighting account, dimension-specific ITF is characteristic of the weight-shifting process. The pattern of dimensional weighting generated when one is responding to targets persists across consecutive trials, so that the RT to a target on a given trial is dependent on the dimensional identity of the target on the preceding trial. If a target is preceded by a target defined along the same dimension, the preexisting

weight pattern will reduce the amount of weight shifting required to produce an above-threshold peak of activation on the master map, in comparison with a preceding target defined on a different dimension. Importantly, weight shifting produces dimension-specific, rather than feature-specific, ITF.

To illustrate, when trial *n* contains a left-tilted bar, a dimension-specific saliency signal will be generated (at the target location) and transferred to the master map. If the master map saliency signal produced by the target is insufficient to elicit a response, weight must be shifted to the orientation dimension to amplify the signal until the response threshold is exceeded. The pattern of dimensional weighting thus generated will persist into the next trial. If the next display also contains an oriented-bar target (no matter whether left or right tilted), it should immediately produce an above-threshold peak of activation on the master map (since the orientation dimension is already weighted). However, if the target on the next trial is a color target, its master map saliency signal will not be sufficiently large to generate a response. To do so, dimensional weight must be shifted from the orientation dimension to the color dimension, incurring an RT cost.

The predicted dimension-specific intertrial costs and benefits in RT were observed in Experiments 1 and 2. Indeed, the largest and most robust ITF was dimension specific, attributable to the dimensional identity of consecutive targets. However, additional facilitation due to the featural identity of consecutive targets was also evident, though only for the color, but not for the orientation, dimension (in Experiment 2). The implications of this feature-specific ITF with color targets will be discussed in the General Discussion.

EXPERIMENT 3

Experiment 3 was designed to extend the dimensional weighting account elaborated above to a subitization task. It is known that subjects can enumerate (count) up to four items in a display rapidly (e.g., Sagi & Julesz, 1985a), even when the items are featurally distinct targets among a set of distractor items (Trick & Pylyshyn, 1993). One way of accounting for this finding is to assume that subitization performance is based on the generation of peaks of activation on some master (saliency) map which can be counted rapidly (at least up to four).

In Experiment 3, subjects had to decide whether three or four such targets were present in 22 or 21 homogeneous distractor items (displays contained 25 items in total). Targets could be defined in the same dimension—all color or all orientation targets—or else, they could be specified in two dimensions: mixture targets (furthermore, potential targets could take on one of two feature values in each dimension: color—red and/or blue; orientation—left and/or right tilted). Pure and mixture trials were presented in random order (as were trials with three and four targets).

Theories according to which dimensions contribute equally to the computation of the master map saliency representation predict no difference between pure and mix-

ture trials. Also, "4" (four targets present) responses are generally found to take longer than "3" responses, which suggests that there is some small, but nonnegligible cost associated with counting each item (though see Folk, Egeth, & Kwak, 1988).

In contrast, according to the dimension-weighting account, a different pattern of RTs should be evident, depending on the number and dimension of targets in the display. More specifically, the account predicts that (pure) trials containing either four color or four orientation targets will produce the fastest RTs. This is because displays containing a mixture of targets defined on both dimensions, or displays containing only three targets, will require that information from both dimensions be sampled. This involves the shifting of weight from one dimension to the other, incurring an RT cost.

For example, when the display contains three orientation targets, weighting the orientation dimension will produce three peaks of activation on the master map. However, this is not sufficient to permit a response, since it may be the case that the display also contains a color target which, because of the orientation dimension being weighted, did not produce a peak on the master map (or only a small peak). Thus, if subjects detect three master map peaks, they must still check that the other dimension does not contain a target as well. However, if the subjects detect four peaks, they can respond without having to check that the other dimension contains a target—since they know there cannot be more than four targets in the display.

In summary, according to the dimension-weighting account, responding "4" should be faster than responding "3" when all of the targets are present on the same dimension (all color or all orientation). However, on mixture trials, RTs should not differ between three and three targets in the display, since both dimensions must be checked to determine the total number of targets. The RT cost associated with this is essentially the same as that with across-dimension search in the study of Müller et al. (1995).

