Precise attention filters for Weber contrast derived from centroid estimations

Stefanie A. Drew Department of Cognitive Sciences, University of California at Irvine, Irvine, CA, USA

Charles F. Chubb Department of Cognitive Sciences, University of California at Irvine, Irvine, CA, USA

George Sperling Department of Cognitive Sciences, University of California at Irvine, Irvine, CA, USA

How well can observers selectively attend only to dots that are lighter or darker than the background when all dot intensities are present? Observers estimated centroids of briefly flashed, sparse clouds of 8 or 16 dots, ranging in intensity from dark black to bright white on a gray background. Attention instructions were to equally weight: (i) dots brighter than the background, assigning zero weight to others; (ii) dots darker than the background, assigning zero weight to others; (iii) all dots. For each observer, a quantitative estimate of the operational attention filter (the weight exerted in the centroid estimates as a function of dot intensity) was derived for each attention instruction in each dot condition. Attended dots typically have 4× the weights of unattended dots. Whereas observers performed remarkably well in estimating centroids and achieving the three required attention filters, they achieved higher accuracy when equally weighing all dots than when selectively attending to dots of only one contrast polarity. Although their attention filters are similar, individual observers use significantly different parameters in their centroid computations. The complete model of performance enables perceptual measurements of observers' attention filters for shades of gray that are as accurate as physical measurements of color filters.

Keywords: attention, search, model, filter, centroid

Citation: Drew, S. A., Chubb, C. F., & Sperling, G. (2010). Precise attention filters for Weber contrast derived from centroid estimations. *Journal of Vision, 10*(10):20, 1–16, http://www.journalofvision.org/content/10/10/20, doi:10.1167/10.10.20.

Introduction

Our perception of the world around us depends not only on the external world but on the elements that we choose to process. In regard to this selective attention, James (1918) notes "*My experience is what I agree to attend to*. Only those items which I *notice* shape my mind—without selective interest, experience is an utter chaos."

We are concerned here with a particular aspect of selective attention, that is, attentional selection as a filter that enables relevant display information to pass to higher mental processes and that blocks the passage of less useful information. A critical element of this aspect is the voluntary nature of the attentional selection. Voluntariness is demonstrated by showing that different attention instructions cause equivalent stimuli to be processed differently.

The modern experimental study of the filtering aspect of selective attention is often traced to the work of Broadbent (1954) and Cherry (1953; for a review, see, e.g., Pashler, 1997) who studied auditory attention to one of two concurrent speakers while ignoring the other speaker. These studies determined that some physical cues, such as the spatial localization of the speaker and the pitch of the voices, were sufficient to enable the listener to separate concurrent messages and attend to just one, i.e., to select one voice stream for higher level processing. In this demonstration, it was important that given precisely the same voices and locations, observers could (1) selectively attend to either one or the other voice (to exclude masking of one voice by the other), and (2) they could not process both voices simultaneously (to demonstrate that selective attention is required).

An early demonstration of visual selective attention was Sperling's (1960) partial-report paradigm. His studies illustrated the ability of observers to selectively attend to one of three briefly flashed rows of letters, enabling the attended row to be recalled whereas virtually no memory for the unattended rows was expressed. This example also illustrates selective attention as a *spatial filter*. Short-term memory cannot hold more than 4 or 5 letters for recall. When 3 × 3 or 3 × 4 letter arrays (i.e., arrays of 3 rows with either 3 or 4 letters in a row) are flashed briefly, the process of spatial selective attention determines which one of the three rows of letters will ultimately gain access to the limited-capacity short-term memory.

Selective attention has been proposed to work on levels beyond early sensory processing (e.g., for review, see Driver, 2001; Yantis, 2000); our focus here is on the early perceptual filtering of visual information. Early perceptual

doi: 10.1167/10.10.20 Received September 30, 2009; published August 19, 2010 ISSN 1534-7362

filtering has been demonstrated and studied experimentally in many paradigms (e.g., for review, see Wolfe, 1998), though of particular relevance for the present experiments, because the interpretations involved attention filters, are studies in which equivalent stimuli were used with different selective attention instructions in visual search. Though a number of these studies have proposed means of applying attention filters to visual search (e.g., Bergen & Adelson, 1988; Foster & Ward, 1991), none actually describes a filter in mathematical detail, i.e., none defines the physical parameters of selective attention filters. Here, we introduce a methodology that permits the precise measurement of the operational properties of attention filters.

