



The psychophysics of visual search

John Palmer ^{a,*}, Preeti Verghese ^b, Misha Pavel ^c

^a *Department of Psychology, University of Washington, Box 351525, Seattle, WA 98195-1525, USA*

^b *Smith-Kettlewell Eye Research Institute, 2318 Fillmore St., San Francisco, CA 94115, USA*

^c *Department of Electrical and Computer Engineering, Oregon Graduate Institute, 20000 Walker Road, Beaverton, OR 97006, USA*

Received 4 December 1998; received in revised form 10 December 1999

Abstract

Most theories of visual search emphasize issues of limited versus unlimited capacity and serial versus parallel processing. In the present article, we suggest a broader framework based on two principles, one empirical and one theoretical. The empirical principle is to focus on conditions at the intersection of visual search and the simple detection and discrimination paradigms of spatial vision. Such simple search conditions avoid artifacts and phenomena specific to more complex stimuli and tasks. The theoretical principle is to focus on the distinction between high and low threshold theory. While high threshold theory is largely discredited for simple detection and discrimination, it persists in the search literature. Furthermore, a low threshold theory such as signal detection theory can account for some of the phenomena attributed to limited capacity or serial processing. In the body of this article, we compare the predictions of high threshold theory and three versions of signal detection theory to the observed effects of manipulating set size, discriminability, number of targets, response bias, external noise, and distractor heterogeneity. For almost all cases, the results are inconsistent with high threshold theory and are consistent with all three versions of signal detection theory. In the Discussion, these simple theories are generalized to a larger domain that includes search asymmetry, multidimensional judgements including conjunction search, response time, search with multiple eye fixations and more general stimulus conditions. We conclude that low threshold theories can account for simple visual search without invoking mechanisms such as limited capacity or serial processing. © 2000 Elsevier Science Ltd. All rights reserved.

Keywords: Search; Attention; Conspicuity; Psychophysics; Signal detection theory

1. The psychophysics of visual search

Looking for a friend's face, foraging for food, and detecting predators are all examples of visual search. These search behaviors are accomplished by processes and representations common to nearly any visual task. These behaviors depend upon the representation of visual stimuli, the limits of attention, and the integration of information for decision. Thus, an understanding of representation, attention, and decision goes hand in hand with an understanding of search. Here, we develop a theoretical framework for visual search that encompasses these issues at the heart of vision. We argue that this framework provides a foundation that is missing in many current theories of visual search.

Our first guiding principle is to develop a theory of search based upon the theories of simple detection and discrimination developed within visual psychophysics. This choice has a number of advantages. Visual psychophysics provides theories of the representation of simple attributes such as color or orientation, and for the processes that yield a response in simple detection and discrimination tasks. Thus it provides hypotheses for both the stimulus representation and the decision process. In addition, psychophysics inspires both the theoretical and empirical foci of this article. Theoretically, we focus on applying signal detection theory to search (e.g. Shaw, 1980, 1982, 1984) because of the theory's success in combining internal processes such as bias with the representation of the stimulus. Empirically, we focus on applying the experimental methods of psychophysics to search (e.g. Geisler & Chou, 1995) because these methods are most likely to reveal a common account of both search and simple detection and dis-

* Corresponding author. Tel.: +1-206-5430706.

E-mail address: jpalmer@u.washington.edu (J. Palmer)

crimination. In short, we argue for building a theory of visual search based on earlier work in psychophysics.

Our second guiding principle is to emphasize the distinction between theories assuming a high versus a low threshold. This distinction has an important place in the history of psychophysics and was one of the motivations for the development of signal detection theory. In simple detection and discrimination, the consensus of research supports the need for some kind of low threshold. For search, the situation is different. Many search response time theories ignore error, a move which can be justified only by assuming a high threshold. As a consequence, these theories must include additional mechanisms such as serial processes or limited capacity to account for phenomena that arise naturally from a low threshold theory. Thus, while the high threshold theory is largely discredited within the domain of simple detection and discrimination, it has survived within the domain of visual search. We extend the previous analysis of high and low threshold from simple detection and discrimination to visual search.

Here is a preview of the article. In the remainder of the introduction, we develop alternative theories of visual search and specify a simplified domain in which to compare the theories. The theories are then tested using six distinct phenomena of visual search. Next is a discussion of alternative interpretations and of generalizations from the simplified domain to the larger domain of visual search. We conclude that a high threshold is untenable for search just as it is untenable for simple detection and discrimination.

2. Theoretical analysis

2.1. High versus low threshold

2.1.1. High threshold

High threshold theories begin from two assumptions (for reviews see Green & Swets, 1966; Graham, 1989). First, the relevant representation of a stimulus is discrete. One either detects a target or one does not. Second, the absence of a target never produces a target-detect state. This means that distractors can never produce enough misleading evidence to surmount the ‘high threshold’ and produce a false alarm. False alarms arise in behavior only as a consequence of guessing.

Consider the example of searching for a slightly pink target disk among white distractor disks. Typical displays contain many distractors and only occasionally contain one or more targets. The high threshold theory presumes that the perceptual system never categorizes a distractor as a target. Thus white distractors are never seen as pink. Instead, the mistakes in the representation are always misses: a pinkish target is categorized as a

white distractor. This asymmetry in the representation presumably arises because of the numerous distractors relative to the occasional target.

2.1.2. Low threshold

Low threshold theories¹ make a wide variety of assumptions about stimulus representation and decision. They include the well known signal detection theory (Green & Swets, 1966), and less known alternatives such as biased-choice theory (Luce, 1959, 1963a; Macmillan & Creelman, 1991) that are not pursued here. What they share is the premise that a distractor can produce the target-detect state. Specifically, in signal detection theory there is always a chance that a distractor can yield enough evidence to pass the ‘low threshold’ criterion. Thus, distractors can produce misleading evidence that results in a false alarm.

Consider again the pinkish target and white distractor example. Any low threshold theory allows both kinds of errors in the representation. Now white distractors can be mistakenly categorized as pinkish targets and pinkish targets can be mistakenly categorized as white distractors.

2.1.3. Consequences for visual search

This distinction has an interesting consequence for visual search tasks because of the multiple distractors. If one assumes a high threshold, then additional distractors remain safely irrelevant to the decision process. In contrast, if one assumes a low threshold, then each additional distractor has the possibility of introducing enough evidence to result in an error. Thus, for low threshold theory, distractors affect performance even if processing of each stimulus is completely independent. This point was perhaps first made by Tanner (1961) and has been discussed by many authors since (e.g. Green, 1961; Cohn & Lasley, 1974; Shaw, 1980; Pelli, 1985; Graham, 1989). It holds for response time tasks as well as accuracy tasks (see Discussion). Because of this effect of noisy distractors, it is a mistake to infer limited capacity from a set-size effect.

2.2. A theoretical framework based upon signal detection theory

There is no single low threshold theory to contrast with the classic high threshold theory. Instead, we describe a framework for constructing alternative low threshold theories based upon signal detection theory and describe three specific theories based upon the

¹ In this article, we use the term *low threshold theory* to denote the large class of theories that complements high threshold theory. However, the term has also been used to denote a specific theory developed by Luce (1963b) based upon discrete representations. We refer to Luce’s specific theory as the ‘two-state, low threshold theory’.

framework (for reviews see Swets, 1984; Sperling & Doshier, 1986). The framework provides a method for constructing theories appropriate to specific domains but it is not a testable theory in itself.

Theories of visual search that are based upon signal detection theory have several common elements. The first element is that they all suppose a representation of the individual stimulus that corresponds to a continuous random variable. Thus, they always assume a *noisy representation* which naturally leads to a low threshold. On the other hand, the framework says nothing about what aspects of the stimulus are represented. For example, one of the specific theories below takes comparisons among neighboring stimuli to be the relevant representation rather than attributes of individual stimuli. The specific theories developed in Appendix A are general for any distribution of distractor and target representations. However, to make more specific predictions, we follow two paths. First, we present the predictions of the usual default assumption of equal variance, normal distributions. Second, we present analyses of more general families of distributions. In particular, these generalizations include a possible non-linear relation between the stimulus and the representation and unequal variance between targets and distractors.

A second element shared by theories of signal detection is a *decision process* distinct from initial perception that incorporates effects of bias in a specific task. In particular, we concentrate on the yes–no task (and related rating tasks) in which the decision process is a comparison between the internal representation of a test stimulus and an internal representation of a comparison standard. This comparison standard can be either a decision criterion or a random variable representing a comparison stimulus. The common feature is that the representation of the comparison standard is a function of the biases of the observer and the history of previous decisions and not a function of the test stimulus.

The third element of these theories is specific to the search situation with multiple, simultaneous stimuli. Some assumption must be made about how information from individual stimuli is integrated to make a decision. The specific theories described below explore three different assumptions about such *information integration*. These include the ideal observer that assumes a very detailed integration rule and the independent decisions model that assumes independent decisions for each stimulus.

The fourth element is also specific to search. Some assumption must be made about the *effect of attention* on the representation of individual stimuli. For all of the theories detailed here, we assume independence among the individual stimuli. This subsumes three distinct properties: unlimited capacity, parallel processing,

and statistical independence. Unlimited capacity is the independence of each individual stimulus representation from the number of other stimuli; parallel processing is the simultaneous processing of all stimuli; and, statistical independence is the lack of trial-to-trial correlations among the representations of multiple stimuli (for more detail on these properties see Busey & Townsend, in press; Townsend, 1981; Townsend & Ashby, 1983). Our intent is to develop elaborations of a simple ‘independent channels’ theory (e.g. Graham, 1989) that assumes all three properties.

In summary, our theoretical framework allows a variety of choices about representation, decision, information integration, and attention. What the theories share are an independent noisy representations of the stimuli and a decision making process that integrates information and incorporates biases.

2.3. Specific theories

We now introduce four specific theories of visual search. The first is the familiar high threshold theory. The other three are low threshold theories that are special cases within the signal detection framework. These three theories make a range of assumptions about information integration and the stimulus representation. They are presented in order from the theory that assumes the most knowledge of the stimulus on the part of the observer to the theory that assumes the least knowledge.

The theories can be described using the information flow diagrams of Fig. 1. Each panel depicts a different theory. These flow diagrams show how information from a set of stimuli illustrated on the left side is transformed into a single yes/no response illustrated on the right side. The boxes indicate the nature of the initial stimulus representation, the rule used to integrate information from the multiple stimuli, and the nature of the decision process. Each theory is briefly described below and is formally defined in Appendix A.

2.3.1. High threshold theory

The two key assumptions of the high threshold theory have already been introduced. First, the representation of the stimulus is one of two discrete states, detecting the target or not detecting the target. Second, the target-detect state never arises from a distractor. This theory is represented in the top panel of the figure with the left box representing the stimulus processing that produces a discrete representation for each stimulus. Next, in the middle box it is assumed the observers introduce guessing responses on some proportion of the trials in which no target is detected. Finally, in the rightmost box the information from all of the stimuli and the guessing processes are integrated by a simple OR rule. If any of the representations indicates a target

detection or a guess that the target was present, then the response is ‘yes;’ otherwise, the response is ‘no.’ Thus, guesses are the only way this theory produces a false alarm. In addition, we follow other recent authors (e.g. Quick, 1974; Watson, 1979; Graham, 1989) in combining these assumptions of the classic high threshold theory with the assumption that the probability of the target-detect state is a Weibull function of the relevant stimulus parameter (e.g. contrast difference between target and distractor). This formulation has several unique properties such as a consistent shape of the psychometric function under probability summation (Green & Luce, 1975; Watson, 1979) and it has proven to be quite successful in describing detection threshold experiments, particularly summation experiments (Graham, 1989). With this additional assumption, one can more easily compare the predictions of the high threshold theory to signal detection theories with specific distributional assumptions. Some of the tested predictions require this assumption while others test the high threshold theory more generally.

2.3.2. Ideal observer

The ideal observer theory has been an important reference point in psychophysical theories of detection and discrimination (Green & Swets, 1966; Geisler, 1989). As with all of the signal detection theories, the initial stimulus representation is more general than it is with the high threshold theory. The representation of

each stimulus is assumed to be continuous and noisy. The unique feature of the theory lies in the middle panel that represents how information is integrated from the different stimuli. It is assumed that information from all stimuli is combined optimally to maximize performance in the particular task. Hence, the ideal observer theory includes an information integration rule that is particular to the task and specific to the set of possible stimuli in the experiment. In Appendix A, the ideal combination rule is presented for the well known case of one target among m stimuli and for the less known case of h targets among m stimuli. As the last step in the theory, the integrated representation is compared to a criterion. If the criterion is surpassed, then the observer responds ‘yes’, otherwise ‘no’.


2.3.3. Maximum of outputs

The third theory is perhaps the most common in visual psychophysics, the maximum-of-outputs theory (see Graham, 1989 for more complete references). As discussed in Appendix A, this theory is equivalent to the independent decisions model of Shaw (1980, 1982) and previously has been referred to as the decision theoretic model (Green & Swets, 1966; Green & Birdsall, 1978), the extreme detection model (Swenson & Judy, 1981), the separable decision model (Smith, 1998) and perhaps was first introduced by Tanner (1961) in his descriptions of the effects of uncertainty predicted by signal detection theory. The sensory representation is continuous and noisy as with the ideal observer theory. The difference arises in the middle panel in which the different sources of information are combined according to a maximum rule: The stimulus with the maximum amount of evidence favoring a target is used in a subsequent decision process. The final decision process is common for all of the versions of signal detection theory. If the criterion is surpassed, the observer responds ‘yes’.


2.3.4. Relative coding: maximum of differences

The last theory is motivated by ideas from robust statistics (e.g. Hampel, Ronchetti, Rousseeuw & Stahel, 1986). By way of introduction, consider again ideal observer theory. It requires complete knowledge of the stimulus including the joint distributions of all possible stimulus representations. By comparison, the maximum-of-outputs theory requires knowledge of only the marginal distributions for each stimulus. It makes no use of information from the joint distribution. What decision rules might require even less knowledge of the underlying distributions? Our candidate is to consider differences in the attributes of neighboring stimuli rather than the value of the individual attributes. This theory is referred to as the *maximum-of-differences theory*. To illustrate, consider the problem of comparing the brightness of stimuli in different parts of the


High Threshold



Ideal Observer



Maximum of Outputs



Maximum of Differences




Fig. 1. A schematic illustration of the flow of information in four theories of visual search. The three arrows on the left represent information input from three stimuli. The one arrow on the right represents the information output to determine a single response. At some point in each model, the multiple sources of information from the multiple stimuli are reduced to a single representation indicated by a single arrow.

visual field. Suppose this task is complicated by an inherent variation of brightness with eccentricity: Peripheral stimuli appear dimmer than central stimuli. The use of local differences is robust to such eccentricity effects compared to a direct representation of brightness.

The maximum-of-differences theory is represented in the bottom panel of Fig. 1. It is similar to the maximum-of-outputs theory but with the maximum taken of the differences between neighboring stimuli as is represented in the second from the left panel of the figure. This idea is very similar to the feature contrast theory of Nothdurft (1991, 1992, 1993). He has demonstrated for certain search tasks that feature differences rather than feature values determine performance. For example, the orientation of lines in a texture pattern may vary systematically across the scene. Even if all orientation values are in a scene, one can easily discriminate an odd line if its orientation differs sharply from its neighbors.

2.3.5. Summary of theories

The four theories can be thought of as varying in two ways. The first distinction is between high and low threshold. The second distinction is among the low threshold theories that assume different amounts of knowledge about the stimulus representation. Comparisons of the predictions of these three theories to the high threshold theory are intended to give an indication of how sensitive different signal detection theories are to the particular assumptions about the details of the representation, information integration, and decision process. To preview the outcome, the predictions for the three signal detection theories are quite similar and all are distinct from the predictions of the high threshold theory.

3. The domain of study

3.1. A simplified domain

The primary principle guiding this article is to integrate a theory of search into the better established theories of simple detection and discrimination. By simple detection and discrimination, we are referring to experiments in spatial vision with one (or two) simple stimuli and a homogeneous surround. The domain of possible search experiments overlaps with such simple detection and discrimination experiments for the case of a single stimulus (set size 1). We argue that focusing on the methods of such simple detection and discrimination experiments is likely to reveal the common mechanisms responsible for both the simple tasks and visual search (cf. Palmer, 1995; Verghese & Stone, 1995). This strategy sacrifices the breadth of previous visual search

research in exchange for the more controlled conditions found in psychophysical studies of spatial vision. If simple theories are adequate for these restricted conditions, then such theories are a plausible starting point for generalizing to a broader range of conditions. The *simplified domain* is defined by the following four principles.

3.1.1. Unidimensional judgments

Judgments of differences along a single dimension such as contrast, hue, size, orientation, etc. are the core of visual psychophysics. We focus on visual search experiments that require judgments along a single dimension.

3.1.2. Accuracy measures

The bulk of visual psychophysics has been conducted with matches and discrimination thresholds that are based upon response accuracy rather than response time. Therefore, we focus on response accuracy experiments.

3.1.3. Single eye fixation

Most visual psychophysics is conducted under conditions in which the retinal stimulus can be precisely specified and eye movements are minimized. Similarly, we focus on experiments that control fixation and use brief displays to minimize eye movements.

3.1.4. Plausibly distinct and independent stimuli

In typical spatial vision experiments, the relevant aspects of a stimulus are manipulated in the presence of as simple a display as possible. For example, the contrast of a disk or grating patch is manipulated within an otherwise homogeneous visual field. Presenting other stimuli within a field can affect performance for a variety of reasons. To minimize these extraneous effects, we emphasize experiments with the following properties:

1. distinct stimuli that are well above detection threshold (e.g. pedestal experiments);
2. widely separated stimuli that are arranged to minimize configural cues (e.g. randomly placed);
3. single displays with no masks.

The use of suprathreshold stimuli allows one to focus on the effects of the number of stimuli rather than considering both the number and the spatial uncertainty of the stimuli as is necessary for near-threshold stimuli. The use of widely separated stimuli minimizes spatial interactions. And, the use of single displays without masks minimizes the possibility of temporal interactions.

In summary, our simplified domain excludes many interesting studies with complex stimuli, multiple eye fixations, crowded stimuli, etc. No simple theory can hope to account for all of the phenomena of visual

search. In this article, our goal is to test simple theories within the simplified domain. To preview the results, restricting oneself to this simplified domain yields a much more homogenous set of results than is otherwise evident. In the discussion, we consider how to elaborate the simple theories to account for more complicated phenomena.

3.2. Overview of phenomena

The body of this article is an analysis of six phenomena of visual search. All can be studied under conditions satisfying the simplified domain inspired by simple detection and discrimination experiments. Three of the phenomena arise from the multiple stimuli present with visual search: *Set-size effects* are the most studied and measure the possible cost of attending to multiple stimuli; *multiple target effects* reveal the possible benefits of multiple stimuli; and, *distractor heterogeneity effects* are diverse and have received little quantitative analysis. We also consider three other phenomena common to search and simple detection and discrimination: *target-distractor discriminability effects* are very large and are sometimes ignored in studies of search; *response bias effects* can also be large and vividly show the difference between high and low threshold theories; and, *external noise effects* reveal the characteristics of the internal noise.

4. Effects of set size

Perhaps the most fundamental phenomenon of search is the effect of multiple stimuli, referred to as a set-size effect. The predicted magnitude of a set-size effect depends critically on whether a theory assumes a high or a low threshold. The high threshold theory predicts that variation in set size will have no effect, while any theory with a low threshold predicts that variations in set size will matter. Furthermore, the magnitude of the set-size effect predicted by low threshold theories depends on the discriminability of the stimuli. Thus, to meaningfully compare set-size effects across experiments, one must equate the discriminability of the stimuli used in the different experiments. This can be accomplished by the common psychophysical technique of estimating the stimulus difference that yields a given level of discrimination performance as a function of the variable of interest (here set size). By estimating such a *difference threshold* at each set size, one can quantify set-size effects at the level of discriminability that defined the threshold measure. While numerous studies have measured performance as a function of set size (e.g. Estes & Taylor, 1964; Shaw, 1980, 1984; for reviews see Teichner & Krebs, 1974; Palmer, 1995; Pashler, 1998), fewer have varied discrim-

inability to estimate a threshold as a function of set size (Bergen & Julesz, 1983; Nagy & Sanchez, 1990; Poirson & Wandell, 1990; Palmer, Ames & Lindsey, 1993; Zacks & Zacks, 1993; Vergheese & Nakayama, 1994). There have been also a few that measured the related uncertainty effects in detection (e.g. Cohn & Wardlaw, 1985) and in simulated medical images (e.g. Burgess & Ghandeharian, 1984). Here, we present details for one example experiment that measures difference thresholds as a function of set size.

4.1. An example contrast increment experiment

Palmer (1994, experiment 1), measured the contrast increment threshold for set sizes 1, 2, 4, and 8. The stimuli were white, 0.5° diameter disks presented for 100 ms to minimize eye movements. To control sensory effects such as those of eccentricity and lateral masking, the disks were arranged randomly with eccentricities between 5° and 8° and the center-to-center spacing was at least 3°. A yes-no procedure was used with equal frequencies of Target Present and Target Absent trials. All distractors had 20% contrast and the single target had a contrast increment that ranged from 8 to 24%, thus the target contrast itself ranged from 28 to 44%. A psychometric function was fit to each condition for each observer and the difference threshold was estimated for 75% correct performance.

Fig. 2 shows the estimated thresholds as a function of set size. Individual observer estimates are indicated by the light dotted lines; the mean thresholds for each set size are indicated by the solid points; and, the best fitting regression line on this log–log plot is shown by the solid line. Because these *threshold versus set size* functions (TvS functions) have logarithmic scales, the linear regression function on the graph is given by,

$$\log t = \log a + b_{\text{TVS}} \log m \quad (1)$$

where m is the set size and t is the contrast increment threshold. This equation is, of course, equivalent to the threshold being a power function of set size. The parameter b_{TVS} is the slope on this log–log plot and $\log a$ is the y -intercept. Because the abscissa is logarithmically scaled, the y -intercept is defined at $\log m = 0$ which implies $m = 1$. Thus, the parameter a is the estimated threshold for set size 1. Because this function is not meaningful for fractional set sizes, it is intended as a familiar continuous description of this discrete effect.

The four observers show similar results with TvS slopes ranging from 0.27 to 0.40. The mean slope was 0.30 ± 0.08 . (Throughout this article, standard errors are the standard error of the mean of a sample from multiple subjects or replications of a single subject.) Palmer (1994) contains two additional exper-

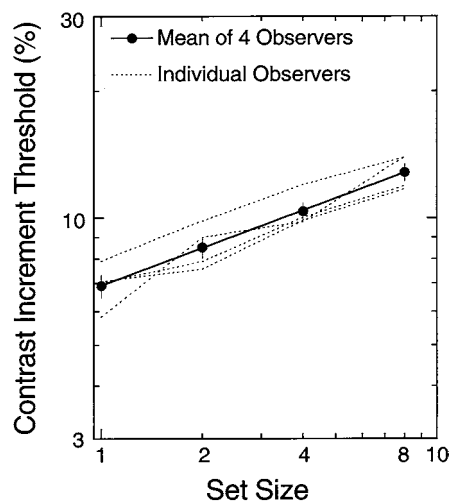


Fig. 2. The contrast increment threshold is plotted as a function of set size (Palmer, 1994, experiment 1). The dotted lines indicate individual observer thresholds; the solid points indicate the mean thresholds of four observers; and, the solid line is the best fitting linear regression line on this log-log plot. Error bars throughout the article are standard errors of the mean.

iments measuring the effect of set size on the contrast increment threshold. The mean TvS slopes in these experiments for set sizes 2 and 8 are 0.26 ± 0.02 (experiment 2, display set size) and 0.25 ± 0.02 (experiment 2, relevant set size). In summary, these three experiments show average set-size effects on threshold with TvS slopes in the 0.25–0.30 range.

