

Visual Search and Attention: A Signal Detection Theory Approach

Review

Preeti Verghese¹

Smith Kettlewell Eye Research Institute
2318 Fillmore Street
San Francisco, California 94115

Background

Visual search is an everyday task. One might search for a familiar face in a crowd, or a memo on a cluttered desktop. The ease of search depends chiefly on how distinguishable the target of search is from the background. For instance, it is hard to find a memo on a desktop cluttered with paper. The task becomes much easier if the memo is a different color than all the other pieces of paper on the desktop. Attention can also improve visual search. Even if the memo is not distinguishable by color, it is easier to find if attention is drawn to its location. For instance, the telephone might ring and as I turn to answer it, I might discover the memo sitting next to the phone. In the course of this review, I will consider visual search and attention, in turn, and will present converging psychophysical and physiological evidence on the mechanisms that underlie these two phenomena.

The task of finding a target among distractors is in general referred to as visual search. Typically there is a single target that differs from the distractors along one or more dimensions. Let us start with the simple case in which the target differs from homogenous distractors along a single dimension. Earlier search literature has typically measured search performance as the response time required to find the target. Search is called serial when search times increase with the number of elements in the display, or parallel when search times are independent of the number of items in the display. Consider the case of searching for a tilted target among vertical distractors. Search times are roughly independent of the number of elements in the display when the tilt is as large as 45°, but increase with the number of elements when the target tilt is small (say 5°). According to the now classic feature integration theory (Treisman and Gelade, 1980), search performance results from a two-stage visual process. The first stage of processing is thought to occur before the influence of attention and is called preattentive. It is made up of maps specialized for various “features”—dimensions such as color, orientation, motion, and spatial frequency. These feature maps are thought to correlate to neurons selective for a restricted range of values across each of these dimensions. Elements in this first stage operate in parallel and allow simultaneous processing across the display, resulting in search times that are independent of the number of items in the display. If the first stage does not isolate the target, then a second limited-capacity serial stage is required, which focuses attention on single items or groups of items in turn. This serial stage is thought to account for search times that increase with the number of elements in the display.

The attentive, serial second stage has been invoked especially to explain conjunction search. In conjunction search, the target differs from the distractors by a unique combination of features; the target cannot be characterized by a single unique attribute. An example of a conjunction search task is a target that is a bright vertical line among distractors that are dark vertical lines and bright tilted lines. Neither brightness nor orientation alone defines the target uniquely. Treisman and Gelade (1980) suggest that the second limited-capacity stage, which depends on attention, integrates the features that define the target. Attention is therefore thought to be indispensable for successful conjunction search performance.

I will present a different approach to understanding visual search and attention. This approach shows that a large body of the psychophysical and physiological visual search results can be explained without invoking a second limited-capacity attentive stage. In the scope of this review, visual attention plays a role when observers are explicitly asked to attend to a location or feature, or when they are implicitly cued to a location or feature. There is a growing body of literature that suggests that visual attention selects objects rather than locations, although the independence of attention to objects and to locations is still being debated (see Scholl, 2001 for a review). In some cases, attention has been shown to select objects as complex as faces and surfaces (O’Craven et al., 1999; He and Nakayama, 1995). These representations are more complex than a simple conjunction of features and involve knowledge, classification, and recognition of stimulus categories. It is unknown whether these higher order processes can be included in the signal detection framework outlined below, but certainly this complex topic is outside the scope of this review. Nevertheless, the signal detection theory approach is a general framework that applies to a broad range of results involving visual search and attention that generates quantitative predictions consistent with experimental data.

Visual Search and Signal Detection Theory

More recently, a few researchers have proposed that signal detection theory (SDT), the standard model used to predict human psychophysical thresholds, can be applied to visual search. SDT has successfully predicted detection and discrimination in a search paradigm, without additional assumptions about limited-capacity. Instead of a two-stage process with a parallel, unlimited-capacity front end, followed by a limited-capacity attention stage, they have proposed a parallel stage followed by a simple decision rule. SDT describes search accuracy with many fewer assumptions than are required to predict response time to find the target. So in the interests of a simpler exposition, I will first present SDT as it applies to search accuracy, and then present SDT in the context of response time. According to SDT, the elements in a visual search display are internally represented as independent, noisy random variables. Consider for in-

¹Correspondence: preeti@ski.org

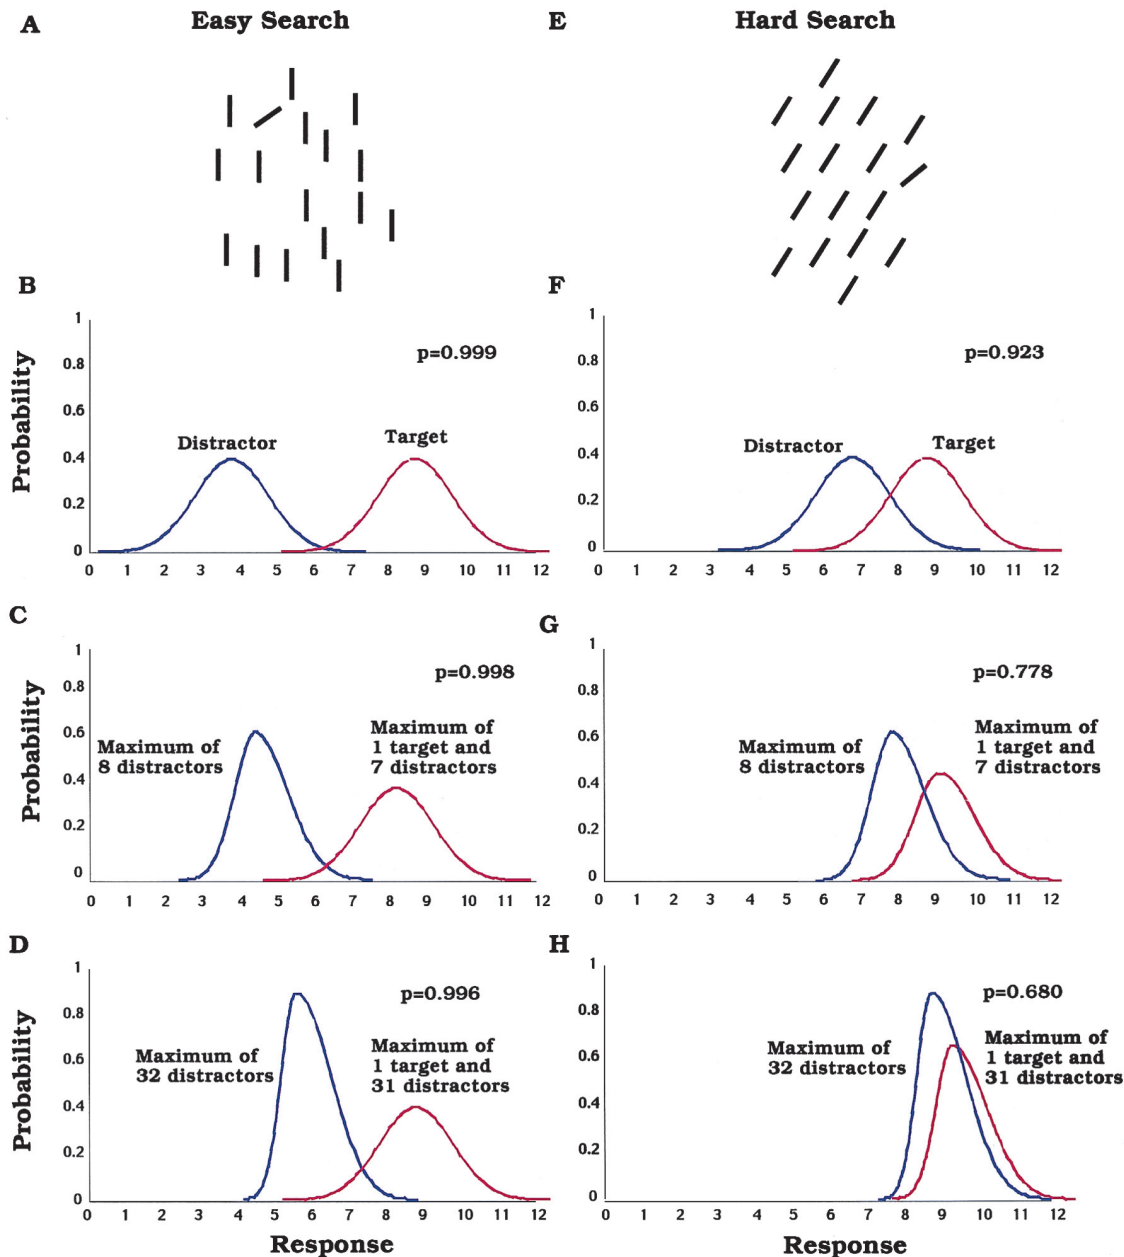


Figure 1. Examples of Easy and Difficult Search and a Graphical Representation of the Maximum Rule