Stimuli that consist of spatially extended targets have been utilized in investigations of attention and visual processing (e.g., Cohen, Schnitzer, Gersh, Singh, & Kowler, 2007), and subjects have revealed the ability to accurately locate the center of mass of a target (e.g., Baud-Bovy & Soechting, 2001; Friedenberg & Liby, 2002). In our task, observers estimate the centroid (the center of gravity) of a cloud of dots flashed on a screen. The displays are very brief and contain a great deal more information than can be remembered (Rubin, Chubb, Wong, & Sperling, 2008). Therefore, it seems likely that only spatially parallel, bottom-up selection mechanisms are useful in making the centroid judgments. The Weber contrasts of the dots vary from deep black to bright white on a neutral gray background.

Attentional selection comes into play when we ask observers to judge the centroid of only those dots that are lighter than the neutral background and to ignore darker dots or, in other conditions, to judge the centroids of only those dots that are darker than the background. Our approach is to vary the attention instructions to our observers and to measure the effects of these changes on observers' responses. From their responses, we can infer the selection properties of their implicit attention filters, i.e., the influence of each dot contrast on their centroid judgments. We will show that observers can indeed attend selectively either to dots of positive or of negative contrast polarity relative to the gray background. We will measure the costs in efficiency that are incurred by attending selectively to dots that are brighter (or darker) than the background versus attending equally to all dot contrasts. To enable more accurate estimation of the attention filters, we also determine a characteristic parameter of each observer's centroid-computing algorithm, i.e., we reduce this idiosyncratic source of uncontrolled variance.

Methods

Observers

Eleven observers volunteered to be a part of this experiment (three females and eight males). All observers had normal or corrected-to-normal vision. All observers gave informed written consent approved by the Institutional Review Board at the University of California, Irvine.

Apparatus and stimuli

The stimuli were displayed on a 21″ NANAO Flex Scan 6600 monochrome monitor connected to a Macintosh computer. The display screen had a resolution of 1250 × 850 pixels, with a refresh rate of 75 Hz. The stimulus field in which the dots could occur comprised the central 512 × 512 pixels. This region was surrounded by a thin black border. At the viewing distance of 69 cm, this stimulus field subtended 14.897 degrees. A single stimulus dot (7 × 7 pixels) subtended 0.204 deg. Black dots had a luminance of 0.38 cd/m^2 and white dots had a luminance of 68.1 cd/m^2. The mean gray background had a luminance of 34.2 cd/m^2.

We use the term "Weber contrast" to refer to the normalized deviation of luminance from the mean. In all of our stimuli, the mean was 34.2 cd/m^2, and each display consisted entirely of this value aside from a sparse cloud of dots. In any given display, there were either 8 or 16 dots. The dots in a given 8-dot display had Weber contrasts of −1, −0.75 −0.5, −0.25, 0.25, 0.5, 0.75, and 1 (Figure 1). In 16-dot stimuli, there were two dots of each of these eight Weber contrasts (Figure 2). The locations of the dots in a given display were randomly drawn from a circular bivariate Gaussian density with a standard deviation of 60 pixels (1.70 deg).

The first trial was initiated by a key press, and each subsequent trial began immediately after the feedback for the previous trial. Prior to the onset of a given trial, the observer viewed a blank field circumscribed by a thin black frame inside which the stimulus would appear. The observer initiated the first trial with a button press. Thereafter, each stimulus was presented 200 ms after the offset of the feedback from the previous trial (Figure 3).

