



Segmentation of objects from backgrounds in visual search tasks

Jeremy M. Wolfe^{a,b,*}, Aude Oliva^{a,b}, Todd S. Horowitz^{a,b}, Serena J. Butcher^c,
Aline Bompas^d

^a Center for Ophthalmic Research, Brigham and Women's Hospital, USA

^b Department of Ophthalmology, Harvard Medical School, USA

^c Department of Psychology, Harvard University, USA

^d Laboratoire de Psychologie Expérimentale, Université René Descartes, Paris, France

Received 2 April 2002; received in revised form 12 August 2002

Abstract

In most visual search experiments in the laboratory, objects are presented on an isolated, blank background. In most real world search tasks, however, the background is continuous and can be complex. In six experiments, we examine the ability of the visual system to separate search items from a background. The results support a view in which objects are separated from backgrounds in a single, preattentive step. This is followed by a limited-capacity search process that selects objects that might be targets for further identification. Identity information regarding the object's status (target or distractor) then accumulates through a limited capacity parallel process. The main effect of background complexity is to slow the accumulation of information in this later recognition stage. It may be that recognition is slowed because background noise causes the preattentive segmentation stage to deliver less effectively segmented objects to later stages. Only when backgrounds become nearly identical to the search objects does the background have the effect of slowing item-by-item selection.

© 2002 Elsevier Science Ltd. All rights reserved.

1. Introduction

In a typical laboratory visual search task, observers look for a target that may or may not be present among some number of distractor items. The experimenter typically measures the accuracy of response and the time required to make the response (reaction time, or RT). The number of items (the “set size”) is often varied, allowing an RT × set size function to be measured. The slope of this function represents the added cost of each additional item and can be seen as a measure of the efficiency of the search. The intercept of this function represents the “fixed costs” of processes such as those involved in the motor response. A lot has been learned from the large body of research using this paradigm (reviewed in Driver & Frackowiak, 2001; Sanders & Donk, 1996; Wolfe, 1998a). However, the laboratory task is necessarily artificial, leading to concerns about ecological validity.

Real world search tasks are ubiquitous—from the natural (Where is the raspberry on this bush?) to the artificial (Is there a weapon in this carry-on bag?)—and they differ from the laboratory versions in ways that need to be investigated if we wish to generalize from the lab to the world. This paper concentrates on one difference. In laboratory search tasks, items are usually scattered over a uniform background. Very little work needs to be done to identify the set of task-relevant items. In real world searches, however, this is not the case. Targets and potential distractors are spread over a continuous and usually heterogeneous background. Before an item can be identified as a target or rejected as a distractor, it must first be distinguished from the background. How is this done?

Broadly speaking, background complexity might make search more difficult in four different ways. Our framework for this discussion begins with the two-stage conceptualization of search put forward by Neisser (1967) and developed by Treisman and Gelade (1980). They describe a “preattentive” stage of processing in which the entire image is processed in parallel. As shown in Fig. 1, when an observer searches for a target, it could be that segmentation of the image into background and

* Corresponding author. Address: Center for Ophthalmic Research, Brigham and Women's Hospital, USA. Fax: +1-617-277-6717.

E-mail address: wolfe@search.bwh.harvard.edu (J.M. Wolfe).




Fig. 1. The effect of adding “noise” to the background could be to lengthen the time it takes for preattentive mechanisms to separate the candidate objects from the background.

search items is a purely preattentive operation that takes longer when the background is more complex (or noisier).

The second stage in the classic two-stage account is a limited capacity stage in which items or groups of items are selected for further processing. If selection cannot begin until the initial, preattentive clean-up is completed, then the addition of noise to the background should only produce an additive increase in RT, as cartooned in Fig. 2a. A second possibility is that the initial “clean-up” of noisier backgrounds might be less effective. In that case, the preattentive stage might permit some pieces of the background to be selected as search items as well as the “official” search items. These items, misinterpreted as potential targets, would effectively increase the set size (and, thus the mean RT) by a constant amount. This would produce an additive increase in RT which should be greater for target-absent trials, since a greater pro-

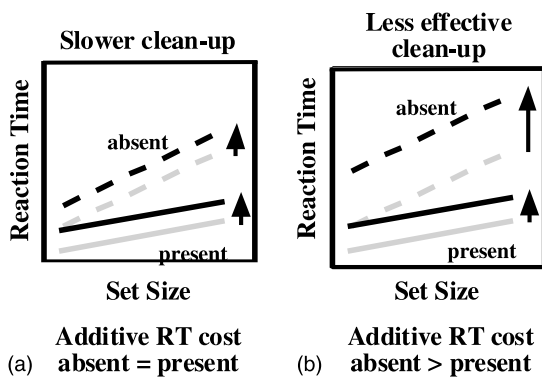


Fig. 2. A harder (noisier, more complex) background might either take longer to separate from search items (2a) or might be imperfectly separated (adding an average of N extra items to each trial). In either case, the effect of background on RT would be additive with the effects of set size. (a) Slower clean-up predicts the same effect on target-present and target-absent trials. (b) Less effective clean-up predicts a larger additive effect on target-absent trials.

portion of the search items must be examined on target-absent trials (Fig. 2b).

Explaining the third possibility requires some more theoretical background. The two-stage architecture of Neisser and Treisman has required some modification over the years. For example, it has proven useful to recognize that these are not independent stages; the preattentive stage provides information that guides the selection of items in the next stage (Egeth, Virzi, & Garbart, 1984; Hoffman, 1979; Wolfe, 1994a; Wolfe, Cave, & Franzel, 1989). It is also becoming clear that the selection stage represents a bottleneck between a parallel, preattentive stage and another parallel, if limited-capacity, recognition stage. Evidence from visual search experiments indicates that items are selected for processing at an average rate of one every 25–50 ms (reviewed in Wolfe, 1998b). However, no credible evidence suggests that items can be processed to the point of identification in that amount of time (e.g. Duncan, Ward, & Shapiro, 1994; Raymond, Shapiro, & Arnell, 1992; Thorpe, Fize, & Marlot, 1996; VanRullen & Thorpe, 2001; Ward, Duncan, & Shapiro, 1996). This strongly suggests that the slope of the RT \times set size function reflects the rate at which items can be fed into some sort of pipeline process (e.g. Harris, Shaw, & Altom, 1985) where several hundred ms might be required to accumulate enough information to identify an item as target or distractor. Details of this argument can be found elsewhere (Moore & Wolfe, 2001). A carwash can serve as a metaphor. Cars go in and come out, say, once a minute but it might take 5 min to fully “process” each car. As a consequence, multiple items/cars can be in the pipeline at the same time even if they are selected sequentially. Thus, this stage can be considered a limited-capacity parallel stage. One way to model such a stage is as a diffusion process (Palmer & McLean, 1995; Ratcliff, 1978) as illustrated in Fig. 3.

As each item is selected, information regarding its status (target or distractor) begins accumulating. An identification decision is made when that information reaches the target threshold (upper bound in Fig. 3) or the distractor threshold (lower bound). The rate of information accumulation can be described by the slope of the function showing the information accumulating over time (indicated by the tilted arrows in Fig. 3). A steeper




Fig. 3. It could take longer for a parallel recognition device to do its work if the initial segmentation of the image is less effective when there is more noise in the background.

slope indicates a more rapid progression from selection to an identification decision. The effect of a noisier/more complex background might be to reduce that rate of accumulation. This is shown in the second half of Fig. 3 as a reduction in the slopes of the accumulation functions. Since this reduction occurs in a parallel stage, the effect of the background on RT would be additive with the effect of set size; changing the slope of the accumulation function does not change the slope of the RT \times set size function. Successful search ends when the target is found. As long as items are selected at the same rate from easy and hard displays, any extra RT cost would reflect only the added time to make a decision about the target item. The effects on target-absent trial RTs are a more complex issue that will be taken up later in this paper.

This slowed accumulation could result from a less effective segmentation process. If, rather than adding virtual items to the set size (the second possibility described above), the segmentation process delivers up candidate objects which carried with them some extraneous background information, such objects would take more time and effort to recognize than “cleaner” objects. It is important to reiterate that this slowing of recognition is assumed to be separate from any slowing of item-by-item selection. We are assuming that attentional selection is a bottleneck between initial parallel processing of visual input that includes parallel segmentation of the scene into component objects and a subsequent, parallel recognition stage in which several items can progress toward recognition at the same time.

The fourth possibility is that the background might exert its effect on the selection stage. Perhaps it takes longer to select each item because that item must be more laboriously and individually separated from the background as illustrated in Fig. 4. This would result in an increase in slopes as shown in Fig. 5.

To summarize, then, there are three hypothetical factors that could produce additive effects on search RT as the background becomes noisier. (1) Preattentive segmentation or “clean-up” could take longer, (2) segmentation might be imperfect and might add extra ob-




Fig. 5. The effect of slowing the segmentation of each item in an attentive stage of processing would be to increase the RT versus set size slope for both target-absent and target-present trials.

jects to the search, (3) segmentation might be imperfect, lengthening the time required to identify objects in a later, parallel recognition stage. The fourth alternative is that noise in the background could have its effect on the time required to select individual items for further identification. In this case, the slopes of the RT \times set size function should become steeper.