(A) A search task with a single 45° target among vertical distractors. The orientation difference between target and distractors is large. (B) The response of a detector tuned to the target orientation (red curve) is a noisy distribution centered about a mean value. The response of the same detector to the distractor is a similar noisy distribution shifted to a lower mean value. For a 45° target and vertical distractors, these distributions barely overlap. (C) A graphical illustration of the distribution of the maximum value of eight samples from the distractor distribution (blue curve) as opposed to seven samples from the distractor distribution and one sample from the target distribution (red curve). The distribution of eight distractors is shifted slightly rightward, whereas the corresponding distribution for seven distractors and one target is largely unchanged. (D) The maximum rule for 32 distractors and for 31 distractors and 1 target. The distractors-only curve shifts further rightward. (E) A search task with the same 45° target, but with distractors very slightly tilted counterclockwise from 45° . (F) The target distribution (red curve) is identical to (B), but the distractor distribution (blue curve) is shifted rightward, so that the two distributions corresponding to the target and distractor overlap. (G) The distribution of the maximum of eight distractors overlaps significantly with the distribution for seven distractors and one target. (H) This degree of overlap increases further as the number of elements in the display increases to 32. Proportion correct for each of these distributions is plotted on the top right of each graph.

stance an easy search task such as searching for a line tilted 45° clockwise among vertical lines as shown in Figure 1A. SDT assumes that the observer is monitoring the output of a matched filter—a detector matched to

the properties of the signal to be detected or discriminated. In the case of the 45° target, the matched filter is a detector at the location of the target with a preferred orientation of 45° , and a receptive field size that matches

the target size; repeated presentations of the same stimulus generate responses that vary about a mean value. The probability of a given response strength is idealized as a bell-shaped Gaussian distribution that is centered on this mean response value. The variability about the mean is due to noise within the visual system, at the level of the detector. A similar response profile is generated physiologically when a neuron is presented with multiple trials of the same stimulus (Bradley et al., 1987; Newsome et al., 1989). The physiological equivalent of the SDT framework is a neuron that is well suited to detecting the target, i.e., one that has a receptive field that includes the target location and a preferred orientation that matches the target tilt of 45° . The neural response to multiple trials is a distribution of spike counts with a mean spike rate and some variability about the mean. Thus, the red distribution in Figure 1B represents both the abstract SDT concept of a noisy internal representation generated by the tilted target, as well as the probability of different spike counts in response to the tilted target. SDT assumes that the vertical distractors generate a smaller response from the filters selective to the tilted target. Thus, the mean response level to a distractor is smaller than the mean response to the target (blue versus red distribution in Figure 1B). In the physiological domain, neurons tuned to the tilted target will have a smaller response to the vertical distractors, which also can be reasonably represented by the blue distribution in Figure 1B. For simplicity, I assume that the variances of the two distributions (red and blue) are the same.

Discriminating Target from Distractor

Before I discuss how these distributions explain search performance, let us consider the simple task of discriminating which one of two sequentially presented stimuli is more tilted. The responses to the target and distractor are represented by the red and blue distributions, respectively, of Figure 1B. In the SDT framework, the decision about whether a stimulus is a target or distractor is reasonably based on which one produces a greater response in a filter selective for the tilted target. Equivalently, a physiological decision might be based on which stimulus evokes a greater response in a neuron selective for the tilted target orientation. When the target and distractor are widely separated in orientation as in this case, the target will almost always generate a larger response. In this example, it is easy to discriminate the target from the distractor. The probability of correctly identifying the 45° target is 0.999, and is shown in the top right corner of the graph.

Consider a harder search task in which there is a much smaller difference in orientation between target and distractor. In Figure 1F, the target has the same orientation as before, but the distractors are now tilted clockwise. If a decision is based on the outputs of filters or neurons selective for the 45° target orientation, then the distributions corresponding to the distractor and target will overlap considerably. Note that the response to the 45° target (red distribution) is in the same location as in Figure 1B, but the response to the tilted distractors (blue distribution) is shifted rightward, and is much closer to the red distribution. This is because the tilted distractors (as opposed to the vertical distractors) are more likely to evoke a response in the filters or neurons selective

for the 45° target. Now a sample from the distractor distribution may be larger than a sample from the target distribution, leading to an incorrect response. Similarly, in the physiology domain, the distractor could evoke a greater spike count than the target. The probability of errors is thus inversely related to the orientation difference between the target and distractor lines. In this case of the 45° target and the tilted distractors, the probability of correctly identifying the target interval is 0.923. The ease of the discrimination task depends on the separation between the means of these distributions, which is directly related to the orientation difference between target and distractor.

Discrimination also depends on the variability of these distributions. In this example, the variability of the distribution representing the vertical and tilted lines is assumed to be the same. In the SDT framework, equal variance implies that the internal representations are equally noisy, while in the physiological framework it implies that the variability of the response of the neuron responding to the vertical and tilted lines is the same. (This is a simplifying assumption; experiments show that the variability of spike rates is typically proportional to the mean firing rate.) Both the mean separation between target and distractors and the intrinsic variability of these representations determine how discriminable the target is from the distractor. The discriminability of the target from the background is defined as $(\mu_t - \mu_d)/\sigma$, where μ_t and μ_d refer to the means of the target and distractor distributions, respectively, and σ is the common standard deviation of these distributions. The distributions can be of any form, although a Gaussian distribution has been historically assumed. Thus for a given orientation difference between target and distractor, distributions with larger variance (noisier representations/larger bandwidth) are less discriminable than distributions with smaller variances. As the overlap between the distributions increases, the probability that a sample (response) from the distractor distribution will be more tilted than a sample (response) from the target distribution also increases, leading to a higher probability of errors. Thus, signal detection theory makes straightforward predictions about accuracy as a function of discriminability.

The abstract SDT distributions that I have introduced are similar to those that have been measured from cortical neurons. Bradley, Skottun, Ozhawa, Sclar, and Freeman (1987) used the signal detection theory framework to measure the ability of single striate cortical neurons to discriminate two orientations. They measured the number of trials that produced firing of different strengths. This histogram of spike counts was measured separately for the two orientations that were to be discriminated. Orientations close to the preferred orientation of the cell had higher mean spike rates, while orientations far from the preferred orientation had much lower mean spike rates. The separation between the histograms corresponding to two different orientations (on the same side of the preferred orientation) increased as a function of orientation difference. The variance of the firing rate in response to a given orientation increased with the mean firing rate. Newsome, Britten, Movshon, and Shadlen (1989) also used a similar technique to measure the ability of a single neuron in area MT to

signal motion direction. From the distribution of responses generated by each of the stimuli, Bradley et al. (1987) and Newsome et al. (1989) determined the probability that a decision based on the neuronal response would correctly detect the target of interest.

Searching for a Target among Distractors

Now let us return to the problem of visual search. A search task is essentially a discrimination task with a single target and multiple distractors. The effect of multiple distractors on a search task is inherently related to discriminability between target and distractor distributions. Let us assume that there are n elements in the display, and that there is a detector (neuron) selective for the target attribute at each of the element locations. Based on the response of these units, the observer has to decide whether the display contained a target. It turns out that for finding a single target, a decision based on choosing the largest response across the units (the maximum rule) is close to the best use of the available information, provided that the responses from each of the units is independent (Nolte and Jaarsma, 1967). Under these conditions, measuring the response of the units to each of the distractors and the target is equivalent to taking samples from the distractor and target distributions, respectively.