After the dot cloud disappeared, a crosshair cursor appeared on the screen that could be controlled by moving the mouse. Observers were asked to move the cursor to the perceived center of the dots and to click to indicate their centroid judgment. After making a judgment, the observer received a feedback display, which displayed the

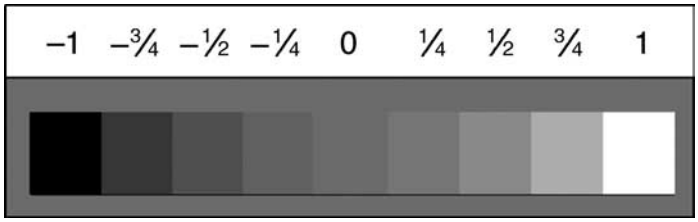

Figure 1. Weber contrasts of dots used in the experiment. The Weber contrast −1 appears black and +1 appears white. The mean-luminance background has a Weber contrast value of 0.

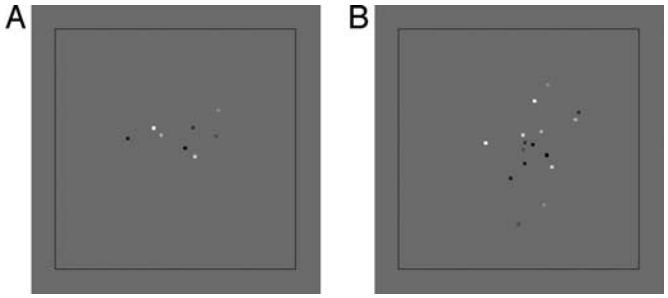

Figure 2. Sample stimuli containing dot clouds of eight or sixteen dots. (A) An 8-dot display. (B) A 16-dot display. The frequency at each contrast value is 1 in (A) and 2 in (B).

location of the target along with the location of the response. Figure 4 shows a display and the three centroids: dark, light, all.

Procedure

Each observer completed three blocks of trials: Attend-to-Dark, Attend-to-Light, and Attend-to-All. Each block consisted of 200 trials, a mixed list of 100 trials of the 8-dot clusters, and 100 trials of the 16-dot clusters. In the Attend-to-Dark condition, observers were instructed to click the cursor on the centroid of the subset of dots darker than the background, to weight all of darker-than-the-background dots equally, and to ignore all dots brighter than the background. Similarly, in the Attend-to-Light condition, observers strove to click on the centroid of just the dots lighter than the background, to weight all lighter-than-the-background dots equally, and to ignore all dots darker than the background. In the blocks of the Attend-to-All condition, observers attempted to click on the centroid of the entire cloud of dots, weighting all dots equally. The order of the three kinds of attention blocks (instructions) was balanced across observers. Observers were given the opportunity for a break between each of the blocks.

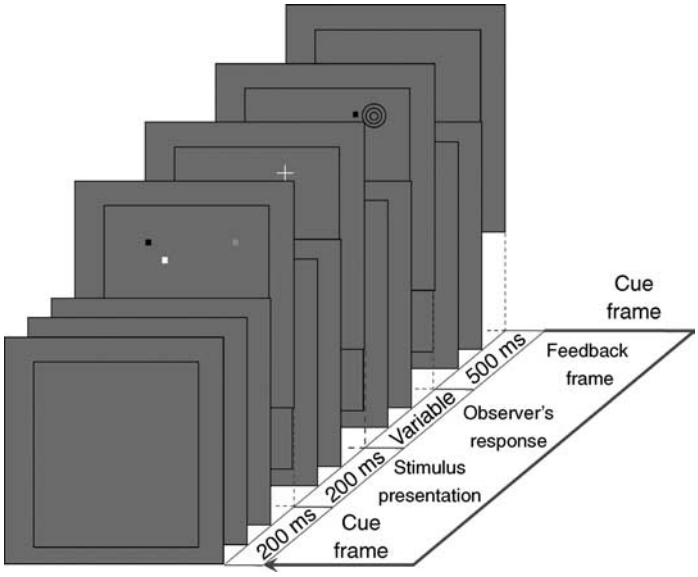

Figure 3. Experimental design. Between trials, observers viewed a 200-ms blank frame with a black outline square. This was followed by a display of dots for 200 ms. Observers then moved the cursor to the location of the perceived centroid and clicked. A 500-ms feedback frame showed the observer's response as a black dot relative to the correct location indicated as a bull's-eye. The 200-ms blank frame initiating the next trial followed immediately upon the offset of the feedback display.