To test these hypotheses, we had observers search for targets on a variety of backgrounds, embodying several different operational definitions of background complexity. In Experiments 1 and 2, complexity is varied by changing the number of additional background objects in a realistic scene. Complexity is manipulated by varying the similarity of background features to features of the search items in Experiments 3 and 6. Similarity between background and search item spatial frequency spectra is varied in Experiments 4 and 5. In general, we found that more complex backgrounds reliably imposed only additive RT costs. The slopes of RT \times set size functions were largely unaffected by an increase in background complexity. This result is inconsistent with that predicted by a slowed selection rate (Option 4: Figs. 4 and 5). That leaves three parallel-stage possibilities. Two of these are preattentive: (1) a longer initial clean-up stage (see Fig. 1); or (2) a less effective clean-up stage that adds virtual items to the search set. The added search items of the latter would be stimuli that the visual system treats as candidate targets even if they were not placed in the image by the experimenter. The remaining possibility is that a less effective initial clean-up stage causes slower accumulation of information in the identification stage that follows attentional selection (see Fig. 5). The details of the additive RT cost further constrain the possibilities. If the cost were due to a single preattentive clean-up step, then that cost should be the same on target-present and target-absent trials. It is not: the cost is greater on target-absent trials. We will use this fact to argue against a strong version of Parallel Option 1. In Experiment 5, subjects distinguish between the presence of one or two targets in a display. The results of this experiment can be used to distinguish between the remaining two sloppy clean-up hypotheses. Does sloppy clean-up add extra candidates for attentional selection (Option 2, Fig. 2b) or does it entail the




Fig. 4. The effect of background noise might be to lengthen the time required for each deployment of attention.

selection of less well-segmented objects, thus slowing accumulation of information in the limited capacity “carwash” (Option 3, Fig. 3)? We will argue that the data are best explained by assuming Option 3, that imperfect segmentation of more complex backgrounds slows the rate at which information accumulates in a limited-capacity, parallel stage following attentional selection. Only when backgrounds truly camouflage items do we observe effects on $RT \times$ set size slope (Experiment 6 and, perhaps, in Experiment 3).

These experiments comparing the effects of backgrounds of different complexity are similar in their logic to experiments from a number of labs in which the search items were degraded, generally by superimposed noise of some sort (e.g. Egeth & Dagenbach, 1991; Logan, 1975; Lu & Doshier, 1998; Mewhort, Johns, & Coble, 1991; Pashler, 1984; Pashler & Badgio, 1985; Swensson & Judy, 1981). For example, Becker and Pashler (2001) placed a noise mask over their stimuli and found that moderate amounts of noise produced an additive change in RT while greater amounts of noise increased the slope of the $RT \times$ set size function. These masking experiments ask a somewhat different question from the question asked here. The degraded stimulus experiments can be seen as part of the larger project to understand visual search in terms of signal-detection theory (e.g. Eckstein, Thomas, Palmer, & Shimozaki, 1996; Graham, Kramer, & Haber, 1985; Kinchla, 1977; Palmer, Verghese, & Pavel, 2000; Verghese, 2001). Many standard search tasks involve discriminations that are trivially easy if the observer is confronted with just a single item (e.g. Is it red or green, a T or an L?). These discriminations can be made more difficult by reducing the difference between target attributes and distractor attributes (e.g. Nagy & Sanchez, 1990, for color stimuli or Foster & Ward, 1991a,b for orientation). Alternatively, discriminations can be made difficult by adding noise to the items as in the studies cited above. Once the discriminability drops below some critical level, it takes more time to decide if a given item is a target or a distractor and the slope of the $RT \times$ set size function increases. Note that, in these masking experiments, there is no particular problem in identifying the set of items. The difficulty lies in discriminating target from distractors.

In our experiments, in contrast to the masking experiments, the items are not degraded. Given a single item, target versus distractor discriminability will be very high. We are adding noise *between* the objects, in the background. We are increasing the possibility that attention will be allocated to a region where there are no task relevant objects at all. Both masking and background effects can co-occur in real search tasks. Consider the problem of screening an X-ray image for tumors (e.g. Kundel, 1991; Nodine, Krupinski, & Kundel, 1993; Nodine, Kundel, Lauver, & Toto, 1996;

Samuel, Kundel, Nodine, & Toto, 1995; Swensson, 1980). The radiologist needs to determine the loci that should be selected for attention. This is a problem of segmenting candidate targets from the background image. Only then can the radiologist determine if a specific item is a tumor. The latter is a version of the signal detection problem described above. Most prior work has dealt with this second step. We are addressing the first.

2. General methods

Our experimental strategy was the same across the six experiments of this paper. We measured $RT \times$ set size functions for a standard visual search task (search for a T among L's) and varied the background complexity. All of the experiments described hereafter were programmed with Matlab 5.1 (MathWorks) using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Stimuli were presented on a 21-in. Mitsubishi monitor running at a refresh rate of 75 Hz, and controlled by a Power Macintosh G4. Targets and distractors subtended, on average, 1.5° of visual angle from a 50 cm viewing distance (except Experiment 3). Background images subtended up to 30° of visual angle. In each experiment, subjects performed a standard visual search task with target and distractors superimposed on a background of different levels of complexity. Targets were present on 50% of trials. Observers were fully informed about the identity of targets and distractors. Observers were instructed to answer as rapidly and accurately as possible and they were given feedback about their response accuracy after each trial. All observers were between the ages of 18 and 50 and had vision of 20/25 or better with correction as-needed. All participants passed the Ishihara test for color blindness. They gave informed consent and were paid \$9 per hour for their participation. Departures from these general methods will be noted as-needed.

3. Experiment 1: refrigerator magnets on messy desks

In Experiment 1, we varied the complexity of the background by manipulating the “messiness” of a desk scene (see Fig. 6). This experiment was intended to simulate the common, realistic situation of a search for a specific object (here, a refrigerator magnet) among similar objects (other refrigerator magnets) lying on a desk.

3.1. Method

Three background desks of different levels of complexity were composed using the 3D scene synthesis




Fig. 6. Front view of the three desk backgrounds: (a) empty desks, (b) neat desk and (c) messy desk.

software *Home Designer 3.0* (from Data Becker, Inc.). These backgrounds will be referred to as *empty*, *neat*, and *messy* desks (see Fig. 6). Here, greater complexity means a greater number of irrelevant objects on the desk (with resulting increases in density and occlusion). Eight rotated versions (every 45°) of the letter T (target) and the letter L (distractors) were designed to look like beveled magnets that you might put on a refrigerator (see Fig. 7). Twelve participants performed 1080 experimental trials preceded by 54 practice trials, searching for the T among 4, 8, or 12 letters. A central fixation cross was presented for 500 ms prior to each trial. The desk background and search stimuli were displayed until the observer responded.

3.2. Results

RTs less than 200 ms and greater than 4000 ms were labeled as errors. Error rates averaged 3%, including the out-of-range RT trials. Error rates did not differ significantly across conditions. Mean RTs for correct trials are plotted against set size, for the three backgrounds, in Fig. 8. The figure shows an additive RT effect of desk complexity. There are substantial changes in mean RT and only minor differences in slopes across the three background desk conditions. Moreover, it is clear that the effect of our complexity manipulations on mean RT is greater for the target-absent trials than for the target-present trials.

These impressions are supported by ANOVA. An ANOVA combining target-present and target-absent data reveals a significant interaction of target presence/absence with background ($F(2, 22) = 23.3$, $p < 0.0001$), reflecting the larger effect of background complexity on target-absent trials' RTs. It is more informative to perform separate ANOVAs on the target-present and target-absent trials. Analysis of correct, target-present trials reveals significant main effects of background ($F(2, 22) = 35.48$, $p < 0.0001$) and set size ($F(2, 22) = 141.83$, $p < 0.0001$) on mean RTs, but no differences in the RT \times set size slopes between conditions (i.e., background \times set size interaction: $F < 1$). The same pattern is seen with ANOVAs performed on correct target-absent trials: a large effect of background complexity on mean RTs, ($F(2, 22) = 32.2$, $p < 0.0001$), as well as a significant set size effect ($F(2, 22) = 145.6$, $p < 0.0001$), but no slope differences between conditions ($F(4, 44) = 1.6$).

3.3. Discussion

The results of Experiment 1 show that observers handled these backgrounds in a parallel manner, independent of set size effects. There have been a variety of previous studies involving visual search in approximations of “real” scenes (Biederman, Glass, & Stacy, 1973; Carmody, Scanlon, & Dasaro, 1990; Kingsley, 1932; Wolfe, 1994b). These include studies of eye movements



Fig. 7. This figure shows examples of the search task in the complex desk background condition. The *beveled* refrigerator magnet stimuli of Experiment 1 are shown on the left and the *Post-It*® stimuli of Experiment 2 are shown on the right.




Fig. 8. RT \times set size functions for the three background “desks” of Experiment 1. (—■—) Empty-present, (---□---) empty-absent, (—●—) neat-present, (---○---) neat-absent, (—◆—) messy-present, (---◇---) messy-absent.

in search (e.g. Ballard, Hayhoe, Pook, & Rao, 1997; Kroll, 1992; Melcher, 2001; Melcher & Kowler, 2001; Rayner & Pollatsek, 1992), as well as the current interest in search for changes in scenes (Rensink, O’Regan, & Clark, 1997; Simons & Levin, 1997). As a group, these studies have shown that the patterns of results seen with artificial search tasks can also be obtained with more realistic stimuli. But despite these substantial gains in ecological validity, the question of how a set of search items is extracted from a continuous scene has remained unanswered. Indeed, one of the difficulties with real scenes as search displays has been the problem of defining set size.

As mentioned above, the present laboratory task is somewhat different from those that have preceded it. We may not know exactly how many items an observer considers to be candidate targets for the search task. However, as experimenters, we can control the number of items that we have placed in the scene and we can conclude that adding four magnet distractors to the display adds four items to the set considered by the visual system (e.g. the difference between nominal set size 4 and 8). This allows us to use the slope of the RT \times set size function to estimate the cost of each additional item even if the intercept of that function might be inflated by the effects of what we, as experimenters, consider to be the background.

Our search task can be seen as a form of “Guided Search” (Wolfe, 1994a, 2001; Wolfe et al., 1989; Wolfe & Gancarz, 1996). The Guided Search model builds on the two-stage architecture of Neisser (1967) and Treisman and Gelade (1980). The two stages consist of a “preattentive” processing stage in which a limited number of basic features (color, size, several depth cues, etc.) can be extracted in parallel across the visual field and an “attentive” stage in which individual items (or

perhaps groups of items) are selected for further analysis by limited-capacity processes. Guided Search argues that information from the preattentive stage guides the deployment of attention in the second stage. The more effective the guidance, the more efficient the search (see also Egeth et al., 1984; Hoffman, 1979; Tsotsos et al., 1995).