Method

Subjects. Eight subjects (including one of the experimenters) took part in Experiment 3. Most of the subjects had experience with visual search tasks (5 had taken part in Experiments 1 and 2). But (with the exception of the experimenter) they were naive as to the purpose of the experiment. Subjects' ages ranged from 20 to 32 years; they all had normal or corrected-to-normal vision and were all right-handed.

Stimuli. The stimuli were the same as in Experiments 1 and 2. Displays contained 25 items arranged in a jittered 5×5 grid pattern, subtending an area of approximately $4.80^\circ \times 4.80^\circ$. There were no constraints on the positioning of three or four target items; they were randomly allocated to any of the 25 display positions. Subjects viewed the CRT from a distance of 90 cm.

Procedure. Prior to the actual experiment, subjects performed a short practice run of 40 trials. Experimental trials were presented randomly in blocks of 50 (with the last block consisting of 38 trials). Subjects initiated a block of trials by pressing the space bar of a keyboard placed in front of them. An experimental run consisted of 288 trials presented in a random order. Subjects performed six such runs, giving a total of 1,728 trials.

In half of the trials, the targets were all specified within the same dimension (*pure* trials); in the other half, displays consisted of a mix-

ture of targets from both the color and orientation dimensions (*mixture* trials). Twenty-five percent of the pure (all color or all orientation) trials contained targets that were featurally identical (all red, all blue, all left tilted, or all right tilted). When displays contained three targets, the mixture trials contained two targets from one dimension and one from the other dimension. When displays contained four targets, the mixture trials could consist of either two targets from each dimension, or three targets from one dimension and one from the other dimension.

Subjects had to respond "3" or "4" (as quickly and accurately as possible) to indicate the number of targets in the display. Half of the subjects responded "4" with their left hand and "3" with their right hand, and vice versa for the other half.

Results

Trials were grouped according to the number of targets they contained (three or four) and whether targets were all orientation-defined items, all color-defined items, or a mixture of both. The correct group mean RTs and associated error rates for each of these types of trials are shown in Table 4. (RTs less than 150 msec or greater than 2,200 msec—overall less than 1% of the trials—were excluded from the analysis.)

Table 4 shows that, when target items were either all color or all orientation defined, subjects were faster to respond "4" than "3" (by about 55 msec, on the average). However, this difference was not evident when targets were a mixture of color- and orientation-defined items (the difference was only 12 msec, on the average).

The RT data were subjected to a two-way ANOVA with the factors target number (three or four) and trial type (all orientation, all color, or mixture). This ANOVA revealed significant main effects of target number [$F(1,7) = 5.60$, $MS_e = 3,507.04$, $p < .05$] and trial type [$F(2,14) = 8.09$, $MS_e = 1,468.73$, $p < .01$]. The target number \times trial category interaction was also significant [$F(2,14) = 5.09$, $MS_e = 491.1$, $p < .05$]. "3" (targets present) responses were overall slower than "4" responses, and mixture trial RTs were slower than pure (all color and all orientation) trial RTs. Furthermore, "4" responses were significantly faster than "3" responses for all color and all orientation trials (simple main effects: $p < .001$ for both comparisons), but not for mixture trials (*n.s.*).

The error rates were subjected to a similar target number \times trial type ANOVA. The main effects were not significant nor was the interaction. Thus, the faster RTs for "4" than for "3" all color and all orientation targets were not caused by speed-accuracy tradeoffs.

In Experiments 1 and 2, the RT to a target (trial n) was facilitated when preceded by a target on the same dimen-

Table 4
Group Mean Reaction Times (RTs, in Milliseconds)
and Error Percentages in Experiment 3

	All Color		All Orientation		Mixture	
	3	4	3	4	3	4
RT	726.2	674.8	754.5	696.3	760.8	749.0
Error	4.86	6.16	6.04	3.94	5.54	6.84

Note—The data are grouped according to the number (three or four) and composition of the targets (all color, all orientation, mixture) on a trial.

sion (trial $n-1$). In addition, there was some evidence of extra feature-specific facilitation, in particular for color targets (see Experiment 2 above). To examine the effect of the featural identity of targets in Experiment 3, the effect on RT of targets' being featurally homogenous (e.g., all red) or featurally heterogeneous (e.g., reds and blues) was examined for all color and all orientation trials. Only trials containing four targets were analyzed, since the responses could be based on information available within a single dimension (trials containing three targets would involve the added component of checking the alternative dimension, according to the dimensional weighting account). The results of this analysis are shown in Table 5.