4.2. Predictions

The predicted effect of set size on the threshold is defined relative to the threshold predicted for set size 2 because the maximum-of-differences theory makes no prediction for set size 1. For the high threshold theory, set size is predicted to have no effect. For each of the other theories, the relative threshold depends on distributional assumptions. In this section, the predictions are based on independent, equal-variance, normal distributions for all distractors and targets; in following sections, we consider more general assumptions. For the maximum-of-outputs theory, one can derive an expression for the relative threshold as a function of set size with no free parameters (Palmer, Ames & Lindsey, 1993; for a more general form see Palmer, 1998). For the ideal observer and maximum-of-differences theories, no known expression exists for the predicted thresholds, but one can simulate the results to any degree




Fig. 3. For each of the four theories, the predicted relative threshold is plotted as a function of set size. All thresholds are relative to that predicted for set size 2.

of precision from the expressions given in Appendix A.²

Fig. 3 shows the predicted relative thresholds as a function of set size for each of the four theories. The flat line represents the constant threshold predicted by the high threshold theory. The ideal observer and maximum-of-outputs predictions are the two finely dashed curves and they fall essentially on top of one another (cf. Nolte & Jaarsma, 1966). For this range of set sizes, the predicted set-size effect is about 1% smaller for the ideal observer theory than the maximum-of-outputs theory. Finally, the maximum-of-differences theory predicts slightly smaller effects for set sizes larger than 2 but makes no prediction for set size 1. The predictions of the signal detection theories deviate slightly from linearity on this log–log graph in the concave downward direction. The small differences between the predictions can be quantified by calculating the TvS slope predicted at set size 2. The ideal observer theory predicts 0.25, the maximum-of-outputs theory predicts 0.26, and the maximum-of-differences theory predicts 0.19. In summary, the high threshold theory predicts a slope of zero and the three signal detection theories all predict slopes in the range of 0.19–0.26 at set size 2.

² The predictions given here were simulated from 40 000 trials per condition assuming a rating experiment with three levels of discriminability centered around the threshold value. Thus it used 120 000 trials to estimate one threshold. These simulations followed the design of the speed discrimination experiment described in Appendix B. Repeated simulations showed that these conditions resulted in predictions of the relative thresholds that were reliable to at least two decimal places. Simulations in the remainder of this article also use these methods.

4.3. Discussion

The first point to emphasize is the consistency of the observed set-size effects. Set-size effects with similar magnitudes have been measured for disk color, disk size, ellipse orientation (Palmer, 1994), line length, line orientation, rectangle shape (Palmer et al., 1993), speed discrimination (Verghese & Stone, 1995), vernier acuity judgments (Fahle, 1991, experiment 3), detection of Landolt C's (Davis & Peterson, 1998), and letter discrimination (Bennett & Jaye, 1995; McLean, Palmer & Loftus, 1997). Restricting ourselves to the simplified domain, there are only a few exceptions to these results. For example, larger set-size effects have been measured for more complex tasks such as the orientation of pairs of widely separated objects (Palmer, 1994; Pöder, 1999). Such cases may require some kind of limited capacity processing specific to the particular complex task (cf. Broadbent, 1958, 1971; Hoffman, 1979). Such limited capacity results are also found with tasks that make explicit memory demands (Palmer, 1988, 1990; Palmer & Ames, 1992; see also Scott-Brown & Orbach, 1998). In summary, set-size effects of similar magnitude are reported by all studies of simple search tasks that restrict themselves to our simplified domain and measure set-size effects at threshold. This consistency is in striking contrast to the variability of results under more general conditions.

The second point to emphasize is the failure of the high threshold theory. The results are inconsistent with the predictions of a high threshold and are consistent with the predictions of all three versions of low threshold theory. This failure of the high threshold theory has led some to propose additional mechanisms with limited capacity or serial processing to account for the set-size effects. Here, we stress that set-size effects do not imply limited capacity or serial processing. Indeed, they are predicted by any low threshold theory.

5. Effects of target-distractor discriminability

The manipulation of discriminability has a large effect on search performance just as it does with other discrimination tasks (e.g. Neisser, 1973; Pashler, 1987b; Duncan & Humphreys, 1989). By manipulating the difference between the target and the distractor, one can move performance from chance to perfect accuracy. This large effect must be accounted for within any comprehensive theory of search. Moreover, discriminability takes on even more importance because comparisons between experiments have often been made with little or no effort to equate discriminability. For example, Duncan and Humphreys have criticized the early evidence supporting feature integration theory (Treisman & Gelade, 1980) because the feature and

conjunction search conditions were not equated in terms of discriminability. This is a crucial control because the magnitude of set-size effects decreases with increasing discriminability as is documented below.

Psychophysical methods provide a natural method for controlling discriminability by the use of threshold estimates as has already been introduced in this article. In addition, there is a body of work that analyzes discriminability effects using a psychometric function that plots the probability of a correct response as a function of a stimulus manipulation such as contrast (for an introduction see Gescheider, 1985). Psychometric functions have been characterized by several analytic functions such as the cumulative normal distribution or the Weibull. Such psychometric functions have two parameters: A threshold parameter specifying the stimulus magnitude necessary to yield a given performance such as 75% correct, and a steepness parameter that specifies the shape (or slope) of the curve (for details see Pelli, 1985, 1987; for examples see Nachmias, 1981; Maloney, 1990). Of interest here, the steepness of the psychometric function has been predicted to increase with increases in uncertainty (Tanner, 1961; Nachmias & Kocher, 1970; Cohn, 1981). This is an interaction between uncertainty and discriminability where the uncertainty effect is larger for less discriminable stimuli. Previous work on uncertainty effects in detection have demonstrated this effect under some conditions (Cohn, 1981; Cohn & Wardlaw, 1985; Pelli, 1985). Here, we examine whether or not set size changes the shapes of psychometric functions in search tasks in a similar fashion. In the process, we also begin exploring more general distributional assumptions for the three versions of signal detection theory.

5.1. An example contrast increment experiment

Consider again the experiment from Palmer (1994) in which contrast increment thresholds were estimated as a function of set size. These estimates were based on psychometric functions for each set size and each observer. In particular, this yes/no experiment measured the probability of a hit and a false alarm for three contrast increments. The hits and false alarms were used to calculate d' using the assumptions of equal-variance normal distributions.

The observed psychometric functions are shown in Fig. 4 with each panel showing a separate observer and the two curves representing the extreme set sizes of 1 and 8. The psychometric functions are presented as d' versus the contrast increment with both logarithmically scaled. Following Tanner & Swets (1954; see also Nachmias & Kocher, 1970; Pelli, 1985, 1987), the psychometric functions are summarized by linear functions on $\log-d'$ versus \log -contrast. In other words, d' is assumed to be a power function of the contrast. By this d' power




Fig. 4. Psychometric functions are plotted for two set sizes with separate observers in each panel. Performance is shown in terms of d' as a function of the contrast increment with both axes logarithmically scaled (Palmer, 1994, experiment 1).

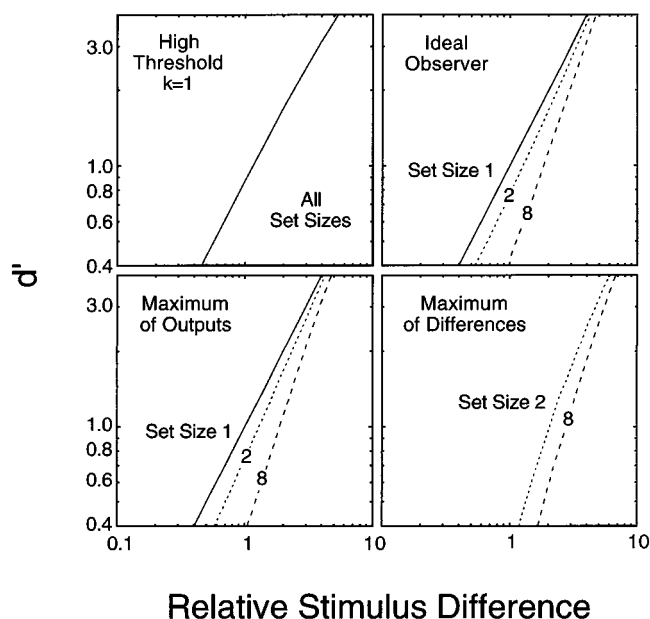


Fig. 5. Each panel shows the psychometric functions predicted by a different theory. Performance is shown in terms of d' as a function of the stimulus difference for several set sizes with both axes logarithmically scaled.

law, variations in the shape of the psychometric function are characterized by changes in slope on this log-log plot. We parameterize the function in terms of the slope b_{pf} , the threshold t , and the performance criterion defining threshold $d'_{criterion}$,

$$\log d' = \log d'_{criterion} + b_{pf} \log (x/t). \quad (2)$$

The two free parameters of these psychometric functions are the threshold, t , and the psychometric function slope, b_{pf} .

In analyzing the psychometric function, the d' power law is similar to, but not identical to the Weibull function which plays the analogous role in high threshold theory (see the last section of Appendix A for a definition of the Weibull function). The Weibull function also has a threshold parameter, t , and a differently defined steepness parameter k . Pelli (1987) has pointed out that these two functions are nearly indistinguishable and one can estimate the parameters of one from the other. Focusing our analyses at threshold performance, we find that the steepness parameters are proportional with a coefficient³ of roughly 0.9 ($b_{pf} = 0.9 k$). For consistency in this article, all of the analyses use the d' power law.

The manipulation of increasing the contrast increment between targets and distractors increases performance from near chance ($d' = 0$) for small increments to near perfect ($d' > 3$) for large increments. The observed data are reasonably approximated by a line the log-log graph. But with only three points, this experiment cannot thoroughly address the quality of this approximation.

Set size has two effects on these functions. First, increasing the set size increases the threshold as was described in detail in the section on set-size effects. Second, increasing set size increases the steepness of the psychometric function: the less discriminable the stimulus, the larger the effect of set size on d' . This effect is summarized in Fig. 6 which shows the psychometric function slope as a function of set size. There is a consistent increase in the slope of the psychometric function as set size increases. The slope parameter is 0.59 ± 0.03 for set size 1 and is 1.2 ± 0.1 for set size 8.

5.2. Predictions

5.2.1. Set size and psychometric functions

Fig. 5 shows the predicted psychometric functions for set sizes 1, 2, and 8. The functions are shown in terms of d' versus the stimulus difference (e.g. contrast increment) with each axis logarithmically scaled. The predictions of the four different theories are shown in the four panels. For the three signal detection theories, we assume equal-variance, normal distributions as before. For all of the theories, the shape of the psychometric functions is well approximated by a linear function on log d' versus log of the stimulus difference.

³ This 0.9 estimate is only slightly different than the 0.8 estimated by Pelli (1987) for his conditions. He used a 2AFC rather than a Yes/No procedure, and he fit the entire function rather than matching the functions at threshold. Both estimates are approximate because, relative to the normal, the Weibull is steeper at low contrasts and shallower at high contrasts.

The differences among the theories are made clear by examining the parameters of predicted psychometric functions. The section on set-size effects has already summarized that increases in set size result in increases in threshold for all of the predictions except that of the high threshold theory. The other parameter is the slope of the psychometric function estimated at the 0.75 probability correct point ($d' = 1.35$). In Fig. 7, the predicted slope of the psychometric function is shown as a function of set size. The flat solid lines show the predictions of the high threshold theory for k equal to 1 and 3; the finely dashed curve shows the predictions of the ideal observer theory; the coarse dashed curve shows the prediction of the maximum-of-outputs theory; and the coarsest dashed curve shows the prediction of the maximum-of-differences theory. To summarize the predictions, the high threshold theory predicts no




Fig. 6. The slope of the psychometric function is plotted as a function of set size for the contrast increment experiment. Points represent the mean of all of the observers in each experiment.

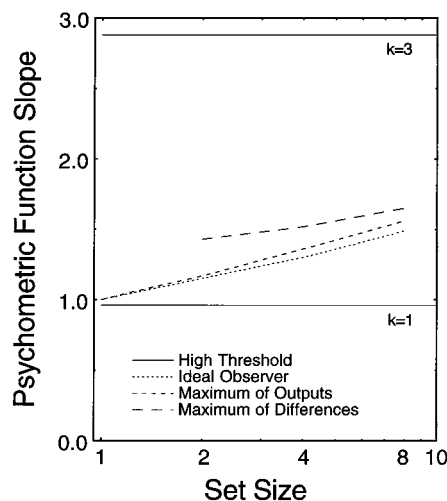


Fig. 7. For each theory, the predicted slope of the psychometric function is plotted as a function of set size.

effect of set size on the psychometric function slope while all three low threshold theories do predict an effect on the slope.

5.2.2. Generalized distributional assumptions

The predictions discussed above are for equal-variance, normal distributions where the mean representation is linearly related to the relevant stimulus parameter (e.g. contrast increment). This is clearly too restrictive for many cases including the example because it predicts too stereotyped a psychometric function. Of particular relevance here is the linearity between stimulus and representation. This linearity assumption can be generalized by introducing representations that are a power function of the stimulus. For instance, Pelli (1985) has considered such an assumption for contrast detection where the psychometric functions are much steeper than they are for contrast discrimination (e.g. Leshowitz, Taub & Raab, 1968). This generalized theory makes the same assumptions about the internal representations and decision processes as conventional signal detection theories. The only change is to relate the value of the mean internal representation to the stimulus by a power function rather than a linear function.

McLean et al. (1997) have applied this generalization to search experiments. In particular, they combined the power function and the maximum-of-outputs theory. The predictions of this combination are illustrated in Table 1 and Fig. 8. Table 1 lists several predicted values as a function of the steepness parameter (b_{pf}) in the leftmost column. The predicted thresholds are given for set sizes 1 and 8. The thresholds are defined relative to the assumed unit variance of the distractor distribution. Together, these two set sizes yield a log–log TvS slope that is given in the fourth column. The predicted psychometric function slopes are described in the rightmost three columns. For set size 1, the predicted psychometric function slope is equal to the steepness parameter; for set size 8, the predicted psychometric function slope is proportionally higher than the steepness parameter. Finally, we also include the mean of the predicted psychometric function slopes for both set sizes. Such an aggregate measure has the advantage of being better estimated in an experiment because one can use all of the available data to estimate the mean psychometric function slope. The ultimate prediction is illustrated in Fig. 8. It shows the TvS slope as a function of the mean psychometric functions slope. The contour is the prediction of the generalized maximum-of-outputs theory. As the psychometric function steepens, the predicted set-size effect decreases. With linear relations ($b_{pf} = 1$), the mean psychometric function slope is predicted to be 1.3 and the TvS slope is predicted to be about 0.25. For conditions with an accelerating nonlinearity ($b_{pf} = 3$), the mean psychometric function slope is predicted to be

Table 1
Predicted effect of set size and the steepness parameter on thresholds and psychometric function slopes

Steepness parameter	Relative threshold ($m = 1$)	Relative threshold ($m = 8$)	TvS slope	PF slope ^a ($m = 1$)	PF slope ($m = 8$)	PF slope mean
0.4	2.11	7.63	0.617	0.4	0.620	0.510
0.5	1.82	5.08	0.494	0.5	0.775	0.638
0.6	1.65	3.88	0.412	0.6	0.93	0.765
0.7	1.53	3.19	0.353	0.7	1.09	0.893
0.8	1.45	2.76	0.309	0.8	1.24	1.02
1.0	1.35	2.25	0.247	1.0	1.55	1.28
1.5	1.22	1.72	0.165	1.5	2.32	1.91
2.0	1.16	1.50	0.124	2.0	3.1	2.55
3.0	1.10	1.31	0.0826	3.0	4.64	3.82
5.0	1.06	1.18	0.0499	5.0	7.69	6.34

^a PF slope is short for psychometric function slope.

3.8 and the TvS slope is predicted to be 0.08. More generally, for these set sizes, the predicted TvS slope is given by $0.247/b_{\text{PF}}$. Thus, the predicted set-size effects are sharply attenuated with steep psychometric functions. By comparison, the high threshold theory predicts no effect of set size, thus its predictions are off the bottom of the graph in Fig. 8. The predictions of the ideal observer and maximum-of-differences theories are essentially indistinguishable from the predictions of the maximum-of-outputs theory.

5.3. Discussion

The observed psychometric functions show an effect of set size on the slope of the psychometric function. This result is consistent with all of the low threshold theories based on signal detection theory and is inconsistent with the high threshold theory. Looking more closely at the example, the observed psychometric functions slopes are too low for the standard assumptions. For set size 1, they predict a slope of 1.0 while 0.6 is observed. The power function generalization can describe such low slope values. Furthermore, this generalization still predicts an increase in the psychometric function slopes. Another prediction of this generalized theory is how the TvS slope depends on the psychometric function slope. Fig. 8 shows that the observed pair of values is consistent with the prediction of all three versions of signal detection theory. A further test this prediction is found in McLean et al. (1997) that describes data with steeper psychometric functions and shallower TvS functions for letter detection tasks. In summary, the signal detection theories describe well the effects of set size on the psychometric functions if more general distributional assumptions are made.

To conclude, discriminability and set size interact. Here, this interaction is documented in terms of d' but the same interaction is evident with probability correct. By either measure, set-size effects are smaller for more discriminable stimuli. Because of this interaction, one

must question any comparison of set-size effects that does not equate for the effects of discriminability.

6. Effects of multiple targets

In visual search, the introduction of multiple stimuli allows for the presence of multiple targets as well as multiple distractors. The processing of multiple targets is an instance of the larger topic of integrating information from multiple sources. Shaw (1982) has reviewed information integration theories. She distinguishes between *first-order integration models* that integrate information about the relevant stimulus attributes and then make a single decision, and *second-order integration models* that make separate decisions for each attribute and integrate the decisions. In the response time litera-

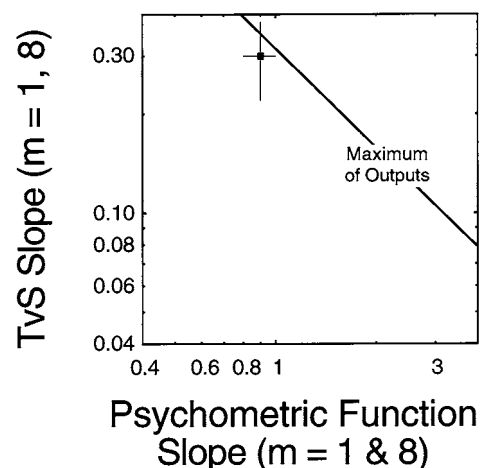


Fig. 8. This scatterplot illustrates the interaction between the slope of the psychometric function and set-size effects. The log-log slope of the TvS function is plotted against the average slope of the psychometric functions. The contour indicates the predicted values for the maximum-of-outputs theory and the symbol indicates the observed value in the contrast increment experiment.

ture, the corresponding models have been termed *coactivation* and *race* models, respectively (e.g. Raab, 1961; Miller, 1982; Mordkoff & Egeth, 1993). Of the theories considered here, the ideal observer theory is the best example of first-order integration whereas the maximum-of-outputs theory is the best example of second-order integration. The latter is because the maximum-of-outputs theory is equivalent to Shaw's (1980) independent decisions model which is the prototype second-order model (see Appendix A). In summary, the ideal observer and maximum-of-outputs (independent decisions) theories span the range of information integration theories.

Information integration has also been analyzed using detection experiments in the summation paradigm (Graham, 1989). In this paradigm, the information is integrated from separate stimulus attributes rather than separate stimuli. Examples include summation across spatial frequencies (e.g. Graham, Robson & Nachmias, 1978), summation across spatial positions (e.g. Robson & Graham, 1981), and summation across color directions (e.g. Poirson, Wandell, Varner & Brainard, 1990). The analysis of the summation paradigm includes detailed predictions for both high and low threshold theories. In addition, this work includes controls for discriminability using thresholds as was discussed above for set-size effects. Surprisingly, there have been almost no studies that exploit these methods for the effect of multiple targets in visual search. Fortunately, there is one example for speed discrimination where threshold was estimated as a function of the number of targets.

6.1. An example speed discrimination experiment

In Verghese and Stone (1995), speed discrimination was measured for multiple patches of grating. One, two, four or six patches were presented moving at an identical speed. This experiment maximized the effect of multiple targets by making the number of targets equal to the number of stimuli. The patches were presented in two intervals with all of the patches in one interval moving faster than the other interval. The observer's task was to choose the interval with the faster motion. This was a 2IFC task rather than the Yes/No tasks that we are focusing on in this paper. The theory relating 2IFC and Yes/No tasks has been extensively developed and we choose to focus on Yes/No only to simplify the presentation (Green & Swets, 1966). The relative threshold predictions for 2IFC have been given in Verghese and Stone (1995) and for this case are essentially indistinguishable from the predictions for Yes/No.

In Fig. 9, speed discrimination thresholds are shown as a function of the number of targets. The dotted lines indicate the thresholds for each of the four observers,

the bold points indicate the mean threshold, and the solid line indicates the best fitting linear regression on this log–log plot.⁴ A linear function on this plot can be characterized by Eq. (1) with the number of targets substituted for the set size. The individuals all show a similar effect of the number of targets with slopes on these log–log plots ranging from -0.20 to -0.44 . The mean *threshold versus number of targets* (TvNT) slope is -0.30 ± 0.05 . All four subjects also show signs of the function asymptoting at four to six targets. We do not try to account for this effect here.

6.2. Predictions

For the maximum-of-outputs and high threshold theories, analytic expressions are given in Appendix A for the predicted effect of the number of targets on threshold. For the ideal observer theory, for this particular situation, the predictions are the same as the analytic predictions of the sum-of-outputs theory (Green & Swets, 1966). For the maximum-of-differences theory, performance is predicted to be at chance for the case in which the number of targets equals the set size because there are no differences between the stimuli (all targets or all distractors) to yield information about the presence or absence of a target. Thus, this theory is inappropriate for this example. The predictions for the ideal observer and maximum-of-outputs theories depend on the choice of distributions and their statistical independence. As before, we first present predictions for independent, equal-variance normal distributions.

In Fig. 10, the predicted relative thresholds are shown as a function of the number of targets for each of the theories. As before, all predictions are scaled relative to the threshold predicted for set size 2. The three solid lines represent the predictions of the high threshold theory for three values of the steepness parameter, k , of the Weibull psychometric function. When k is infinite, then the high threshold theory predicts no summation and there is no effect of multiple targets. When k is 1, the theory predicts linear summation and there is a large effect of multiple targets. In that case, the predicted function on this log–log plot is linear with a TvNT slope of -1 . More generally, the high threshold theory predicts a linear relation on a log–log plot with a TvNT slope of $-1/k$. Thus, when $k = 3$, the predicted TvNT slope is $-1/3$ as is shown in the figure.

⁴ This presentation of the results differs slightly from that published in Verghese and Stone (1995). In order to maintain consistency with the other examples in this paper, we define the threshold at 75% correct performance instead of 82% and fit the psychometric functions to a d' power law rather than a Weibull.




Fig. 9. The speed difference threshold is plotted as a function of the number of targets (Vergheese & Stone, 1995, experiment 1). The dotted lines indicate the individual observer thresholds; the solid points indicate the mean thresholds of four observers; and, the solid line is the best fitting linear regression on this log–log plot.