With the maximum rule, it is easy to illustrate the effect of the number of elements in the display on search performance. Consider an easy search task in which the target and distractor distributions are far apart, a situation sometimes referred to as "popout" (Figure 1A). The distributions of the responses to a single distractor and a single target overlap very little in this case (Figure 1B). The task is to decide which of two presentations (intervals) has the target. Both intervals have n elements: one of the intervals has 1 target and $n-1$ distractors, the other has n distractors. Let us assume that the observer looks for the largest value or maximum of the samples in each presentation and then chooses the presentation interval that has the larger of the two maximum values; this decision rule is called the "Max Rule." Consider the case when there are only eight elements in the display. In the distractor-only interval, the responses to the eight distractors are represented by eight samples from the distractor distribution. The distribution of the maximum of these eight samples is shifted toward the right tail of the distribution for the single distractor (compare blue curves in Figures 1B and 1C). When the display has one target and seven distractors, the responses are represented by seven samples from the distractor distribution and one sample from the target distribution. The distribution of the maximum of these samples (max distribution) is virtually identical to the original target distribution (compare red curves in Figures 1B and 1C). The observer makes a decision based on these two max distributions. Note that even for larger values of n , say 32, the max distribution corresponding to distractors does not overlap significantly with the max distribution of the target plus distractors (Figure 1D). Therefore, there is hardly any effect of the number of distractors on search performance, and finding the target is effortless. Proportion correct with eight distractors and with 32 distractors is 0.998 and 0.996, respectively.

When the target and distractor distributions are close together (Figure 1F), search performance degrades

quickly with the number of distractors (Figures 1G and 1H). When there are eight elements in the display, the max distribution for the distractor-only case has the same shape as the distribution described above (compare blue curves in Figures 1G and 1C), but it now overlaps significantly with the max distribution of one target and seven distractors (red curve in Figure 1G). The degree of overlap increases with the number of distractors so that the effective separation between the two max distributions decreases with set size (Figure 1H). The proportion correct for eight distractors has now fallen to 0.778, and for 32 distractors is 0.680.

Figure 2A plots the probability of a correct response as a function of set size, for both the easy and hard search tasks. These values are taken from Figure 1. Performance is almost unchanged with increasing set size in the easy search task, whereas performance falls dramatically with set size for the hard search task. Figure 2B plots performance at the same task when the more familiar response time measure is plotted versus set size. The response time measures are approximated from Palmer's (1998) model, which predicts response time as a function of set size, for a fixed accuracy level. I will defer a more detailed discussion of how SDT can predict response time measures to a later section. Both accuracy and response time measures show that increasing set size has little effect on performance when search is easy and a significant effect when search is hard. Previously, the differences between these two kinds of search tasks were attributed to parallel and serial processing. Here I show that the strategy of basing a decision on the largest sample, along with the assumption of inherently noisy representations, accounts for the effect of set size on visual search performance over a large range of target-distractor difference. Large set size effects are not necessarily due to an attentive bottleneck. In the SDT framework, set size effects depend on the discriminability of the target from the distractors; set size effects are small when discriminability is high and large when discriminability is low. This is not a serial-parallel dichotomy, but a whole continuum that depends on discriminability. Several psychophysical studies have successfully used this framework to predict the effect of discriminability on search performance (Shaw, 1982, 1984; Graham et al., 1987; Pavel et al., 1992; Palmer et al., 1993; Verghese and Nakayama, 1994; Verghese and Stone, 1995; Eckstein et al., 2000). In these studies, calculations based on the largest response have done a good job of predicting psychophysical search accuracy as a function of the number of elements in the display.

Physiology of Search and Texture Discrimination

The SDT approach has not been applied to physiological studies of visual search so far. The neural responses to a target and to distractors can be approximated by sampling the distribution of responses of a prototypical cell selective for the target attribute to multiple presentations of the target and distractor, respectively. This approximation avoids the need to measure neural responses to an array of target and distractors at different locations in the display simultaneously. Recent physiological evidence suggests that this simplifying assumption may not be justified as it ignores the modulation of a neuron's response by the presence of surrounding

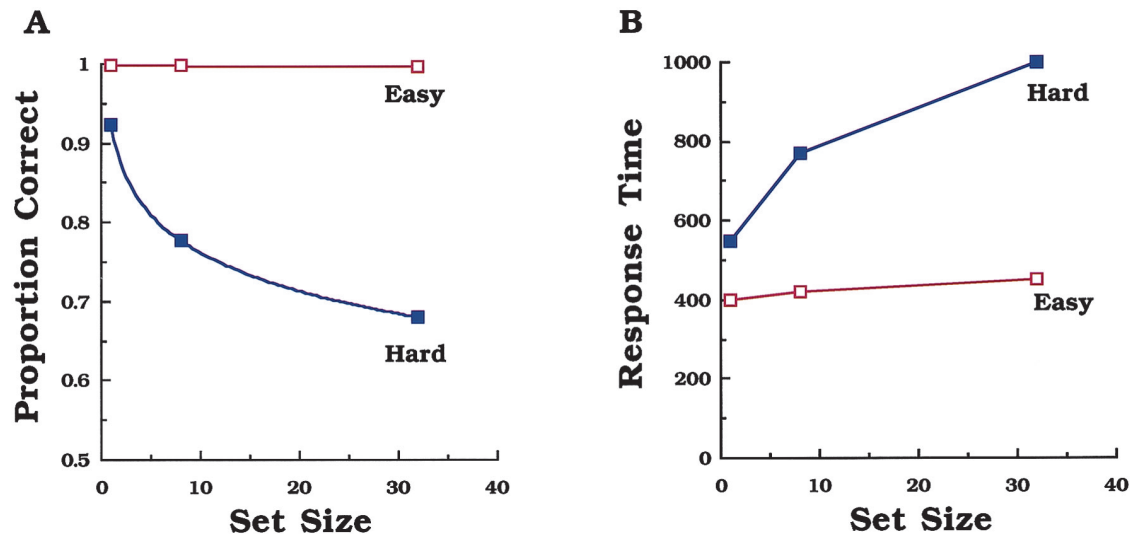


Figure 2. Two Measures of Search Performance as a Function of Set Size

(A) Proportion correct is plotted as a function of the number of elements in the display (set size). The proportion correct values are taken from Figure 1. The red and blue lines represent the easy and hard search tasks, respectively.

(B) Response times for the same tasks estimated from the model of Palmer (1998), assuming that accuracy was held constant.

elements (Knierim and van Essen, 1992; Levitt and Lund, 1997; Kastner et al., 1999). These studies showed that responses of striate cortical cells in monkeys and cats were often suppressed by the presence of surrounding texture. On average, responses were suppressed by 34% compared to the response when the stimulus was presented alone. This finding suggests that the responses to all the elements in the search display, both targets and distractors, are suppressed. SDT can accommodate this general suppressive effect by including a masking or normalizing effect of the surround on the response to the target. If the response of the cell is normalized by the activity of neighboring cells (Legge and Foley, 1980; Heeger, 1992), then the response in the presence of surrounding elements will be lower than when a stimulus element is presented alone.

A subset of these suppressed cells (roughly one-third of the cells in the sample) were more suppressed when the surround elements had the same orientation than when the surround elements had an orthogonal orientation. These cells have been called orientation contrast selective because they respond better when the stimulus within their receptive field has a different orientation than the background. This context-dependent response can also be incorporated into the SDT framework by assuming that orientations similar to the stimulus orientation within the receptive field have a greater weight in the normalization pool than do very dissimilar orientations, as suggested by the data of Cavanaugh, Bair, and Movshon (1997). According to Knierim and van Essen (1992), the difference in response to surrounds of opposite orientation as compared to surrounds with the same orientation, averaged across all the cells in the sample, was 9%. If similar orientation contrast effects occur in humans, this difference by itself is too small to account for the improvement in human search performance with orthogonal distractors as compared to similar orientation distractors. However, an SDT model that incorpo-

rates these orientation contrast effects can predict the magnitude of improvement in going from hard to easy search.

SDT Explains Other Aspects of Visual Search

Other psychophysical search results are also accounted for by the parallel SDT approach. For example, it is known that making the distractor sample more heterogeneous degrades search performance. In the orientation search example, increased distractor heterogeneity implies a larger range of orientations for the distractors. In the SDT framework, increasing the distractor heterogeneity corresponds to increasing the variance of the distractor distribution, so that the distractor distribution overlaps more with the target distribution, thus reducing the effective discriminability of the signal.