The data were subjected to a two-way ANOVA with the factors dimension and homogeneity. This ANOVA revealed a significant main effect of dimension [$F(1,7) = 7.62$, $MS_e = 952.91$, $p < .05$], due to overall faster color RTs, and a significant dimension \times homogeneity interaction [$F(1,7) = 20.62$, $MS_e = 269.91$, $p < .005$]. This interaction was caused by responses to homogeneous color targets' being faster than responses to heterogeneous color targets, while responses to orientation targets were little affected by within-dimensional target homogeneity. This pattern is consistent with the tendency for an extra same-dimension/same-feature (i.e., feature-specific) advantage for color targets in Experiment 2.

To determine whether the RT differences were compromised by speed-accuracy tradeoffs, the error rates were subjected to an equivalent ANOVA. This revealed significant main effects of dimension [$F(1,7) = 5.90$, $MS_e = 9.38$, $p < .05$] and homogeneity [$F(1,7) = 9.69$, $MS_e = 6.85$, $p < .025$], the interaction was not significant [$F(1,7) = 0.95$]. More errors were made with orientation targets and with homogeneous displays. The insignificant interaction leaves the possibility that the RT advantage for homogeneous over heterogeneous color displays was due to an error difference between these conditions. This possibility was ruled out by a t test which failed to reveal the error difference between homogeneous and heterogeneous color displays to be significant [$t(7) = 0.87$, n.s.]; in fact, 3 of the 8 subjects made more errors with heterogeneous displays.

Discussion

In summary, Experiment 3 showed that, when targets were all defined within a single dimension (all color or all orientation), enumerating four such targets was significantly faster than enumerating three. Importantly, there

was no difference between enumerating three and four targets on trials containing a mixture of targets from each of the two dimensions. This pattern of effects is as predicted by the dimension-weighting account.

According to this account, search for singleton feature targets across several dimensions involves a cost due to the differential weighting of dimension-specific saliency signals onto an overall (master) saliency map on which detection responses are based (see, e.g., Experiment 1 of Müller et al., 1995). Experiment 3 extended this proposal to a subitization task that required subjects to count (either three or four) feature targets against a homogeneous background of distractors (color targets, red or blue; orientation targets, left or right tilted). The dimension-weighting account assumes that subitization performance, too, is based on the master map representation. Strong peaks of activation on this map produced by salient feature differences can be used to determine the presence and location of targets in the display. Counting the number of such peaks can operate rapidly for up to about four items. If subitization performance does indeed rely on the same representation as target detection, the enumeration of multiple feature targets defined in different dimensions should show a similar (cross-dimensional) cost in comparison with the enumeration of targets all defined within the same dimension. Such a cost was indeed observed in Experiment 3. RTs were significantly faster to displays containing four targets rather than three, provided the four targets were all defined within a single dimension (all color or all orientation).

Subjects in Experiment 3 knew that displays would contain either three or four targets. Thus, when four peaks of activation were detected on the master map, subjects could immediately respond "4" (as this was the maximum number of targets). However, when three (or fewer) peaks were detected on the master map, subjects could not respond "3" without shifting weight to the other task-relevant dimension to determine whether it, too, contained any targets. This redistribution of dimensional weight produced the RT cost for three-target and mixed-target trials. Only displays containing four targets all specified within a single stimulus dimension did not require weight redistribution, permitting fast enumeration responses.

Trials that contained a mixture of color and orientation targets would produce one, two, or three peaks of activation on the master map, depending on the dimensional composition of the targets and on which dimension was weighted. Similarly, trials that contained three color or three orientation targets would generate three peaks of activation on the master map, provided that the appropriate dimension was weighted. Since mixture trials containing four targets could also produce three master map peaks, the detection and enumeration of three peaks alone was not sufficient to permit a correct response to be made (the same applied to the detection of one or two peaks). Rather, the response had to wait until the other task-relevant dimension was weighted to determine whether or not it generated any other master map peak(s).