The predictions of the ideal observer theory are shown by the finely dashed line. It predicts a linear function on a log–log graph with a TvNT slope of -0.5 . The prediction of the maximum-of-outputs theory are shown by the coarsely dashed line. It predicts a curve that deviates from a straight line on a log–log plot. The TvNT slope decreases slightly with an increasing number of targets. For set size 2, the TvNT slope is -0.25 . The specific values given here are for a yes/no task but the relative threshold predictions for a 2IFC tasks are essentially indistinguishable. For example, the predicted slope at set size 2 is -0.28 for 2IFC and -0.25 for Yes/No.

The predicted combinations of TvNT slope and psychometric function slope are shown in Fig. 11. The ordinate is the absolute value of the TvNT slope to avoid negative values and allow logarithmic scaling. The abscissa is the slope of the psychometric function in terms of the d' power law. As described before, this slope is about 0.9 of the Weibull k parameter. For the high threshold theory, the predicted values fall on a line with slope of -1 and pass through the point $(0.9, 1.0)$. The predictions of different versions of signal detection theory also form lines with slopes of -1 , but differ in their vertical position. For a given psychometric function, the high threshold theory predicts the largest multiple target effect, the ideal observer theory predicts the next largest effect, and the maximum-of-outputs theory predicts the least effect. In summary, all three theories when generalized predict a similar relation between of the effect of multiple targets and the slope of the psychometric function. However, they differ in the overall magnitude of the effect.

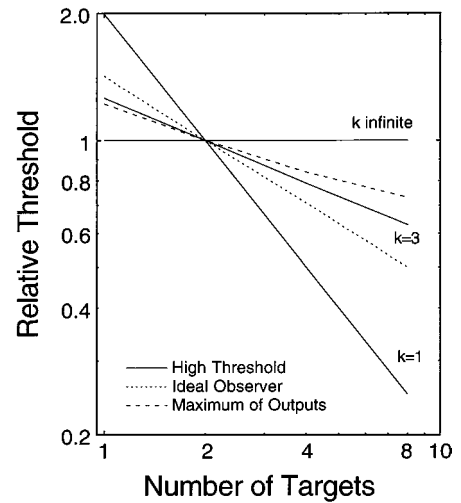


Fig. 10. Predicted relative thresholds are plotted as a function of the number of targets for several theories. All thresholds are relative to that predicted for set size 2.

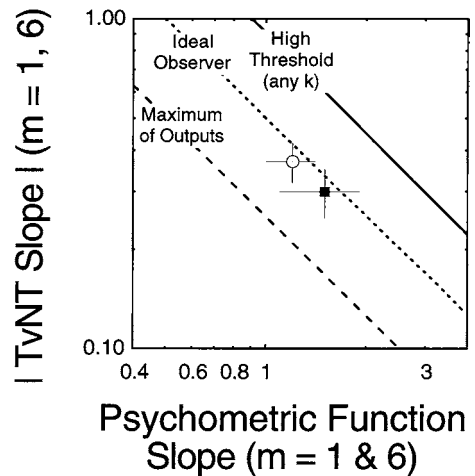


Fig. 11. An illustration of the interaction between the effect of multiple targets and the slope of the psychometric function. The absolute value of the log–log slope of the TvNT function is plotted against the slope of the psychometric function. The contours indicate the predicted values for each theory and the symbols indicates the observed value for the no-noise (solid) and noise (open) conditions of the Vergheese and Stone (1995) speed discrimination experiment.

6.3. Discussion

The results of Vergheese and Stone (1995) are shown as the solid symbol in Fig. 11. The TvNT slope is -0.30 ± 0.05 and the psychometric function slope is 1.5 ± 0.4 . This result falls closest to the prediction of the ideal observer theory. However, the large variability in the psychometric function slope makes it hard to reject the alternative theories. This conclusion differs from Vergheese and Stone (1995) because they only considered the linear model rather than the power law generalization. Further evidence using this analysis (the open symbol) is presented in the section on external

noise. To foreshadow that result, it also favors the ideal observer over the other theories. In closing, the evidence from this one experiment on multiple targets slightly favors the ideal observer over the alternatives.

7. Effects of response bias

Response bias is the tendency to use one response more than another and it can have a large effect when uncontrolled. Extreme biases in a yes/no task can push the probability of responding yes down to 0 or up to 1. An important success story of psychophysics has been both the reduction of biases using the forced-choice procedure and the analysis of bias apart from sensitivity for the yes/no procedure (Green & Swets, 1966; Swets, 1986a,b). One can also exploit response bias to reveal information about how decision processes combine bias and stimulus information. The relevant data are usually plotted in a *receiver operating characteristic function* (ROC function): the probability of a hit is plotted against the probability of a false alarm. In detection, such functions were important in distinguishing high threshold from low threshold theories (Tanner & Swets, 1954; for reviews see Luce, 1963a; Green & Swets, 1966; Swets, 1986a,b). The ROC function has also been used to study the effects of uncertainty in detection (Nachmias & Kocher, 1970; Graham, 1989) and the effects of set size in search tasks (Cohn & Lasley, 1974; Swensson & Judy, 1981; Swensson, 1996). The last line of work is perhaps the best evidence against the high threshold theory in visual search. Here we describe the ROC functions measured for another example search task.

7.1. An example speed discrimination experiment

In a new experiment, Verghese measured speed discrimination as a function of set size. The experiment was similar to the search experiment described in Verghese and Stone (1995, experiment 2) and is described in Appendix B. The observer's task was to indicate whether or not the display contained a single grating patch that moved faster than the other patches. Stimuli were briefly presented (195 ms), 20% contrast drifting grating patches in a stationary Gaussian window. Two experienced observers participated, BB and PV (an author). The major difference from Verghese and Stone (1995) was the use of a rating procedure rather than 2IFC. The observers responded on a 4 point scale: 1 indicated certainty that the target was present and 4 indicated certainty that it was absent. These ratings were used to construct ROC functions and the d'_e was estimated for each of several speed differences.

The results are shown as ROC functions in Fig. 12: the upper panels are for set size 2 and the lower panels

are for set size 6. Panels on the left are for Observer BB; those on the right are for Observer PV. Each panel plots the ROC using a double probability plot. The double probability plot is the z -transform of the hit probability plotted against the z -transform of the false alarm probability. (The z -transform is the inverse of the cumulative normal distribution function.) The lines are the best fitting linear functions on these axes for each condition (using maximum likelihood, Dorfman & Alf, 1969) and are given by

$$z(\text{hit}) = a + b_{\text{roc}} z(\text{false alarm}) \quad (3)$$

where a is the y -intercept and b_{roc} is the slope in this double probability plot. This function is more usefully parameterized in terms of d'_e and the slope

$$z(\text{hit}) = d'_e (b_{\text{roc}} + 1)/2 + b_{\text{roc}} z(\text{false alarm}). \quad (4)$$

The value of d'_e is equal to d' with equal bias defined as by where the ROC function intersects the negative diagonal. Within each ROC plot, functions for three discriminability values are plotted with the solid lines fit to the smallest speed difference and the coarsest dashed lines fit to the largest speed difference. Thus, we estimate both the sensitivity parameter d'_e and the ROC slope parameter b_{roc} for each condition.

Consider first the top left graph that shows the set size 2 condition for Observer BB. The observer has succeeded in using the rating scale to yield responses that show very different biases. The central point is near the negative diagonal which indicates an equal bias between misses and false alarms. In contrast, the points to the upper right show higher false alarms accompanied by fewer misses and the points to the lower left show lower false alarms accompanied by more misses. For BB, these effects are well described by the linear functions on these axes. Such linear functions are predicted by a signal detection theory which assumes normal distributions and are often found in detection experiments (Green & Swets, 1966; Nachmias & Kocher, 1970). The fits are particularly good for the two lower speed differences but less good for the largest speed difference. Next consider the rest of the data. Many of the conditions are better described by a curve that is concave down. This deviation is minor for Observer BB but quite consistent for Observer PV. Examining both set sizes, BB has only one function that is concave down. In contrast, for PV all six functions are concave down. The concave downward functions may be fit by further generalizing the distributional assumptions. One way is to modify the normal distributions to have heavier tails (e.g. Kassam, 1988). Another way is to introduce a small percentage of trials ($\sim 5\%$) with pure guessing. Because only a single observer shows this effect, we do not pursue these possibilities here.

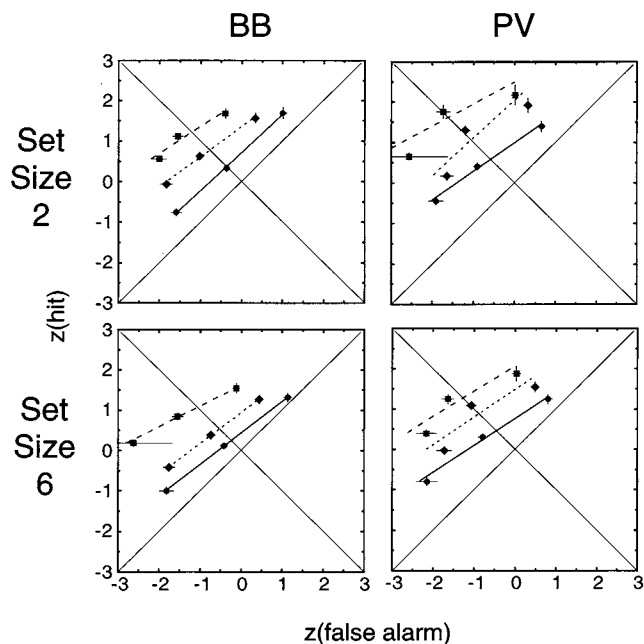


Fig. 12. ROC functions are plotted for set size 2 in the top two panels and for set size 6 in the bottom two panels. Separate observers are shown in the left and right panels. Within each panel, the transformed probability of a hit is shown as a function of the transformed probability of a false alarm. Each triplet of points is for a separate level of discriminability. Each of these discriminability levels is fit by a separate linear contour on this plot. The positive diagonal represents chance performance and the negative diagonal represents performance with an equal probability of a 'yes' and a 'no' response.

To quantify the slope of the ROC functions, we focus on the middle discriminability condition because it is nearest threshold. For BB, the ROC slope is 0.74 ± 0.06 and 0.77 ± 0.05 for set sizes 2 and 6, respectively. For PV, the ROC slope is 0.94 ± 0.10 and 0.69 ± 0.06 for set sizes 2 and 6, respectively. The interpretation of PV's data is hindered by the poor fit of the linear functions to her concave down data. In summary, the ROC curves are well fit by linear functions for one observer but require a concave down function for the other. The ROC slopes are in the 0.7–0.8 range and are little affected by set size.

7.2. Predictions

Fig. 13 shows in separate panels the predictions of the high threshold, ideal observer, maximum-of-outputs, and maximum-of-differences theory. Each panel is a double probability plot identical to that used to plot the data. The plots include the positive diagonal which results from chance responses and the negative diagonal which indicates an equal bias response. Each plot shows the predicted ROC functions for several set sizes. For set size 1, the discriminability is chosen that yields 75% correct at the equal bias point (negative diagonal).

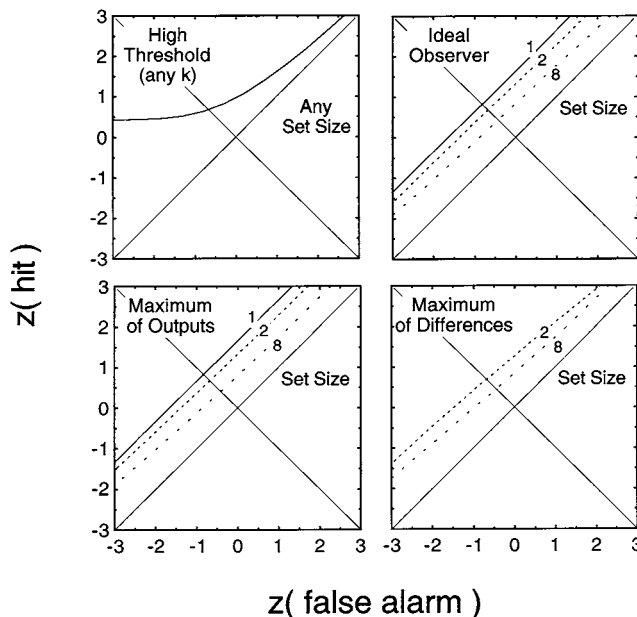


Fig. 13. Each panel has the ROC curves predicted by a different theory. Performance as the transformed probability of a hit is shown as a function of the transformed probability of a false alarm. Separate curves are shown for different set sizes.

The same stimulus is used for the other set sizes to illustrate the effect of set size at a given level of discriminability.

The upper left panel contains the predictions for the high threshold theory. This theory predicts no effect of set size, so a single ROC function is predicted for all set sizes. The predicted function is concave upward with the curve asymptotically approaching a minimum hit rate. This results in an ROC function that is very shallow near the equal-bias point described by the negative diagonal. Such ROC functions are rarely observed.

Next we turn to the three versions of signal detection theory. For all three, we show the predictions assuming equal-variance, normal distributions. The upper right panel contains the ideal observer predictions and the lower left panel contains the maximum-of-outputs predictions. These two theories have nearly identical predictions. For set size 1, the predictions are identical straight lines on this plot with a slope of one. For larger set sizes, performance is reduced but the curves remain very close to straight lines with the slopes decreasing slightly with increasing set size. Under these conditions, the curves are visibly indistinguishable from straight lines. The bottom right panel contains the predictions of the maximum-of-differences theory. Again, the predictions are nearly straight lines but the slopes are uniformly less than one and are less affected by set size.

We calculated the ROC slope as a function of set size for each of the theories. The slope is specified at the negative diagonal (equal bias) for a stimulus with

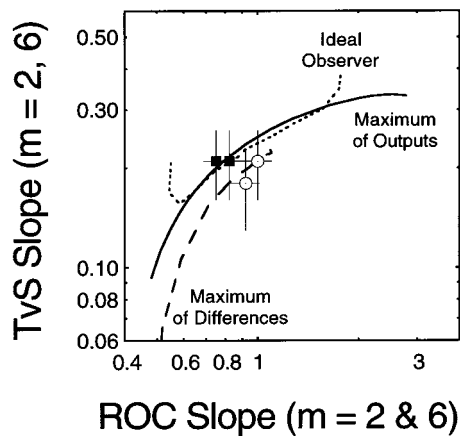


Fig. 14. An illustration of the interaction between the slope of the ROC function and set-size effects. The log-log slope of the TvS function is plotted against the average slope of the ROC functions. The contours indicate the predicted values for each theory and the symbols indicate the observed values for the no-noise (solid) and the noise (open) conditions of the example speed discrimination experiment.

threshold discriminability. At this point on the ROC function, the high threshold theory predicts a ROC slope of 0.33 for all set sizes; the ideal observer and maximum-of-outputs theories predict a gradual decline in slope from 1.0 at set size 1 to around 0.9 at set size 8. The maximum-of-differences theory predicts a relatively constant ROC slope of around 0.84. Thus, the signal detection based theories make a prediction quite distinct from that of the high threshold theory.

7.2.1. Generalized distributional assumptions

The predictions presented above are for equal-variance, normal distributions. This is clearly too restrictive because for set size 1 it predicts an ROC slope of 1 for any stimulus. Introducing unequal-variance, normal distributions yields ROC slopes that are equal to the ratio of the standard deviation of the distractor distribution relative to the standard deviation of the target distribution (*sigma ratio*, Swets, Tanner & Birdsall, 1961; Green & Swets, 1966; Nachmias & Kocher, 1970). Thus, this generalization can account for any 'linear' ROC function.

One of the predictions of this generalized theory is illustrated in Fig. 14. The magnitude of the predicted set-size effect depends on the ROC slope. For set sizes 2 and 6, this Figure shows the TvS slope as a function of the mean ROC slope for the two set sizes. Both threshold and ROC slope are estimated for the threshold level of performance (probability correct = 0.75). The prediction of the maximum-of-outputs theory is shown by the solid curve, the prediction of the ideal observer theory is shown by the fine dashed curve, and the prediction of the maximum-of-differences theory is shown by the coarse dashed curve. Consider first

the prediction of the maximum-of-outputs theory. For ROC slopes greater than 1, the TvS slope increases from 0.24 to 0.33 and flattens out. For ROC slopes less than 1, the TvS slope decreases sharply. For example, a mean ROC slope of 0.5 yields a TvS slope of only 0.11. Reducing the distractor variability makes this model more like a high threshold model and thus decreases the magnitude of the set-size effect. In fact, one can think of the high threshold model as a limiting case where the distractor variability is zero.

The other two theories yield similar predictions. The prediction of the ideal observer theory follows closely the prediction of the maximum-of-outputs theory over the middle range of sigma ratios. At extreme sigma ratios, the theories diverge. The prediction of the maximum-of-differences theory has the same form as the maximum-of-outputs theory, but it is shifted to lower TvS slope values. It also has the special characteristic that even an extremely high sigma ratio does not result in a high ROC slope.

7.3. Discussion

First consider the predictions of the high threshold theory in Fig. 13. It predicts concave upward ROC functions that differ from the observed linear or concave downward functions and it predicts values of ROC slope that are near 0.33 while those observed are around 0.75. Thus, one can reject the high threshold theory. This reinforces the similar evidence against the high threshold theory found in the previous search experiments measuring the ROC function (Cohn & Lasley, 1974; Swensson & Judy, 1981; Swensson, 1996).

Next consider the predictions of the three signal detection theories as a group. When assuming equal-variance, normal distributions, all three predict nearly linear ROC functions. This is a good approximation for one observer but not for the other. The theories also predict ROC slopes that range from 1.00 to 0.85. These predictions are higher than the observed values of around 0.75. In sum, with the equal variance assumption, the theories capture most of the qualitative features of the data, but not their quantitative details.

The quantitative details are better captured by the theories with more general distributional assumptions. Allowing unequal variance allows one to accurately describe the ROC slopes. The set-size effects and ROC slopes for the example experiment are shown in Fig. 14. The solid points indicate the values observed for the two observers and the open points are similar data from a second part of this experiment to be discussed in the section on distractor heterogeneity. For both conditions, the observed results are consistent with the similar predictions of the three versions of signal detection theory. In contrast, the prediction of the high threshold theory is for zero TvS slope which is off the bottom of this graph.

In summary, the observed ROC functions for search resemble those found for detection. This is expected if a common set of mechanisms underlies both tasks. In addition, multiple lines of evidence reject the high threshold theory for search and are consistent with signal detection theory.

8. Effects of external noise

External noise is a useful tool for testing alternative theories of visual mechanisms. It highlights the limitations of high threshold theory and provides a quantitative approach to the analysis of distractor heterogeneity as will be discussed in the following section. The analysis of noise in the visual system began with statistical theories of light detection (e.g. Rose, 1948; Tanner & Swets, 1954; Barlow, 1956). The manipulation of external noise is common in some areas of engineering (e.g. Mumford & Scheibe, 1968) and was introduced to contrast detection experiments by Nagaraja (1964). In this approach, noise is added by perturbing the luminance of each pixel in a raster display by an amount determined by a random sample from a normal distribution with a mean of zero and a given variance. The detection threshold is measured as a function of the variability of this *pixel noise*. The effects on contrast threshold are roughly proportional to the amount of noise, which results in very large effects for large amounts of noise. This work has been extended by many including Burgess, Wagner, Jennings, and Barlow (1981), Pelli (1981), Legge, Kersten and Burgess (1987) and Swenson and Judy (1996). The use of external noise has also been applied to a variety of other discriminations including color discrimination (Gegenfurtner & Kiper, 1992) and dot patterns (Burgess & Barlow, 1983). For a review see Pelli (1990), Pelli and Farell (1999); for a more general model see Lu and Doshier (1998, 1999).

The manipulation of external noise allows one to determine the characteristics of the internal noise that limits performance. In the most basic model, one assumes that external and internal noise sources are additive and independent. Assuming this additive model and normal distributions for both sources of noise, then the combined noise is also normal with a variance σ_t^2 that is the sum of the external noise variance σ_e^2 and the internal noise variance σ_i^2 :

$$\sigma_t^2 = \sigma_e^2 + \sigma_i^2. \quad (5)$$

In this formulation, the internal noise is represented by the *equivalent input noise* defined as follows: The σ_i is the standard deviation of the external noise that produces the same effect on performance as the internal noise. This allows the internal noise to be expressed in the same units as the external noise and the stimulus (e.g. contrast).

Combining these assumptions with signal detection theory, several authors have derived the relationship between threshold performance and the external noise. Here we follow the notation of Legge et al. (1987); for a more general formulation, see Lu and Doshier (1998, 1999). The main result is that a linear relation is predicted between the squared threshold t^2 and the external noise variance σ_e^2 :

$$t^2 = (d'_{\text{criterion}}/e)(\sigma_e^2 + \sigma_i^2) \quad (6)$$

where $d'_{\text{criterion}}$ is the d' value used to define threshold and e is an efficiency parameter called by a variety of names by various authors in different contexts. We follow Pelli and Farell (1999) and use the term *high-noise efficiency* for the e parameter. It estimates the efficiency of the observer relative to the ideal under conditions of high external noise. This formulation allows one to test the predicted linearity and estimate the internal noise from the intercept and the high-noise efficiency from the slope. When the external noise is large relative to the internal noise, this relationship simplifies to the threshold being proportional to the external noise. Under ideal conditions, the internal noise is zero and the high-noise efficiency is one. For contrast detection experiments with pixel noise, the high-noise efficiency is usually well below one, but does reach values as high as 0.8 for carefully selected stimuli (Burgess et al., 1981). An interpretation of high-noise efficiency can be found in models of multiplicative noise or gain control (e.g. Sperling, 1989; Lu & Doshier, 1998, 1999). We next consider one of the few studies of external noise and visual search. This example is from Verghese and Stone (1995) whose speed discrimination experiment was described in the section on multiple targets.

8.1. An example speed discrimination experiment

8.1.1. Method

Verghese and Stone (1995) conducted a control experiment that added external noise to the target and distractor stimuli. In the no-noise conditions already described, the distractor patches all moved at the reference speed of 5.3°/s; in the noise conditions, the speed of the patches was modified by an amount drawn from a normal distribution with mean of zero and a standard deviation of 1.8°/s. The no-noise and noise conditions were presented in separate blocks of trials. Note that the noise here was added to the speed of the stimulus rather than to the luminance of individual pixels. This kind of noise has been discussed by a number of researchers but has been used in only a few experiments (e.g. Ashby & Gott, 1988). Barbara Doshier suggested to us the name *dimensional noise* to distinguish it from pixel noise.

8.1.2. Results

Fig. 15 shows the mean threshold as a function of set size for both the noise and the no-noise conditions. The no-noise data have been shown above and illustrate a multiple target effect with a TvNT slope of -0.30 ± 0.05 . For the noise conditions, the set-size effects are quite similar with a slope of -0.35 ± 0.02 . Thus, one obtains a multiple target effect in a high noise condition and when measured as a slope on a log–log graph, there is little sign of an interaction between the level of the noise and the number of targets. Such an additive effect on logarithmic coordinates is equivalent to a multiplicative effect on linear coordinates.




Fig. 15. The speed difference threshold is plotted as a function of the number of targets for two levels of noise (Verghese & Stone, 1995, Figs. 3 and 7a). The lines are the best fitting linear regression on this log–log plot.