So far in the discussion it has been implicitly assumed that the variance of the target and distractor distributions is the same. In reality, this might not be the case. The variance of the two distributions can be different leading to quite different search outcomes depending on whether the elements with the larger intrinsic variance are assigned to the target or distractor. When the target has the larger variance, search is easier than when the distractor is drawn from the large-variance distribution. This explains why search performance sometimes changes dramatically if the identities of the target and distractor are reversed. Treisman and coworkers (Treisman and Souther, 1985; Treisman and Gormican, 1988) have called this class of findings "search asymmetries." These findings are also predicted by signal detection theory. Palmer et al. (2000) have shown that accuracy in finding the target decreases as a function of sigma ratio, which is the ratio of the standard deviations of the distractor to the target distribution. This trend accounts for the change in performance when the target and distractor identities are interchanged. For example, it is much easier to find a tilted line among vertical lines than to find a vertical line among tilted lines. Sutter et al.

(2000) showed that this asymmetry was due to the higher intrinsic variability of the tilted line.

The idea that a serial capacity-limited process is required to find a target defined by a conjunction of basic features has been directly challenged by the studies of Eckstein et al. (2000). He showed that when the task was to find a bright tilted ellipse among dark tilted ellipses and bright vertical ellipses, human performance was well predicted by a model that defined the target by the sum of its discriminabilities along the orientation and contrast dimensions. His results suggest that the signal detection theory approach is applicable to conjunction search as well as to other types of search.

Response Time Measurements and Signal Detection Theory

The vast majority of the behavioral studies of visual search measure the response time to report the presence or absence of the target. The key here is recognizing that the speed of response time is related to the difficulty of discrimination. In the SDT framework, it is slightly less straightforward to predict response times from the discriminability of the target and distractor. To make this relationship quantitative, response accuracy has to be held constant, a constraint that is not true of most of the response time data. Palmer and McLean (1995; Palmer, 1998) actually measured response time as a function of discriminability, while asking subjects to maintain error levels at about 10%. Their study shows a clear relationship between response time and discriminability: response times increase as a power function of decreasing discriminability. Response times are also monotonically related to set size, as has been shown in many earlier search studies. Response times increase slowly with set size when discriminability is high and search is easy, whereas they rise much more steeply when discriminability is low and search is hard (Figure 2b). Furthermore, Palmer and McLean (1995) were able to predict the response time trends with an unlimited-capacity, parallel model that accumulates information over time in a noisy manner. Evidence accumulates to support the presence of a target or a distractor at two different rates. Palmer and McLean (1995) assumed that the difference between these rates was proportional to the stimulus difference between target and distractor. Other studies that measure both response time as well as accuracy support parallel processing in both simple search tasks as well as in conjunctions (McElree and Carrasco, 1999).

Search with Cluttered Displays

Recent search experiments that relate search performance to the predictions of a signal detection theory model have taken great care to avoid confounds due to extraneous factors. Display duration is typically brief to avoid eye movement contamination. The stimuli are placed at about equal eccentricity, and are presented without masks; stimuli are not too closely spaced so that they do not violate the independence assumption of signal detection theory (Palmer, 1994; Verghese and Stone, 1995; Eckstein et al., 2000).

Recently, we have applied the visual search approach to a much less constrained display that approximates the clutter in real world situations (Verghese et al., 1999;

Verghese and McKee, 1999). Our displays were crowded with randomly placed dots (Figure 3A). The target was a string of five evenly spaced dots oriented at 45° and the distractors were pairs of dots with the same spacing and variable orientation. The target could appear anywhere in the central 2° of the display while the distractors were placed randomly throughout the circular display that was 4° in diameter. The display duration was brief (200 ms) to minimize eye movements and shifts of attention. To reduce performance below 100% correct responses, the displays had to be dense (greater than 10 dots/degree²). Rather than dealing with individual distractors at these high densities, we assumed that the entire stimulus was seen by a biologically plausible array of orientation-selective detectors. Considering the output of realistic detectors had an added advantage; it took into account the local masking of the signal by noise dots that happened to fall in the vicinity of the signal string.

The detectors in the model had properties that were consistent with both physiology and psychophysics. They had a 2-dimensional Gabor profile with an aspect ratio of 3 elongated along the preferred direction, and an orientation bandwidth of 30°–40° (full width at half height). The preferred orientation of the detectors matched that of the target and their size was chosen to be optimal for the target string. They were arranged so as to tile the display area (Figure 3B). Their center-to-center spacing was two times the space constant of the spatial profile of the detector. This spacing allowed a complete coverage of the display area without too much overlap between neighboring filters. This degree of overlap allowed us to assume that the outputs of neighboring detectors were independent. The output of each detector was normalized by the mean response of the detectors in the array and then fed to the decision stage. Since the observer's task was to choose which of two intervals contained the target, we assumed that the model chose the largest sample in each interval and then picked the interval with the larger response. As the detectors were matched to the properties of the target, it was quite likely that the detector that saw the target would have the largest response. Errors arose when a detector that saw only noise dots had a stronger response than one that saw the signal. Choosing the largest response was in essence the maximum rule applied to the outputs of the oriented detectors, rather than to an idealized response distribution of each element. Thus, this approach took into account both the competition from filters responding to noise and local masking.

This simple model successfully predicted that the probability of finding the target decreased with the number of possible locations in which the target could appear (the set size effect). It also did a reasonable job of predicting search performance as a function of the orientation difference between the target and the distractor pairs (Figure 3C), as well as the difference in length (number of dots in target string) between the target and distractors. It did underestimate human performance when the orientation difference between target and distractors is small, less than 5°. But given its simplicity (for example, individual detectors were only tuned for the orientation of the target and not for the

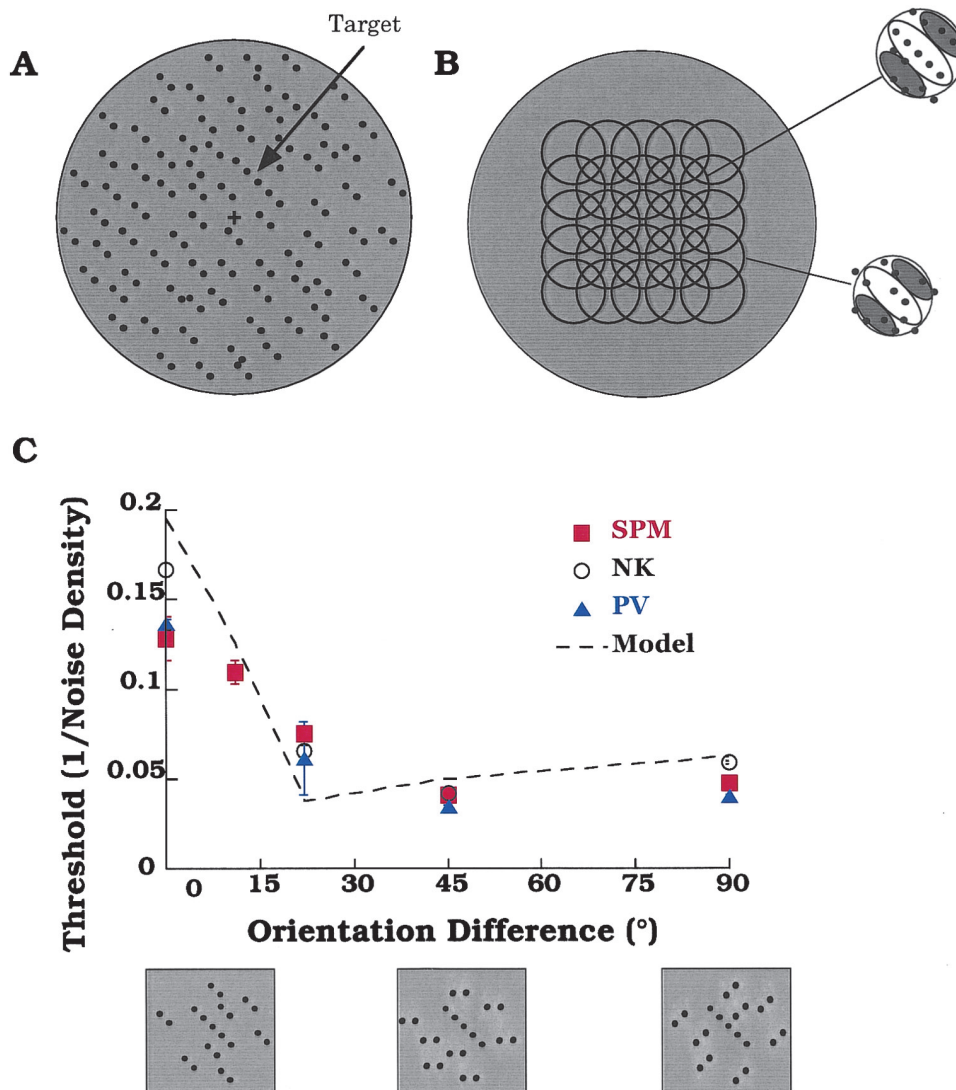


Figure 3. Search in a Dense Display

(A) The stimulus contained a 5-dot target string among oriented distractor pairs. In this example, the target and distractors have the same orientation. The arrow points to the location of the target string.