In Experiment 1, attention can be guided to the T’s and L’s by their basic feature attributes (probably color, size, and some still ill-defined form features). Within the set of T’s and L’s, no further guidance is possible and search probably proceeds in a serial, item-by-item manner (Horowitz & Wolfe, 1998, 2001; Kwak, Dagenbach, & Egeth, 1991; Wolfe, 1998a; Woodman & Luck, 1999). The effect of background complexity was shown to be additive with increasing set size. This argues against any hypothesis that holds that rate of attentional selection of items is slowed. Also, the additive effect is larger for the target-absent trials. This fact argues against the proposal that observers wait longer to start searching when the background is more complex (Figs. 1 and 2). As discussed earlier, two accounts predict an additive increase in RT that would be greater for target-absent than for target-present trials. Those are the “less-effective clean-up” and the “slowed accumulation” accounts. In the next three experiments, we explore the generality of the above result. Experiment 5 will distinguish between these two options.

4. Experiment 2: a messy desk with Post-It® notes

The primary goal of Experiment 2 was to replicate Experiment 1 with stimuli presenting a slightly different segmentation challenge for the visual system. In this experiment, the T’s and L’s were drawn on pieces of paper resembling Post-It® notes scattered across the desk backgrounds used in the previous experiment (see Fig. 7b). These Post-It® stimuli added right angle corners similar to the books and journals found as objects in the background. This similarity might be expected to make the search task more difficult. On the other hand, the T’s and L’s were now placed on locally blank backgrounds, which might be expected to improve performance. Methods were otherwise identical to those of Experiment 1. Twelve observers participated in this study.

4.1. Results and discussion

RTs less than 200 ms and greater than 4000 ms were coded as errors. Error rates averaged 3.7% including the out-of-range RT trials. These rates did not significantly differ across conditions. Mean RTs for correct trials are plotted against set size, for the three backgrounds, in Fig. 9. The overall pattern of results is similar to the




Fig. 9. RT × set size functions for the three background desks of Experiment 2. In this experiment, there is a large main effect of background complexity on mean RT and a reliable effect on the slope of the RT × set size functions. (—■) Empty-present, (-□-) empty-absent, (—●) neat-present, (-○-) neat-absent, (—◆) messy-present, (-◇-) messy-absent.

pattern of results in Experiment 1. The primary effect of background complexity appears to be an additive increase in RT with a substantially larger effect on target-absent trial RTs. There are large main effects of background and set size in the target-present data ($F(2, 22) = 79.52, p < 0.0001$ and $F(2, 22) = 118.83, p < 0.0001$, respectively). However, in contrast to Experiment 1, Experiment 2 shows a statistically reliable interaction of background and set size ($F(2, 22) = 3.75, p < 0.02$). Slopes are slightly shallower with an empty desk background than with a messy desk background. The same pattern of results was observed for the target-absent trials sets (Main effects: background $F(2, 22) = 88, p < 0.0001$, set size $F(2, 22) = 103, p < 0.0001$; Interaction: $F(4, 44) = 7.76, p < 0.0001$).

The increases in slope with background complexity may indicate observers' decreased ability to discriminate targets from distractors. Since the backgrounds remain the same across the two experiments, it is the nature of the search items that must make the difference. A likely candidate is the structure of those items. In Experiment 1, the T's and L's are objects in their own right. In Experiment 2, the squares of yellow paper are the objects and the letters might be seen as surface markings on those objects; i.e., parts of a larger whole. Part structure is known to have an effect on search (Bilsky & Wolfe, 1995; Enns & Kingstone, 1995; Humphreys, Cinel, Wolfe, Olson, & Klempe, 2000; Wolfe, Friedman-Hill, & Bilsky, 1994; Xu & Singh, in press) with properties of the whole object usually being more accessible than properties of the part (Navon, 1977). Either this, or some other factor appears to be modestly slowing the attentive stage in its effort to select each potential target.

As a consequence, slopes increase modestly. The next experiments move away from naturalistic stimuli in an effort to isolate factors that might hinder the segmentation of scenes into a set of search items and a background to be ignored.

5. Experiment 3: “Brick Wall” backgrounds—effects of overt features

With natural stimuli like those in the first two experiments, it is difficult to know how to manipulate the similarity of search items and background. In Experiment 3, this issue was addressed by making the backgrounds out of the same picture elements as the search items (T's among L's). We manipulated the type of junctions and the local terminators in the backgrounds. Six of the eight backgrounds are shown in Fig. 10.

The background of Fig. 10a contains T-junctions but no line terminators. Fig. 10b contains terminators but no actual T-junctions. Fig. 10c contains both. Fig. 10d contains X-junctions but no terminators, Fig. 10e contains terminators but no actual X-junctions, and Fig. 10f contains both. In addition, two control backgrounds were used: a blank background and a background containing only the horizontal lines of the backgrounds shown in Fig. 10. Note that the overall complexity of the wall backgrounds is about the same in all conditions except for the controls. The total amount of background contour and the number of “cells” in which to place a T or L remains constant. Only the junctions and terminators change.

5.1. Method

Target letter T and distractor letters L were presented in one of the four cardinal orientations. They subtended $2.5^\circ \times 2.5^\circ$ of visual angle. Set sizes were 3, 6 and 9. The backgrounds were composed of a 4×4 grid that




Fig. 10. Six of the eight backgrounds used in Experiment 3.

subtended $19^\circ \times 19^\circ$ of visual angle. Search stimuli were randomly assigned to the 16 locations of the grid.

On each trial, a fixation cross appeared and was followed, after 200 ms, by one of the backgrounds containing search items. This display remained visible until the observer made a response regarding the presence of a T. Participants performed 96 practice trials, followed by 2400 experimental trials (two sessions of 1200 trials). Background, set size, and target presence/absence were all randomized within the same experimental block. Subjects were told to respond as quickly and accurately as possible and they were given feedback about their accuracy. Thirteen observers were tested in all the conditions of this experiment.

5.2. Ranking control experiment

To assess the subjective complexity of the backgrounds used in this experiment, we asked two separate groups of observers to rank order paper copies of the stimuli according to perceived visual complexity. Specifically, one group was asked to “order the pictures according to their complexity”. The other was asked “How hard do you think it would be to find a T on this background?” Observers responded by physically lining up the images from least to most complex.

5.3. Results

Results are shown in Table 1. Complexity ratings shown in the table were obtained from the ranking control experiment by assigning the value “1” to the image rated least complex by an observer, “2” to the next and so on up to 8. These values were then averaged

across observers and rescaled from 0 to 1. Since this is an ordinal scale, these averages should be treated with caution.

The data were analyzed for effects of “junction” (“X”, “T”, or “none”) and of terminators (“present” or “absent”) on RT and RT \times set size slope. Inspection of Fig. 11 suggests that, as with Experiments 1 and 2, the main effect of background complexity is primarily an additive one on mean RT. Again, the slope, or rate of the search, did not appear to vary with backgrounds, indicating that the background junctions and terminators did not alter search efficiency. Moreover, it is clear that the largest impact on mean RTs comes from the presence of local terminators while the nature of background intersections (X or T) has a more modest effect, largely limited to the conditions without line terminators. These impressions are borne out by statistical analysis.

Looking first at the slope data, Fig. 11 shows that there are only modest differences in slopes, with some tendency for steeper slopes with terminators present and with junctions. An ANOVA on target-present slope data bears out the impression that these are not large effects. There is no significant effect on slopes of terminator presence/absence ($F(1, 12) = 2.9$, $p = 0.11$) or junction ($F(2, 24) = 2.1$, $p = 0.14$), or their interaction ($F(1, 12) = 0.011$, $p = 0.99$). In the target-absent slope data, there is a significant effect of junction ($F(2, 24) = 8.6$, $p = 0.0015$). The presence of X-junctions seems to make subjects a bit more cautious about giving up on a search. There is no significant effect on slope of terminator presence/absence ($F(1, 12) = 1.5$, $p = 0.25$) or an interaction between terminator presence/absence and junction ($F(1, 12) = 2.4$, $p = 0.11$).

Table 1
Data for Experiment 3

Condition	Perceived complexity ratings	Ease of search ratings	Slope (ms/item)	Mean RT (s.e.m.)	Intercept (ms)	Error (%)
<i>Target present</i>						
1. T-junctions w/terminators	0.82	0.89	24.3	637 (36)	491	5.94
2. T-junctions no terminators	0.65	0.74	21.8	596 (34)	465	3.5
3. X-junctions w/terminators	0.61	0.53	21	628 (38)	501	5.49
4. X-junctions no terminators	0.36	0.29	17.5	587 (33)	482	4.12
5. Terminators with broken T-junctions	0.58	0.75	21	664 (42)	538	5.48
6. Terminators with broken X-junctions	0.35	0.65	18.3	618 (35)	508	5.70
7. Control horizontal lines	0.13	0.14	16.2	580 (30)	482	4.98
8. Control blank field	0	0	17.5	561 (29)	456	3.64
<i>Target absent</i>						
1. T-junctions w/terminators	0.82	0.89	45.7	815 (65)	541	2.43
2. T-junctions no terminators	0.65	0.74	48.2	763 (58)	474	2.29
3. X-junctions w/terminators	0.61	0.53	47.5	794 (63)	509	1.59
4. X-junctions no terminators	0.36	0.29	39.7	702 (49)	467	1.48
5. Terminators with broken T-junctions	0.58	0.75	45.6	888 (81)	615	1.99
6. Terminators with broken X-junctions	0.35	0.65	41.7	767 (56)	517	2.5
7. Control horizontal lines	0.13	0.14	38.8	698 (50)	465	1.97
8. Control blank field	0	0	36.2	672 (42)	455	2.08

Ratings are rescaled subjective ratings of the stimulus complexity (see text for details of the two versions). Other columns are self-explanatory.