Normally, the enumeration of items up to four incurs a cost of some 40–120 msec for each item (see, e.g., Trick

Table 5
Group Mean Reaction Times (RTs, in Milliseconds)
and Error Percentages for Trials With Four Featurally
Homogeneous or Heterogeneous All-Color and
All-Orientation Targets in Experiment 3

	Color		Orientation	
	Homogeneous	Heterogeneous	Homogeneous	Heterogeneous
RT	647.1	688.8	703.6	692.5
Error	5.05	3.33	8.95	4.80

& Pylyshyn, 1993). However, Folk et al. (1988) have shown that, when subjects make a two-alternative forced-choice response (e.g., three or four targets, as in Experiment 3), responding to the larger number of targets, "4," can be faster (if they are presented within a texture array of distractor items). Folk et al. attributed this RT advantage to the uncertainty experienced by subjects upon detecting three targets that an additional target may still lurk in the display; there is no such uncertainty upon detecting four targets (since this is the maximum number). It is important to note that this does not compromise the interpretation of the results of Experiment 3 in terms of dimensional weighting (see above). The critical result is that the RT advantage for "4" responses was dependent on the dimensional composition of the targets (significant interaction). When the targets were dimensionally mixed, the "4" response advantage was only 12 msec; but when they were dimensionally homogeneous, the advantage was 55 msec. This interaction cannot be explained simply in terms of greater uncertainty associated with "3" responses.

One further result of Experiment 3 was that, when the all-color and all-orientation trials were analyzed according to whether they were featurally homogenous or heterogeneous, there was an asymmetry between targets defined in the color and orientation dimensions. For orientation targets, no RT advantage was evident for homogeneous trials (e.g., all right tilted) relative to heterogeneous trials (some right and some left tilted). However, homogeneous color targets (e.g., all blue) exhibited a significant RT advantage over heterogeneous targets (some blue and some red) (although this may be somewhat compromised by a speed-accuracy tradeoff). This asymmetry had an analogue in the patterns of ITF observed in Experiment 2, namely: there was significant feature-specific facilitation (in addition to dimension-specific facilitation) only for the color, not for the orientation, dimension.

GENERAL DISCUSSION

The Dimension-Weighting Account

According to the dimension-weighting account, the transmission of dimension-specific saliency signals to a master saliency map is modulated by a process that weights (scales) the dimension-specific signals. Responses can be based on the master map representation of weighted and integrated dimension-specific saliency signals (see Koch & Ullman, 1985, and Wolfe, 1994, for a similar proposal). Above-threshold peaks of activation on this map can be used to determine the presence and location of salient feature differences. The master map signal indicates *that* there is a difference in the display and *where* it is, but not *what* it is. This assumption accounts for the dissociation between target detection/localization and target identification (e.g., Nothdurft, 1992; Sagi & Julesz, 1985a, 1985b). The master map representation underlies the rapid responses to single feature targets and the efficient enumeration of up to four feature targets presented simultaneously (e.g., Trick & Pylyshyn, 1993). This proposal is broadly in line with GS (e.g., Wolfe, 1994). However, it makes the

additional assumption that master map peaks of activation can lead "directly" to response, without involving "focal attention" and the object recognition system (see below).

The master map sums dimension-specific saliency signals in parallel, but greater weight is assigned to input from the dimension that is known (or very likely; see Experiment 3 of Müller et al., 1995) to contain a target. As a result, saliency signals from this dimension are amplified relative to signals from other dimensions, generating a strong peak of activation on the master map which forms the basis of a rapid response. But when a target can be present on one of several dimensions, the target dimension cannot be weighted in advance with any degree of certainty. In this case, weight must be shifted between dimensions until signals from the target dimension reach sufficient strength at the master map level to induce a response. This weight shift causes the cost in the search for singleton feature targets across dimensions (Müller et al., 1995).

The weight pattern generated in the weight-shifting process persists across consecutive trials, producing characteristic intertrial dependencies. If a target is preceded by a target defined along the same dimension, the preexisting weight pattern will reduce the amount of weight shifting required for response compared with a preceding target defined along a different dimension (ITF). Experiments 1 and 2 permitted dimension-specific ITF to be dissociated from feature-specific facilitation. The analysis of the intertrial effects showed that, to a large extent, the ITF is indeed dimension specific (rather than feature specific) in nature, as predicted by the dimensional weighting account.