(B) The overlapping circles represent the array of oriented filters that were used to model performance in this search task. The enlarged filters on the right show a filter responding to target and distractors, and another responding to distractors alone.

(C) Comparison of data for three observers (symbols) and the model (dashed line). Thresholds, in terms of reciprocal noise density to reach criterion performance, are plotted as a function of orientation difference between target and distractor.

spacing of the dots in the string), it is impressive how well it predicted human performance. Our results indicate that the signal detection approach, combined with biologically plausible detectors, can predict performance in more realistic search tasks with crowded displays.

The Role of Visual Attention in Search *Psychophysics and Modeling*

Although many aspects of search performance can be explained with SDT as shown above, visual attention can influence performance in a search task. This section of the review is devoted to understanding how attention works. I assume that attention can be manipulated by informing the observer of the location or feature of inter-

est, or by cueing these attributes. Attention can also be divided by having the observer perform a concurrent task. Some studies suggest that features and locations may not always be the units of attention. Attention may be directed to parts of objects (Vecera et al., 2000), to single objects (Duncan, 1984; see also review by Scholl, 2001), to groups of objects (Driver et al., 1994), or to surfaces (He and Nakayama, 1995). The complexity of the attentional unit might evolve with increasing levels of visual processing. For example, the phenomenon of object attention, which involves attention to all features that define an object, might depend on how well the object is segmented from the rest of the visual scene, as well as how familiar the observer is with that category of objects. For the purposes of this review, we consider

the effect of attention at an early level, at the level of features and locations.

How does attention work? It might increase the gain of the responses at the attended location (enhancement), or it might effectively turn down the gain at the unattended locations (exclusion), or a combination of the two. The left panels of Figure 4 plot responses as a function of orientation for each of these possible models. The left panels show orientation tuning curves, where response strength in arbitrary units is plotted versus stimulus orientation. The units have a preferred orientation of 45° that is centered on the target orientation shown in red. The distractor orientation is in blue. The curves in the right panels of Figure 4 represent probability distributions of response strength for the target (red) and distractor (blue). In both panels, the continuous and dashed lines represent the unattended and attended cases, respectively. How are the curves in the left and right panels related? For instance, in the unattended case, the tuning curve in the left panels of Figure 4 shows that the target and distractor evoke mean responses of 5 and 3 units, respectively. So their probability distributions shown in the right panels of Figure 4 are centered about these values. The variance of the probability distributions in the right panel reflects the variability in the response to a given orientation in the left panel. The degree of overlap between the target and distractor distributions is inversely related to the discriminability of these distributions. For all the types of attention mechanisms that I will consider, attention causes less overlap between the target and distractor distributions (higher discriminability) than in the unattended case.

Before we consider each of the attention mechanisms, there are a few points to note. The curves in the left and right panels are Gaussian, but they refer to different measures. The left panels show orientation tuning curves, while the right panels show probability distributions. So, the area under each of the curves on the right must be one. If attention causes a change in response strength at a given orientation in the tuning curve on the left, it leads to a corresponding change in the mean value of the response of the probability distribution on the right. Since changes in mean response level are usually accompanied by proportional changes in variance, the attended curves show this altered variance. Furthermore, since the area under each of these probability distributions is one, an increase in variance is associated with a concomitant decrease in the probability associated with the mean spike count. Thus if attention causes an overall increase in spike count (and variance) in the tuning curve on the left, then the mean of the associated probability distribution curve on the right moves to a higher value. Because of a proportional increase in variance, the width of the curve increases, and consequently its height is lower.

Let us consider each of these attention mechanisms in turn. If the target location is cued, and attention works by enhancing response strength, then there is a proportional increase in response for all orientations on the tuning curve (Figure 4A). If each trial has two presentations (intervals) with the target randomly occurring in one of the intervals, then attention to the cued location will enhance responses to the target in one interval and to the distractor in the other. This enhancement results

in a shift of the response distributions representing both target and distractors (Figure 4B) toward larger values (rightward). For instance, the target response in the attended case now has a mean value of 6 units, and the distractor response has a mean value of 3.6 units. As defined earlier, discriminability is the ratio of the differences in the means to the standard deviation, $d' = \Delta\mu / \sigma$. If signal enhancement increases both mean μ and variance σ^2 by a factor k , then discriminability increases by a factor \sqrt{k} .

$$d' = \frac{k\Delta\mu}{\sqrt{k}\sigma} = \frac{\sqrt{k}\Delta\mu}{\sigma}$$

This is an approximation. I have assumed that the variances of the target and distractor distributions are equal in the unattended case and that the variance of both distributions increases by a factor k in the attended case. However, even if the target and distractor distributions in a given attentional state have unequal variances, as they might if the variance is proportional to the mean response level, enhancement still results in an increase in discriminability.

If attention works to exclude elements that differ significantly from the distractor by decreasing the bandwidth of the filter or neuron responding to the target (Figure 4C), then the response to the target is unchanged, whereas the response to the distractor decreases. This change is depicted in Figure 4D. The target distributions in the attended and unattended cases are identical, but the distractor distribution in the attended case now shifts to a smaller mean response value, consistent with the narrowing of the tuning curve in Figure 4C. The variability of the distractor distribution also decreases in proportion to the lower mean response. Discriminability increases because of the larger separation between the means of the target and distractor distributions.

Attention might also work to increase discriminability by a combination of these two mechanisms, i.e., enhancement and exclusion. This effectively results in an enhancement of responses at the preferred (target) orientation and a suppression of responses at the distractor orientation (Figure 4E). In the corresponding probability distributions in Figure 4F, the distractor distribution shifts to a lower mean value, while the target distribution shifts to a higher value, thus increasing the discriminability of these two distributions. In the next two sections, I will consider psychophysical and physiological evidence in support for each of these mechanisms.

Psychophysics

The most definitive psychophysical work that distinguishes between different attention mechanisms has come from recent experiments by Doshier and Lu (Lu and Doshier, 1998, 1999; Doshier and Lu, 2000a, 2000b; Lu et al., 2000). They manipulated attention to one of two or eight locations and compared performance in attended and unattended conditions. The novelty in their approach was the use of added external noise. They studied the effect of attention as a function of added external noise (pixel noise) and tried to distinguish between three different possibilities: that attention enhanced the stimulus, that it excluded extraneous infor-