While the emphasis in this experiment is on the effect of multiple targets, one can also estimate the main effect of noise. Because the modeling is simplest for set size 1, we focus on the main effect of noise in that condition. The mean threshold is $0.82 \pm 0.07^\circ/\text{s}$ for no noise and $1.90 \pm 0.13^\circ/\text{s}$ for the noise condition. This is an increase of 133%. Fitting the data (only two points) to the predicted linear function between the squared threshold and the external noise variance (Eq. (6)) allows one to estimate the internal noise in terms of the equivalent input noise and the high-noise efficiency. The mean of the estimated standard deviation of the internal noise is $0.9 \pm 0.1^\circ/\text{s}$ and the mean of the estimated high-noise efficiency is 1.1 ± 0.2 . Thus, the internal noise is about half the size of the external noise used in this experiment and the high-noise efficiency is essentially perfect at 1. An experiment with multiple noise levels is needed to elaborate this result and, specifically, to test its linearity.

8.2. Predictions

The traditionally defined high threshold theory predicts no effect of noise because distractors always map into the no-detect state and targets map into the detect state by a deterministic function. It is possible to more generously interpret the theory by mapping the noisy distractors into the detect state according to the probability assigned to the particular noisy value of a stimulus. This, however, fundamentally changes the theory into a two-state, low threshold theory. Thus, it is unclear how, if at all, external noise can be integrated into high threshold theory.

Next consider the prediction for the three signal detection theories. Adding noise to both the target and the distractors simply scales all thresholds proportionally to the standard deviation of the total noise. This is because performance depends on the signal-to-noise ratio and the effect of multiple targets is a function of the total noise rather than the separate components of the noise. Thus, the interaction between noise and either multiple targets or set size should be multiplicative or, when on log–log graphs, a vertical displacement. Fig. 16 shows this relationship for the maximum-of-outputs and ideal observer theories. In each panel, the predicted threshold is shown as a function of the number of targets. The curve parameter is the amount of external noise normalized as a multiple of the internal noise. For the Verghese and Stone experiment, the appropriate comparison is to a normalized external noise of 2. The maximum-of-differences theory predicts chance performance for the case where the stimuli in the display are either all targets or all distractors. Thus, it is inappropriate for the case considered here. Another prediction of these theories was presented in Fig. 11 with the analysis of multiple target




Fig. 16. Each panel shows the predicted effects of noise for a different theory. The curves indicate the relative threshold as a function of the number of targets for different levels of external noise. The external noise magnitude is measured by the standard deviation of the external noise normalized to the assumed standard deviation of the internal noise.

effects. It shows that the predicted multiple target effect depends on the slope of the psychometric function.

8.3. Discussion

The high threshold theory predicts little or no effect of noise and extensions to allow it to do so turn it into a low threshold theory. Thus the high threshold theory cannot predict the large effects of noise such as observed here and elsewhere. The theories based upon signal detection theory all predict noise effects and independence between the noise level and the number of targets. To make quantitative comparisons, return to Fig. 11 which showed the predicted multiple target effect as a function of the slope of the psychometric function. On this graph, the solid points indicate the results for the no-noise condition and the open points indicate the results for the noise condition. Both results are closest to the predictions of the ideal observer theory. Taken together, they weigh against both the high threshold and maximum-of-outputs theory.

Next consider the interaction between external noise and multiple target effects. Both the ideal observer and maximum-of-output theories predict a multiplicative effect of the number of targets and noise. Thus, the parallel functions on the log–log plot of Fig. 15 are consistent with both theories. In contrast, such a result is inconsistent with other theories of set size or multiple target effects that pose some modification of the internal noise alone, without also affecting the consequences of the external noise (see Lu & Doshier, 1998, 1999).

Finally, consider the application of the external noise model to these search data. The model distinguishes between internal noise and high-noise efficiency. The high-noise efficiency estimates are remarkably high, near the value of 1 expected for an ideal observer. Such a high value puts a limit on the loss of information from any decision process such as criterion variability. This high value is almost certainly due to the use of dimensional noise in the speed of the stimulus rather than the pixel noise of previous experiments. Manipulating the variability in the dimension being judged avoids any inefficiencies in extracting that representation from a noisy image. For example, a less than optimal receptive field might include pixel noise from a larger region of the image than is relevant to a given task. The same receptive field could be perfectly efficient in conveying the variability among the different speeds of the stimuli.

The external noise model also allows one to estimate the internal noise. For this experiment, the internal noise has an equivalent input noise of about $0.9^\circ/s$ which is about half the $1.8^\circ/s$ external noise used in this experiment. The 1:2 ratio of standard deviations between these noise sources becomes a 1:4 ratio of variances. Thus, the total noise variance is composed of

20% internal noise and 80% external noise. This means that performance in the noise condition is largely limited by the external noise that had known independent distributions for each stimulus compared to the unknown internal noise distribution. Despite this change, the effects of the number of targets and of discriminability are very similar for the no noise and noise conditions. This is consistent with the internal noise being similar to the external noise in both distribution and independence.

In summary, the effects of external noise are inconsistent with the usual assumptions of high threshold theory and illustrate its limitations. In contrast, the signal detection theories do predict an effect of external noise which is consistent with the observed results.

9. Effects of distractor heterogeneity

Another consequence of introducing multiple distractors in a search task is the potential for introducing distractor heterogeneity. Thus, distractor heterogeneity ranks with set size and multiple targets as a manipulation unique to search. Furthermore, distractor heterogeneity often has a major effect on visual search tasks as demonstrated by Eriksen (1953), Gordon (1968), Farmer and Taylor (1980). More recently, Duncan and Humphreys (1989), Duncan (1989) have emphasized the potential role of distractor heterogeneity and how it interacts with target-distractor discriminability. They have suggested that the difference between conjunction and feature search may be accounted for in part by a confounding of the conjunction task with an increase in distractor heterogeneity.

More detailed work on distractor heterogeneity can be found in the texture, color, and medical imaging literature. An example of the studies in texture is the analysis of whether heterogeneity of irrelevant attributes has an effect on performance (e.g. Callaghan, Lasaga & Garner, 1986). Pashler (1988) has extended this work arguing that distractor heterogeneity of irrelevant dimensions usually has little effect, but there are exceptions that depend on the task as well as the stimulus. The color literature also contains several detailed studies (e.g. Bauer, Jolicoeur & Cowan, 1996a,b; Nagy, 1997). For example, Nagy has shown substantial effects of distractor heterogeneity unless the heterogeneity is specifically along a dimension of color that is both irrelevant to the task and independently represented from the dimension being judged. Finally, the medical imaging literature includes studies that introduce heterogeneity into the structured noise backgrounds of simulated medical images. Again there is often a large effect of this kind of heterogeneity (e.g. Rolland & Barrett, 1992; Eckstein, Ahumada & Watson, 1997).




Fig. 17. The orientation threshold is plotted as a function of set size for two levels of distractor heterogeneity. The no-noise condition has homogenous distractors and the noise condition has distractors sampled from a normal distribution with a standard deviation of 4° . The lines are the best fitting linear regression on this log–log plot.

The existence of distractor heterogeneity effects is well established. However, previous work has provided relatively little quantitative analysis of distractor heterogeneity effects. To this end, we next describe two studies that introduce distractor heterogeneity by adding noise to the distractors but not the targets. Both studies allow one to quantify the effect of distractor heterogeneity in a one-dimensional search task.

Adding noise to only the distractors results in a somewhat different prediction about the effect of noise. The distractor noise manipulation introduces a marked shift in the ratio of the variability of the target and distractor distributions. This should change the slope of the ROC function. Thus one needs a d' based measure that is defined over changes in ROC slope. (The usual d' measure assumes a ROC slope = 1.) The best candidate appears to be d_a rather than the d'_e measure used earlier in this article. The d_a is similar to the others in measuring the difference between the target and distractor distributions in terms of the standard deviation of the distributions (Simpson & Fitter, 1973; Macmillan & Creelman, 1991). It differs in scaling the difference with respect to the root mean squared deviation

$$d_a = \frac{\Delta m}{\sqrt{(\sigma_s^2 + \sigma_n^2)/2}} \quad (7)$$

rather than assuming equal variance as in the case of d' or the average of the standard deviations as in the case of d'_e . The nice feature of d_a is that it allows one to derive an analog of Eq. (6) for the case of noise added to distractors only. The prediction is again for the squared threshold to be linearly related to the external noise variance, but the parameters have changed:

$$t^2 = (d_{a-criterion}^2 / e) (\sigma_e^2 / 2 + \sigma_i^2), \quad (8)$$

In particular, the external noise variance is only half as effective in increasing the squared threshold.

We now consider the effect of distractor heterogeneity using ideas from external noise experiments. This analysis again shows the limitations of high threshold theory. Furthermore, it suggests an account of distractor heterogeneity based upon the variability already built into signal detection theory without requiring the introduction of any novel perceptual mechanisms.

9.1. An example orientation discrimination experiment

Pavel (unpublished study related to Pavel, Econopoulou & Landy, 1992) measured the threshold for orientation discrimination at set sizes 1, 2, 4, 6, and 12. The stimuli were 1° dark bars on a gray background. A 2IFC procedure was used with a display duration of 150 ms and an interstimulus interval of 750 ms. There were both a no-noise and a noise condition. The no-noise condition replicated previous set size experiments and the noise condition introduced distractor heterogeneity. For the no-noise condition, the distractors were always vertical and the target orientation was varied to determine a psychometric function and estimate a threshold. For the noise condition, the distractors were drawn from a normal distribution with a mean of 0° (vertical) and a standard deviation of 4° .

Fig. 17 shows the mean orientation threshold for 6 observers as a function of set size for both the no-noise and noise conditions. There are effects of both set size and distractor heterogeneity. For the no-noise condition, the TvS slope is 0.37 which is somewhat higher than the set-size effects described in the earlier section. For the noise condition, the TvS slope is 0.40. Thus, the magnitude of the set-size effect in terms of the log–log slope is little changed by the addition of distractor heterogeneity.

The effect of noise for set size 1 is to increase the threshold from 1.5° to 3.1° . Using the noise model just defined, this increase is consistent with a high-noise efficiency of 0.98 and an equivalent input noise of 1.6° . Thus, the high-noise efficiency is near 1 as with the previous noise experiment and the external noise standard deviation of 4° was much larger than the 1.6° that was estimated for the standard deviation of the internal noise. In summary, both set size and distractor heterogeneity have the expected main effects and these effects combine approximately multiplicatively resulting in a simple vertical displacement on the log–log graph.

9.2. An example speed discrimination experiment

The second example is again from our own speed discrimination and set size experiment (Appendix B). In addition to the no-noise condition that was described in the response bias section, there was another condition that added external noise to the speed of the distractors. In the no-noise condition, the distractor patches

all moved at the reference speed of 5.3°/s, while in the noise condition, the speed of the distractor patches was drawn from a normal distribution centered about the reference speed with a standard deviation of 0.89°/s. For both, the target speed was constant. The no-noise and noise conditions were presented in separate blocks of trials.

Fig. 18 shows the threshold as a function of set size for both the noise and the no-noise conditions. Data from two observers are plotted in the separate panels. The no-noise data illustrate the expected set-size effect with a TvS slope of 0.21 for both BB and PV. For the noise conditions, the set-size effects are similar with

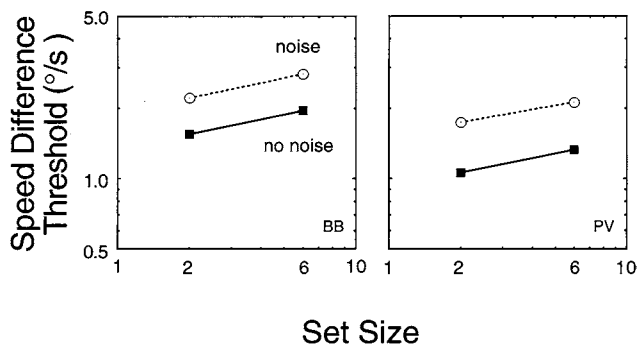


Fig. 18. The speed difference threshold is plotted as a function of set size for two levels of distractor heterogeneity. The no-noise condition has homogenous distractors and the noise condition has distractors sampled from a normal distribution with a standard deviation of 0.89°/s. The two panels contain data from separate observers.

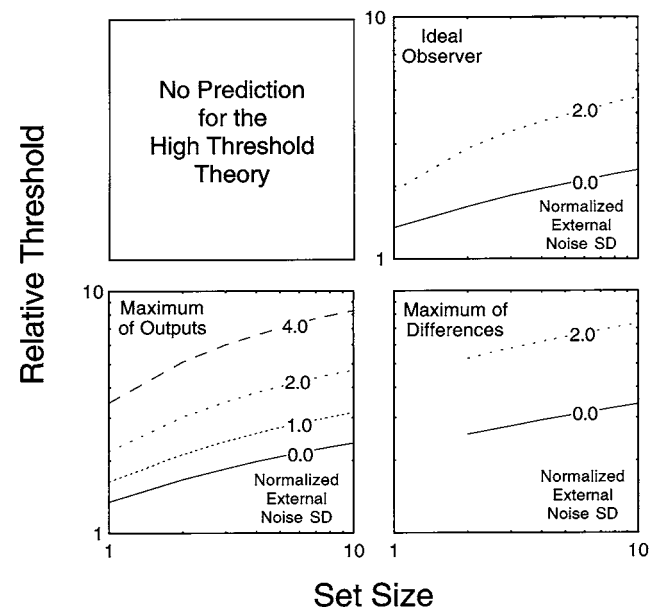


Fig. 19. Each panel shows the predicted effects of distractor heterogeneity for a different theory. The curves indicate the relative threshold as a function of the set size for different levels of distractor heterogeneity. The degree of distractor heterogeneity is indicated by the standard deviation of the external noise added to distractors alone normalized to the assumed standard deviation of the internal noise.

TvS slopes of 0.18 and 0.21 for BB and PV, respectively. There is no sign of an interaction between the level of the noise and set size.

To measure the effect of noise, we extrapolate the thresholds to set size 1 using a power function approximation to the set-size effects. For the no-noise condition, the extrapolated threshold is 1.34 and 0.92 for BB and PV, respectively; for the noise condition, the extrapolated threshold is 1.91 and 1.54, for BB and PV, respectively. From these values, the high-noise efficiency is estimated to be 0.43 and the equivalent input noise is estimated to be 0.54°/s. Surprisingly, the high-noise efficiency is about half that observed in the two external noise and distractor heterogeneity experiments described above. Task differences probably account for this difference as will be discussed below.

Distractor heterogeneity also has an effect on the slope of the ROC function. All of the noise conditions have ROC functions with higher slopes than the corresponding no-noise conditions. For BB, adding noise steepens the ROC slope from about 0.7 to 0.9. For PV, adding noise steepens the ROC slope from about 0.8 to 1.0. Overall, the noise increases the slope by an average of 40% which is nearly as large as the 50% effect on threshold. The observed combination of the ROC slope and the TvS slope is shown by the open symbols of Fig. 14.

In summary, adding variability to the single relevant attribute of these experiments decreases performance and raises the difference thresholds. The effect of set size in terms of log–log slope does not interact with distractor heterogeneity. On the other hand, the effect of noise in terms of high-noise efficiency is different for the orientation and speed experiments. It will be interesting to see how this quantification of heterogeneity effects extends to other attributes and to attributes irrelevant to the immediate judgment.

9.3. Predictions

Consider next the predicted effects of manipulating distractor heterogeneity. As with the external noise experiments, the high threshold theory makes no prediction. Predictions for the other three theories are shown in Fig. 19. Each panel shows the predicted relative threshold as a function of set size. The curve parameter is the standard deviation of the external noise normalized to the internal noise. The 0.0 curve corresponds to the no-noise condition and the 1.0 curve corresponds to the noise condition in the speed discrimination experiment. For the ideal observer theory, the set-size effect is larger with more external noise. This interaction is slightly reduced for the prediction of maximum-of-outputs theory and reduced further for the maximum-of-differences theory. One also notices that all of the predictions are concave down. However,

Table 2
Coarse summary of results

Phenomena	High threshold theory	Low threshold theories		
		Ideal	Max	Difference
Set-size effects on the observed threshold	Fail	Pass	Pass	Pass
Discriminability effects: interaction with set size	Fail	Pass	Pass	Pass
Multiple target effects on the observed threshold	Pass?	Pass	Pass?	n/a
Response bias effects: shape of the ROC	Fail	Pass ^a	Pass ^a	Pass ^a
External noise effects: interaction with number of targets	Fail	Pass	Pass	n/a
Distractor heterogeneity effects: interaction with set size	Fail	Pass	Pass	Pass

^a Passes for some observers but not others.

because we only measured performance at two set sizes, we ignore this aspect of the data and instead focus on the TvS slopes. To quantify this effect, we calculate the predicted TvS slope between set sizes 2 and 6 for both the no-noise (0.0) and noise condition (1.0). This pair of predicted TvS slopes are 0.23 and 0.34 for the ideal observer theory, 0.23 and 0.30 for the maximum-of-outputs theory, and 0.18 and 0.22 for the maximum-of-differences theory. The data are most consistent with the nearly equal TvS slopes predicted by the last theory but probably cannot reject any of the three theories. Hence, we emphasize the similarities rather than the differences.

9.4. Discussion

The high threshold theory as it exists cannot account for the effects of distractor heterogeneity. The heart of the high threshold theory is that distractors never result in a false alarm. Thus, adding noise to the distractors cannot affect performance. The three versions of signal detection theory are more flexible. Consider three points of comparison.

First, consider the main effect of set size. The analysis of Fig. 14 gives the predicted combinations of ROC slope and TvS slope. The solid symbols are for the no-noise condition and the open symbols are for the noise conditions. All of the points fall near the contour predicted by the maximum-of-outputs theory. The other versions of signal detection yield similar predictions.

Second, consider the interaction between distractor heterogeneity and set size. On these log–log plots, set-size effects are similar for the no-noise and noise conditions in both experiments. This is consistent with the predictions of all three theories based on signal detection theory. They all predict an increase in set-size effects for larger degrees of distractor heterogeneity with the largest predicted by the ideal observer and the smallest predicted by the maximum-of-differences theory.

Third, consider the main effect of distractor heterogeneity on thresholds. For both of the experiments, the modified noise model can be used to fit the parameters of equivalent input noise and high-noise efficiency. Surprisingly, the two experiments yield quite different values of high-noise efficiency. The 2IFC orientation experiment yields efficiencies near 1 similar to the 2IFC speed discrimination experiment with noise added to both targets and distractors. In contrast, the yes–no speed discrimination experiment with noise added to just distractors yields lower efficiencies of around 0.5. It remains to be seen whether or not this difference is idiosyncratic of these two experiments or if it reflects a more general phenomenon such as larger criterion variability in yes–no tasks.

In summary, the effect of distractor heterogeneity is not consistent with high threshold theory. In contrast, the examples of low threshold theory are all consistent with several aspects of the data. This analysis provides a new quantitative treatment of distractor heterogeneity that follows naturally from the ideas of signal detection theory and requires no additional perceptual mechanisms.

10. General discussion

The discussion begins with a brief summary of the results for the six phenomena. This is followed by a discussion of the theoretical implications and a consideration of alternative theories. Then the discussion turns to generalizing our analysis. We address search asymmetry, multiple dimensions and conjunction search, response time paradigms, search with multiple eye fixations, and more general stimulus conditions.

10.1. Summary of results

The results from all of the reviewed studies are summarized in Table 2. It specifies a coarse pass–fail summary of each theory for each of the major results.

10.1.1. Set size

Many studies have established the occurrence of a set-size effect of a particular magnitude. All of the versions of the signal detection theory predict the right magnitude of the effect. The high threshold theory does not predict such an effect.

10.1.2. Discriminability

The few available studies demonstrate the steepening of the psychometric function with increasing set size. This effect is predicted by all three versions of signal detection theory and is not predicted by the high threshold theory.

10.1.3. Multiple targets

The one available study yields modest multiple target effects. It favors the ideal observer theory over both the high threshold and the maximum-of-outputs theory. However, this result is not as conclusive as the others on this list. So the alternatives are given a marginal pass in the summary of Table 2.

10.1.4. Response bias

The few available studies show the shape of the ROC function as linear or concave down. This shape is partly consistent with the nearly linear predictions of the three signal detection theories and it is inconsistent with high threshold theory.

10.1.5. External noise

The one existing study shows that the multiple target effect interacts multiplicatively with external noise (a vertical shift on the log–log graphs). Such a relation is predicted by the ideal observer and maximum-of-output theories. (The maximum-of-difference theory predicts the effect of external noise but not the effect of multiple targets in the paradigm reviewed here). External noise effects cannot be accounted for by the high threshold theory.

10.1.6. Distractor heterogeneity

Two studies show that distractor heterogeneity interacts multiplicatively with set size (a vertical shift on the log–log graphs). Similar interactions are predicted by all three signal detection theories. An effect of distractor heterogeneity cannot be accounted for by the high threshold theory.

10.2. Theoretical implications

10.2.1. High threshold theory

The high threshold theory fails on almost every test of its predictions. It can be clearly rejected as a theory of visual search accuracy as it was as a theory of simple detection and discrimination. See the section below on response time for a discussion of the role of a high threshold in many response time theories of search.

10.2.2. Low threshold theory

The three versions of signal detection theory fail in only a few comparisons. In particular, the ideal observer theory with generalized distributional assumptions passed all of the six tests. The other two signal detection theories do nearly as well. The maximum-of-outputs theory may fail in predicting too small an effect for multiple target experiments. The maximum-of-differences theory does not fail but could not be applied to the multiple target effects as reviewed here. It also cannot be applied to set size 1. It needs further generalization to be as widely applicable as the other theories considered here. For example, one could combine its comparison to neighbors with a comparison to a memory standard using a Bayesian rule for combining different sources of evidence. One could use comparisons to memory or comparisons to neighboring stimuli depending on which provided the best source of information. This is a hybrid of the maximum-of-outputs and maximum-of-differences theory.

In summary, we emphasize the similarities among the predictions of the three versions of signal detection theory. This approach seems to be robust across a variety of assumptions about the representation and decision process.

10.2.3. Alternative high threshold theories

Thus far in this paper, we have focused on the traditional high threshold theory without any additional mechanisms. To better account for some of the results, one could add mechanisms such as a limited capacity process (Broadbent, 1958, 1971). An example of one such limited capacity process is the fixed capacity theory described in Palmer et al. (1993; also known as the sample size model, Taylor, Lindsay & Forbes, 1967; Lindsay, Taylor & Forbes, 1968; Shaw, 1980). In this theory, perception is assumed to be the result of a sampling process in which one can take a fixed number of samples of the entire visual scene. If only one attribute of the scene is relevant, then all of the samples can be concentrated on that attribute to result in a very precise representation of that attribute. If more attributes are relevant, then the samples are distributed across the relevant attributes which results in less precise representations of each attribute. For equally distributed and independent samples, the result of this process is that the standard deviation of the representation is proportional to the square root of the number of relevant attributes.

If this fixed-capacity sampling perceptual process is combined with the high threshold theory, the result is a theory that predicts larger set-size effects and smaller multiple target effects than the high threshold theory alone. Both of these changes make the fixed-capacity, high threshold theory more consistent with the observed results than is the high threshold theory in its

original form. Indeed, if one adopts a yet more general limited-capacity, high threshold theory, then one can adjust the degree of limited capacity to predict any set-size effect. Such general versions of the high threshold theory make the magnitude of set-size effects no longer a test of the theory. Thus, one must focus on other phenomena to test this generalized theory. Adding limited-capacity mechanisms does not change the predicted ROC functions or the steepness of the psychometric functions. Thus they remain a problem for this theory. In addition, limited capacity does not make it any easier to account for the external noise and distractor heterogeneity effects. Thus, generalizing the theory to include a limited-capacity process does not alleviate the limitations of the high threshold theory.