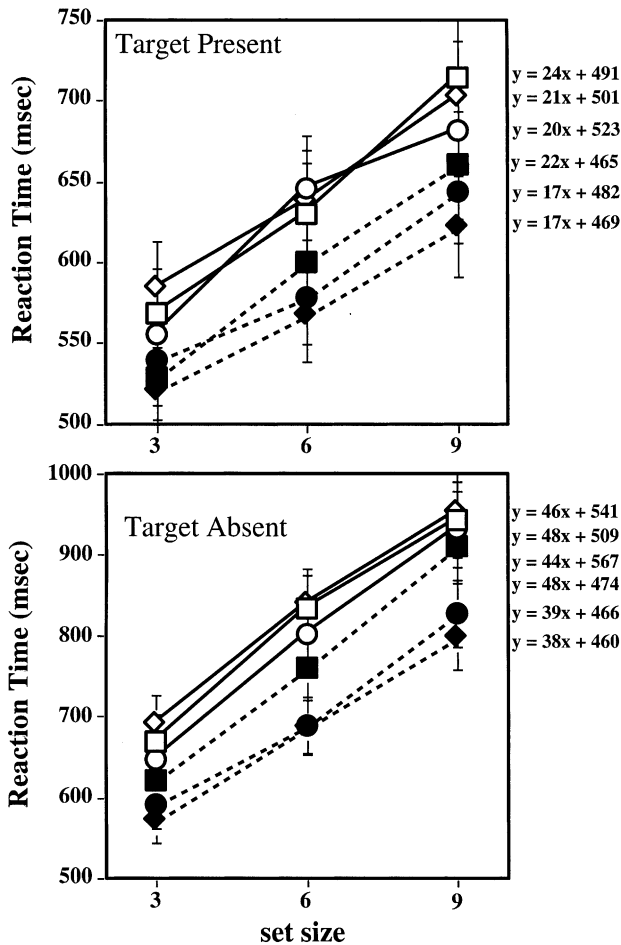


Fig. 11. Results for Experiment 3. Note that the scale for target-absent trials is double that for target-present trials. Term: (—□—) T-junction; (—○—) X-junction; (—◇—) no junctions. No term: (—■—) T-junction; (—●—) X-junction; (—◆—) no junctions.

Turning to the mean RT data, in a combined ANOVA, the interaction of background and target presence/absence is significant (for the terminator variable: $F(1, 12) = 10.5, p < 0.01$; for the junction variable: $F(2, 24) = 8.1, p < 0.01$) reflecting the larger effect of background on target-absent RTs. Again, it is more informative to perform separate ANOVAs for target-present and target-absent data. Here, we see a large effect of terminator presence/absence on RT in the target-present data ($F(1, 12) = 65.0, p < 0.0001$). The effect of junction does not quite reach the 5% significance level ($F(2, 24) = 2.7, p = 0.08$), but the interaction of terminator presence/absence and junction does ($F(2, 24) = 5.1, p = 0.01$).

The effect of set size on mean RT is unsurprisingly significant ($F(2, 24) = 51.6, p < 0.0001$). The small slope effects are reflected in modestly significant interactions of terminator presence/absence and set size ($p = 0.05$) and terminator presence/absence, junction, and set size ($p = 0.03$). The target-absent data are comparable to the target-present data described above. All of the main

effects are highly significant (all $p < 0.0001$). The interaction of terminator presence/absence and junction is also significant ($p < 0.0001$). Of the interactions with set size, reflecting effects on the slope, only the set size/X-junction interaction is modestly significant ($p = 0.04$).

RTs less than 200 ms and greater than 4000 ms were labeled as errors. Error rates averaged 3.5%, including the out-of-range RTs. An ANOVA showed that errors did not differ across junction types, but were reliably, if only slightly, higher for backgrounds containing terminators (3.9%) than those without terminators (3%), ($F(2, 24) = 10.2, p < 0.01$).

Turning to the subjective rating tasks, subjects have a reasonably accurate notion of the relative effects of the backgrounds. When asked about the perceived complexity of the images, average subject rankings have a Spearman rank correlation of 0.71 with the means of the RT data. When asked the more directed question about the difficulty of finding a “T” on a specific background, the correlation increases to 0.88. In agreement with the results, observers judged that the general complexity of the image increased with the presence of local terminators. The two sets of rankings are strongly related to each other (Spearman’s $\rho = 0.83$). Note that the correlations of rankings and RTs are correlations between the average data from two different groups of subjects. One group performed complexity ratings. Another group performed the search task.

Interestingly, even though the effects of background on slope are marginal at best, perceived complexity and perceived difficulty are strongly related to the slope (Spearman rank correlation of 0.92 and 0.88, respectively). This suggests that there might be a real effect of background on slope and that subjects are introspectively sensitive to it. Is it possible that the main effect of background characteristics on RT which we observed is in fact due to a real effect on the slope? If the slope increases by X ms per item and there are N search items, then mean RT will increase by XN ms on average. The resulting change in mean RT might be statistically reliable while the slope change was not, because slope measurements are inherently less sensitive than measurements of RT. This seems unlikely. Analysis of the intercept data (based on the same number of data points as the slope data) reveals a powerful effect of terminator presence/absence ($F(1, 12) = 39.12, p < 0.0001$), and an interaction of junction type with terminator presence/absence ($F(2, 24) = 4.15, p < 0.05$). There is also a strong trend towards an interaction between target presence/absence and terminator presence/absence ($F(1, 12) = 4.74, p = 0.05$). As with the mean RT, the effects of the background are stronger on target-absent data, a separate analysis of target-present data reveals only a powerful effect of terminator presence/absence ($F(1, 12) = 93.58, p < 0.0001$). However, the target-absent data yield a significant main effect of junctions ($F(1, 12) = 17.98,$

$p < 0.005$) and a junction by terminator presence/absence interaction ($F(2,24) = 4.90$, $p < 0.05$). Thus, the background effects we observe on mean RT seem to be driven primarily by statistically reliable changes to the intercept of the RT \times set size function, rather than by subtle changes in slope “leaking” into the mean RT measure.

5.4. Discussion

As in the previous experiments, these results show that the main effect of background is an additive effect on mean RT. Of course, we cannot conclude that there is no effect at all on search rate. Given that perceived difficulty and perceived complexity both correlate strongly with the RT \times set size slope, it is likely that there is in fact a subtle change to the slopes which we could not detect statistically. The slope effect may be relatively weak here because the backgrounds do not differ sufficiently to produce a statistically reliable effect. The issue of slope effects will be revisited in Experiment 6, where we are able to induce a detectable change in search rate. Nevertheless, the fundamental result of this experiment is that the additive effects of these backgrounds on mean RT are much more powerful than any effects on mean RT \times set size slope. Also, line terminators have a more substantial effect on RT than intersections.

These additive effects implicate parallel processes either before or after item-by-item selection. As with the desk scene experiments, it is possible to entertain several accounts of the added time required with more complex backgrounds. It could be that it simply takes longer to preattentively separate items from a complex background. However, as discussed before, the larger effect of background complexity on target-absent RTs argues against a fixed, preattentive waiting period that gets longer as backgrounds become more complex. That leaves two other potential sources of the additive effect on mean RT with increasing set size. It could be that each complex background adds some number of candidate items to the set of search items—increasing each set size by a constant amount. Alternatively, it could be that the presence of the background makes it harder to accumulate the information needed to identify each item, especially when there are features such as line terminators in the background (cf. “crowding effects” Intriligator & Cavanagh, 2001; Toet & Levi, 1992). Subsequent experiments, notably Experiment 5, seek to differentiate between these accounts.

6. Experiment 4: the role of spatial frequency

In the first three experiments, the primary effect of the background was an additive effect on RT rather than a

change in the slope of RT \times set size functions. These experiments (especially, Experiment 3) used backgrounds that contained figural elements that could be confused with search items. Experiment 4 introduces a different sort of background complexity. Here the search items and the background share common spatial frequency content. It has been shown that stimuli with similar frequency spectra interfere with each other (e.g. Blake & Holopigian, 1985; Regan, 1985). One could imagine that a T or an L on a background of the same component frequencies would be harder to segment and might increase the slope of an RT \times set size function in visual search.

6.1. Method

6.1.1. Stimuli

This experiment used the same T versus L search as used in the previous experiments. To create the backgrounds, the Fourier transform was computed for a black letter T of size 64×64 pixels, placed in a white image of size 1024×1024 pixels. We randomized the phase spectrum, keeping the amplitude spectrum intact. By simply rescaling this background spectrum, background textures were created with spectra of different ratios to the target ratio (1:8, 1:2, 1:1, 2:1, 8:1). Portions of the textures are shown in Fig. 12. All the textures had a gaussian distribution of gray levels (0–256), centered on 128.

Search items were T’s (target) and L’s (distractors) that subtended $1.5^\circ \times 1.5^\circ$ of visual angle. Items could be presented in any of four orientations (0° , 90° , 180° , or 270°). Search items of two contrast levels were tested in separate blocks. Measured against the average luminance of the background, the contrast of the *high contrast* T’s and L’s was 75% and the contrast of the *low contrast* letters was 23%. Examples of search displays are shown in Fig. 13.

6.1.2. Subjects

Twenty-two observers participated in the study, 11 participants per contrast group (low versus high).

6.1.3. Procedure

Set sizes were 1, 4, 7 and 10, and the experimental procedures of the search task were identical to those of the previous experiments. Different groups of subjects participated in the high contrast and low contrast conditions. Each background was paired with each set size to produce 56 target-present and 56 target-absent trials. Background, set size, and target presence/absence were randomly distributed across trials. There were a total of 2240 trials per contrast condition, split into four blocks with rest periods in between.




Fig. 12. Portions of the five background textures used in Experiment 4: from left to right, the background to target frequency ratios are 1:8 (coarse), 1:2, 1:1, 2:1, and 8:1 (fine).