Müller et al. (1995) considered the possibility that, when an unknown feature target can be present on one of several (unpredictable) dimensions, pop-out requires knowledge of the source dimension of critical difference. However, Experiments 1 and 2 showed evidence of a difference between responding to target presence and target dimension ("present" responses being faster than "dimension" responses), which cautions against a strong interpretation of the data of Müller et al. The faster RTs to target presence suggest that (at least explicit) knowledge of the dimensional identity of a feature target is not a necessary condition for its detection. Although responding to the dimension of a singleton feature target involves some additional process compared with simply responding to its presence (see note 6), the similar patterns of intertrial effects (in particular, the dimension-specific facilitation) in Experiments 1 and 2 argue that dimensional weighting underlies performance in both tasks.

Experiment 3 provided a strong test of the dimensional weighting account in a subitization task (determining whether a display contained three or four salient feature targets). According to this account, the master map representation forms the basis not only of rapid responses to single feature targets, but also of performance in the subitization task. If so, enumeration RTs depend on whether or not the targets are specified within a single dimension. In particular and counterintuitively, when all targets are specified within a single dimension, responding "3" should be slower than responding "4." The results of Experiment 3

did indeed show the expected cost (of approximately 50 msec) for “3” responses on dimensionally homogeneous target trials. No such cost was evident on trials containing a mixture of color and orientation targets, also as predicted by the dimension-weighting account.

Current models of visual search have difficulty accounting for this pattern of results (as well as the across-dimension search cost observed by Müller et al., 1995) unless they are augmented to incorporate dimensional weighting. Models such as revised FIT and GS, which postulate an overall saliency map with dimensionally segregated input, could easily be modified to account for these results. In fact, the revised version of GS (Wolfe, 1994) already assumes scaling of dimension-specific saliency signals prior to their summation onto the master map. The dimension-specific ITF observed in Experiments 1 and 2 would additionally require retention of the coefficients for scaling dimension-specific saliency signals into the next trial.

Physiological Evidence of Dimensional Weighting

It is interesting that the results of a recent PET (positron emission tomography) study by Corbetta, Miezin, Dobmeyer, Shulman, and Petersen (1990, 1991) are consistent with the dimensional weighting account advocated here. The subjects of Corbetta et al. had to perform discrimination tasks, either within a single prespecified dimension or across several dimensions, while their cerebral blood flow patterns were scanned by means of PET. Subjects were presented with two frames, each lasting 400 msec, separated by a 200-msec blank interval. The first frame displayed a randomly arranged field of small identical bars moving coherently either from left to right or from right to left. In the second frame, the shape, color, or velocity of all the elements might independently change. The subjects were required to respond “same” or “different,” depending on whether a change had or had not occurred between the two frames. The subjects performed two conditions: the selective- and divided-attention conditions. In the selective-attention condition, they were required to respond to a change only in a single prespecified dimension (e.g., color). On some trials, changes occurred on the other dimensions, but the subjects were instructed to ignore these changes and respond “same.” In the divided-attention condition, changes in any of the dimensions required a different response. Corbetta et al. found that subjects were better at performing the *same/different* discrimination task in the selective-attention condition than in the divided-attention condition. The results of PET scans taken during task performance showed that performance in the selective-attention condition was accompanied by an increase in cerebral blood flow to the task-relevant cortical area (e.g., the inferior parietal cortex in the case of velocity). This (blood flow pattern) is consistent with the idea proposed above that, when known, the task-relevant dimension can be weighted in advance, allowing rapid responses in comparison with when the task-relevant dimension is not known.

Interestingly, Corbetta et al. (1991) also found that the two conditions involved “the use of different premotor or

response selection pathways. In the divided task, dorso-lateral frontal activation was coupled with activation of the anterior cingulate, while in the selective task, extrastriate enhancement was coupled with activity in insular cortex and inferior premotor cortex.” In other words, the task in the selective-attention condition can be performed efficiently, on the basis of “the use of nonfrontal [“early” posterior] processors and a [direct] premotor-insular output pathway” (p. 2397). In contrast, the divided-attention task was more complex, “involving a coupling of frontal and anterior cingulate activity,” which has also been observed in, for example, the conflict condition of a Stroop task (see, e.g., Pardo, Pardo, Janer, & Raichle, 1990). This is also consistent with the idea proposed above that performance of the across-dimension task involves a top-down component that actively redistributes weight among the various task-relevant dimensions.