10.2.4. *Other low threshold theories*

In this article, we focus entirely on signal detection theory as an instance of a low threshold theory. A more complete treatment requires consideration of other low threshold theories (for reviews see Luce, 1963a). These theories include modifications of the high threshold theory such as two-state, low threshold theory (e.g. Luce, 1963b), three-state, low- and high-threshold theory (Krantz, 1969), and various extensions of Luce's choice theory (1959). For example, Bundesen (1990) develops choice theory and applies it to a variety of attention phenomena particularly those using identification paradigms. We do not pursue these alternatives here and instead focus on signal detection theory as a familiar and plausible example.

We also focus on the 'independent channels' version of the various signal detection theories. One could certainly consider limited capacity versions of signal detection theory. For example, Palmer et al. (1993) described a combination of the maximum-of-outputs integration rule and a fixed capacity process that is based on the sample size model (e.g. Shaw, 1980). They showed that such a theory can be rejected for simple search because it predicts set-size effects of much greater magnitude than observed (TvS slopes of 0.75 instead of 0.25). Thus, at least this specific combination can be rejected. More generally, there seems little reason to consider limited capacity models when an unlimited capacity model already accounts for the data of simple visual search.

10.3. *Generalizations for search asymmetries*

We now turn the discussion to related phenomena in visual search. Asymmetries in discrimination are interesting because they are not predicted by most simple theories. Such asymmetries are rarely reported for simple discrimination tasks but may be more common in visual search. Treisman and Souther (1985) provided an early report of search asymmetry. An example of their

task was to discriminate between a circle and a circle intersected with a small line segment similar to the letter Q. They measured response time as a function of set size and found larger set-size effects for discriminating an O from Qs than for discriminating a Q from Os. A variety of other stimulus situations have also shown such asymmetries (e.g. Treisman & Gormican, 1988; Driver & McLeod, 1992; Ivry & Cohen, 1992; von Grünau & Dubé, 1994; Carrasco, McLean, Katz & Frieder, 1998). Here we discuss an empirical example of an asymmetry and then consider how asymmetry can be accounted for using the theories developed in this article. In particular, signal detection theories provide a natural account of asymmetries due to potentially asymmetric variances of targets and distractors.

10.3.1. *Empirical examples*

Perhaps the most detailed study of search asymmetry is that of Nagy and Cone (1996). They compared a variety of color pairs in a search task in which one had to localize a small target disk of one color among 53 small disks of the other color. For each color pair, they systematically varied the target-distractor discriminability and measured response time as a function of the color difference. They found that search asymmetries occurred, but only for certain pairs of colors. In particular, searching for saturated among desaturated colors resulted in faster performance than searching for desaturated among saturated colors. In addition, by measuring response time as a function of the color difference, Nagy and Cone estimated the difference threshold for a given response time for each condition. In terms of these difference thresholds, the asymmetry increased the threshold by 10–30% depending on the color pair. Thus, this was a modest but reliable effect.

Palmer and Teller (unpublished) have also demonstrated asymmetries in search accuracy experiments modeled after early reports of Nagy and Cone. They measured color thresholds for a white target and pink distractors (and vice versa) using methods similar to Palmer (1994). The chromatic contrast thresholds were lower for discriminating between a pink target and white distractors than a white target and pink distractors. The change in threshold was $13 \pm 4\%$ of the smaller threshold. Thus, these asymmetries appear in accuracy as well as response time paradigms.

10.3.2. *Theoretical accounts*

Asymmetries raise at least two questions. The first is why do they occur for some stimuli and not others? The second is what mechanism accounts for them when they do occur? We pursue the second of these two questions.

Treisman and Souther (1985; see also Williams & Julesz, 1992) have suggested that differences in the representation of different stimuli produced the asymmetries. In particular, some stimulus characteristics

might be directly represented in unlimited-capacity feature maps while others might need to be inferred from more complete representations only available to a limited-capacity process. Alternatively, Nagy and Cone (1996) have suggested that their asymmetries between saturated and unsaturated colors were linked to processing time differences between saturated and unsaturated colors. Similar differences have been observed for simple discrimination tasks (Nissen, Pokorny & Smith, 1979).

Another possible explanation arises from supposing differential variability in the representation of different stimuli (Rubenstein & Sagi, 1990). For example, nearly all intensive dimensions show decreasing discriminability for increasing values (e.g. Weber's Law). These effects may be due to increasing noise with increasing stimulus values (cf. Graham, Kramer & Yager, 1987). For example, the representation of the length of a long line is more variable than the representation of the length of a short line. Combining this idea with signal detection theory predicts that searching for a long line among short lines yields different performance than vice versa. More generally, differences in the variability of targets and distractors might account for other asymmetries such as between saturated and unsaturated colors or between prototypical and non-prototypical stimuli.

These ideas can be made more specific by combining them with the theory presented here. The size of the predicted asymmetry depends on the distributional assumption about the representation of the target and distractor stimuli. Following the analysis presented in the response bias section, suppose one assumes un-




Fig. 20. An illustration of how an unequal variance for targets and distractors predicts asymmetries in performance. The predicted relative threshold is plotted as a function of the ratio of target and distractor standard deviations (sigma ratio). This relation is shown for two set sizes. The relative magnitude of the set-size effect grows with larger values of the sigma ratio.

equal-variance, normal distributions and the maximum-of-outputs theory. The relative thresholds for various conditions depend on the ratio of the standard deviation of the distractor distribution to the standard deviation of the target distribution. This *sigma ratio* is equal to the predicted ROC slope for set size 1. Fig. 20 shows the predicted relative threshold as a function of the sigma ratio for set sizes 1 and 8. The threshold rises with increasing sigma ratio and rises yet more quickly for larger set sizes. Compared to a sigma ratio of 1, increasing the sigma ratio results in larger set-size effects and decreasing the sigma ratio results in smaller set-size effects. To compute the predicted asymmetry, one can compare the predicted performance for a given sigma ratio and its inverse which represents reversing the roles of targets and distractors. Asymmetry effects of 10% (threshold ratio = 1.1) require a sigma ratio of 1.17 and effects of 30% require a sigma ratio of 1.57. Such ratios are compatible with those observed in the speed discrimination experiment presented in the response bias section. To conclude, the relative variability of the representations of targets and distractors may mediate some instances of search asymmetries.

10.4. Generalization to multiple dimensions and conjunction search

The body of this article focuses on experiments within a simplified domain chosen at the intersection of visual search and simple detection and discrimination. Now we begin to consider how to generalize to a larger variety of visual search experiments including multiple dimensions, response time, eye movements, and more general stimulus conditions.

Our first generalization is to tasks that depend on multiple dimensions rather than a single dimension. For example, in conjunction search the target is distinguished from the distractors by a conjunction of properties. No one property is sufficient as is the case with feature search. We compare two approaches to this generalization. The first is multidimensional signal detection theory as it has been applied in psychophysics. The second is feature integration theory (Treisman & Gelade, 1980) that assumes separate one-dimensional representations that are often called feature maps. This alternative theory has been proposed to account for several phenomena including the relative difficulty of conjunction and feature search.

10.4.1. Multidimensional signal detection theory

Multidimensional generalizations of signal detection theory have been used in engineering (e.g. van Trees, 1968; Duda & Hart, 1973) and have found some application in psychology and vision. The concepts can be found in selected chapters of Graham (1989), Macmillan and Creelman (1991), Ashby (1992).

In brief, there are three parts to such a theory. First, stimuli are represented in a space of multiple dimensions with the simplest case being a two-dimensional space defined by orthogonal dimensions. Representations with more dimensions and nonorthogonal dimensions are possible. Second, the representation is characterized by a random variable in this space. Typically this variable is assumed to be multivariate normal with unit variance in each dimension and no correlation between the dimensions. Again, other distributions, variance scales, and correlational assumptions are possible. Third, the observer's decision is based on a criterion contour defined in the space. In two dimensions, the simplest case is a line parallel to one of the axes which indicates that the value on a single dimension determines the response. However, information from multiple dimensions can be used by comparing a weighted sum of the observed dimensional values to a criterion. In two dimensions, this is equivalent to dividing the space by a tilted line. Alternatively, independent criteria can be applied to each dimension resulting in a decision criterion that is equivalent to an AND or an OR rule. Or, more complex functions for integrating information between dimensions result in more complex contours such as the case for the ideal observer theory (e.g. Green & Birdsall, 1978). Here we consider only theories with orthogonal dimensions, unit normal distributions with statistical independence between the dimensions, and either linear or independent decision criteria.

To apply this theory to search, one must integrate the analysis of multiple stimuli with the analysis of multiple dimensions. Some of the possible integrated theories are in Eckstein, Thomas, Whiting, Palmer and Shimozaki (in press). They combine a maximum-of-outputs rule for combining information across stimuli with either the optimal linear decision criterion (they refer to as the max-linear rule) or with an independent decisions rule (AND) for each dimension (they refer to as the max-min rule). For these two rules, predictions are calculated for a variety of tasks including feature search, conjunction search, triple conjunctions, and disjunction search. Performance is predicted to vary with the task. For example, assuming two dimensions, the predicted d' for a conjunction task is smaller than the corresponding feature tasks by a factor of one over square root of two (0.71). If one adjusts for this performance difference by measuring set-size effects at threshold, then the various tasks are predicted to have similar set-size effects.

In summary, the simple one-dimensional theories can be generalized to multiple dimensions and to the variety of tasks possible with multidimensional stimuli. For the simplest such generalizations, the choice of task is predicted to have an effect. However, at least for the cases discussed, the choice of task has little interaction with set size if discriminability is controlled.

10.4.2. Feature integration theory

An interesting contrast to multidimensional signal detection theory is provided by feature integration theory (Treisman & Gelade, 1980; Treisman & Sato, 1990). It is an instance of a two-stage search theory (Hoffman, 1979). In such theories, the first stage is typically assumed to process all stimuli in a parallel, unlimited-capacity fashion. The second stage is assumed to process only a single stimulus (or perceptual object) at a time. In feature integration theory, the first stage has separate one dimensional representations (feature maps) rather than a combined representation. These representations are separate in the sense that one cannot base a decision on combinations of dimensions. This is equivalent to a multidimensional representation with orthogonal dimensions and decision criteria that must depend on only a single dimension at a time. Thus, in this respect, the feature map representation is a special case of the representations discussed above. The second stage has a full multidimensional representation of one stimulus at a time. Thus, it can use decision rules that depend on more than one dimension but only as part of a serial comparison process. Feature integration theory makes several predictions as a result of being able to combine feature information in only the second stage. Other things being equal, the theory predicts that conjunction search must result in qualitatively larger set-size effects than feature search. This prediction has found mixed support in a variety of studies (e.g. Treisman & Gelade, 1980; Egeth, Virzi & Garbart, 1984; Nakayama & Silverman, 1986; Wolfe, Cave & Franzel, 1989; Duncan & Humphreys, 1989; Mordkoff, Yantis & Egeth, 1990).

In the simple form described thus far, feature integration theory has a fundamental deficit that was perhaps most forcefully pointed out by Duncan and Humphreys (1989). The theory does not explicitly account for the large effects of target-distractor discriminability and distractor heterogeneity. Thus, comparisons between different stimuli and different tasks cannot be critical tests of the various predictions unless discriminability is also controlled. This control has been addressed in more recent studies with mixed results (e.g. Treisman, 1991; Duncan & Humphreys, 1992; and see below). In short, one cannot know if the poor performance found for a conjunction search is due to the conjunction per se or due to the set of stimuli being less discriminable. This question can be naturally addressed by the psychophysical methods described here because they place discriminability at the heart of the theory and measure performance at a constant threshold level of performance.

10.4.3. Relevant experiments

A number of experiments have pursued multidimensional search tasks and compared conjunction search to simpler search tasks (e.g. Cave & Wolfe, 1990; Duncan

& Humphreys, 1992; Wolfe, 1994), but there have been only a few that come close to the proposed simplified domain of accuracy measures, single fixation displays, and plausibly distinct and independent stimuli.

A small study of Aiken (1992) compared conjunction and feature search using ellipses of differing contrast and orientation. The conjunction target was a right-tilted, dark ellipse among vertical, dark and right-tilted, medium-gray ellipses. The feature target was a left-tilted, light ellipse that differed from the distractors in both orientation and contrast. Thus it can also be described as a disjunction target. Her experiment had two novel characteristics. First, she designed her two tasks so that when the target was absent the displays of distractors were identical for conjunction and disjunction conditions. This design provided a complete control for distractor heterogeneity effects that probably contribute to differences between feature and conjunction displays. Second, while she did not measure thresholds, she did adjust the discriminability of the two conditions to be roughly similar. Under these conditions, she found similar effects of set size for disjunction and conjunction search. These data were analyzed quantitatively by Eckstein et al. (in press). The predictions of their multidimensional signal detection theory were generally consistent with the observed set-size effects for the disjunction and conjunction tasks. In sum, these results were inconsistent with a qualitative difference between the two tasks as predicted by feature integration theory.

Eckstein (1998) has also conducted a comparison of feature and conjunction tasks using the discrimination of contrast and orientation. The displays were crowded compared to most of the displays reviewed here, but the effects of crowding were minimized by the use of a cueing procedure to manipulate set size (cf. Palmer, 1994). As expected, the overall performance was less for the conjunction condition than for the feature conditions. Measured by the difference in percent correct, set-size effects were larger for the less discriminable conjunction conditions. These results were well fit by the predictions based upon multidimensional signal detection theory. For comparison, a simple serial theory developed from Bergen and Julesz (1983) was also fit to the data. This theory predicted set-size effects that were much larger than those observed. Thus, even for conjunction search, the set-size effects were consistent with a theory based upon signal detection with no special mechanisms such as serial search or limited capacity.

The examples thus far suggest that the predicted larger set-size effects for conjunction search have not been found for the simplified domain. However, there is at least one instance of a task inspired by feature integration theory that does yield large set-size effects even under the simplified domain. The instance followed a study of O'Connell and Treisman (1990) who

examined search for targets defined by relations between separate objects. They and several authors (Steinman, 1987; Enns & Rensink, 1991; Logan, 1994; Moore, Marrara & Elsinger, 1998) have suggested that tasks that depend upon relations among objects may be processed differently than those that depend on attributes of a single object. The key to such an experiment is to assure that observers cannot recode the relational information between the objects as an attribute of a composite object. In other words, one must prevent grouping of the separate objects. One way to minimize the grouping is to use opposite contrast polarities: a white dot does not easily group with a black dot when seen on a gray background (Zucker & Davis, 1988). In Palmer (1994), set-size effects were measured for such pairs of opposite polarity dots. The task was to detect a tilted pair of dots among distractors of horizontally arranged pairs. Orientation thresholds were measured as a function of set size and the TvS slope was found to be twice as large as that for simple orientation tasks. Thus, this task produces a large set-size effect even within the simplified domain.

In summary, the methods and theories reviewed here for one-dimensional tasks can be generalized to multiple dimensions. The experiments reviewed suggest that this generalization may be adequate for combinations of simple dimensions such as contrast and orientation but probably fails for more complex judgments such as the relative position of distinct objects. This pattern of results is different from that expected for feature integration theory but probably does require some kind of two-stage search theory.

10.5. Generalization to response time

The generalization of our analysis from accuracy to response time is a serious challenge. The issues go well beyond visual search to the heart of theories of response accuracy and response time (cf. Luce, 1986). Fortunately, there has been progress developing appropriate paradigms and theories. Here we describe paradigms and theories that extend the analyses of accuracy paradigms presented in this article.

10.5.1. Empirical issues

To generalize the analysis described here to response time one needs corresponding paradigms. For an analog to the psychometric function, one can examine response time as a function of stimulus intensity or a similar variable such as contrast.

Response time is a very consistent function of stimulus intensity (see Luce 1986 for a review). Piéron (1914, 1920) was probably the first to suggest that the effect of stimulus intensity on response time could be characterized as a generalized power function,

$$y = a x^b + t_0 \quad (9)$$

where y is the response time, x is the stimulus variable, a and b are the power function parameters, and t_0 is an additive time constant. This equation can be reparameterized (Palmer, 1998) as,

$$y = c (x/d)^b + t_0 \quad (10)$$

where d is the response time threshold and c is the increase in response time from t_0 that defines threshold performance. Compared to Eq. (9), d is equal to $(c/a)^{1/b}$. This parameterization specifies a threshold parameter d in the units of the stimulus, an additive response time parameter t_0 in units of time, and a curve shape parameter b that is dimensionless. This function has been measured for simple response time of a variety of attributes (for a review see Luce, 1986). It has also been applied to choice response time (Schweickert, Dahn & McGuigan, 1988; Pins & Bonnet, 1996) and to visual search (Palmer, 1998). Thus, the generalized power function appears to be a promising candidate for the analog of the accuracy psychometric function.

The next step is to use the generalized power function (or similar functions) to estimate thresholds based on response time. In other words, what stimulus value results in a given response time. For light detection, Mansfield (1973) has measured such thresholds to obtain spectral sensitivity functions and measurements of the effect of eccentricity and flash duration. In visual search, Nagy and Sanchez (1990) have measured such thresholds for different color directions to compare color sensitivity in discrimination and matching experiments. More recently, Palmer (1998) has measured such thresholds as a function of set size for both response accuracy and response time. In summary, a few initial studies have measured the analog of psychometric functions and thresholds based upon response time. Such paradigms are necessary to address the control of discriminability effects in response time.

10.5.2. Theoretical issues

As with accuracy, the distinction between high and low threshold is critical for response time theories. Consider first theories with a high threshold. A high threshold greatly simplifies a theory of response time because the 'yes' response is only a function of the target and not the distractors. Thus, if bias, speed/accuracy and guessing are constant with set size, increasing the number of distractors cannot affect 'yes' responses. The predictions for 'no' responses depend on the details of the theory. In typical theories, one assumes a set of independent processes and a decision rule of responding 'no' when all processes indicate a 'no' response (e.g. parallel, unlimited capacity, self-terminating search, Townsend & Ashby, 1983). This results in a response time that is the maximum of the component response

times. Thus, an effect of set size is predicted for correct rejections but not hits. In summary, for the appropriate conditions, high threshold theories predict no effect of set size on the 'yes' response time as well as no effect on accuracy.

Many response time theories do not make explicit whether they assume a high or a low threshold. To determine if a theory implicitly assumes a low threshold, one must determine if false alarms can arise from a distractor. Theories that ignore errors (e.g. Treisman & Gelade, 1980) or that assume all false alarms are due to guessing (e.g. Wolfe, 1994)⁵ must be assuming a high threshold because any low threshold theory would predict more errors with larger set sizes. Such implicit high threshold theories are common. For example, the theories forming the parallel-serial equivalence relations assume a high threshold (Townsend, 1974; Townsend & Ashby, 1983; Townsend, 1990). One way to generalize these theories is to develop a separate model of error that goes along with the theory of response time (e.g. Schweickert, 1985, 1989). But this alternative is limited for most cases because it does not include a mechanism for the speed/accuracy tradeoff. One needs an integrated theory of accuracy and response time.

Now consider low threshold theories. Here the story is more complex because the response depends on all of the stimuli. Consequently, the distractors affect the 'yes' as well as the 'no' response. As the number of stimuli increases, one of two things must happen. False alarms may increase due to the additional distractors providing misleading evidence. Or, the observers may adjust their bias and speed/accuracy criteria to maintain a constant error rate, in which case the response time must increase as a function of set size.

This general idea has been pointed out by several authors (Lappin, 1978; Schweickert, 1985; Broadbent, 1987; Ward & McClelland, 1989; Zenger & Fahle, 1997) and has been illustrated in a specific theory by Palmer and McLean (1995). The authors applied the diffusion model (Feller, 1957) to response time as applied to (memory) search by Ratcliff (1978). The diffusion model is a continuous version of the random walk (see Luce 1986). In this theory, each stimulus is assumed to be analyzed by an independent and parallel diffusion process and the rate of accumulating information is assumed to be linearly related to the target-distractor difference. The bias parameters are fixed to maintain a constant error rate and to accumulate infor-

⁵ Wolfe's guided search model is a hybrid with respect to the high versus low threshold distinction. The first stage of guidance has a low threshold and is explicitly based on signal detection theory. However, the second stage of deciding if a particular stimulus is a target is explicitly a high threshold theory. False alarms arise only from guessing.

mation symmetrically for ‘yes’ and ‘no’ responses. The resulting theory has only two parameters for all set size and discriminability conditions. One is a sensitivity parameter and the other is the constant response time parameter that describes the contribution of the processing times that are independent of the target-distractor difference.

This two-parameter theory captures all of the qualitative properties of Palmer’s (1998) search time experiments. For example, this theory predicts a reasonable value for the curve shape parameter of the RT versus stimulus difference plot described above. Most importantly, even though the component processing is independent, the theory can predict arbitrarily large set-size effects on response time. The less discriminable the stimuli, the larger the predicted set-size effects. Furthermore, the theory predicts a threshold-versus-set-size function with a TvS slope of 0.26 which is similar to that observed for the corresponding experiment for short threshold criteria such as 100 ms. The theory also has several shortcomings: for example, it always predicts response time versus set size functions that are concave downward whereas the observed functions are often close to linear.

It appears likely that some low threshold theory such as just described can account for the results found with typical search response time experiments. Properly constructed, these theories can be relatively simple extensions of the signal detection theories of search accuracy. They stand in contrast to the implicitly high threshold theories found in much of the search literature. Given the failure of the high threshold theory for detection and search accuracy, it seems unlikely that it will be useful for search response time.

10.6. Generalization to multiple fixations

The studies reviewed thus far have been restricted to single eye fixations. This is a special case of the larger search problem in which an observer can move his or her eyes and perhaps head and body as well (for reviews see Kowler, 1990; Viviani, 1990). The rationale for focusing on the single fixation special case is that it greatly simplifies the stimulus at the eye. To generalize this analysis, we describe two additional phenomena and present the visual lobe theories of multiple fixation visual search. Then we sketch a theory that combines the maximum-of-outputs and the visual lobe theories.

10.6.1. Two additional phenomena

Search with multiple eye fixations introduces at least two additional phenomena. The first is variability in eccentricity of the stimulus. Eccentricity can be controlled with brief displays allowing only a single fixation display but cannot be as easily controlled with longer displays allowing multiple fixations. The impact

of eccentricity effects can sometimes be large. The most detailed work is with simple detection tasks but there are a few studies using search tasks themselves (e.g. Carrasco & Frieder, 1997). At one extreme are detection tasks that usually show large effects of eccentricity. For example, the acuity threshold for a fine grating can double with eccentricities of just 1° and increase ten fold by 10°. Such large effects have been reported for a variety of tasks (e.g. Yager & Davis, 1987). For such detection tasks, appropriate eye movements can sharply improve performance. In contrast, a second group of tasks is not limited by eccentricity over at least the central 10° or so. These tasks are typically suprathreshold discriminations such as the contrast increment discrimination task considered earlier in this article. This task appears to be limited by contrast phenomena (i.e. Weber’s law for contrast) rather than eccentricity (Legge & Kersten, 1987). For such stimuli, the effects of eccentricities of 10° can be less than a 10% increase in threshold. In summary, many but not all tasks are limited by eccentricity effects.