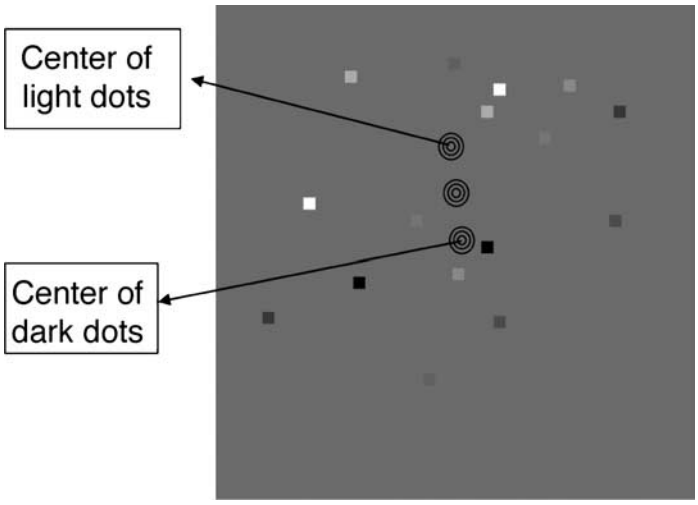

Figure 4. An enlarged sample of a 16-dot display showing the centroid of the dark dots, the centroid of the light dots, and the centroid of all the dots, which is directly between the first two.

Model for interpreting the results

Data from each observer were obtained for two dot-cloud sizes, N = 8, 16, with each of the three attention instructions A = Attend-to-Dark, Attend-to-Light, Attend-to-All. We refer to the condition in which N dots were displayed and the attention instruction was A as condition (N, A). The raw data from condition (N, A) consists of the x- and y-locations, $X_{N,A}(t)$ and $Y_{N,A}(t)$ of the observer's response for each trial t. Although these raw centroid-estimation data are interesting in and of themselves, it is of even greater interest to interpret the data in terms of the process of selective attention that enables observers, in their centroid computation, to give differential weights to attended dots and to ignore unattended dots. The process of selective attention is encapsulated in a model in which filters, determined by attention instructions, selectively weight the dots according to the dots' Weber contrasts, and these weights are what is incorporated into the centroid computation. The model has filters; observers have selective attention. Because of the intimate connection between

the two, we will speak of the "observers' filters" even though the attention filter is a model construct.

Model overview

Models are used for many different purposes in psychophysical research. Often the aim of a model is to describe a process that might plausibly operate in the brain to produce the observed results of an experiment; we emphasize, however, that this is *not* the case for the model introduced below.[1] Our focus in this study is on the attentional selection of input information, not on how the brain computes centroids. We seek to determine which sorts of attention filters people can achieve and to describe these filters quantitatively. The point of our model is to enable a more accurate characterization of attention filters across different attention conditions (Attend-to-Dark, Attend-to-Light, Attend-to-All) and different numbers of stimulus dots (8 vs. 16) by taking into account the effects of imperfect centroid computations and individual differences.

Causes of response errors

Across different attention conditions and different numbers of dots, we expect response errors to be determined in varying degrees by three distinct factors:

1. The attention filter used by the observer may differ from the ideal (target) filter.
2. The centroid computation may be imperfect. We model the imperfection by a single parameter, a bias to weight peripheral dots in the cloud differently from more central dots.
3. The resulting centroid computation may be corrupted by internal response noise.

Note that factors (1) and (2) produce *systematic* errors in observers' responses whereas factor (3) leads to *random* errors in observers' responses. The model enables us to measure variations in each of these three factors for each observer in each attention condition and in each different numbers-of-dots condition.