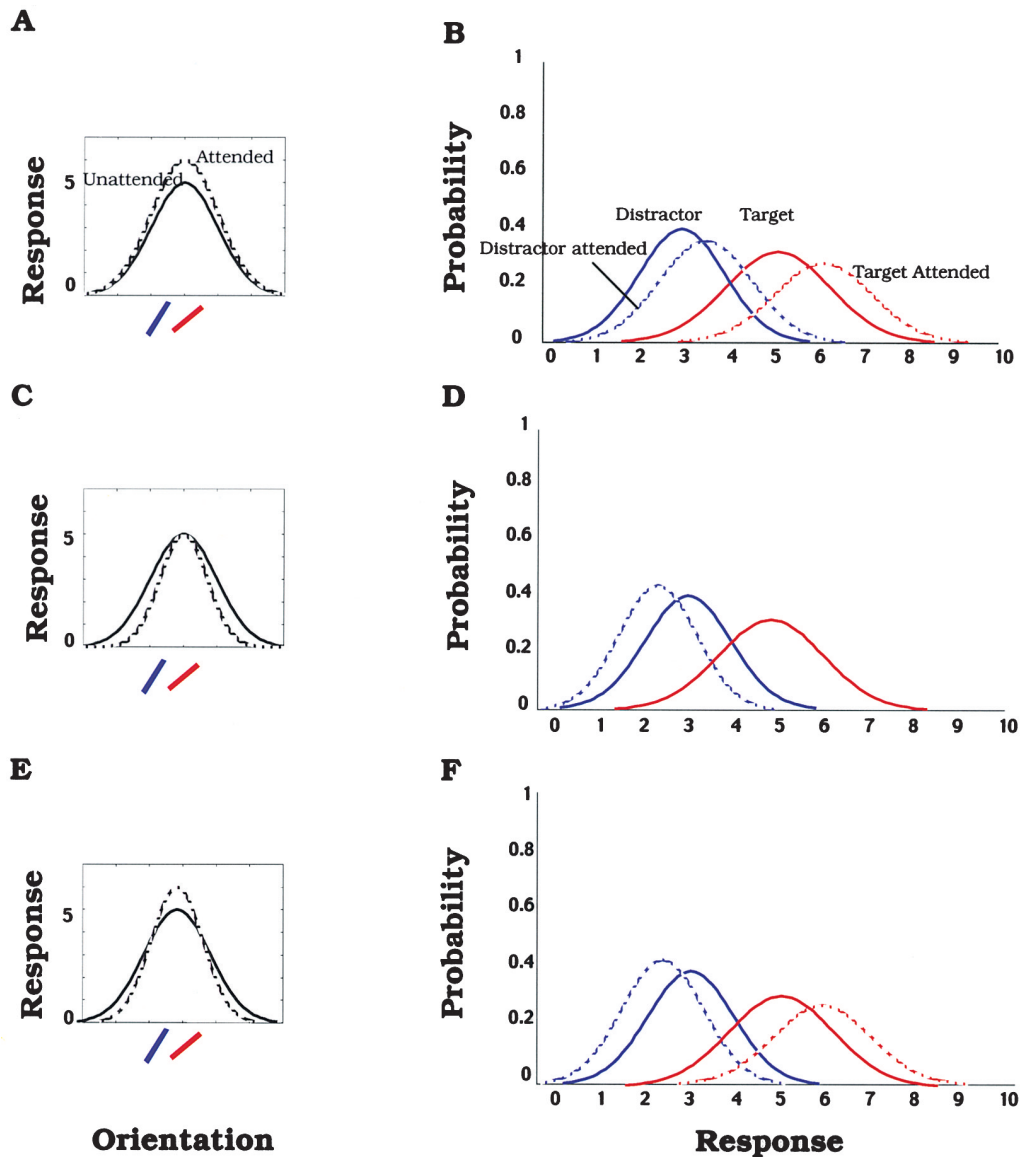


Figure 4. Three Attention Models

The panels on the left show the tuning curves of the neuron; response is plotted as a function of stimulus orientation. The target orientation is in red and the distractor orientation is in blue. The panels on the right show the probability distributions of response strength for the target and distractor. In both the left and right panels, the solid and dashed lines represent the unattended and attended cases, respectively. This figure (unlike Figure 1) depicts the variance of the distribution as proportional to the mean response. Because the area under a probability distribution curve is constrained to be 1, a distribution with a larger variance also has a lower height. (A) Attention increases the gain of the detector selective for the target. The solid line plots response versus orientation for a unit tuned to the target orientation in the unattended case. When the observer attends to the target, the response (dashed line) is enhanced across all orientations by a multiplicative factor. (B) The probability distributions of response strength for target and distractor are shown in red and blue, respectively. The solid lines represent the distributions in the unattended case and the dashed lines represent the attended case. In the attended case, the target and distractor distributions shift rightward. The shift is proportional to the mean response level so that the target distribution shifts more than the distractor distribution. This causes a smaller overlap between the distributions in the attended case, and results in a smaller proportion of errors. (C) Attention narrows the bandwidth of the detector selective for the target. The response to the target is unchanged, but the response to orientations away from the preferred orientation is reduced. (D) The probability distribution of distractor responses shifts to a lower mean level, thus increasing the separation between target and distractor responses. (E) A combination of signal enhancement and noise exclusion. Attention increases the response at the preferred (target) orientation, and also reduces the bandwidth of the detector. (F) In the attended case, the probability distribution of responses to the target is shifted to a higher mean value, whereas the response to the distractor is shifted to a lower value. Both these changes lead to a smaller overlap between the distributions, and to a smaller proportion of errors.

mation (distractors), or that it caused a decrease in the multiplicative noise that is inherent in their model.

Doshier and Lu (2000a, 2000b) compared performance for validly and invalidly cued trials, as well as neutral

trials when no precue was available. Their data support the hypothesis that attention improves performance by sharpening the selectivity of filter to exclude the added noise (Figure 4C). This is the attention mechanism that

Shiu and Pashler (1994) identified in their single stimulus cueing paradigm. Under low noise conditions, some of the Doshier and Lu data are modeled by assuming that attention enhances the stimulus, i.e., turns up the gain of the filter (Figure 4A). More recently, Carrasco, Penpeci-Talgar, and Eckstein (2000) have shown that attention increases constant sensitivity across all spatial frequencies, providing support for signal enhancement.

Lee, Itti, Koch, and Braun (1999) have also investigated mechanisms of attention in humans. The observer's attention was divided by having them perform a same-different letter task at fixation in addition to the discrimination task in the periphery. In the fully attended condition, they performed the discrimination task alone. Their data show dramatic improvements in threshold in the fully attended condition—evidence that they claimed was consistent with both higher gain and a more selective filter shape. Lee et al. fit these data with a filter similar to the one shown in Figure 4E.

The psychophysical studies to date do not suggest a single attention mechanism. Attention could improve performance in a more general way, by changing how relevant and irrelevant stimuli are weighted. This can be done by enhancing relevant stimuli or by excluding irrelevant stimuli, or by a combination of the two. The psychophysical results described here are all well fit by this general reweighting strategy.

Physiology

Psychophysics and modeling have used the term filter to represent a population of neurons tuned to a particular property. Are the psychophysical inferences about attention mechanisms consistent with physiology? Several investigators starting with Goldberg and Wurtz (1972) have studied the effect of attention on the responses of neurons in awake behaving monkeys. All of the physiological studies have measured the response and sensitivity of a cell to particular stimuli in the absence of noise. For instance, McAdams and Maunsell (1999a) have investigated the effect of attention on cells' orientation tuning. Treue and Maunsell (1996) and Treue and Trujillo (1999) have measured direction-tuning curves when attention is directed inside or outside the receptive field. All of these studies show an enhancement in response, in the form of increased spike count when the animal is attending to the cell's preferred orientation or direction within its receptive field. The increased response is evident over the entire tuning curve of the cell (see Figure 4A). Maunsell and coworkers have fit separate Gaussian functions to these response profiles for the conditions when attention is directed inside or outside the receptive field. A simple increase in the gain of the Gaussian, without a change in orientation bandwidth, best describes the effect of the animal attending to a preferred stimulus within the receptive field of the cell. In the absence of added noise, there is converging evidence from physiology and psychophysics that attention increases the gain of the stimulus. In fact, Reynolds and Desimone (1997) show that at low contrasts, the increase in response caused by attention is equivalent to increasing the stimulus contrast by about 50% on average.

McAdams and Maunsell (1999b) sought to investigate an important consequence of this increased gain—whether attention improved discriminability. To this end,

they measured both the spike count and the variability in the spike count in the presence of attention. The variability increased in proportion to the mean spike count, suggesting that the increase in gain occurred without a change in the proportionality constant between mean response and noise variance. The resulting standard deviation of the spike count grew more slowly (as the square root of the variance), so the discriminability measure defined as the ratio of the mean spike count to its standard deviation increased with attention.

The first studies investigating the role of attention in increasing the selectivity of cortical neurons were aimed at studying spatial selectivity (Bushnell et al., 1981). Several physiological studies in areas V2 and V4 of the ventral pathway showed that when there were two or more stimuli within a neuron's receptive field, attention to one of the stimuli filtered out the effect of the unattended stimulus (Moran and Desimone, 1985; Motter, 1993; Luck et al., 1997; Reynolds et al., 1999). These results were consistent with a spatial selectivity to the attended stimulus rather than a sharpening of the cell's tuning curve. For instance, Reynolds et al. (1999) measured the response of cells in areas V2 and V4 of macaque to stimuli of two different orientations. Spike counts were measured both when the animal was attending inside or outside the receptive field. In the absence of attention, the response to the two stimuli was part way between the responses to the individual stimuli. When the animal attended one of the stimuli within the receptive field, the response was closer to the response evoked by the stimulus when it was presented alone. Other studies with multiple stimuli within a receptive field (Moran and Desimone, 1985; Motter, 1993; Luck et al., 1997) also support a role for attention in increasing the spatial selectivity of the cell.