Fig. 13. Examples of low contrast items and high contrast items superimposed on a 1:1 background texture.

6.2. Results

Two subjects (one per group) were discarded from the analysis due to high error rates (>15%). Error rates otherwise averaged 3.7% and 4.9%, including out-of-range RT trials (RT < 200 ms and RT > 7000 ms), for the high contrast and low contrast groups, respectively.

Table 2 shows the average RT × set size slopes in ms/item for the target-present trials. Fig. 14 shows the mean RT as a function of the ratio of background to search item spatial frequency. As can be seen, the background had no effect on the high contrast letters. Accordingly, Fig. 15 shows only the RT × set size functions for the low contrast letters.

Visual examination of the data suggests that, as in the previous experiments, the important effects of background are only seen in mean RT measures and not in




Fig. 14. Mean RTs as a function of log ratio of spatial frequencies between background and search items (coarser background on the left). (—■—) High-present; (-□-) high-absent, (-●-) low-present, (-○-) low-absent.

Table 2
Slope in ms/item as a function of frequency ratio between background and search items (1:8 = coarse, 8:1 = fine) for the target-present trials of Experiment 4

Frequency ratio	High contrast	Low contrast
1:8	34	64
1:2	40	58
1:1	40	60
2:1	41	70
8:1	39	74

RT × set size slope measures. This is borne out by statistical analysis. An ANOVA performed on the target-present trials, shows a clear effect of the background on mean RT ($F(4, 72) = 51.67, p < 0.0001$). The interaction of background with set size, which would reflect a slope effect, is not significant ($F(12, 216) = 1.34, p = 19.7$), nor is the triple interaction of background × set size × item contrast ($F \sim 1$).

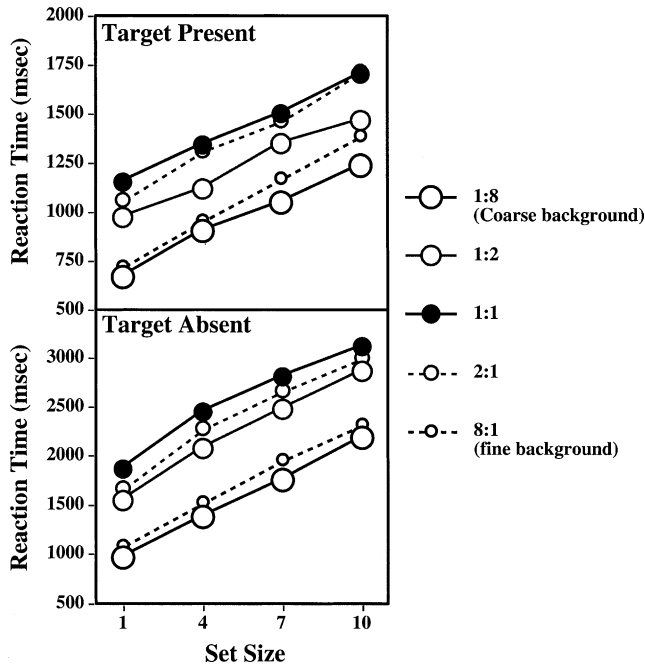


Fig. 15. RT × set size functions for low contrast letters on the five different backgrounds of Experiment 4. Note that the y-axis scale is twice as great for the absent trial as for the target present trials.

Unsurprisingly, there are significant main effects of set size ($F(3, 54) = 255.93, p < 0.0001$) and item contrast ($F(1, 18) = 27.33, p < 0.0001$) on RT as well as an interaction of background and item contrast ($F(4, 72) = 30.88, p < 0.0001$) that reflects the fact that background has an effect only on the low contrast items. The high contrast items are uniformly easy to segment from the background.

An ANOVA including target-present and target-absent trials for the low contrast letters reveals a significant interaction of background with target presence/absence ($F(4, 36) = 27.6, p < 0.0001$). This reflects the larger effect of background on the target-absent trials as seen in the preceding experiments.

Note that the slope effect, even if it were reliable, runs opposite to the RT effect. The largest slope is found with the low-contrast 8:1 ratio of background to search items. Yet, this is the second fastest background condition. If the 1:1 ratio produced the same slope as the 8:1, the mean RT difference would be greater than what is seen here. The slight reduction in slope in the 1:1 condition probably reflects a small speed-accuracy tradeoff as subjects abandon a few of the longest searches.

6.3. Discussion

Variation in the spatial frequency content of the background produces a tuning curve function (Fig. 14) reminiscent of the curves produced in detection experiments when a target of one spatial frequency is masked

by another (e.g. Legge & Foley, 1980). In our case, target and distractors are “camouflaged” within the background. Note that, in this case, the “mask” is the background and is not superimposed on the target item. As with the previous experiments, these data argue against the hypothesis that the background exerts its primary effect at the selection stage (Figs. 3 and 4). These data can also be used to argue against the hypothesis that it takes longer to segment items from harder/noisier backgrounds and that search is delayed until the segmentation process is complete (Figs. 1 and 2). This is most obvious if we consider the result for a set size of 1. In this case, there is just a single “real” item in the display. All the observer needs to do is to identify the item as either a “T” or an “L”. Relative to the coarsest (1:8) background, it takes about 500 ms longer to confirm that the item is a “T” with the 1:1 background. However, it takes nearly 1000 ms longer to confirm that the item is an “L” with the same 1:1 background! This is very hard to explain if the preattentive processes are presenting the limited-capacity stage with a single item. Why should it take 500 ms longer to deliver an “L” to the selection stage? Moreover, the putative delay is so long that it should be detectable by inspection of the RT distributions, which should appear to be shifted noticeably to longer RTs as background spatial frequency approaches that of the search item.

Fig. 16 shows that this is not the case. These data are still consistent with the hypothesis that a sloppy preattentive stage could add extra virtual items (about 8) to the one “official” item when the background frequency matched that of the target and distractors. If the effective set size is nine, then it would be a standard result to find that it takes longer to determine that none of these nine possible items are “Ts” than it does to confirm that one of those items is a “T”. The data are also consistent




Fig. 16. RT distributions for two backgrounds and a set size of 1. Note that that the distribution of the similar (1:1) background RTs (solid black) does not look like the dissimilar (1:8) RT distribution shifted by 500 ms (dashed line) as predicted by Option 1.

with the hypothesis that accumulation of information is slowed in a limited-capacity parallel stage following selection (Fig. 5). Experiment 5 provides data to distinguish between these two hypotheses.

7. Experiment 5: search for one versus two targets

Experiment 5 has two conditions. One is a simple replication of Experiment 4 with subjects searching for a T among L's on backgrounds defined by the similarity of their spatial frequency spectrum to that of the target item. In the other condition, subjects distinguished displays containing two targets from those containing only one. By this point, we can be reasonably sure that the effect of background on mean RT will be additive with set size, but how will that additive cost change if subjects must search for one versus two targets? Subjects must search through an average of N items in order to find a single target in a set of M items. The value of N will vary depending upon your model of search. In a standard serial-self-terminating search model (Sternberg, 1969), $N = (M + 1)/2$. In an "amnesic" model (Horowitz & Wolfe, 1998, 2001) $N = M + 1$. The choice of model, however, is not critical to the present argument. In order to find the second of two targets, the subject must conduct two searches. Under either model, more items must be examined to find two targets than to find one target. Call the number of items which must be examined to find two targets kN , where k is a value greater than 1.0.

Now suppose that increasing the similarity of the background to the search items has the effect of adding " J " virtual items to the set of items that must be searched. Instead of having to search through N items to find a target, the subject must search through $N + J$ items. In order to find the second target, the subject will need to search through $k(N + J)$ items rather than just kN items. If the additive cost in the one target case is created by the need to search J extra items, then the additive cost of the background in the two-target case should be ' kJ '—i.e., greater than the one-target cost.

Suppose, on the other hand, that the effect of a more difficult background is to slow the accumulation of target-identifying information in a limited-capacity parallel stage after selection. This situation is shown in Fig. 17 (modeled in Fig. 5).

On one-target trials, the target is selected for analysis and recognition after the same average amount of time, whether it is on an *easy* or *hard* background. If it is on an easy background, identity information accumulates quickly (dashed, rising lines in Fig. 17). If the target is on a hard background, the information accumulates slowly (solid rising lines in Fig. 17). The cost of the change in background is captured by the angle, alpha, between the lines. In the two-target task, the RT is changed by the increase in time required to select the




Fig. 17. Schematic rendering of the accumulation of information for target identification. Dashed, rising lines represent the faster rate of accumulation for items on easy backgrounds (e.g., a 1:8 or 8:1 spatial frequency ratio between background and search items). Thicker lines show a slower accumulation rate for items on harder/noisier background (e.g., 1:1). 1T = one-target trial starting point. 2T = selection of the second target in a two-target trial (see text for further details).

second target. The starting point of accumulation shifts from the "1T" point to the "2T" point. This simply shifts the critical accumulation (the one that determines the RT) later in time. Recall that the angle, alpha, reflects the effects of the background on the accumulation. This effect will be the same on Target 2 as it was on Target 1. Thus, alpha remains the same and the cost of the *harder* background remains the same.

Negative trials are more complex to model (Chun & Wolfe, 1996). In the one-target case, a negative trial is characterized by the complete absence of a target. In the two-target case, subjects make a negative response to the presence of only one target. Both the less efficient clean-up and the slowed accumulation model predict that the effects on negative trials should be proportional to the effects on target-present trials. That is, the additive cost of a harder background should be greater for negative trials than for positive trials. Recall that the less efficient clean-up model—with its added items—predicts that the additive cost will be greater for two-target trials than for one-target *positive* trials. Thus, it follows that the additive cost on negative trials should be greater for the two-target negative trials than for the one-target negative trials. The increase should be proportional. If the additive costs for positive and negative one-target RTs differ by a factor of 2, then the greater two-target costs should also differ by a factor of 2. On the other hand, the slower accumulation model predicts that the switch from one-target to two-target trials will produce no increase in the additive cost of background complexity on positive-trial RTs. It follows from this reasoning that there should be no difference between the costs for one-target and two-target negative trials.