The second phenomenon is the pattern of eye fixations in search tasks. Unless efforts are taken to prevent eye movements, they almost always occur in search tasks. For easy discriminations, the first fixation is often on the target (Eckstein, Beutter & Stone, submitted; Findlay, 1997). For more difficult discriminations, individual stimuli are sometimes fixated (e.g. Motter & Belky, 1998b), but other times groups of relevant stimuli are fixated (He & Kowler, 1989, 1991; Zelinsky, Rao, Hayhoe & Ballard, 1997). Thus, the pattern of eye fixations shows clear signs of a systematic and purposeful behavior that is likely to contribute to search performance.

10.6.2. Visual lobe theories

Many theories of multiple fixation search have been called *visual lobe theories* (for review see Overington 1976; for examples see Engel, 1977; Geisler & Chou, 1995). These theories combine the constraint of eccentricity effects with the freedom to move one’s eyes to overcome this constraint. They begin with the assumption that the probability of detecting a target depends largely on the eccentricity of the target. Hence, for each fixation, there is a region of visual field, called a *visual lobe*, within which a target is likely to be detected. Search proceeds by moving one’s eyes to different locations in the visual field. Detecting the target in any one fixation is assumed to be independent of prior fixations and thus depends on only the current fixation. These theories make several predictions that have had some support. First, for random search patterns, the probability of detection as a function of time is $1 - e^{-at}$ (Koopman, 1956; Williams, 1966b). Second, the size of the visual lobe measured with a single fixation is negatively correlated with the search time (Engel, 1977;

Bloomfield, 1979; Geisler & Chou, 1995). The bigger the visual lobe, the fewer fixations are needed to ‘cover’ the relevant part of the visual field. This theory has interesting variations such as Motter and Belky’s (1998a) argument that under their conditions it is the stimulus density rather than the eccentricity that determines the size of the visual lobe.

The visual lobe theories become more complex when more specific information about eye movements is considered. First, there is good evidence (e.g. Rayner, McConkie & Ehrlich, 1978; Rayner & Fisher, 1987; Irwin, 1991) that memory for information from previous fixations is sharply limited. This is consistent with the fixation by fixation account of the visual lobe models. But, it need not mean that the processing of stimulus information is concurrent with the fixation. Hooze and Erkelens (1996) make the case that processing continues for some time following the beginning of the next fixation. Processing may ultimately be interrupted by the new saccade but only after some lag in time. Second, there is good evidence that peripheral information can guide eye movements at least when the stimulus information is very discriminable (Williams, 1966a; Luria & Strauss, 1975; Hooze & Erkelens, 1999). But there is also evidence that eye movements are programmed ahead of time (Zingale & Kowler, 1987). The latter suggests that the pattern of fixations is constrained over a series of fixations rather than being decided fixation by fixation as implied by a simple version of the visual lobe theories. Thus, describing the accumulation of stimulus information at the molar level may be different than describing the control of eye movements at the molecular level.

10.6.3. A sequential independent decisions theory

The ideas from the visual lobe theories can be combined with the signal detection framework given here. Indeed, Geisler & Chou (1995) explicitly incorporate many ideas from signal detection theory into their version of a visual lobe theory. Here, we sketch a theory that combines the visual lobe concept with the independent decision variant of the maximum-of-outputs theory (Shaw, 1980). The resulting hybrid theory illustrates how one can generalize the theories presented here from a single fixation to multiple fixations.

This theory follows closely that of Geisler & Chou (1995) but applied to discrete, clearly visible stimuli rather than a target texture patch on a surrounding texture. The size of the visual lobe is defined in terms of the number of stimuli that are analyzed in parallel during each fixation. Each stimulus is assumed to have an independent perceptual representation that is compared to a criterion for an independent decision. The quality of the representation depends on all the factors determining the discriminability of the stimulus including eccentricity. If any of the individual decisions are

positive, then the observer responds ‘yes’. The next eye movement is directed to a new, possibly overlapping visual lobe. Overlap is allowed so that stimuli can be processed repeatedly to gain the advantage of ‘multiple looks’. The choice of the location of the next fixation is probably a function of the information available in the periphery and the observer’s strategy. To simplify the possibilities, we follow Geisler & Chou (1995) among others in assuming that observers can be instructed to ‘tile’ the display systematically with their visual lobes until all stimuli are included within a lobe for at least one fixation. If all of the stimuli have been scanned without finding a target, then the observer responds ‘no’. For m stimuli, this model can allow either parallel processing of up to m stimuli or serial processing of up to m fixations. The size of the visual lobe is used to compensate for eccentricity effects and the degree of overlap is used to compensate for overall low discriminability (even in the fovea).

In summary, multiple fixation search highlights the importance of limited peripheral vision and the sequential contribution of eye movements. These effects can be described by combining ideas from signal detection applied to single fixation search and visual lobe theories developed for multiple fixation search.

10.7. More general stimulus conditions

The experiments reviewed here are limited to idealized stimulus conditions typical of simple detection and discrimination experiments. These conditions include: distinct stimuli well above detection threshold, widely separated stimuli, and single displays with no mask. The intent of this simplified domain is to minimize stimulus specific phenomena such as crowding that might be confounded with set size or other manipulations of interest. Consider efforts to generalize this analysis to situations including each of these conditions.

10.7.1. Indistinct stimuli

Laboratory search experiments typically use stimuli that are distinct from the surround. This allows one to manipulate the number of distractors unambiguously. If the stimuli are indistinct such as they would be if they were near contrast threshold, then the spatial uncertainty of the target becomes critical rather than the number of distractors. In effect, the observer’s task changes from discriminating targets from distractors to detecting the target against the background. There is still uncertainty, but it becomes a function of the possible target locations rather than the number of distractors. Such spatial uncertainty has been much studied in its own right (e.g. Davis, Kramer and Graham, 1983).

How can one generalize the theories considered here with distinct stimuli to more realistic cases involving stimuli that may or may not be distinct? Such situations

occur in naturalistic settings where targets must often be distinguished from a structured background as well as from specific distractors. Fortunately, this problem has been addressed by two recent papers for the case of discriminating increments (or decrements) in contrast for simple stimuli such as disks, Gaussian blobs, or Gabor patches (Solomon, Lavie & Morgan, 1997; Foley & Schwarz, 1998). For these stimuli, the threshold contrast increment (or decrement) was measured as a function of the contrast of the standard stimulus, often called the pedestal contrast. For zero or near zero pedestal contrast, the task reduced to simple detection and for large pedestal contrast, the task was to discriminate changes in a stimulus that is quite distinct from the surround. In Foley and Schwarz (1998), they manipulated the pedestal contrast, set size, and spatial uncertainty in a factorial design. Within this large experiment, they replicated the expected set-size effects such as reported here for the high contrast pedestals. In contrast, set size had no effect for low contrast pedestals. Similarly, they replicated the expected spatial uncertainty effect for low contrast pedestals but found no effect of spatial uncertainty for high contrast pedestals. Thus, this experiment spanned the range of conditions that yield set size and spatial uncertainty effects. Both Foley and Schwarz (1998) and Solomon et al. (1997) independently constructed a three part theory to account for these effects. First, they assumed a particular representation of contrast (nonlinear transduction to account for the dipper effect). Second, they selected the inputs to the decision process on the basis of both the uncertainty manipulation and set size. Only information from relevant locations or clearly visible stimuli are considered in the decision process. Third, they chose the maximum-of-outputs rule to determine the response. This theory accounted well for the effects of the factorial combination of pedestal contrast, set size and spatial uncertainty. The key feature was to select only the appropriate information for the decision process. This selection predicted the particular interaction between these three effects. In summary, these two papers successfully generalized the kind of theory described here to a case with mixed distinct and indistinct stimuli.

10.7.2. Crowding effects

Verghese and Nakayama (1994) conducted a detailed study of visual search manipulating the set size and discriminability of three attributes: orientation, color and spatial frequency. Two aspects of their experiments were different from the simplified domain described here. They used a mask and the spacing between stimuli was approximately two times the size of the individual stimuli rather than five times as typical of the studies reviewed here. Verghese and Nakayama found that the effects of set size on threshold were different for the

three different attributes. In terms of difference thresholds, color showed very small set-size effects, spatial frequency intermediate effects and orientation the largest effects. In a threshold-versus-set-size graph, the three attributes had TvS slopes of 0.1–0.6 depending on conditions. This is a larger range of effects than the slopes of 0.2–0.3 typically reported in experiments within the simplified domain.

Two control experiments in Verghese and Nakayama (1994) imply that these results may arise from crowding. For their orientation task, they compared performance between the usual conditions in which the display set size varied and a relevant set size manipulation in which the same number of stimuli were always displayed but a subset of them were cued as relevant for each block of trials. The relevant-set-size manipulation produced smaller set-size effects (TvS slope = 0.42) compared to the display-set-size manipulation (TvS slope = 0.65). In a second control using the color task, they manipulated the space between stimuli directly as well as the presence of irrelevant stimuli. For less discriminable colors they found effects of both decreasing the spacing between the stimuli and adding irrelevant distractors. Thus, both control experiments indicated that the spacing of the stimuli interacts with set-size effects. It is likely that crowding resulted in stimulus specific phenomena that are inconsistent with the simple theories described here.

In the next few paragraphs, we speculate on alternative accounts of the crowding effects. Perhaps the first possibility to consider is *lateral masking* and similar spatial interactions (e.g. Eriksen, 1980; Breitmeyer, 1984; Chubb, Sperling & Solomon, 1989; Cannon & Fullenkamp, 1991; Spillmann & Werner, 1996; Zelinsky, 1999). While some researchers suggest a relatively narrow spatial range with these masking effects, there are a variety of situations where spatial interactions operate at larger distances, particularly in the periphery. For example, Bouma (1970), Andriessen & Bouma (1976), Kröse & Burbeck (1989) suggests that to eliminate masking effects, one must have stimulus separations of approximately half the eccentricity of the stimulus. To predict the masking due to crowded stimuli, one can explicitly model the relevant channels in early vision (e.g. Verghese, Watamaniuk, McKee & Grzywacz, 1999).

Another possibility is that perceptual grouping creates different relevant representations at different set sizes and stimulus separations. For wide separations, stimuli may be represented individually; but for small separations, stimuli may be represented in groups that integrate information from the individual stimuli. By this *group scanning* hypothesis, the number of groups determines the set-size effect rather than the number of individual stimuli (Williams, 1966b; Treisman, 1982; Bundesen & Pedersen, 1983; Egeth et al., 1984; Pashler,

1987a; Humphreys, Quinlan & Riddoch, 1989; Macquistan, 1994). In the extreme, the entire field may be represented as a single texture and a target detected as a deviation in an otherwise homogeneous texture.

Texture may enter into the process in another way as well. The above hypothesis emphasized the grouping of similar stimuli into a single texture. Alternatively, differences between nearby stimuli might introduce *texture gradients* (Nothdurft, 1985, 1991, 1992, 1993; Sagi & Julesz, 1987). By this hypothesis, it is the contrast between nearby stimuli that is critical rather than the stimuli themselves. A simple version of this idea is embodied in the maximum-of-differences theory. In summary, there are a number of possible accounts of crowding effects in visual search.

10.7.3. Masking effects

Morgan, Ward and Castet (1998) investigated the effect of having a mask in addition to crowding. They compared a masked, long duration display to a masked, brief display. With the long duration display and mask, they found set-size effects of the same magnitude as the studies reviewed here without a mask. With short duration display and mask, the set-size effects were several fold larger. This result was also found for a cueing paradigm in which the relevant set size was manipulated and the display set size was held constant. Thus, this study suggested that temporal masking as well as spatial crowding interacts with set-size effects.

What kind of mechanism is responsible for these effects of masking? One possibility is that masking makes information from individual stimuli unavailable and instead the observer must depend on information from perceptual groups. Morgan et al. (1998) went on to conduct an experiment to test this idea by perturbing the distractors in the short duration, mask condition. They showed that nearby distractors affect the perception of the target as one might expect if the individual stimuli could not be processed separately. Of course, there are other possibilities such as the mask interrupting serial processing. Distinguishing among these possibilities awaits further measurements with and without masks.

10.7.4. Summary

Relative to the simplified domain of this article, new phenomena are introduced by indistinct stimuli, crowded stimuli, or masks. For the case of indistinct stimuli, one must account for the effects of spatial uncertainty. Foley and Schwarz (1998) and Solomon et al. (1997) describe simple extensions of the current theory that successfully account for these effects and their interactions. For crowding and masking effects, the evidence of new phenomena is clear, but there remain a variety of possible accounts of these effects.

Indeed, generalizing the theory to these cases may prove difficult if the new phenomena are specific to the particular stimulus.

11. Conclusions

The departure point for this paper was that it would be profitable to study visual search using laboratory idealizations that make visual search as similar as possible to the simple detection and discrimination tasks long studied in spatial vision. With this approach, we reviewed the quantitative results from six visual search paradigms.

We next compared high threshold theory to low threshold theories based upon signal detection theory. In almost all testable cases, the high threshold theory failed to describe the results observed in the six paradigms. In contrast, the three versions of low threshold theory succeeded with few exceptions. It seems likely that a successful theory of visual search must be based on a low threshold. This argument is perhaps even more important for the domain of response time search where implicit high threshold theories are common.

In the discussion, we explored generalizations of the simple signal detection theories that might account for a wider range of phenomena such as search asymmetry, conjunction search, response time measures, multiple eye fixations, and less ideal stimulus conditions. These generalizations sketched a comprehensive theory of visual search that is based on the same psychophysical concepts that have been used to describe simple detection and discrimination in spatial vision.

12. Summary of notation

a	intercept parameter of the linear models used in data analysis
b	slope parameter of the linear models used in data analysis
c	criterion for decision to respond yes
d'	discriminability measure defined for equal variance distributions
d'_e	discriminability measure defined at equal bias
d_a	discriminability measure defined for RMS weighting of variance
e	high-noise efficiency of external noise models
f	probability density function
F	cumulative distribution function
F_{normal}	cumulative function of the normal distribution
f_U	joint probability density function for the random variables of U

g	guessing parameter used with high threshold theory
$G(x)$	general function used with high threshold theory
h	maximum number of stimuli that can be a target on a single trial
I	index for individual stimuli within a trial
I_h	set of indexes specifying which stimuli are the h targets
j	secondary stimulus index
k	index or miscellaneous constant
$L(X)$	likelihood ratio of an event X for hypotheses s relative to n
$L_h(X)$	likelihood ratio for the case of exactly h targets
m	number of stimuli (targets and distractors)
n	noise trial (all stimuli are distractors)
$P(x)$	probability of x
P_h	set of probabilities of the number of targets being exactly $1, 2, \dots, h$
q	a priori probability of a target stimulus
Q	set of a priori probabilities of a target stimulus
R	relative coding function
s	signal trial (at least one stimulus is a target)
t	threshold stimulus value
u	internal representation of a stimulus on a single trial
U	vector of random variables corresponding to a representation
v	physical stimulus value
V	vector of physical stimulus values
x	a value specifying the domain of some functions
X	a vector specifying the domain of some functions
y	range of some functions
z	z -transform (inverse of the Normal cumulative function)
σ	standard deviation of a random variable
μ	mean of a random variable
Δm	difference in the means of two random variables

Acknowledgements

Some of this article was first presented in Pavel (1990) and Pavel et al. (1992). We thank Jennifer McLean for helping with many of the ideas in this article, particularly in her close reading of Appendix A and her work on generalizing the distributional assumptions. We also thank Davida Teller for many useful discussions and suggestions. In addition, we thank Beth Davis, Barbara Doshier, Miguel Eckstein, Bill Geisler, Eileen Kowler, Zhong-Lin Lu, Michael

Morgan, Hal Pashler, George Sperling, Jim Townsend and Greg Zelinsky for criticisms and comments on earlier versions of the manuscript. This work was supported in part by NASA-Ames grant NAG 2-1202 and RTOP 199-16-12-37 to PV.

Appendix A

This appendix presents descriptions of the four theories discussed in the article. First, a detailed notation is introduced that is sufficient for all four theories. Then each theory is defined and its predictions derived. Specifically, ROC curves are derived as a function of set size and the number of targets. These predictions are for arbitrary distributions. With more specific assumptions of the distribution and the relative variability of the target and distractor distributions, one can calculate all of the results discussed in this article.

A.1. Background

A.1.1. Notation

The notation has four parts. The first part defines the task. A yes–no search task has two kinds of trials. *Signal trials* contain at least one target stimulus and are denoted s . *Noise trials* contain only distractor stimuli and are denoted n . On each trial, observers respond either ‘yes’ or ‘no’ yielding a two-by-two contingency table (Table 3) of stimulus–response pairings: hit, false alarm, miss, and correct rejection. In a typical experiment, the specification of signal versus noise is under experimental control. Thus, the responses can be characterized by two conditional probabilities, one for each kind of stimulus: the probability of a hit, $P(\text{‘yes’}|s)$, and the probability of a false alarm, $P(\text{‘yes’}|n)$.

The second part of the notation defines the stimuli in more detail. On any trial, m stimuli are presented at different locations. On a noise trial, all m stimuli are distractors; on a signal trial, at least one and as many as h stimuli can be targets with the remaining stimuli being distractors. The m stimuli presented on a noise trial are denoted by $V_n = \{v_{n1}, v_{n2}, \dots, v_{nm}\}$ with the i th stimulus denoted by v_{ni} . If a signal is included as the i th stimulus, then the trial is a signal trial and the m stimuli are denoted by $V_s = \{v_{n1}, v_{n2}, \dots, v_{si}, \dots, v_{nm}\}$. More generally, h targets can be included among the distractors. For example, if $h = 2$ and the 1st and the i th stimuli are targets, they are denoted by $V_s = \{v_{s1}, v_{n2}, \dots, v_{si}, \dots, v_{nm}\}$.

The third part of the notation defines the internal representations that are assumed to correspond to each stimulus. Specifically, each stimulus results in an internal representation that corresponds to a random variable. The random variables corresponding to noise stimuli V_n , are denoted by $U_n = \{u_{n1}, u_{n2}, \dots, u_{nm}\}$. The

Table 3
The two-by-two contingency table for the yes–no task

Stimulus	Response	
	Yes	No
Signal (s)	Hit	Miss
Noise (n)	False alarm	Correct rejection

random variables corresponding to signal stimuli V_s , with the i th stimulus a target, are denoted $U_s = \{u_{n1}, u_{n2}, \dots, u_{si}, \dots, u_{nm}\}$. As in the previous example, if $h = 2$ and the 1st and the i th stimuli are targets, then the representation is denoted by $U_s = \{u_{s1}, u_{n2}, \dots, u_{si}, \dots, u_{nm}\}$. When quantitative stimulus attributes are used, v_{ni} and v_{si} correspond to specific physical values (e.g. contrast) while u_{ni} and u_{si} correspond to internal representations of those values (e.g. signal-to-noise ratio or subjective contrast). Sometimes it is convenient to drop the n and s subscripts for noise and signal to focus on the observer's internal representation $U = \{u_1, u_2, \dots, u_m\}$. The need for this modified notation arises because observers in the yes–no task do not know ahead of time that a given trial is a noise or a signal trial (even though the experimenter does). Consequently, to model their decision process, one must describe a representation that can be from either a noise or signal trial.

The fourth part of the notation specifies the density and distribution functions for the random variables that correspond to the internal representations. For each random variable u_{ni} or u_{si} , the probability density function is denoted by f_{ni} or f_{si} , respectively, and the cumulative distribution function is denoted by F_{ni} or F_{si} , respectively. The joint probability density of all m random variables is denoted by f_{Un} for noise trials and f_{Us} for signal trials. When all m noise representations, u_{ni} , are identically and independently distributed, then the density is denoted by f_n which is equal to f_{ni} for all i . Similarly, when all h signal stimuli are identically and independently distributed, then the density of those h representations is denoted by f_s . In the example above with $h = 2$ and the 1st and the i th stimuli are independently and identically distributed targets, then $f_{s1} = f_{si} = f_s$.

A.1.2. Signal detection theory for a single stimulus

The next part of the background is a brief description of how signal detection theory (Green & Swets, 1966) accounts for performance in the yes–no task for a single stimulus ($m = 1$). A 'yes' response is made whenever a unidimensional internal representation, u_{s1} or u_{n1} , exceeds a decision criterion, c . The probability of a false alarm and a hit are, respectively,

$$P(\text{'yes'}|n) = P(u_{n1} > c), \text{ and } P(\text{'yes'}|s) = P(u_{s1} > c). \quad (\text{A1})$$

By the definition of the cumulative distribution function, this can be rewritten as:

$$P(\text{'yes'}|n) = 1 - F_{n1}(c), \text{ and } P(\text{'yes'}|s) = 1 - F_{s1}(c). \quad (\text{A2})$$

These equations predict how the false alarms and hit probabilities depend upon the criterion. The resulting relation between the false alarm and hit probabilities is known as the Receiver Operating Characteristic (ROC) function. Eq. (A2) predicts the form of the ROC function.

A.1.3. Independence properties

The last part of the background is to define the independence properties that are to be assumed in the theories presented here. By *statistical independence*, we mean that for the representation of either noise or signal trials, all pairs of distinct random variables are independent. Thus, for a noise trial, the random variables u_{ni} and u_{nj} , $i \neq j$, must be independent. Similarly for the representations of a signal trial. Further discussions of this independence property can be found in Pelli (1985), Graham (1989), Ashby (1992). Generalizations described by these authors can be applied to the theories presented here, but are beyond the scope of this article. We consider statistical independence as an idealization appropriate for widely separated stimuli.

A second kind of independence property is commonly referred to as unlimited capacity (Townsend, 1974). By *unlimited capacity*, we mean that the representations are independent of the number of stimuli, m . Thus, the distribution of u_{ni} or u_{si} for all i is unchanged with a change in the number of stimuli. This property can be further broken down to distinguish between unlimited capacity for additional distractors and for additional targets such as present in multiple target experiments. Duncan (1980) has suggested that for some tasks, unlimited capacity may hold for additional distractors but not for targets. In the development here, we assume all of the independence properties.

A.2. The ideal observer theory

The ideal observer theory is next described in several steps. We begin by stating the general definition of the ideal observer, its application to a single stimulus task, and its application to a task with one target among m stimuli. Then we present generalizations of this result to a task with exactly h targets among m stimuli and to a task with h targets or less among m stimuli.

A.2.1. General definition

Several authors have shown that the optimal decision rule depends on the likelihood ratio (Peterson, Birdsall & Fox, 1954; Green & Swets, 1966). The likelihood ratio is defined as the ratio of the conditional probability of the representation given the signal relative to the conditional probability of the representation given noise,

$$L(X) = \frac{P(U = X|s)}{P(U = X|n)} = \frac{f_{Us}(X)}{f_{Un}(X)}. \tag{A3}$$

It has been shown that adopting the likelihood ratio as the decision variable yields the optimal decisions in terms of probability correct as well as other criteria. For a yes–no task, this is done by adapting Eq. (A1) to use the likelihood ratio rather than the internal representation itself. Thus, the probability of a false alarm and a hit are, respectively,

$$P(\text{'yes'}|n) = P(L(U_n) > c), \text{ and} \\ P(\text{'yes'}|s) = P(L(U_s) > c). \tag{A4}$$

In words, the probability of responding ‘yes’ given a noise trial is equal to the probability that the likelihood ratio on a noise trial is greater than c . Similarly, the probability of responding ‘yes’ given a signal trial is equal to the probability that the likelihood ratio on a signal trial is greater than c .