How is this spatial selectivity achieved? It is possible that attention narrows the spatial receptive field of the cell (similar to Figure 4C, but with space rather than orientation on the x axis). More likely attention to a location causes an increase of the gain of neurons tuned to that location, while possibly suppressing the response of neurons tuned to other locations. The reweighting of gains probably occurs at a stage prior to areas V2 and V4 where attentional modulation of the spatial selectivity is evident. In fact, the competitive scheme outlined by Reynolds et al. (1999) and Desimone and Duncan (1995) has similarities with the signal detection theory approach outlined above. In the competitive framework, attention affects the competition between the inputs of the cell. In the absence of attention, each of the inputs has an excitatory and inhibitory effect on the cell so that the response to a pair of stimuli is intermediate between the responses to each presented on its own. Attention to one of the inputs increases its weighting so that it overwhelms the effect of the other input. This competitive scheme has a similar outcome to signal detection theory, which bases its decision on the largest response. If attention increases the gain of the response to a particular stimulus, then according to signal detection theory, there is an increased probability that the largest response comes from the attended stimulus.

Studies that require attention to a stimulus attribute, rather than to a spatial location, do not suggest a clear role for attention in narrowing the tuning profile of the

cell to attributes such as orientation or motion direction (see Figure 4C). Early studies by Spitzer, Moran, and Desimone (1988) showed that neurons in V4 responded more vigorously and more selectively while the animal performed a hard orientation discrimination task than with an easy task. While they did not explicitly manipulate attention, they inferred that the animal was paying more attention in the hard condition than in the easy condition. Their measure of selectivity was the orientation tuning bandwidth of the cells. Although it does appear that the bandwidth of the population of their cells decreased on average, a more careful analysis by McAdams and Maunsell (1999a) suggests that when the baseline activity is subtracted from the cell's response, the decrease in bandwidth is much less clear. Treue and Trujillo (1999) explicitly set out to measure the effect of attention to a stimulus attribute rather than to a location in space. They recorded the response of MT neurons to stimuli that moved in the cells' preferred direction. The animal was always attending to a location outside the receptive field, either to a stimulus moving in the cell's preferred direction or to a stimulus moving in the null direction. Attending to the preferred-direction stimulus increased the response from the recorded cell, whereas attending to the null direction stimulus decreased its response. Treue and Trujillo characterize these effects as changes in gain, rather than changes in the direction tuning of the cell. While individual cells do not show evidence for narrowing of their tuning functions, attention might achieve filter sharpening by modulating the response of a population of cells tuned to the attended direction. Attention does appear to increase the gain of neurons tuned to the attended attribute while suppressing neurons tuned to a very different attribute. This modulation is analogous to the stimulus enhancement and distractor exclusion mechanisms that have been proposed to explain the psychophysical data.

Physiological evidence for attention acting through a combination of signal enhancement and possibly narrower bandwidth (as shown in Figure 4E) comes from a study by Croner and Albright (1999). They had macaque monkeys discriminate the direction of motion of coherently moving signals in dynamic noise (as in Newsome and Pare, 1988). In addition to this behavioral measure, they recorded from cells in area MT. In the control condition, the signal dots were the same color as the noise. In the "attended" condition, the signal dots were colored differently, thus potentially cueing the animal to their locations. Croner and Albright (1999) report improved neurometric and psychometric functions in the attended case. They attribute the improvement in the attended case to both a larger mean response motion in the preferred direction of the cell, as well as a reduced variability of this response. The larger mean response is consistent with signal enhancement. The decrease in the variability of the cell's response (a narrowing of the probability distributions in the right of Figure 4) is potentially related to a decreased bandwidth. Recall that several dots typically fall within the cells receptive field in area MT. A few of these are signal dots moving in the preferred direction of the cell, amidst several noise dots moving in random directions. Thus the variability of the cell's response is partly due to the heterogeneity of the motion direction of the noise dots within its receptive

field. The reduced variance of the response might be a reflection of the narrower tuning (increased direction selectivity) of the cell.

At what stage of cortical processing does attention modulate neuronal responses? Attention clearly affects processing in extrastriate areas V2 and V4, and in area MT. Areas V2 and MT receive direct projections from area V1, so it is important to determine the role of attention in modulating responses of V1 cells. Evidence from physiology is mixed. Motter (1993) shows that in the presence of competing stimuli, V1 cells show a modulation of responses if attention is directed within the receptive field. However, Luck et al. (1997) failed to find an effect of attention on the response of V1 cells when a single stimulus was presented to the cell. Evidence from functional magnetic resonance imaging (fMRI) studies suggest that attention modulates the response of V1 neurons. Gandhi, Heeger, and Boynton (1999) showed that cueing the observer to attend to one of two locations while performing a motion discrimination task increased the response at a site corresponding to the cued location in area V1. Brefczynski and DeYoe (1999) measured the pattern of activation when human observers were required to shift attention from one location to the next in a dense stimulus array. The pattern of MRI activation in area V1 corresponded to the pattern of attention shifts. A study by Martinez et al. (1999) resolves the discrepancy between some of the physiology studies that show no effect of attention in area V1, and the fMRI studies that do. They compared human fMRI responses and event related potentials (ERPs) in the same attention task. While the fMRI response (averaged over multiple 20 s stimulus presentations) showed a modulation of responses in area V1 by attention, the ERP data indicate that area V1 is not the first area to be modulated by attention. Rather, the ERP shows that the earliest attention-modulated responses occur in extrastriate areas at a latency of about 70–75 ms, and only appear in area V1 subsequently. A likely possibility for the discrepancy between the physiology and fMRI studies is that response in V1 is due to feedback from higher areas that are directly modulated by attention.

The other attention mechanism, filtering out the effect of unattended stimuli, has been observed only at higher levels starting at V2 and MT. V1 receptive fields are typically too small to accommodate the multiple stimuli that have been considered in these studies. It is likely that the attention-modulated sharpening of the neuronal response profile observed in these higher areas is due to a selection of a subset of V1 responses that feed into these higher-level neurons, similar to the competitive scheme suggested by Reynolds et al. (1999).

Conclusion

This review has introduced the concept of discriminability in search and has shown how the signal detection theory approach, based on noisy internal representations of these stimuli, accounts for the psychophysical data. Furthermore, adding a biologically plausible array of filters to this signal detection framework predicts search performance in dense visual displays. Visual attention has been implicated in searching for targets among distractors, but it is only recently that converging

evidence from physiology and psychophysics has clarified the mechanisms by which attention influences search. Attention acts mainly by enhancing the response to the attended stimulus, and by restricting the range of units responding to the stimulus, so as to exclude distractors and noise. The response gain associated with signal enhancement seems to occur in a way that increases the discriminability of the signal. Thus attention improves visual search by increasing the response to the target and by excluding distractors.

Acknowledgments

I would like to thank Suzanne McKee, John Palmer, Tony Movshon, John Maunsell, and an anonymous referee for their insightful comments on this review. This work was supported by NEI grant RO1 EY12038 and by NASA grant NAG 9-1163.