7.1. Method

The method was essentially the same as that of Experiment 4, with the following exceptions. Item contrast was set to 45% of the mean background luminance. Ten

subjects were tested for two blocks of 1680 trials; one block in each task (one-target, two-target), with the order counterbalanced across subjects. Trials were evenly divided among three set sizes (4, 7, and 10) and five backgrounds (1:8x, 1:4x, 1:1x, 2:1x, 8:1x where “x” is the frequency spectrum of the target T). Within a block, trial types were randomly intermixed.

7.2. Results and discussion

The general pattern of results was similar to the results of the earlier experiments. For the positive trials, there were large, additive effects of background difficulty ($F(4, 36) = 47, p < 0.0001$) on mean RT and no significant effect of background on the slope of the RT \times set size functions ($F \sim 1$). Slopes for the two-target conditions were steeper than the slopes for the one-target conditions ($F(2, 18) = 58, p < 0.0001$). The comparison of the additive costs of the background between the one-target and two-target conditions is central to this experiment. To illustrate this comparison, the average RT for all one-target, correct positive trials was subtracted from the average one-target, correct positive trial RTs for each background. The same calculation was done for the two-target correct positive trials and for the one-target and two-target, correct negative trials. The results are shown in Fig. 18.

The relative effects of background difficulty on mean RT are strikingly similar in the one-target and two-target tasks of Experiment 5. Thus, for example, there is a roughly 200 ms difference between the easiest and




Fig. 18. Relative mean RT as a function of background in Experiment 5. Squares represent the one-target condition. Circles represent the two-target condition. Clearly, one-target and two-target conditions show essentially identical effects of background similarity on relative mean RT.

hardest average one-target positive trial RTs. The difference between the easiest and hardest average two-target positive trial RTs is the same 200 ms. Note that, as in the previous experiments, the RT cost of harder/noisier backgrounds is markedly larger on target-absent than on target-present trials. However, here the mean RT cost during negative trials is effectively the same for the one-target and two-target tasks.

This is *not* the pattern of results one would expect if search through more complex backgrounds is slowed by the introduction of “virtual” items into the search set. If there are more items, then there should be a greater mean RT cost in the two-target condition because subjects must make two searches through those items (with or without memory for prior deployments of attention). The results of Experiment 5 do fit the RT pattern predicted by the slower accumulation model (see Fig. 17). In that model, mean RT cost attributable to background complexity is driven entirely by the rate at which information is accumulated about the last target selected. The predicted mean RT cost is not dependent on the number of items selected prior to the last item, so the one-target and two-target costs should be the same.

We began with four hypotheses about the possible effects of background complexity:

1. The selection stage hypothesis (Figs. 3 and 4) is ruled out by the failure to find a reliable increase in mean RT \times set size slopes with more complex backgrounds, regardless of whether natural scene complexity (Experiment 1) or spatial frequency similarity (Experiments 4 and 5) were manipulated. There might be an effect of local feature similarity on slope in Experiment 3. However, it is relatively small when compared to the additive RT effects.
2. The purely preattentive clean-up hypothesis is ruled out by the difference in target-present and target-absent RT costs observed in all of the above experiments (Fig. 1). Additionally, the distribution of set size 1 RTs (Experiment 4, Fig. 16) fails to show the lateral shift predicted by this hypothesis. It is simply implausible that subjects are delaying their search for up to a full second when backgrounds are difficult.
3. The less effective clean-up hypothesis (i.e., the addition of virtual items) is ruled out by the failure to find a different RT cost for one-target and two-target search tasks in Experiment 5 (see Fig. 17).
4. The slower accumulation hypothesis (see Figs. 5 and 17) fares the best with the present set of results. It correctly predicts an RT cost that is
 - (a) additive with set size,
 - (b) greater on target-absent than on target-present trials,
 - (c) the same for one-target and two-target search tasks.




Fig. 19. Search difficulty was manipulated by varying the ratio of check size in the background to check size in search items. Panels (a–d) show examples of 1:8 (a), 1:4 (b), 1:1 (c), and 2:1 (d) ratios. Panels (e) and (f) show examples of three item search arrays. Search items would be red and black in panel (e) and yellow and black in (f).

The only models considered here have been situated within the two-stage, preattentive–attentive architecture growing out of the work of Neisser and Treisman. There are, of course, other classes of models (e.g. Palmer et al., 2000; Vergheze, 2001). Any other contender would need to be able to satisfy the constraints laid down by these data.

8. Experiment 6: checkerboards: altering selection rate

In the five experiments to this point, background complexity and similarity to search items have been varied and the effects of these backgrounds on RT have been consistent with the view that a one-step preattentive process separates the display into search items and background. When backgrounds are more complex or more similar to search items, there is an RT cost that is additive with set size. We have attributed that cost to a slowing in the rate of accumulation of information about object identity. Is there ever a reliable influence of the background on the rate at which items are selected for recognition? It seems there should be conditions under which the processing of *each individual* item can be made slower by the presence of a sufficiently complex or similar background. Recall from the Introduction that such an effect should be seen as an increase in slope of the $RT \times$ set size function. In this final experiment, we present a condition wherein backgrounds do modu-

late search slopes. As illustrated in Fig. 19, to get this effect, we need to approximate the search for the proverbial “needle in a haystack”.

8.1. Method

8.1.1. Stimuli

The backgrounds in this experiment were checkerboard patterns of yellow and black squares as shown in Fig. 19. The search items were 2×3 check pieces of the checkerboard shown in Fig. 19c. Targets were oriented vertically while distractors were oriented horizontally. Nine ratios of background to search item checks sizes were used: 1:16, 1:8, 1:4, 1:2, 1:1, 2:1, 4:1, 8:1, and 16:1. The 1:16 coarse image corresponded to a uniform yellow background. All other backgrounds were yellow and black. Search items were either the same color as the background (yellow and black) or different (red and black). The search items were not aligned with the background and were slightly beveled in appearance, otherwise they would have been invisible when background and search items had the same check size and color. Nevertheless, as can be seen most clearly in Fig. 19c, these items are hard to segment from similar backgrounds. Set sizes of one, three, and five were tested.

8.1.2. Procedure

On each trial, observers saw one, three, or five elements pasted at random locations onto one of the nine

backgrounds. A vertical target was present on 50% of trials. Observers were tested for 1620 trials each with red & black and yellow & black search items. The display remained on-screen until the subject responded. Subjects were told to answer as rapidly and accurately as possible and they were given feedback about their accuracy following each trial. Subjects took breaks after blocks of several hundred trials.

8.1.3. Subjects

Twelve subjects were paid for their participation in the study.

8.2. Results

RTs less than 200 ms and greater than 4000 ms were labeled as errors. One subject was discarded from the analysis because of a high error rate. Otherwise, average error rates were very low: 1.5% and 2.2% for the colour-cue and no-colour-cue conditions, respectively, including the out-of-range RTs.

Fig. 20 shows RTs averaged across set size. As in Experiment 4, the results take the form of a tuning curve with a peak at the 1:1 ratio. Also consistent with the previous experiments, the effect of background is greater for target-absent than for target-present trials.

Fig. 21 shows the slopes of RT \times set size functions for target-present trials (target-absent results are similar but with slopes that are about twice as steep). Here, for the first time, there is a clear effect of background on RT \times set size slope, even with items differing in color from the background. Slopes are markedly elevated when the background check size is similar to the search item check size. This result suggests that the processing

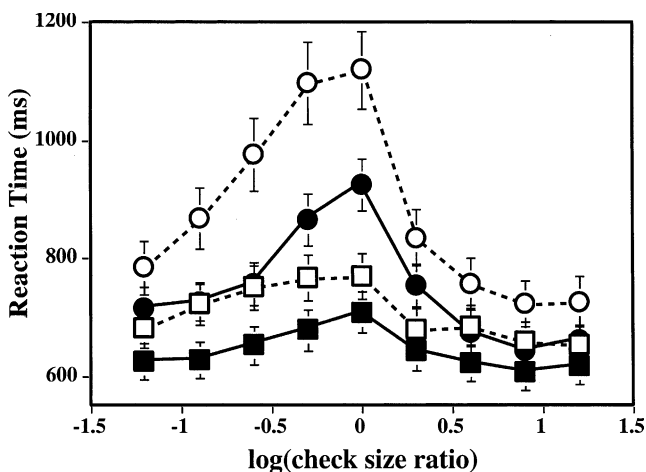


Fig. 20. RTs as a function of background to search item check size ratio for yellow & black (circle symbols) and red & black (squares) items. Target-present data are shown by solid lines and filled symbols; target-absent data are represented by dashed lines and open symbols.




Fig. 21. Mean RT \times set size slopes as a function of background to search item check size ratio for yellow & black (open circles) and red & black (closed squares) items. Target-present data are shown. Target-absent data are comparable.

of each item has been slowed by its placement on this background.

Statistical analysis bears out these conclusions. In this experiment, literally every effect and interaction is statistically reliable ($p < 0.001$ in *all* cases). Restricting discussion to those effects of interest, an ANOVA on mean RTs shows that there are main effects of background, search item color, set size, and target presence/absence. The interaction of background with target presence/absence is significant, showing the usual greater effect of background on target-absent trials. The interaction of background with set size is significant, indicating a slope difference—unlike what was seen in the preceding experiments. This is confirmed by an ANOVA on slope data. Here there is a significant main effect of background on slope ($F(1, 10) = 7.46$, $p < 0.001$) and a main effect of search item color reflecting the greater effect of background when background and search items are of the same color ($F(1, 10) = 52.16$, $p < 0.001$).