Responses Based on the Master Map

In all three experiments, it was assumed that subjects can base responses more or less directly on some representation such as the “master map of activations/locations” (coding the locations of objects to be attended *and* responded to) without waiting for complete object knowledge to become available. This proposal is at variance with GS and, perhaps, with revised FIT, which mainly focus on the role of attention for object recognition. However, it is consistent with the view there are two main visual processing streams: the occipitoparietal *where* pathway, which mediates knowledge of the whereabouts of objects, and the occipitotemporal *what* pathway, which is concerned with object recognition (see, e.g., Desimone & Ungerleider, 1989; Ungerleider & Mishkin, 1982). Indeed, the *where* system is thought to play a role in the selection of objects for further processing by the capacity-limited recognition system (see, e.g., LaBerge, 1990; LaBerge & Brown, 1989). Some form of master map is likely to reside in the parietal cortex, which is also known to mediate directed motor actions such as eye and hand movements (see Goodale & Milner, 1992, for a review). It is therefore plausible that at least some types of response to objects (e.g., to the presence of an odd-one-out item) can be initiated directly from the master map representation, while further processing of the object of interest (*what is it?*) is still going on. This would be ecologically desirable—for example, in situations in which an object approaches the observer at speed, requiring him/her to take evasive action and/or orient toward the object (head and eye movements) before knowing precisely what that object is. Although the suggestion of responses’ being directly driven from the master map representation is at variance with GS and revised FIT, it is not fundamentally incompatible with these models. They could easily be extended to provide direct links to the response output systems.

The assumption of direct links to response would not necessarily conflict with GS and revised FIT. It is possible that some kind of object file system (Kahneman, Treisman, & Gibbs, 1992) is involved in setting the-trigger conditions under which direct responses are released. Object

files code information such as the time, location, spatial and temporal relations, identity, name, and so forth, of objects in the display. There is now ample evidence indicating that some types of information (such as object location) may become available before others (such as object identity) (e.g., Atkinson & Braddick, 1989; Johnston & Pashler, 1990; Müller & Rabbitt, 1989; Sagi & Julesz, 1985a, 1985b). It is therefore possible that direct responses are released (disinhibited) as soon as the relevant information (transmitted via the attentional system) is recorded in the object file. This way, not all responses have to wait until the information that is derived slowest becomes available.

Differential Preattentive Processing of Color and Orientation

Both Experiments 2 and 3 exhibited a significant feature-specific effect confined to targets specified in the color dimension. In Experiment 2, significant feature-specific ITF was observed for color targets, but not for orientation targets. This differential effect is consistent with Müller et al. (1995), who found no ITF with orientation targets (within-dimension condition). In Experiment 3, "4" responses were significantly faster when targets on all color trials were featurally homogeneous (all red or all blue) rather than heterogeneous (red and blue targets); no such difference was evident on all orientation trials.

Similar feature-specific effects have recently been reported by Nothdurft (1993) and Wolfe, Chun, and Friedman-Hill (1995). For example, Nothdurft's subjects were presented with brief displays (terminated by masks) of an array of 11×11 items which contained either 3 or 4 items that were featurally distinct from the other (background) items. In particular, each of the featurally distinct items had the same local contrast to the surrounding background items, but these distinct items could vary in their actual featural values. Nothdurft separately examined a number of dimensions along which items could differ (orientation, motion direction, and color). For the orientation dimension, the featural similarity of the distinct items had no effect on subjects' detection performance (see also Nothdurft 1991, 1992). However, for the color dimension, performance was improved when the featurally distinct items were featurally identical. Nothdurft (1992) interpreted this finding as indicating that "grouping of color items is not exclusively based on feature contrast but depends to some extent also on the features themselves" (p. 1952).⁷

An alternative interpretation of this feature-specific color effect in terms of dimensional weighting would be to conceptualize the color dimension as further subdivided into subdimensions representing broad categories of color (such as red, green, and blue). That is, the dimensional weighting account would also apply *within* the encompassing color dimension. One possible reason for the semiindependent representation of primary colors is likely to be found in the fact that color is the result of comparisons. The color of a surface is determined not only by the wavelength composition of the light reflected from it, but also by that reflected from surrounding surfaces. This re-

quires the lightness information to be coded in separate (narrow-band) channels before being submitted to a comparison stage (see Land, 1977; Zeki, 1993).