How the model estimates factor 1 (deviation of attention filter from target filter)

The model assumes that the observer applies an attention filter $F_{N,A}$ to each dot in the internally represented stimulus. This process replaces each dot in the stimulus by a point with a weight that depends on the dot's Weber contrast. Specifically, if a dot has Weber contrast c, then this dot gets replaced by a point with weight $F_{N,A}(c)$. It is crucial to realize that although the observer may be trying to achieve a particular target attention filter (e.g., to weight all dots darker than the background equally while giving all dots lighter than the background weight 0), in practice the attention filter $F_{N,A}$ that he/she achieves always deviates from this target filter. This deviation of the attention filter from the target filter is one important source of systematic error.

How the model estimates factor 2 (misweighting of dots depending on how peripheral they are in the cloud)

The model assumes that the observer computes the centroid of the resulting field of weighted points. However, the model admits the possibility that this computation may differ systematically from a veridical centroid computation. Specifically, the model assumes that the observer may give differentially either higher or lower weights to more central versus more peripheral dots in the cloud. We follow common practice in calling a centroid computation that underweights peripheral dots relative to central dots "robust." Conversely, we call a centroid computation that underweights central dots relative to peripheral dots "anti-robust." Any deviation from veridicality of the observer's centroid computation contributes systematic error to his/her responses. We call this source of systematic error "distance distortion."

How the model estimates factor 3 (degree of corruption by internal noise)

The model assumes that the observer's response is corrupted by additive noise. It assumes that these random response perturbations are independent and identically distributed in the vertical and horizontal directions and across trials. To assess the amount of random error corrupting the performance of a given observer in a given condition (N, A), we extract an unbiased estimate of the standard deviation of the observer's internal response noise.

Our treatment of random noise in the model is determined by the fact that if response noise were due to "early" dot-specific noise sources (i.e., uncertainty in each dot's position), and if those dot-specific noise sources had equal variance in the 8- and 16-dot conditions, then judged centroids based on 16 dots would be more accurate (would exhibit lower levels of response noise) than judged centroids based on 8 dots. In fact, the data are opposite; that is, 16-dot judged centroids are more variable than 8-dot judged centroids. Therefore, the model assumes that the main source of random error compromising performance is "late" response noise (versus "early" stimulus noise—error in representing individual dot positions; Figure 5).

For each condition (N, A), we first derive the attention filter and distance-distortion parameter that minimize the sum of squared deviations (across trials) of the model-predicted centroids from the observer's centroid settings. We then derive an unbiased estimate of the standard deviation of the observer's internal response noise[2] from residual sum of squared deviations of predicted from observed centroids (see Appendix A for details).