8.3. Discussion

With the stimuli of Experiment 6, it appears the task of segmentation has not been successfully completed when preattentive processes have done their work. Preattentive processes guide attention to plausible target locations. In the other five experiments, the time required to select individual items for further processing was not altered by the nature of the background. In Experiment 6, it was altered. When the check items are presented on a background of checks of the same or similar size, it appears to be harder to handle each item in turn. It is interesting that the background alters the search slope even when the search items are red and the

background is yellow. Apparently, while the color cue must help in the preattentive segmentation stage, it is not good enough to present a reliably vertical or horizontal “item” to the recognition stage.

This result, indicating that preattentive processes can only do so much, is mirrored elsewhere in the search literature. For example, in simple feature search, it is known that search becomes harder as the target becomes more similar to the distractors or as the distractors become less similar to each other (Duncan & Humphreys, 1989). Over a moderate range of variations in similarity, making the task harder leads to an increase in mean RT without an increase in slope (see Treisman & Gormican, 1988 for examples in color search or Royden, Wolfe, & Klempen, 2001 for examples in motion). However, at some point, the increase in difficulty exceeds the capabilities of the preattentive feature processes and slopes increase (Nagy & Sanchez, 1990).

9. General discussion

How do observers perform visual search tasks when the background is not the blank screen usually employed in such experiments? The results of the experiments presented here indicate that there is a preattentive step in which candidate targets are segmented from the background. Attention is then deployed to the locations of those candidates. This allows the items to be selected for further processing in a limited-capacity stage that leads to their identification. This account fits very well into the broad framework of the Guided Search theory of visual search (Wolfe, 1994a, 2001; Wolfe et al., 1989; Wolfe & Gancarz, 1996). The core idea in Guided Search is that preattentive information is used to guide the deployment of attention. A limited number of basic features can provide guidance. These include color, orientation, size, motion and a variety of cues to the 3D layout of the world. There are perhaps a dozen such features (the details are reviewed in Chun & Wolfe, 2001; Wolfe, 1998a). Thus, in a search for a red vertical target among red horizontal and green vertical distractors, attention can be guided toward the set of red items and toward the set of vertical items. The intersection of those two sets is a likely locus for any red vertical items. Guidance is a somewhat noisy affair and so some red horizontal and/or green vertical items might receive enough guiding activation to attract attention. Hence slopes for conjunctions of features like color and orientation are fairly shallow but tend not to be zero ms/item.

Understood in these terms, “object” is just another feature property. A preattentive process guides attention toward objects and away from other areas of the display. In the example just given, the search would be a search not for “red vertical” but for a red vertical *object*.




Fig. 22. It is easier to find the mid-gray target on the right than on the left. Why?

In the experiments presented in this paper, the search is not for disembodied T-ness but for an object that is a “T”.

This principle can be illustrated using the following thought experiment.

In the left panel of Fig. 22, the target is a mid-gray item among dark-gray and light-gray distractors. This is a relatively inefficient search because it is hard to find a target that lies between distractors in feature space. Thus, it is hard to find a target that is orange among distractors that are red and yellow (Bauer, Jolicœur, & Cowan, 1996; D’Zmura, 1991) or to find a 0° (vertical) target among distractors tilted 20° to the left and right (Wolfe, Friedman-Hill, Stewart, & O’Connell, 1992). In the right panel, search has become much easier even though there is much more of the light-gray color that colored one set of distractors on the left. The reason why is clear enough. On a light-gray background, the light-gray distractors literally disappear. They no longer count in the visual search. It is easy to search for a mid-gray item among dark-gray distractors on a light-gray background.

It might be objected that the set size is cut in half on the right. However, it seems obvious that doubling the number of dark-gray distractors would not do much to this search. It will remain efficient while the search on the left will be inefficient.

If guidance were perfect, search would never be required. Attention would be deployed to the target item, first time, every time. This is roughly what happens in a simple feature search. If the target is a red item and the distractors are green items and all these items are clearly delineated on a blank background, then it will be possible to guide attention to a red target, if it is present, without ever selecting a green distractor. When the target is a T and the distractors are L’s on a blank background, no preattentive feature appears to distinguish between these items. Therefore, attention must be deployed at random among these items until a target is found or the search is abandoned.

The data presented in this paper suggest that, when the T’s and L’s are placed on a moderately complex

background, the search items can still be separated from the background and selected by attention. However, when the background becomes more complex, it seems to take longer to accumulate the information required to identify a selected object. Perhaps the separation of background from item is imperfect and enough background gets included to make identification more difficult.

For Experiments 1 through 5, the selection itself is not impaired by the background. Only in Experiment 6 does a difficult background that exactly matches the features of the target slow the rate of selection. In the discussion of Experiment 6, we suggested that selection could be slowed because of an inability of the preattentive stage to adequately segment the scene into search items and background. There is an alternative possibility. At some point, the background could interfere with the identification stage so severely that the capacity of that stage is reduced. Thus, it might normally be possible to load an item every 50 ms into an identification stage able to handle six items at any one time. If identification became so difficult that only three items could be in the “pipeline” at any one time, then selection might have to be slowed down to accommodate this reduced capacity.

Under either a slowed selection or a slowed identification account, Experiment 6 presents a situation analogous to other difficult feature searches. Search for a red object among green will be efficient. Search will be less efficient if the target is red while the distractors are red with a slight orange tint. An increase in the slope of $RT \times$ set size functions results when preattentive processes can no longer deliver only the target item. Slopes get still steeper when it takes longer to determine the actual color of each item. In Experiment 6, the preattentive object parser delivers some set of candidate objects for subsequent search. When the checks on those objects are similar to the checks in the background, it either becomes harder to select the individual item or to determine if that item is vertical or horizontal and search becomes markedly inefficient. Whether this is a failure of object segmentation or of orientation identification is hard to tell in this experiment.

In summary, the experiments described here help to bridge the gap between laboratory search tasks and search in the real world. It has become clear that attention is generally deployed to objects (Baylis & Driver, 1993; Egly, Driver, & Rafal, 1994; Goldsmith, 1998; Yantis, 1993). As a consequence, most theories of search implicitly accept the idea that some process must segment objects from the visual input. We do not know how to segment complex scenes into objects. However, the data from these experiments allow us to conclude that the object segmenting step is likely to be preattentive, rather than being performed separately for each likely object location. Moreover, the data indicate that

the segmentation process is imperfect. Faced with complex backgrounds or backgrounds similar to the search items, the segmentation process will deliver “objects” that are harder to identify than objects on clean backgrounds. Further experiments of this sort should provide more information about the details of the processes that segment scenes into objects and background.

Acknowledgements

This research was supported by funds from the National Science Foundation (SBR-9710498) and the National Institute of Health (MH/EY56020). We thank Christoph Koch and two anonymous reviewers for useful comments and Antonio Torralba for his assistance with stimulus design.

References

- Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. (1997). Deictic codes for the embodiment of cognition. *Behavioral Brain Science*, 20(4), 723–742, discussion 743–767.
- Bauer, B., Joliceur, P., & Cowan, W. B. (1996). Visual search for colour targets that are or are not linearly-separable from distractors. *Vision Research*, 36(10), 1439–1466.
- Baylis, G. C., & Driver, J. (1993). Visual attention and objects: evidence for hierarchical coding of location. *Journal of Experimental Psychology: Human Perception and Performance*, 19(3), 451–470.
- Becker, M., & Pashler, H. (2001). Information acquisition from scenes: exploring the effects of preview on search. Submitted.
- Biederman, I., Glass, A. L., & Stacy, E. W. (1973). Searching for objects in real-world scenes. *Journal of Experimental Psychology*, 97, 22–27.
- Bilsky, A. A., & Wolfe, J. M. (1995). Part-whole information is useful in size \times size but not in orientation \times orientation conjunction searches. *Perception and Psychophysics*, 57(6), 749–760.
- Blake, R., & Holopigian, K. (1985). Orientation selectivity in cats and humans assessed by masking. *Vision Research*, 25, 1459–1468.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 443–446.
- Carmody, D. P., Scanlon, M. J., & Dasaro, C. R. (1990). Target search in embedding and non-embedding displays. In R. Groner, G. d'Ydewalle, & R. Parham (Eds.), *Studies in visual information processing* (Vol. 1) *From eye to mind: Information acquisition in perception, search, and reading* (pp. 103–117). Amsterdam, Netherlands: North-Holland.
- Chun, M. M., & Wolfe, J. M. (1996). Just say no: how are visual searches terminated when there is no target present? *Cognitive Psychology*, 30, 39–78.
- Chun, M. M., & Wolfe, J. M. (2001). Visual attention. In E. B. Goldstein (Ed.), *Blackwell's Handbook of Perception* (Chapter 9, pp. 272–310). Oxford, UK: Blackwell.
- D'Zmura, M. (1991). Color in visual search. *Vision Research*, 31(6), 951–966.
- Driver, J., & Frackowiak, R. S. J. (2001). Introductory review: neurobiological measures of human selective attention. *Neuropsychologia*, 39, 1257–1262.
- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, 96, 433–458.