REFERENCES

- ATKINSON, J., & BRADDICK, O. (1989). "Where" and "what" in visual search. *Perception*, **18**, 181-189.
- CAVE, K. R., & WOLFE, J. M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology*, **22**, 225-271.
- CORBETTA, M., MIEZIN, F. M., DOBMEYER, S., SHULMAN, G. L., & PETERSEN, S. E. (1990, June). Attentional modulation of neural processing of shape, color, and speed in humans. *Science*, **248**, 1556-1559.
- CORBETTA, M., MIEZIN, F. M., DOBMEYER, S., SHULMAN, G. L., & PETERSEN, S. E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: Functional anatomy by positron emission tomography. *Journal of Neuroscience*, **11**, 2383-2402.
- DESIMONE, R., & UNGERLEIDER, L. G. (1989). Neural mechanisms of visual processing in monkeys. In F. Boller & J. Grafman (Eds.), *Handbook of neuropsychology* (Vol. 2, pp. 267-299). Amsterdam: Elsevier, North-Holland.
- DUNCAN, J., & HUMPHREYS, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, **96**, 433-458.
- FOLK, C. L., EGETH, H., & KWAK, H.-W. (1988). Subitizing: Direct apprehension or serial processing? *Perception & Psychophysics*, **44**, 313-320.
- GOODALE, M. A., & MILNER, A. D. (1992). Separate visual pathways for perception and action. *Trends in Neurosciences*, **15**, 20-25.
- 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.
- KAHNEMAN, D., TREISMAN, A. M., & GIBBS, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, **24**, 175-219.
- KAPTEIN, N. A., THEEUWES, J., & VAN DER HEIJDEN, A. H. C. (in press). Search for a conjunctively defined target can be selectively limited to a color-defined subset of elements. *Journal of Experimental Psychology: Human Perception & Performance*.
- KOCH, C., & ULLMAN, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, **4**, 219-227.
- LABERGE, D. (1990). Thalamic and cortical mechanisms of attention suggested by recent positron emission tomography experiments. *Journal of Cognitive Neuroscience*, **2**, 358-372.
- LABERGE, D., & BROWN, V. (1989). Theory of attentional operations in shape identification. *Psychological Review*, **96**, 101-124.
- LAND, E. H. (1977, December). The retinex theory of color vision. *Scientific American*, **237**, 108-128.
- MALKJOVIC, V., & NAKAYAMA, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, **22**, 657-672.
- MÜLLER, H. J., HELLER, D., & ZIEGLER, J. (1995). Visual search for singleton feature targets within and across feature dimensions. *Perception & Psychophysics*, **57**, 1-17.
- MÜLLER, H. J., & RABBITT, P. M. A. (1989). Spatial cueing and the relation between the accuracy of "where" and "what" decisions in visual search. *Quarterly Journal of Experimental Psychology*, **41A**, 747-773.
- NOTHDURFT, H.-C. (1991). Texture segmentation and pop-out from orientation contrast. *Vision Research*, **31**, 1073-1078.
- NOTHDURFT, H.-C. (1992). Feature analysis and the role of similarity in preattentive vision. *Perception & Psychophysics*, **52**, 355-375.
- NOTHDURFT, H.-C. (1993). The role of features in preattentive vision: Comparison of orientation, motion and color cues. *Vision Research*, **33**, 1937-1958.
- PARDO, J. V., PARDO, P. J., JANER, K. W., & RAICHEL, M. E. (1990). The anterior cingulate cortex mediates processing selection in the Stroop attentional conflict paradigm. *Proceedings of the National Academy of Sciences*, **87**, 256-259.
- SAGI, D., & JULESZ, B. (1985a). Detection versus discrimination of visual orientation. *Perception*, **14**, 619-628.
- SAGI, D., & JULESZ, B. (1985b, June). "Where" and "what" in vision. *Science*, **228**, 1217-1219.
- TREISMAN, A. M. (1988). Features and objects. The fourteenth Bartlett