A.2.2. Single stimuli

To illustrate the ideal observer theory yet more concretely, consider the case of a single stimulus representation $U = \{u_1\}$. The likelihood ratio of Eq. (A3) reduces to:

$$L(x) = \frac{f_{s1}(x)}{f_{n1}(x)}. \tag{A5}$$

The probability of a false alarm and a hit simplify to, $P(\text{'yes'}|n) = P(L(u_{n1}) > c)$, and

$$P(\text{'yes'}|s) = P(L(u_{s1}) > c). \tag{A6}$$

Combining Eqs. (A5) and (A6) yields,

$$P(\text{'yes'}|n) = P\left(\frac{f_{s1}(u_{n1})}{f_{n1}(u_{n1})} > c\right), \text{ and} \\ P(\text{'yes'}|s) = P\left(\frac{f_{s1}(u_{s1})}{f_{n1}(u_{s1})} > c\right). \tag{A7}$$

Written out in this fashion, one can see that the argument for each density function is determined by the trial type, noise or signal. In contrast, the density functions themselves are determined by the definition of the likelihood ratio.

A.2.3. Multiple stimuli: 1 of m

Next consider the case in which one target is pre-

sented among m stimuli in a signal trial (Peterson et al., 1954; Green & Birdsall, 1978). For noise trials, $P(U = X|n)$ is the joint probability of the values $X = \{x_1, x_2, \dots, x_m\}$,

$$P(U = X|n) = f_{Un}(x_1, x_2, \dots, x_m). \tag{A8}$$

Assuming independence, this joint probability is the product of the component densities,

$$P(U = X|n) = \prod_{j=i}^m f_{nj}(x_j). \tag{A9}$$

For signal trials, $P(U = X|s)$ is also a joint probability of the values $\{x_1, x_2, \dots, x_m\}$,

$$P(U = X|s) = f_{Us}(x_1, x_2, \dots, x_m). \tag{A10}$$

This joint probability can be conditionalized upon the location of the target and then, assuming independence, can be written as the product of the component densities,

$$P(U = X|s) = \prod_{j=i}^m q_i f_{si}(x_i) \prod_{j \neq i}^m f_{nj}(x_j), \tag{A11}$$

where q_i is the a priori probability of the i th stimulus being the target. If all stimuli have equal probability of being the target, then q_i equals $1/m$. Combining Eqs A3, A9, and A11 and simplifying yields,

$$L(X) = \sum_{j=i}^m q_i \frac{f_{si}(x_i)}{f_{ni}(x_i)}. \tag{A12}$$

In words, the likelihood is the weighted sum of the local likelihoods for each stimulus. The likelihood can be used with Eq. (A4) to predict the false alarm and hit probabilities for this case of 1 target of m independent stimuli.

A.2.4. Multiple stimuli: exactly h of m

Consider next the case in which exactly h targets are presented among m stimuli. As before we assume statistical independence. By definition, $P(U = X|n)$ is the joint probability as before,

$$P(U = X|n) = f_{Un}(x_1, x_2, \dots, x_m). \tag{A13}$$

Assuming independence, this joint probability is the product of the component densities identical to Eq. (A9),

$$P(U = X|n) = \prod_{j=i}^m f_{nj}(x_j). \tag{A14}$$

By definition, $P(U = X|s)$ is the joint probability of the values $\{x_1, x_2, \dots, x_m\}$,

$$P(U = X|s) = f_{Us}(x_1, x_2, \dots, x_m). \tag{A15}$$

In this case, $P(U = X|s)$ is more complicated because the h targets can occur at many combinations of locations. The joint probability is the weighted sum of the joint probability for each of the possible combinations

of target locations. These locations can be specified by $I_h = \{i_1, i_2, \dots, i_h\}$. For example, if there were three targets and the targets were at locations 1, i , and j , then $I_3 = \{1, i, j\}$. All of the possible locations can be specified by the following nested sum,

$$P(U = X|s) = \sum_{i_1 \neq 1}^m \sum_{i_2 \neq i_1}^m \sum_{i_3 \neq i_1 i_2}^m \dots \sum_{i_h \neq i_1 i_2 \dots i_{h-1}}^m q_{I_h} P(U = X'|s) \quad (\text{A16})$$

where the vector X' contains the values x_{si} as specified in I_h and x_{ni} . The joint probability $P(U = X'|s)$ is the joint density as usual,

$$P(U = X'|s, I_h) = f_{U_s}(X'). \quad (\text{A17})$$

Assuming independence, this becomes the product,

$$P(U = X'|s, I_h) = \prod_{j=1}^h f_{s_j}(x_{i_j}) \prod_{k \neq i_1 i_2 \dots i_h}^m f_{n_k}(x_k). \quad (\text{A18})$$

Eqs. (A11–A13) can be combined into the likelihood ratio as defined by Eq. (A3). Simplifying yields the final result

$$L_h(X) = \sum_{i_1=1}^m \sum_{i_2 \neq i_1}^m \sum_{i_3 \neq i_1 i_2}^m \dots \sum_{i_h \neq i_1 i_2 \dots i_{h-1}}^m q_{I_h} \prod_{j=1}^h \frac{f_{s_j}(x_{i_j})}{f_{n_j}(x_{i_j})}. \quad (\text{A19})$$

As with the 1 of m case, the likelihood ratio simplifies to be a sum of specific likelihood ratios where ‘specific’ refers to one for each possible combination of h target locations. This analysis can be easily extended to cases with mixtures of 1 to h targets in any trial. In summary, the ideal observer theory depends critically on the exact set of possible signals.

A.3. The maximum-of-outputs theory

In the maximum-of-outputs theory, the decision is based on the representation that most favors the target. A ‘yes’ response is made if the maximum of the component representations favors a ‘yes’ response. This theory has been considered by many investigators in several closely related forms. We follow Graham (1989) in using the name, *maximum-of-outputs theory*. The origin of the maximum rule is probably in applications to forced-choice tasks where response is determined by the representation with the maximum value (e.g. Green & Swets, 1966). In our treatment of this maximum-of-outputs theory, we assume that the sign of all internal representations is positive with increasing evidence for a target. A ‘yes’ response is made when the evidence with the maximum value exceeds a criterion

$$\max(U) = \max(u_1, u_2, \dots, u_m). \quad (\text{A20})$$

From this, the probability of a false alarm and a hit are, respectively,

$$P(\text{‘yes’}|n) = P(\text{Max}(U_n) > c), \text{ and}$$

$$P(\text{‘yes’}|s) = P(\text{Max}(U_s) > c). \quad (\text{A21})$$

Assuming the m random variables are identically distributed and statistically independent, the probability of the maximum of m identical random variables not exceeding u is $F(u)^m$. Thus the probability of a false alarm can be rewritten as,

$$P(\text{‘yes’}|n) = 1 - F_n(u)^m, \quad (\text{A22})$$

and by distinguishing between the h target distributions and the $m - h$ distractor distributions, the probability of a hit can be written as

$$P(\text{‘yes’}|s) = 1 - F_s(u)^h F_n(u)^{m-h}. \quad (\text{A23})$$

It is helpful to also describe the independent decisions theory on its own (Shaw, 1980). In the independent decisions theory, the decision rule is to respond ‘yes’ if,

$$u_1 > c, u_2 > c, \dots, \text{ or } u_m > c. \quad (\text{A24})$$

Assuming statistical independence, the probability of a ‘yes’ response is

$$A(U) = 1 - \prod_{i=1}^m [1 - P(u_i > c)]. \quad (\text{A25})$$

The probability of a false alarm and a hit are, respectively,

$$P(\text{‘yes’}|n) = A(U_n), \text{ and } P(\text{‘yes’}|s) = A(U_s). \quad (\text{A26})$$

Assuming identical distributions for all stimuli, one can replace $P(u_{ni} > c)$ with $1 - F_n(u)$. Simplifying yields a false alarm probability of

$$P(\text{‘yes’}|n) = 1 - F_n(u)^m. \quad (\text{A27})$$

Similarly, considering the h targets separately from the $m - h$ distractors yields the hit probability of

$$P(\text{‘yes’}|s) = 1 - F_s(u)^h F_n(u)^{m-h}. \quad (\text{A28})$$

Thus, this version of the maximum-of-outputs theory makes an identical prediction to the independent decisions theory.

To show that the maximum-of-outputs theory is equivalent to the independent decisions theory, consider first a case in which the independent decisions theory holds. If one of the individual representations exceeds the threshold c , then the maximum of all must also exceed c . If none of the individual representations exceeds c , then the maximum cannot either. Thus, the performance predicted by the independent decisions theory is matched by the maximum-of-outputs theory. Similarly, if one assumes the maximum-of-outputs the-

ory, one can show that equivalent performance is predicted by the independent decisions theory. Thus, they are equivalent under the conditions defined here.

A.4. Relative coding models

The third theory to be considered is an instance of a class of models we refer to as *relative coding models*. In these models, each stimulus is always coded relative to its context before contributing to a decision. In general, this idea can be represented by the function $R(u_i, U)$, which depends on both the i th representation and the vector of all representations, U . This function is combined with the maximum combination rule of the independent decisions theory to yield the decision rule:

$$\text{Max}[R(u_1, U), R(u_2, U), \dots, R(u_m, U)]. \tag{A29}$$

Here, we develop a special case referred to as the *maximum-of-differences theory*. This theory is presented as an idealization of more complex theories that can incorporate factors such as the spacing between items in determining which stimuli contribute to the context. For the maximum-of-differences theory, each stimulus is compared to a randomly selected other stimulus. The functions R_d and A_{MaxRd} are defined as,

$$R_d(u_i, U) = u_i - u_j, \text{ and} \\ A_{\text{MaxRd}}(U) = \text{Max}[R_d(u_1, U), R_d(u_2, U), \dots, R_d(u_m, U)], \tag{A30}$$

where u_j is a randomly selected member of the set U such that $j \neq i$. Each stimulus is never compared to itself. From this definition, the probability of a false alarm and a hit is, respectively,

$$P(\text{'yes'}|n) = P[A_{\text{MaxRd}}(U_n) > c], \text{ and} \\ P(\text{'yes'}|s) = P[A_{\text{MaxRd}}(U_s) > c]. \tag{A31}$$

A.5. High threshold theory

High threshold theory has a long history and has been described in a number of recent articles (Quick, 1974; Watson, 1979; Pelli, 1985; Graham, 1989). There are four assumptions to the version of the high threshold theory presented here. First, in a detection task, the internal state of the observer is either a detect-target state or a no-detect state. Second, the detect state results only on signal trials and never on a noise trial. This assumption gives the theory its name because a distractor cannot produce the detect state. Third, observers adjust their bias by guessing ‘yes’ on some proportion of trials when they are in the no-detect state. They never guess when in the detect state because that would only introduce an error. Fourth, the probability of the detect state is a monotonic, single-valued function $G(x)$ of the relevant stimulus attribute x . This

function can be a cumulative distribution function as in signal detection theory, but it need not be in general. Here, we differ from most reviews in presenting this general function in addition to specializing it to the Weibull function, $1 - 2^{-(x/t)^k}$, as is more commonly the case (Quick, 1974; Graham, 1989). In this form, t is the threshold parameter (defined at .75 probability correct) and k is the steepness parameter.

For this general theory, the probability of a false alarm is

$$P(\text{FA}) = g, \tag{A32}$$

where g is the probability of guessing ‘yes’ when in a no-detect state. The probability of a hit is the sum of the probability of detecting the stimulus and the probability of guessing ‘yes’ if the signal is not detected,

$$P(\text{hit}) = P(\text{detect}) + g[1 - P(\text{Detect})]. \tag{A33}$$

This can be rewritten as,

$$P(\text{hit}) = g + (1 - g) P(\text{Detect}). \tag{A34}$$

For m stimuli and h targets, the probability of a detection is

$$P(\text{Detect}) = 1 - \prod_{i=1}^m [1 - G(x_i)]. \tag{A35}$$

For noise stimuli, $G(x_i)$ is zero, so $1 - G(x_i)$ is one and can be dropped from the product. Thus, the product can be rewritten for only the h signals,

$$P(\text{Detect}) = 1 - \prod_{i=1}^h [1 - G(x_i)]. \tag{A36}$$

Assuming these targets are identical and independent, this becomes,

$$P(\text{Detect}) = 1 - [1 - G(x)]^h. \tag{A37}$$

Inserting this into Eq. (A34), one has

$$P(\text{hit}) = g + (1 - g) \{1 - [1 - G(x)]^h\}. \tag{A38}$$

This general form can also be specialized for the Weibull function as

$$P(\text{hit}) = g + (1 - g) \{1 - [1 - (1 - 2^{-(x/t)^k})]^h\}, \tag{A39}$$

which simplifies to

$$P(\text{hit}) = 1 - (1 - g) 2^{-(hx/t)^k}. \tag{A40}$$

Note that performance does not depend on the set size, m , but does depend on the number of targets, h , and the target-distractor discriminability, x .

Appendix B

In a new experiment, Verghese measured speed discrimination as a function of set size with and without noise. The experiment used methods similar to the

search experiment described in Verghese and Stone (1995, experiment 2). Because this study is unpublished, we report here the details of the method. The threshold results for all conditions were described in the distractor heterogeneity section and the ROC results for the no-noise conditions were described in the response bias section.

B.1. Method

The observer's task was to indicate whether or not the display contained a single grating patch that moved faster than the other patches. Stimuli were drifting grating patches in a stationary 2-dimensional Gaussian window, with a horizontal and vertical spatial spread of 0.4° (Gaussian standard deviation). The spatial frequency of the gratings was $1.5\text{ c}/^\circ$, and the reference temporal frequency was 8 Hz, yielding a reference speed of $5.3^\circ/\text{s}$. The peak contrast of the patches was 20%. They were randomly placed at an eccentricity of 4.3° , concentric with the fixation point, with a center-to-center spacing of 2.3° . The stimuli were presented for 195 ms with abrupt onset and offset.

The major difference from Verghese and Stone (1995) was the use of a rating procedure rather than a 2IFC judgment. The observers responded on a 4 point scale: 1 indicated certainty that the target was present and 4 indicated certainty that it was absent. These ratings were used to construct ROC functions and the d'_e (or d_a) was estimated for each of several speed differences. The d'_e (or d_a) values were used to construct psychometric functions that, in turn, were used to estimate thresholds defined by 75% correct performance.

Set size was manipulated using a cueing procedure (cf. Palmer, 1994). The stimulus always consisted of six grating patches, whose locations were cued 500 ms before stimulus onset. The relevant locations were cued by black markers while the other stimulus locations were cued by white markers. For the set size 2 condition, the target could only occur in the two locations cued by black markers. For the set size 6 condition, the target could occur in all six locations and therefore all locations were cued by black markers. This method kept the density of the displayed elements constant despite the increase in relevant set size. In addition, for the set size 2 condition, the two potential target locations were chosen to be opposite each other to minimize eye movements.

The experiment also included a noise condition with external noise added to the speed of the distractors. In the no-noise condition, the distractor patches all moved at the reference speed of $5.3^\circ/\text{s}$; in the noise condition, the speed of the distractor patches was drawn from a normal distribution centered about the reference speed with a standard deviation of $0.89^\circ/\text{s}$. For both, the target speed was constant. The no-noise and noise

conditions were presented in separate blocks of trials. The data reported below are based on 400–600 trials per condition per observer. Two experienced observers participated, BB and PV (one of the authors).

References

- Aiken, D. (1992). *Comparing conjunction and disjunction visual search with equally similar targets*. Unpublished Master's Thesis, University of Washington, Seattle, WA.
- Andriessen, J. J., & Bouma, H. (1976). Eccentric vision: adverse interactions between line segments. *Vision Research*, *16*, 71–78.
- Ashby, G. A. (1992). *Multidimensional models of perception and cognition*. Hillsdale, NJ: Erlbaum.
- Ashby, F. G., & Gott, R. E. (1988). Decision rules in the perception and categorization of multidimensional stimuli. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *14*, 33–53.
- Barlow, H. B. (1956). Retinal noise and absolute threshold. *Journal of the Optical Society of America*, *46*, 634–639.
- Bauer, B., Jolicoeur, P., & Cowan, W. B. (1996a). Distractor heterogeneity versus linear separability in colour visual search. *Perception*, *25*, 1281–1293.
- Bauer, B., Jolicoeur, P., & Cowan, W. B. (1996b). Visual search for colour targets that are or are not linearly separable from distractors. *Vision Research*, *36*, 1439–1465.
- Bennett, P. J., & Jaye, P. D. (1995). Letter localization, not discrimination, is constrained by attention. *Canadian Journal of Experimental Psychology*, *49*, 460–504.
- Bergen, J. R., & Julesz, B. (1983). Rapid discrimination of visual patterns. *IEEE Transactions on Systems, Man and Cybernetics*, *13*, 857–863.
- Bloomfield, J. R. (1979). Visual search with embedded targets: color and texture differences. *Human Factors*, *21*, 317–330.
- Bouma, H. (1970). Interaction effects in parafoveal letter recognition. *Nature*, *226*, 177–178.
- Breitmeyer, B. G. (1984). *Visual masking: an integrative approach*. New York: Oxford University Press.
- Broadbent, D. (1987). Simple models for experimentable situations. In P. Morris, *Modeling cognition* (pp. 169–185). New York: Wiley.
- Broadbent, D. E. (1958). *Perception and communication*. London: Pergamon Press.
- Broadbent, D. E. (1971). *Decision and stress*. NY: Academic Press.
- Bundesen, C., & Pedersen, L. F. (1983). Color segregation and visual search. *Perception & Psychophysics*, *33*, 487–493.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, *97*, 523–547.
- Burgess, A., & Barlow, H. B. (1983). The precision of numerosity discrimination in arrays of random dots. *Vision Research*, *23*, 811–820.
- Burgess, A. E., & Ghandeharian, H. (1984). Visual signal detection. II. Signal-location identification. *Journal of the Optical Society of America A*, *1*, 906–910.
- Burgess, A. E., Wagner, R. F., Jennings, R. J., & Barlow, H. B. (1981). Efficiency of human visual signal discrimination. *Science*, *214*, 93–94.
- Busey, T. A., & Townsend, J. T. Independent sampling versus inter-item dependencies in whole report processing: Contributions of processing architecture and variable attention. *Journal of Mathematical Psychology* (in press).
- Callaghan, T. C., Lasaga, M. I., & Garner, W. R. (1986). Visual texture segregation based on orientation and hue. *Perception & Psychophysics*, *39*, 32–38.

- Cannon, M. W., & Fullenkamp, S. C. (1991). Spatial interactions in apparent contrast: inhibitory effects among grating patterns of different spatial frequencies, spatial positions and orientations. *Vision Research*, 31, 1985–1998.
- Carrasco, M., & Frieder, K. S. (1997). Cortical magnification neutralizes the eccentricity effect in visual search. *Vision Research*, 37, 63–82.
- Carrasco, M., McLean, T. L., Katz, S. M., & Frieder, K. S. (1998). Feature asymmetries in visual search: effects of display duration, target eccentricity, orientation and spatial frequency. *Vision Research*, 38, 347–374.
- Cave, K. R., & Wolfe, J. M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology*, 22, 225–271.
- Chubb, C., Sperling, G., & Solomon, J. A. (1989). Texture interactions determine perceived contrast. *Proceedings of the National Academy of Science*, 86, 9631–9635.
- Cohn, T. E. (1981). Absolute threshold: analysis in terms of uncertainty. *Journal of the Optical Society of America*, 71, 783–785.
- Cohn, T. E., & Lasley, D. J. (1974). Detectability of a luminance increment: Effect of spatial uncertainty. *Journal of the Optical Society of America*, 64, 1715–1719.
- Cohn, T. E., & Wardlaw, J. C. (1985). Effect of large spatial uncertainty on foveal luminance increment detectability. *Journal of the Optical Society of America A*, 2, 820–825.
- Davis, E. T., Kramer, P., & Graham, N. (1983). Uncertainty about spatial frequency, spatial position, or contrast of visual patterns. *Perception & Psychophysics*, 33, 20–28.
- Davis, E. T., & Peterson, S. A. (1998). *Effects of mirror-image distractors on visual search: Data and theory*. Talk presented at the meeting of the Psychonomic Society.
- Driver, J., & McLeod, P. (1992). Reversing visual search asymmetries with conjunctions of movement and orientation. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 22–33.
- Dorfman, D. D., & Alf, E., Jr. (1969). Maximum likelihood estimation of parameters of signal detection theory and determination of confidence intervals-rating method data. *Journal of Mathematical Psychology*, 6, 487–497.
- Duda, R. O., & Hart, P. E. (1973). *Pattern classification and scene analysis*. New York: Wiley.
- Duncan, J. (1980). The locus of interference in the perception of simultaneous stimuli. *Psychological Review*, 87, 272–300.
- Duncan, J. (1989). Boundary conditions on parallel processing in human vision. *Perception*, 18, 457–469.
- Duncan, J., & Humphreys, G. (1989). Visual search and stimulus similarity. *Psychological Review*, 96, 433–458.
- Duncan, J., & Humphreys, G. (1992). Beyond the search surface: visual search and attentional engagement. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 578–588.
- Eckstein, M. P. (1998). The lower visual search efficiency for conjunctions is due to noise and not serial attentional processing. *Psychological Science*, 9, 111–118.
- Eckstein, M.P., Ahumada, A. J., Jr., & Watson, A. B. (1997). Visual signal detection in structured backgrounds. II. Effects of contrast gain control, background variations, and white noise. *Journal of the Optical Society of America A*, 14, 2406–2419.
- Eckstein, M. P., Beutter, B. R., & Stone, L. S. (submitted). Quantifying the performance limits of human saccadic targeting during visual search.
- Eckstein, M. P., Thomas, J. P., Whiting, Palmer, J., & Shimozaki, S. S. A signal detection model predicts the effects of set-size on visual search accuracy for feature, conjunction and disjunction displays. *Perception & Psychophysics* (in press).
- 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.
- Engel, F. L. (1977). Visual conspicuity, visual search and fixation tendencies of the eye. *Vision Research*, 17, 95–108.
- Enns, J. T., & Rensink, R. A. (1991). Preattentive recovery of three-dimensional orientation from line drawings. *Psychological Review*, 98, 335–351.
- Eriksen, C. W. (1953). Object location in a complex perceptual field. *Journal of Experimental Psychology*, 45, 126–132.
- Eriksen, C. W. (1980). The use of a visual mask may seriously confound your experiment. *Perception & Psychophysics*, 28, 89–92.
- Estes, W. K., & Taylor, H. A. (1964). A detection method and probabilistic models for assessing information processing from brief visual displays. *Proceedings of the National Academy of Science*, 52, 446–455.
- Fahle, M. (1991). Parallel perception of vernier offsets, curvature, and chevrons in humans. *Vision Research*, 31, 2148–2149.
- Farmer, E. W., & Taylor, R. M. (1980). Visual search through color displays: effects of target-background similarity and background uniformity. *Perception & Psychophysics*, 27, 267–272.
- Feller, W. (1957). *An introduction to probability theory and its applications* (2nd ed.). New York: Wiley.
- Findlay, J. M. (1997). Saccade target selection during visual search. *Vision Research*, 37, 617–631.
- Foley, J. M., & Schwarz, W. (1998). Spatial attention: effect of position uncertainty and number of distractor patterns on the threshold-versus-contrast function for contrast discrimination. *Journal of the Optical Society of America A*, 15, 1036–1047.
- Gegenfurtner, K. R., & Kiper, D. C. (1992). Contrast detection in luminance and chromatic noise. *Journal of the Optical Society of America A*, 9, 1880–1888.
- Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review*, 96, 267–314.
- Geisler, W. S., & Chou, K. L. (1995). Separation of low-level and high-level factors in complex tasks: visual search. *Psychological Review*, 102, 356–378.
- Gescheider, G. A. (1985). *Psychophysics: method, theory, and application*. Hillsdale, NJ: Erlbaum.
- Gordon, I. E. (1968). Interactions between items in visual search. *Journal of Experimental Psychology*, 76, 348–355.
- Graham, N. (1989). *Visual pattern analyzers*. New York: Oxford University Press.
- Graham, N., Kramer, P., & Yager, D. (1987). Signal-detection models for multidimensional stimuli: Probability distributions and combination rules. *Journal of Mathematical Psychology*, 31, 366–409.
- Graham, N., Robson, J. G., & Nachmias, J. (1978). Grating summation in fovea and periphery. *Vision Research*, 18, 815–826.
- Green, D. M. (1961). Detection of auditory sinusoids of uncertain frequency. *Journal of the Acoustical Society*, 33, 897–903.
- Green, D. M., & Birdsall, T. G. (1978). Detection and recognition. *Psychological Review*, 85, 192–206.
- Green, D. M., & Luce, R. D. (1975). Theoretical note: parallel psychometric functions from a set of independent detectors. *Psychological Review*, 82, 483–486.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Krieger.
- Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J., & Stahel, W. A. (1986). *Robust statistics: the approach based on influence functions*. New York: Wiley.
- He, P., & Kowler, E. (1989). The role of location probability in the programming of saccades: implications for 'center-of-gravity' tendencies. *Vision Research*, 29, 1165–1181.
- He, P., & Kowler, E. (1991). Saccadic localization of eccentric forms. *Journal of the Optical Society of America A*, 8, 440–449.
- Hoffman, J. E. (1979). A two-stage model of visual search. *Perception & Psychophysics*, 25, 319–327.