- Duncan, J., Ward, R., & Shapiro, K. (1994). Direct measurement of attention dwell time in human vision. *Nature*, 369(26 May), 313–314.
- Eckstein, M. P., Thomas, J. P., Palmer, J., & Shimozaki, S. S. (1996). Further predictions of signal detection theory on visual search accuracy: conjunctions, disjunctions, and triple conjunctions. *Investigative Ophthalmology and Visual Science*, 37(3), S15.
- Egeth, H., & Dagenbach, D. (1991). Parallel versus serial processing in visual search: further evidence from subadditive effects of visual quality. *Journal of Experimental Psychology: Human Perception and Performance*, 17(2), 551–560.
- Egeth, H. E., Virzi, R. A., & Garbart, H. (1984). Searching for conjunctively defined targets. *Journal of Experimental Psychology: Human Perception and Performance*, 10, 32–39.
- Egley, R., Driver, J., & Rafal, R. D. (1994). Shifting attention between objects and locations: evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, 123, 161–177.
- Enns, J. T., & Kingstone, A. (1995). Access to global and local properties in visual search for compound stimuli. *Psychological Science*, 6(5), 283–291.
- Foster, D. H., & Ward, P. A. (1991a). Asymmetries in oriented-line detection indicate two orthogonal filters in early vision. *Proceedings of the Royal Society of London B*, 243, 75–81.
- Foster, D. H., & Ward, P. A. (1991b). Horizontal-vertical filters in early vision predict anomalous line-orientation identification frequencies. *Proceedings of the Royal Society of London B*, 243, 83–86.
- Goldsmith, M. (1998). What's in a location? Comparing object-based and space-based models of feature integration in visual search. *Journal of Experimental Psychology: General*, 127(2), 189–219.
- Graham, N., Kramer, P., & Haber, N. (1985). Attending to the spatial frequency and spatial position of near-threshold visual patterns. In M. I. Posner & O. S. M. Marin (Eds.), *Mechanisms of attention: attention and performance XI* (pp. 269–284). Hillsdale, NJ: Lawrence Erlbaum.
- Harris, J. R., Shaw, M. L., & Altom, M. J. (1985). Serial-position curves for reaction time and accuracy in visual search: tests of a model of overlapping processing. *Perception and Psychophysics*, 38(2), 178–187.
- Hoffman, J. E. (1979). A two-stage model of visual search. *Perception and Psychophysics*, 25, 319–327.
- Horowitz, T. S., & Wolfe, J. M. (1998). Visual search has no memory. *Nature*, 394(6 August), 575–577.
- Horowitz, T. S., & Wolfe, J. M. (2001). Search for multiple targets: remember the targets, forget the search. *Perception and Psychophysics*, 63(2), 272–285.
- Humphreys, G. W., Cinel, C., Wolfe, J. M., Olson, A., & Klempen, N. (2000). Fractionating the binding process: neuropsychological evidence distinguishing binding of form from binding of surface features. *Vision Research*, 40(10–12), 1569–1596.
- Intriligator, J., & Cavanagh, P. (2001). The spatial resolution of visual attention. *Cognitive Psychology*, 43(3), 171–216.
- Kinchla, R. A. (1977). The role of structural redundancy in the perception of targets. *Perception and Psychophysics*, 22(1), 19–30.
- Kingsley, H. L. (1932). An experimental study of 'search'. *American Journal of Psychology*, 44, 314–318.
- Kroll, J. F. (1992). Making a scene: the debate about context effects for scenes and sentences. In K. Rayner (Ed.), *Eye movements and visual cognition* (pp. 284–292). New York: Springer-Verlag.
- Kundel, H. L. (1991). Search for lung nodules: the guidance of visual scanning. *Investigative Radiology*, 266, 777–787.
- Kwak, H., Dagenbach, D., & Egeth, H. (1991). Further evidence for a time-independent shift of the focus of attention. *Perception and Psychophysics*, 49(5), 473–480.
- Legge, G. E., & Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America*, 70(12), 1458–1471.
- Logan, G. (1975). On the independence of naming and locating masked targets in visual search. *Canadian Journal of Psychology*, 29(1), 51–58.
- Lu, Z.-L., & Doshier, B. A. (1998). External noise distinguishes attention mechanisms. *Vision Research*, 38(9), 1183–1198.
- Melcher, D. (2001). Persistence of visual memory for scenes. *Nature*, 412(26 July), 401.
- Melcher, D., & Kowler, E. (2001). Visual scene memory and the guidance of saccadic eye movements. *Vision Research*, 41, 3597–3611.
- Mewhort, D. J., Johns, E. E., & Coble, S. (1991). Early and late selection in partial report: evidence from degraded displays. *Perceptions and Psychophysics*, 50(3), 258–266.
- Moore, C. M., & Wolfe, J. M. (2001). Getting beyond the serial/parallel debate in visual search: a hybrid approach. In K. Shapiro (Ed.), *The limits of attention: temporal constraints on human information processing*. Oxford: Oxford University Press.
- Nagy, A. L., & Sanchez, R. R. (1990). Critical color differences determined with a visual search task. *Journal of the Optical Society of America A*, 7(7), 1209–1217.
- Navon, D. (1977). Forest before the trees: the precedence of global features in visual perception. *Cognitive Psychology*, 9, 353–383.
- Neisser, U. (1967). *Cognitive psychology*. New York: Appleton, Century, Crofts.
- Nodine, C. F., Krupinski, E. A., & Kundel, H. L. (1993). Visual processing and decision making in search and recognition of targets. In D. Brogan, A. Gale, & K. Carr (Eds.), *Visual Search 2* (pp. 239–249). London, UK: Taylor & Francis.
- Nodine, C. F., Kundel, H. L., Lauver, S. C., & Toto, L. C. (1996). Nature of expertise in searching mammograms for breast masses. *Academic Radiology*, 3(12), 1000–1006.
- Palmer, J., & McLean, J. (1995). Imperfect, unlimited-capacity, parallel search yields large set-size effects. Paper presented at the Society for Mathematical Psychology, Irvine, CA.
- Palmer, J., Verghese, P., & Pavel, M. (2000). The psychophysics of visual search. *Vision Research*, 40(10–12), 1227–1268.
- Pashler, H. (1984). Evidence against late selection: stimulus quality effects in previewed displays. *Journal of Experimental Psychology: Human Perception and Performance*, 10, 429–448.
- Pashler, H., & Badgio, P. C. (1985). Visual attention and stimulus identification. *Journal of Experimental Psychology: Human Perception and Performance*, 11, 105–121.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108.
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: an attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18(3), 849–860.
- Rayner, K., & Pollatsek, A. (1992). Eye movements and scene perception. *Canadian Journal of Psychology*, 46(3), 342–376.
- Regan, D. (1985). Masking of spatial-frequency discrimination. *Journal of the Optical Society of America—A*, 2, 1153–1159.
- Rensink, R., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: the need for attention to perceive changes in scenes. *Psychological Science*, 8, 368–373.
- Royden, C. S., Wolfe, J. M., & Klempen, N. (2001). Visual search asymmetries in motion and optic flow fields. *Perception and Psychophysics*, 63(3), 436–444.
- Samuel, S., Kundel, H. L., Nodine, C. F., & Toto, L. C. (1995). Mechanism of satisfaction of search: eye position recordings in the reading of chest radiographs. *Radiology*, 194(3), 895–902.
- Sanders, A. F., & Donk, M. (1996). Visual search. In O. Neumann, & A. F. Saunders (Eds.), *Handbook of perception and action, Vol. 3: Attention* (pp. 43–77). London: Academic Press.

- Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, 1(7), 261–267.
- Sternberg, S. (1969). High-speed scanning in human memory. *Science*, 153, 652–654.
- Swenson, R. G. (1980). A two-stage detection model applied to skilled visual search by radiologists. *Perception and Psychophysics*, 27(1), 11–16.
- Swenson, R. G., & Judy, P. F. (1981). Detection of noisy visual targets: models for the effects of spatial uncertainty and signal-to-noise ratio. *Perception and Psychophysics*, 29(6), 521–534.
- Thorpe, S., Fize, D., & Marlot, C. (1996). Speed of processing in the human visual system. *Nature*, 381(6 June), 520–552.
- Toet, A., & Levi, D. M. (1992). The two-dimensional shape of spatial interaction zones in the parafovea. *Vision Research*, 32(7), 1349–1357.
- 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.
- Tsotsos, J. K., Culhane, S. N., Wai, W. Y. K., Lai, Y., Davis, N., & Nuflo, F. (1995). Modeling visual attention via selective tuning. *Artificial Intelligence*, 78, 507–545.
- VanRullen, R., & Thorpe, S. J. (2001). Is it a bird? Is it a plane? Ultra-rapid visual categorisation of natural and artificial objects. *Perception*, 30(6), 655–668.
- Vergheze, P. (2001). Visual search and attention: a signal detection approach. *Neuron*, 31, 523–535.
- Ward, R., Duncan, J., & Shapiro, K. (1996). The slow time-course of visual attention. *Cognitive Psychology*, 30(1), 79–109.
- Wolfe, J. M. (1994a). Guided Search 2.0: a revised model of visual search. *Psychonomic Bulletin and Review*, 1(2), 202–238.
- Wolfe, J. M. (1994b). Visual search in continuous, naturalistic stimuli. *Vision Research*, 34(9), 1187–1195.
- Wolfe, J. M. (1998a). Visual search. In H. Pashler (Ed.), *Attention* (pp. 13–74). Hove, East Sussex, UK: Psychology Press Ltd.
- Wolfe, J. M. (1998b). What do 1,000,000 trials tell us about visual search? *Psychological Science*, 9(1), 33–39.
- Wolfe, J. M. (2001). Guided Search 4.0: a guided search model that does not require memory for rejected distractors. *Journal of Vision*, Abstracts of the 2001 VSS meeting.
- Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided search: an alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 419–433.
- Wolfe, J. M., Friedman-Hill, S. R., & Bilsky, A. B. (1994). Parallel processing of part/whole information in visual search tasks. *Perception and Psychophysics*, 55(5), 537–550.
- Wolfe, J. M., Friedman-Hill, S. R., Stewart, M. I., & O'Connell, K. M. (1992). The role of categorization in visual search for orientation. *Journal of Experimental Psychology: Human Perception and Performance*, 18(1), 34–49.
- Wolfe, J. M., & Gancarz, G. (1996). Guided Search 3.0: a model of visual search catches up with Jay Enoch 40 years later. In V. Lakshminarayanan (Ed.), *Basic and clinical applications of vision science* (pp. 189–192). Dordrecht, Netherlands: Kluwer Academic.
- Woodman, G. F., & Luck, S. J. (1999). Electrophysiological measurement of rapid shifts of attention. *Nature*, 400(6 August), 867–869.
- Xu, Y., & Singh, M. (in press). Early computation of perceptual part structure: Evidence from visual search. *Perception and Psychophysics*.
- Yantis, S. (1993). Stimulus-driven attentional capture. *Current Directions in Psychological Science*, 2(5), 156–161.