- memorial lecture. *Quarterly Journal of Experimental Psychology*, **40A**, 201-237.
- TREISMAN, A. M., & GELADE, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, **12**, 97-136.
- TREISMAN, A. M., & GORMICAN, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychological Review*, **95**, 15-48.
- TREISMAN, A. M., & SATO, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception & Performance*, **16**, 459-478.
- TRICK, L. M., & PYLYSHYN, Z. W. (1993). What enumerations studies can show us about spatial attention: Evidence for limited capacity preattentive processing. *Journal of Experimental Psychology: Human Perception & Performance*, **19**, 331-351.
- UNGERLEIDER, L. G., & MISHKIN, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *Analysis of visual behavior* (pp. 549-586). Cambridge, MA: MIT Press.
- WOLFE, J. M. (1994). Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin & Review*, **1**, 202-238.
- WOLFE, J. M., CHUN, M. M., & FRIEDMAN-HILL, S. R. (1995). Making use of texton gradients: Visual search and perceptual grouping exploit the same parallel processes in different ways. In T. Pappathomas & A. Gorea (Eds.), *Linking psychophysics, neuropsychology, and computational vision* (pp. 189-197). Cambridge, MA: MIT Press.
- ZEKI, S. (1993). *A vision of the brain*. Oxford: Blackwell.

NOTES

1. Treisman and Gormican (1988) defined a dimension as a set of mutually exclusive values for any single stimulus attribute. So, for example, although a line can be both red and vertical (values of different dimensions), it cannot be both vertical and horizontal (mutually exclusive values on the same dimension).
2. Similarity comparisons within a dimension are scaled by the distance between two items. So the saliency signal for two dissimilar items will be greater when they occupy nearby locations than when they are farther apart.
3. A pilot study had shown that search for the five different targets presented in the within- and across-dimension conditions was equally efficient when the target identity was fixed.
4. Experiment 3 of Müller et al. (1995) manipulated the likelihood of the three targets in the across-dimension condition. Subjects were in-

formed that, on *present* trials, the target would be a right-tilted line on 80% of the trials and a large or a black bar on only 20%; there was a similar manipulation of target probabilities in the within-dimension condition. Subjects' RTs to the right-tilted line were as rapid in the across-dimension condition as in the within-dimension condition, with responses to the other (large and black) targets being slower by over 100 msec. The benefit for the likely dimension suggested that determining the target dimension was top-down penetrable. There was no appreciable cuing effect in the within-dimension condition, consistent with the idea that responses in this condition could be given without knowledge of the target's featural identity.

5. Recall that, in Experiment 2 of Müller et al. (1995), which required explicit elimination of potential target dimensions, the across-dimension condition did show a cost of some 20 msec relative to the control condition (and, unusually in visual search experiments, the *present* RTs were slower than the absent RTs). Although small, this cost cautions against a "strong" interpretation of the data of Müller et al.

6. There are several possible reasons for this added cost. One is that target detection requires only minor shifting of dimensional weight, whereas determining the target's dimensional identity requires a greater weight shift to home in on the target dimension. Alternatively, the weight shifts required are equivalent, but the additional cost may be due to the extra need to check the across-dimensional weight pattern to discern the source dimension of the target signal on the master map.

7. The differential homogeneity effect between all-color and all-orientation targets in Experiment 3 might simply be due to the similarity between the color features being lower than that between the orientation features, thereby disrupting their collective selection. However, inconsistent with such an account is the absence of an RT advantage for homogeneous over heterogeneous all-orientation targets (in fact, RTs to homogeneous all-orientation targets were, if anything, slower than RTs to heterogeneous all-orientation targets; see Table 5). Since there is no support for a quantitative, similarity-based, explanation from the data *within* the orientation dimension, it is implausible to provide an account for the differential homogeneity effect *between* the color and orientation dimensions. It is preferable to interpret this effect as reflecting a qualitative processing difference between color and orientation features.

(Manuscript received October 5, 1994;
revision accepted for publication May 23, 1995.)