References

- Bradley, A., Skottun, B.C., Ohzawa, I., Sclar, G., and Freeman, R.D. (1987). Visual orientation and spatial frequency discrimination: a comparison of single neurons and behavior. *J. Neurophysiol.* 57, 755–772.
- Brefczynski, J.A., and DeYoe, E.A. (1999). A physiological correlate of the 'spotlight' of visual attention. *Nat. Neurosci.* 2, 370–374.
- Bushnell, M.C., Goldberg, M.E., and Robinson, D.L. (1981). Behavioral enhancement of visual responses in monkey cortex. I. Modulation in posterior parietal cortex related to selective visual attention. *J. Neurophysiol.* 46, 755–772.
- Carrasco, M., Penpeci-Talgar, C., and Eckstein, M. (2000). Spatial covert attention increases contrast sensitivity across the CSF: support for signal enhancement. *Vision Res.* 40, 1203–1215.
- Cavanaugh, J.R., Bair, W., and Movshon, J.A. (1997). Orientation selective setting of contrast gain by the surrounds of macaque striate cortex neurons. *Abstr. Soc. Neurosci.* 23, 567.
- Croner, L., and Albright, T.D. (1999). Segmentation by color influences responses of motion-sensitive neurons in the cortical middle temporal area. *J. Neurosci.* 19, 3935–3951.
- Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222.
- Dosher, B.A., and Lu, Z. (2000a). Mechanisms of perceptual attention in precuing of location. *Vision Res.* 40, 1269–1292.
- Dosher, B.A., and Lu, Z.L. (2000b). Noise exclusion in spatial attention. *Psychol. Sci.* 11, 139–146.
- Driver, J., Baylis, G., Goodrich, S., and Rafal, R. (1994). Axis-based neglect for visual shapes. *Neuropsychologia* 32, 1353–1365.
- Duncan, J. (1984). Selective attention and the organization of visual information. *J. Exp. Psychol. Gen.* 113, 501–517.
- Eckstein, M.P., Thomas, J.P., Palmer, J., and Shimozaki, S.S. (2000). A signal detection model predicts the effects of set size on visual search accuracy for feature, conjunction, triple conjunction, and disjunction displays. *Percept. Psychophys.* 62, 425–451.
- Gandhi, S.P., Heeger, D.J., and Boynton, G.M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proc. Natl. Acad. Sci. USA* 96, 3314–3319.
- Goldberg, M.E., and Wurtz, R.H. (1972). Activity of superior colliculus in the behaving monkey. II. Effect of attention on neuronal responses. *J. Neurophysiol.* 35, 560–574.
- Graham, N., Kramer, P., and Yager, D. (1987). Signal-detection models for multidimensional stimuli: Probability distributions and combination rules. *J. Math. Psychol.* 31, 366–409.
- He, Z.J., and Nakayama, K. (1995). Visual attention to surfaces in 3-D space. *Proc. Natl. Acad. Sci. USA* 92, 11155–11159.
- Heeger, D.J. (1992). Normalization of cell responses in cat striate cortex. *Vis. Neurosci.* 9, 181–197.
- Kastner, S., Nothdurft, H.C., and Pigarev, I.N. (1999). Neuronal responses to orientation and motion contrast in cat striate cortex. *Vis. Neurosci.* 16, 587–600.
- Knierim, J.J., and van Essen, D.C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J. Neurophysiol.* 67, 961–980.
- Lee, D.K., Itti, L., Koch, C., and Braun, J. (1999). Attention activates winner-take-all competition among visual filters. *Nat. Neurosci.* 2, 375–381.
- Legge, G.E., and Foley, J.M. (1980). Contrast masking in human vision. *J. Opt. Soc. Am. A* 70, 1459–1470.
- Levitt, J.B., and Lund, J.S. (1997). Contrast dependence of contextual effects in primate visual cortex. *Nature* 387, 73–76.
- Lu, Z.L., and Dosher, B.A. (1998). External noise distinguishes attention mechanisms. *Vision Res.* 38, 1183–1198.
- Lu, Z.L., and Dosher, B.A. (1999). Characterizing human perceptual inefficiencies with equivalent internal noise. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 16, 764–778.
- Lu, Z.L., Liu, C.Q., and Dosher, B.A. (2000). Attention mechanisms for multi-location first- and second-order motion perception. *Vision Res.* 40, 173–186.
- Luck, S.J., Chelazzi, L., Hillyard, S.A., and Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J. Neurophysiol.* 77, 24–42.
- Martinez, A., Anillo-Vento, L., Sereno, M.I., Frank, L.R., Buxton, R.B., Dubowitz, D.J., Wong, E.C., Hinrichs, H., Heinze, H.J., and Hillyard, S.A. (1999). Involvement of striate and extrastriate visual cortical areas in spatial attention. *Nat. Neurosci.* 2, 364–369.
- McAdams, C.J., and Maunsell, J.H.R. (1999a). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* 19, 431–441.
- McAdams, C.J., and Maunsell, J.H. (1999b). Effects of attention on the reliability of individual neurons in monkey visual cortex. *Neuron* 23, 765–773.
- McElree, B., and Carrasco, M. (1999). The temporal dynamics of visual search: evidence for parallel processing in feature and conjunction searches. *J. Exp. Psychol. Hum. Percept. Perform.* 25, 1517–1539.
- Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science* 229, 782–784.
- Motter, B.C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J. Neurophysiol.* 70, 909–919.
- Newsome, W.T., and Pare, E.B. (1988). Selective impairment of motion perception following ibotenic acid lesions of the middle temporal visual area of macaque monkey. *J. Neurosci.* 8, 2201–2211.
- Newsome, W.T., Britten, K.R., Movshon, J.A., and Shadlen, M. (1989). Single neurons and perception of visual motion. In *Neural Mechanisms of Visual Perception*, D.M.-K. Lam, and C. D. Gilbert, eds., Proceedings of the Retina Research Foundation, Volume 2. (The Woodlands, TX: Portfolio Publishing Company).
- Nolte, L.W., and Jaarsma, D. (1967). More on the detection of one of M orthogonal signals. *J. Acoust. Soc. Am.* 41, 497–505.
- O'Craven, K., Downing, P., and Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature* 401, 584–587.
- Palmer, J. (1994). Set-size effects in visual search: the effect of attention is independent of the stimulus for simple tasks. *Vision Res.* 34, 1703–1721.
- Palmer, J. (1998). Attentional effects in visual search: relating search accuracy and search time. In *Visual Attention*, R. Wright, ed. (New York: Oxford University Press), pp. 295–306.
- Palmer, J., and McLean, J. (1995). Imperfect, unlimited capacity, parallel search yields large set-size effects. Paper presented at the Society of Mathematical Psychology, Irvine, CA.
- Palmer, J., Ames, C.T., and Lindsey, D.T. (1993). Measuring the effect of attention on simple visual search. *J. Exp. Psychol. Hum. Percept. Perform.* 19, 108–130.
- Palmer, J., Verghese, P., and Pavel, M. (2000). The psychophysics of visual search. *Vision Res.* 40, 1227–1268.

- Pavel, M., Econopouly, J., and Landy, M.S. (1992). Psychophysics of rapid visual search. *Invest. Ophthalmol. Vis. Sci. Suppl.* 33, 1355.
- Reynolds, J.H., and Desimone, R. (1997). Attention and contrast have similar effects on competitive interactions in macaque V4. *Abstr. Soc. Neurosci.* 302.
- Reynolds, J.H., Chelazzi, L., and Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *J. Neurosci.* 19, 1736–1753.
- Scholl, B.J. (2001). Objects and attention: the state of the art. *Cognition* 80, 1–46.
- Shaw, M.L. (1982). Attending to multiple sources of information: I. The integration of information in decision making. *Cognit. Psychol.* 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 *Attention and Performance*, Vol. X, H.B. Bouma and D.G. Bouwhuis, eds. (Lawrence Erlbaum Associates), pp. 109–121.
- Shiu, L.-P., and Pashler, H. (1994). Negligible effect of spatial precuing on identification of single digits. *J. Exp. Psychol. Hum. Percept. Perform.* 20, 1037–1054.
- Spitzer, H., Desimone, R., and Moran, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science* 240, 338–340.
- Sutter, A., dela Cruz, R., and Sheft, S. (2000). Noisy, independent processing of features in visual search explains search asymmetries. *Invest. Ophthalmol. Vis. Sci. Suppl.* 41, 423.
- Treisman, A., and Gormican, S. (1988). Feature analysis in early vision: Evidence from search asymmetries. *Psychol. Rev.* 95, 15–48.
- Treisman, A., and Souter, J. (1985). Search asymmetry: A diagnostic for preattentive processing of separable features. *J. Exp. Psychol. Gen.* 114, 285–310.
- Treisman, A.M., and Gelade, G. (1980). A feature-integration theory of attention. *Cognit. Psychol.* 12, 97–136.
- Treue, S., and Martinez Trujillo, J.C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399, 575–579.
- Treue, S., and Maunsell, J.H. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* 382, 539–541.
- Vecera, S., Behrmann, M., and McGoldrick, J. (2000). Selective attention to parts of an object. *Psychonom. Bull. Rev.* 7, 301–308.
- Vergheze, P., and McKee, S. (1999). Searching for patterns in noise. *Perception Suppl.* 28, 55c.
- Vergheze, P., and Nakayama, K. (1994). Stimulus discriminability in visual search. *Vision Res.* 34, 2453–2467.
- Vergheze, P., and Stone, L.S. (1995). Combination of speed information across space. *Vision Res.* 35, 2811–2823.
- Vergheze, P., Watamaniuk, S.N.J., McKee, S.P., and Grzywacz, N.M. (1999). Local motion detectors cannot account for the detectability of an extended trajectory in noise. *Vision Res.* 39, 19–30.