- Hooge, I. T. C., & Erkelens, C. J. (1996). Control of fixation duration in a simple search task. *Perception & Psychophysics*, *58*, 969–976.
- Hooge, I. T. C., & Erkelens, C. J. (1999). Peripheral vision and oculomotor control during visual search. *Vision Research*, *39*, 1567–1575.
- Humphreys, W., Quinlan, P. T., & Riddoch, M. J. (1989). Grouping processes in visual search: effects with single- and combined-feature targets. *Journal of Experimental Psychology: General*, *118*, 258–279.
- Irwin, D. E. (1991). Information integration across saccadic eye movements. *Cognitive Psychology*, *23*, 420–456.
- Ivry, R. B., & Cohen, A. (1992). Asymmetry in visual search for targets defined by differences in movement speed. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 1045–1057.
- Kassam, S. A. (1988). *Signal detection in non-gaussian noise*. New York: Springer-Verlag.
- Koopman, B. O. (1956). The theory of search. II. Target detection. *The Journal of the Operations Research Society of America*, *4*, 503–531.
- Kowler, E. (1990). The role of visual and cognitive processes in the control of eye movement. In E. Kowler, *Eye movements and their role in visual and cognitive processes*. New York: Elsevier.
- Krantz, D. H. (1969). Threshold theories of signal detection. *Psychological Review*, *76*, 308–324.
- Kröse, B. J., & Burbeck, C. A. (1989). Spatial interactions in rapid pattern discrimination. *Spatial Vision*, *4*, 211–222.
- Lappin, J. S. (1978). The relativity of choice behavior and the effect of prior knowledge on the speed and accuracy of recognition. In N. J. Casteller, & F. Restle, *Cognitive theory*, vol. 3 (pp. 139–168). Hillsdale, NJ: Erlbaum.
- Legge, G. E., & Kersten, D. (1987). Contrast discrimination in peripheral vision. *Journal of the Optical Society of America A*, *4*, 1594–1598.
- Legge, G. E., Kersten, D., & Burgess, A. E. (1987). Contrast discrimination in noise. *Journal of the Optical Society of America A*, *4*, 391–404.
- Leshowitz, B., Taub, H. B., & Raab, D. H. (1968). Visual detection of signals in the presence of continuous and pulsed backgrounds. *Perception & Psychophysics*, *4*, 207–213.
- Lindsay, P. H., Taylor, M. M., & Forbes, S. M. (1968). Attention and multidimensional discrimination. *Perception & Psychophysics*, *4*, 113–117.
- Logan, G. D. (1994). Spatial attention and the apprehension of spatial relations. *Journal of Experimental Psychology: Human Perception and Performance*, *20*, 1015–1036.
- Lu, Z. L., & Doshier, B. A. (1998). External noise distinguishes attention mechanisms. *Vision Research*, *38*, 1183–1198.
- Lu, Z. L., & Doshier, B. A. (1999). Characterizing human perceptual inefficiencies with equivalent internal noise. *Journal of the Optical Society of America A*, *16*, 764–778.
- Luce, R. D. (1959) *Individual choice behavior*. New York: Wiley
- Luce, R. D. (1963a). Detection and recognition. In R. D. Luce, R. R. Bush, & E. Galanter, *Handbook of mathematical psychology* (pp. 105–189). New York: Wiley.
- Luce, R. D. (1963b). A threshold theory for simple detection experiments. *Psychological Review*, *70*, 61–79.
- Luce, R. D. (1986). *Response times*. New York: Oxford University Press.
- Luria, S. M., & Strauss, M. S. (1975). Eye movements during search for coded and uncoded targets. *Perception & Psychophysics*, *17*, 303–308.
- Macmillan, N. A., & Creelman, C. D. (1991). *Detection theory: a user's guide*. New York: Cambridge University Press.
- Macquistan, A. D. (1994). Heterogeneity effects in visual search predicted from the group scanning model. *Canadian Journal of Experimental Psychology*, *48*, 495–514.
- Maloney, L. T. (1990). The slope of the psychometric function at different wavelengths. *Vision Research*, *30*, 129–136.
- Mansfield, R. J. W. (1973). Latency functions in human vision. *Vision Research*, *13*, 2219–2234.
- McLean, J. E., Palmer, J., & Loftus, G. R. (1997). *Comparison of set-size effects in visual search and memory*. Talk presented at the meeting of the Association for Research in Vision and Ophthalmology, Ft. Lauderdale, FL.
- Miller, J. (1982). Divided attention: evidence for coactivation with redundant signals. *Cognitive Psychology*, *14*, 247–279.
- Moore, C. M., Marrara, M. T., & Elsinger, C. L. (1998). *Visual attention and the apprehension of spatial relations*. Talk presented at the meeting of the Psychonomic Society.
- Mordkoff, J. T., & Egeth, H. E. (1993). Response time and accuracy revisited: converging support for the interactive race model. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 981–991.
- Mordkoff, J. T., Yantis, S., & Egeth, H. E. (1990). Detecting conjunctions of color and form in parallel. *Perception & Psychophysics*, *48*, 157–168.
- Morgan, M. J., Ward, R. M., & Castet, E. (1998). Visual search for a tilted target: tests of spatial uncertainty models. *Quarterly Journal of Experimental Psychology*, *51A*, 347–370.
- Motter, B. C., & Belky, E. J. (1998a). The zone of focal attention during active visual search. *Vision Research*, *38*, 1007–1022.
- Motter, B. C., & Belky, E. J. (1998b). The guidance of eye movements during active visual search. *Vision Research*, *38*, 1805–1815.
- Mumford, W. W., & Scheibe, E. H. (1968). *Noise performance factors in communication systems*. Dedham, MA: Horizon House-Microwave Inc.
- Nachmias, J. (1981). On the psychometric function for contrast detection. *Vision Research*, *21*, 215–223.
- Nachmias, J., & Kocher, E. C. (1970). Visual detection and discrimination of luminance increments. *Journal of the Optical Society of America*, *60*, 382–389.
- Nagaraja, N. S. (1964). Effect of luminance noise on contrast thresholds. *Journal of the Optical Society of America*, *54*, 950–955.
- Nagy, A.L. (1997). *Color codes mediating searches for chromatic targets*. Presented at the meeting of ARVO, Ft. Lauderdale, FL.
- Nagy, A. L., & Sanchez, R. R. (1990). Critical color differences determined with a visual search task. *Journal of the Optical Society of America*, *7*, 1209–1225.
- Nagy, A. L., & Cone, S. M. (1996). Asymmetries in simple feature searches for color. *Vision Research*, *36*, 2837–2847.
- Nakayama, K., & Silverman, G. H. (1986). Serial and parallel processing of visual feature conjunctions. *Nature*, *320*, 264–265.
- Neisser, U. (1973). Decision-time without reaction-time: experiments in visual scanning. *American Journal of Psychology*, *76*, 376–385.
- Nissen, M. J., Pokorny, J., & Smith, V. C. (1979). Chromatic information processing. *Journal of Experimental Psychology: Human Perception and Performance*, *5*, 406–419.
- Nolte, L. W., & Jaarsma, D. (1966). More on the detection of one of M orthogonal signals. *Journal of the Acoustical Society of America*, *41*, 497–505.
- Nothdurft, H. C. (1985). Sensitivity for structure gradient in texture discrimination tasks. *Vision Research*, *25*, 1957–1968.
- 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.
- O'Connell, K. M. & Treisman, A. M. (1990). *Is all orientation created equal?* Talk presented at the meeting of the Association for Research in Vision and Ophthalmology, Sarasota, FL.

- Overington, I. (1976). *Vision and acquisition: fundamentals of human visual performance, environmental influences and applications in instrumental optics*. London: Pentech Press.
- Palmer, J. (1988). Very short-term visual memory for size and shape. *Perception & Psychophysics*, *43*, 278–286.
- Palmer, J. (1990). Attentional limits on the perception and memory of visual information. *Journal of Experimental Psychology: Human Perception and Performance*, *16*, 332–350.
- Palmer, J. (1994). Set-size effects in visual search: the effect of attention is independent of the stimulus for simple tasks. *Vision Research*, *34*, 1703–1721.
- Palmer, J. (1995). Attention in visual search: distinguishing four causes of set-size effects. *Current Directions in Psychological Science*, *4*, 118–123.
- Palmer, J. (1998). Attentional effects in visual search: relating search accuracy and search time. In R. Wright, *Visual attention* (pp. 348–388). New York: Oxford University Press.
- Palmer, J., & Ames, C. T. (1992). Measuring the effect of multiple eye fixations on memory for visual attributes. *Perception & Psychophysics*, *52*, 295–306.
- Palmer, J., Ames, C. T., & Lindsey, D. T. (1993). Measuring the effect of attention on simple visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 108–130.
- Palmer, J., & McLean, J. (1995). *Imperfect, unlimited-capacity, parallel search yields large set-size effects*. Talk presented at the 1995 meeting of the Society of Mathematical Psychology, Irvine, CA. Transcript available from the authors.
- Pashler, H. (1987a). Detecting conjunctions of color and form: re-assessing the serial search hypothesis. *Perception & Psychophysics*, *41*, 191–201.
- Pashler, H. (1987b). Target-distractor discriminability in visual search. *Perception & Psychophysics*, *41*, 285–292.
- Pashler, H. (1988). Cross-dimensional interaction and texture segregation. *Perception & Psychophysics*, *43*, 307–318.
- Pashler, H. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.
- Pavel, M. (1990). *A statistical model of preattentive visual search*. Presented at the Meeting of the Psychonomic Society, New Orleans.
- Pavel, M., Econopoulou, J., & Landy, M. S. (1992). *Psychophysics of rapid visual search*. Talk presented at the meeting of the Association for Research in Vision and Ophthalmology, Sarasota, FL.
- Pelli, D. G. (1981). *Effects of visual noise*. Unpublished doctoral dissertation, Cambridge University, Cambridge, UK.
- Pelli, D. G. (1985). Uncertainty explains many aspects of visual contrast detection and discrimination. *Journal of the Optical Society of America*, *2*, 1508–1532.
- Pelli, D. G. (1987). On the relation between summation and facilitation. *Vision Research*, *27*, 119–123.
- Pelli, D. G. (1990). *Vision: coding and efficiency*. Cambridge University Press: New York.
- Pelli, D. G., & Farell, B. (1999). Why use noise? *Journal of the Optical Society of America A*, *16*, 647–653.
- Peterson, W. W., Birdsall, T. G., & Fox, W. C. (1954). The theory of signal detectability. *Transactions of IRE Professional Group on Information Theory*, *4*, 171–212.
- Piéron, H. (1914). Recherches sur les lois de variation de latence sensorielle en fonction des intensités excitations. [On the laws of variation of sensory processing time as a function of excitatory intensity]. *L'Année Psychologique*, *22*, 58–142.
- Piéron, H. (1920). Nouvelles recherches sur l'analyse du temps de latence sensorielle en fonction des intensités excitations. [Further evidence on the laws of sensory processing time as a function of excitatory intensity]. *L'Année Psychologique*, *22*, 58–142.
- Pins, D., & Bonnet, C. (1996). On the relation between stimulus intensity and processing time: Piéron's law and choice reaction time. *Perception & Psychophysics*, *58*, 390–400.
- Pöder, E. (1999). Search for feature and relative position: measurement of capacity limitations. *Vision Research*, *39*, 1321–1327.
- Poirson, A. B., & Wandell, B. A. (1990). Task-dependent color discrimination. *Journal of the Optical Society of America A*, *7*, 776–782.
- Poirson, A. B., Wandell, B. A., Varner, D. C., & Brainard, D. H. (1990). Surface characterizations of color thresholds. *Journal of the Optical Society of America A*, *7*, 783–789.
- Quick, R. F. (1974). A vector-magnitude model of contrast detection. *Kybernetik*, *16*, 65–67.
- Raab, D. H. (1961). Statistical facilitation of simple reaction times. *Transactions of the New York Academy of Sciences*, *24*, 574–590.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*, 59–108.
- Rayner, K., & Fisher, D. L. (1987). Letter processing during eye fixations in visual search. *Perception & Psychophysics*, *42*, 87–100.
- Rayner, K., McConkie, G. W., & Ehrlich, S. (1978). Eye movements and integrating information across fixation. *Journal of Experimental Psychology: Human Perception and Performance*, *4*, 529–544.
- Robson, J. G., & Graham, N. (1981). Probability summation and regional variation in contrast sensitivity across the visual field. *Vision Research*, *21*, 409–418.
- Rolland, J. P., & Barrett, H. H. (1992). Effect of random background inhomogeneity on observer detection performance. *Journal of the Optical Society of America A*, *9*, 649–658.
- Rose, A. (1948). The sensitivity performance of the human eye on an absolute scale. *Journal of the Optical Society of America*, *38*, 196–208.
- Rubenstein, B. S., & Sagi, D. (1990). Spatial variability as a limiting factor in texture-discrimination tasks: implications for performance asymmetries. *Journal of the Optical Society of America A*, *7*, 1632–1643.
- Sagi, D., & Julesz, B. (1987). Short-range limitation on detection of feature differences. *Spatial Vision*, *2*, 39–49.
- Schweickert, R. (1985). Separable effects of factors on speed and accuracy: memory scanning, lexical decision, and choice tasks. *Psychological Bulletin*, *97*, 530–546.
- Schweickert, R. (1989). Separable effects of factors on activation functions in discrete and continuous models: d' and evoked potentials. *Psychological Bulletin*, *106*, 318–328.
- Schweickert, R., Dahn, C., & McGuigan, K. (1988). Intensity and number of alternatives in hue identification: Piéron's law and choice reaction time. *Perception & Psychophysics*, *44*, 383–389.
- Scott-Brown, K., & Orbach, H. S. (1998). Contrast discrimination, non-uniform patterns and change blindness. *Proceedings of the Royal Society, B* *265*, 2159–2166.
- Shaw, M. L. (1980). Identifying attentional and decision-making components in information processing. In R. S. Nickerson, *Attention and performance VIII* (pp. 106–121). Hillsdale, NJ: Erlbaum.
- Shaw, M. L. (1982). Attending to multiple sources of information. *Cognitive Psychology*, *14*, 353–409.
- Shaw, M. L. (1984). Division of attention among spatial locations: A fundamental difference between detection of letters and detection of luminance increments. In H. Bouma, & D. G. Bouwhais, *Attention & performance X*. Hillsdale, NJ: Erlbaum.
- Simpson, A. J., & Fitter, M. J. (1973). What is the best index of detectability? *Psychological Bulletin*, *80*, 481–488.
- Smith, P. L. (1998). Attention and luminance detection: a quantitative analysis. *Journal of Experimental Psychology: Human Perception and Performance*, *24*, 105–133.
- Solomon, J. A., Lavie, N., & Morgan, M. J. (1997). Contrast discrimination function: spatial cueing effects. *Journal of the Optical Society of America A*, *14*, 2443–2448.
- Sperling, G. (1989). Three stages and two systems of visual processing. *Spatial Vision*, *4*, 183–207.
- Sperling, G., & Doshier, B. A. (1986). Strategy and optimization in human information processing. In K. R. Boff, L. Kaufman, & J.

- P. Thomas, *Handbook of perception and human performance* (pp. 2.1–2.65). New York: Wiley.
- Spillmann, L., & Werner, J. S. (1996). Long-range interactions in visual perception. *Trends in Neural Sciences*, 19, 428–434.
- Steinman, S. B. (1987). Serial and parallel search in pattern vision? *Perception*, 16, 389–398.
- Swensson, R. G. (1996). Unified measurement of observer performance in detecting and localizing target objects on images. *Medical Physics*, 23, 1709–1725.
- Swensson, R. G., & Judy, P. F. (1981). Detection of noisy visual targets: models for the effects of spatial uncertainty and signal-to-noise ratio. *Perception & Psychophysics*, 29, 521–534.
- Swensson, R. G., & Judy, P. F. (1996). Measuring performance efficiency and consistency in visual discriminations with noisy images. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 1393–1415.
- Swets, J. A. (1984). Mathematical models of attention. In R. Parasuraman, & D. R. Davis, *Varieties of attention* (pp. 183–242). New York: Academic Press.
- Swets, J. A. (1986a). Form of empirical ROCs in discrimination and diagnostic tasks: implications for theory and measurement of performance. *Psychological Bulletin*, 99, 181–198.
- Swets, J. A. (1986b). Indices of discrimination or diagnostic accuracy: their ROCs and implied models. *Psychological Bulletin*, 99, 100–117.
- Swets, J. A., Tanner, W. P., Jr., & Birdsall, T. G. (1961). Decision processes in perception. *Psychological Review*, 68, 301–340.
- Tanner, W. P., Jr. (1961). Physiological implications of psychophysical data. *Annals of the New York Academy of Science*, 89, 752–765.
- Tanner, W. P., Jr., & Swets, J. A. (1954). A decision-making theory of visual detection. *Psychological Review*, 61, 401–409.
- Taylor, M. M., Lindsay, P. H., & Forbes, S. M. (1967). Quantification of shared capacity processing in auditory and visual discrimination. *Acta Psychologica*, 27, 223–229.
- Teichner, W. H., & Krebs, M. J. (1974). Visual search for simple targets. *Psychological Bulletin*, 81, 15–28.
- Townsend, J. T. (1974). Issues and models concerning the processing of a finite number of inputs. In B. H. Kantowitz, *Human information processing: tutorials in performance and cognition* (pp. 133–185). Hillsdale, NJ: Erlbaum.
- Townsend, J. T. (1981). Some characteristics of visual whole report behavior. *Acta Psychologica*, 47, 149–173.
- Townsend, J. T. (1990). Serial vs. parallel processing: sometimes they look like Tweedledum and Tweedledee but they can (and should) be distinguished. *Psychological Science*, 1, 46–54.
- Townsend, J. T., & Ashby, F. G. (1983). *Stochastic modeling of elementary psychological processes*. New York: Cambridge University Press.
- Treisman, A. (1982). Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 194–214.
- Treisman, A. (1991). Search, similarity, and integration of features between and within dimensions. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 652–676.
- Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Treisman, A., & Gormican, S. (1988). Feature analysis in early vision: evidence from search asymmetries. *Psychological Review*, 95, 15–48.
- Treisman, A., & Sato, S. (1990). Conjunction search revisited. *Journal of Experimental Psychology: Human Perception and Performance*, 16, 459–478.
- Treisman, A., & Souther, J. (1985). Search asymmetry: a diagnostic for preattentive processing of separable features. *Journal of Experimental Psychology: General*, 114, 285–310.
- van Trees, H. L. (1968). *Detection, estimation, and modulation theory*. New York: Wiley.
- Vergheese, P., & Nakayama, K. (1994). Stimulus discriminability in visual search. *Vision Research*, 34, 2453–2467.
- Vergheese, P., & Stone, L. S. (1995). Combining speed information across space. *Vision Research*, 35, 2811–2823.
- Vergheese, P., Watamaniuk, S. N. J., McKee, S. P., & Grzywacz, N. M. (1999). Local motion detectors cannot account for the detectability of an extended trajectory in noise. *Vision Research*, 39, 19–30.
- Viviani, P. (1990). Eye movements in visual search: cognitive, perceptual and motor control aspects. In E. Kowler, *Eye movements and their role in visual and cognitive processes*. New York: Elsevier.
- von Grünau, M., & Dubé, S. (1994). Visual search asymmetry for viewing direction. *Perception & Psychophysics*, 56, 211–220.
- Ward, R., & McClelland, J. L. (1989). Conjunction search for one and two identical targets. *Journal of Experimental Psychology: Human Perception & Performance*, 15, 664–672.
- Watson, A. B. (1979). Probability summation over time. *Vision Research*, 19, 515–522.
- Williams, D., & Julesz, B. (1992). Perceptual asymmetry in texture perception. *Proceedings of the National Academy of Science*, 89, 6531–6534.
- Williams, L. G. (1966a). Target conspicuity and visual search. *Human Factors*, 8, 80–92.
- Williams, L. G. (1966b). The effect of target specification on objects fixated during visual search. *Perception & Psychophysics*, 1, 315–318.
- Wolfe, J. M. (1994). Guided Search 2.0. A revised model of visual search. *Psychonomic Bulletin & Review*, 1, 202–238.
- 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.
- Yager, D., & Davis, E. T. (1987). Special issue on: variations of visual function across the visual field. *Journal of the Optical Society of America*, 4, 1478–1703.
- Zacks, J. L., & Zacks, R. T. (1993). Visual search times assessed without reaction times: a new method and an application to aging. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 798–813.
- Zenger, B., & Fahle, M. (1997). Missed targets are more frequent than false alarms: a model for error rates in visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 1783–1791.
- Zelinsky, G. J. (1999). Precuing target location in variable set size ‘nonsearch’ task: dissociating search-based and interference-based explanations for set-size effects. *Journal of Experimental Psychology: Human Perception and Performance*, 25, 875–903.
- Zelinsky, G. J., Rao, R. P. N., Hayhoe, M. M., & Ballard, D. H. (1997). Eye movements reveal the spatiotemporal dynamics of visual search. *Psychological Science*, 8, 448–453.
- Zingale, C. M., & Kowler, E. (1987). Planning sequences of saccades. *Vision Research*, 27, 1327–1341.
- Zucker, S. W., & Davis, S. (1988). Points and endpoints: a size/spacing constraint for dot grouping. *Perception*, 17, 229–247.

Errata for Palmer, Verghese and Pavel (2000)
Published in Vision Research, 40, 1227-1268
John Palmer
Last updated 4/15/03

1. Page 1258, paragraph 3, sentence 4.

Overstated result in Morgan et al., (1998). Better if "could not be" is replaced by "were not". Morgan et al. also showed selectivity under some cueing conditions.

2. Page 1261, Equation A9.

The index for the product should read $j = 1$.

3. Page 1261, Equation A11.

The first product sign should be a summation sign instead.
The index for the summation sign should read $j = 1$.

4. Page 1261, Equation A12.

The index for the product should read $j = 1$.

5. Page 1261, Equation A14.

The index for the product should read $j = 1$.

6. Page 1262, Equation A16.

The index for the first product should read $j = 1$.