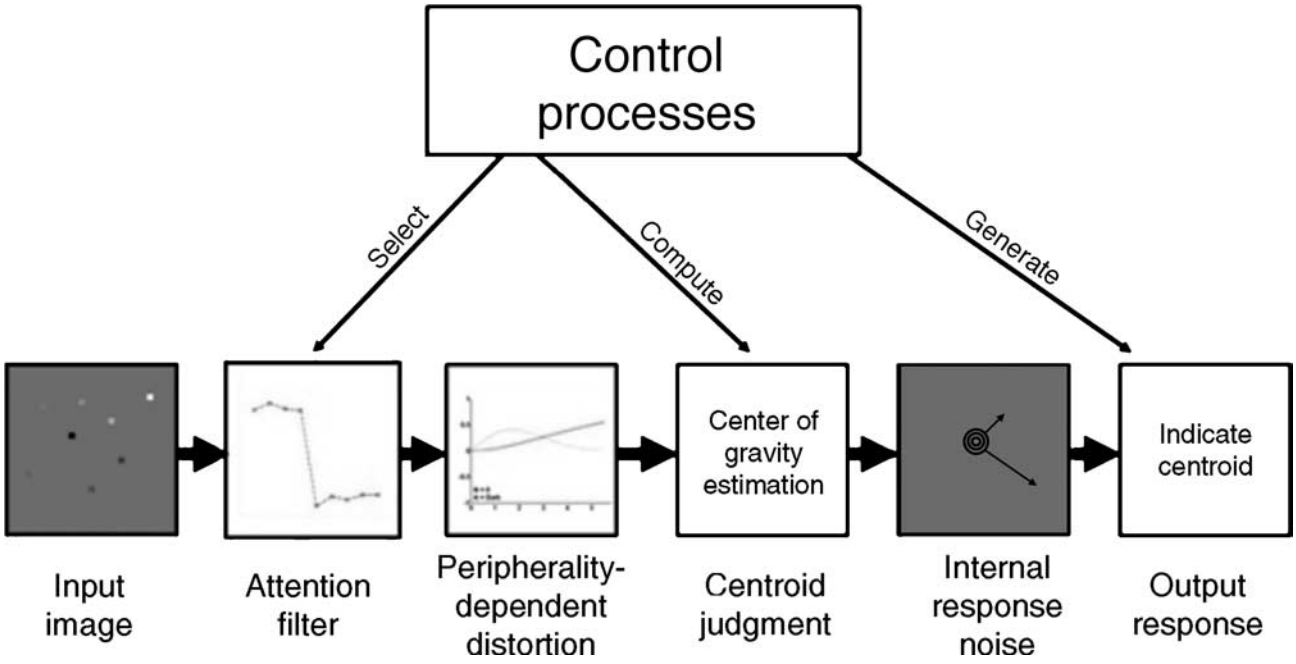

Figure 5. Model of an attention-weighted centroid computation. The model consists of stages in which (1) an attention filter is applied to the input image, (2) the resulting weights of dots are adjusted depending on their peripherality within the cloud, i.e., the reweighting can either overweight central dots (robust) or peripheral dots (anti-robust), (3) the centroid is computed, (4) positional noise is added, and (5) the resulting (x, y) value is the response output.

Model details

To describe the model more precisely, we use the following notation.

Attention filter definition

For N = 8 or 16 and A = Attend-to-Dark, Attend-to-Light, or Attend-to-All, we write $F_{N,A}$ for the attention filter used by the observer in condition (N, A). That is, for any Weber contrast c = −1, −3/4, …, 1, the attention filter $F_{N,A}(c)$ gives the weight assigned to dots of contrast c in centroid estimates in condition (N, A). Target filters for each attention condition (A) are depicted in Figure 6.

Distance distortion (robustness, anti-robustness)

In addition to applying an attention filter to dot contrast, the model allows for a possibly imperfect centroid computation that can vary from subject to subject. A simple computation that captures much of the intersubject variation is weighting dots in the centroid computation differentially depending on how peripheral they are within the cloud of dots. Let z represent a dot in the display, let $c(z)$ be the Weber contrast of z, and let $d(z)$ be the Euclidean distance (a positive real number) of z from the centroid of the $F_{N,A}$-filtered dot cloud. The weight $w(z)$ of a dot z on trial t of condition (N, A) is the product of two factors, an attention filter applied to the dot's contrast and a distance correction function applied to the dot's distance from the centroid of the $F_{N,A}$-filtered dot cloud:

$$w(z) = F_{N,A}(c(z)) \times R_{N,A}(d(z)). \quad (1)$$

The distance correction function $R_{N,A}(d(z))$ is a one-parameter, monotonic function (given by Equation A3 in Appendix A) that can either be a decreasing function of $d(z)$ (in which case the centroid computation is "robust") or an increasing function of $d(z)$ (in which case the centroid computation is "anti-robust").

Model predictions, parameter estimation

The model then predicts that the x- and y-locations of the response on trial t in condition (N, A) are given by

$$X_{N,A}(t) = \frac{\sum w(z)x(z)}{\sum w(z)} + noise_X \quad \text{and}$$

$$Y_{N,A}(t) = \frac{\sum w(z)y(z)}{\sum w(z)} + noise_Y, \quad (2)$$

where each summation ranges over all dots z presented on trial t, $w(z)$ is given by Equation 1, and $x(z)$ and $y(z)$ are

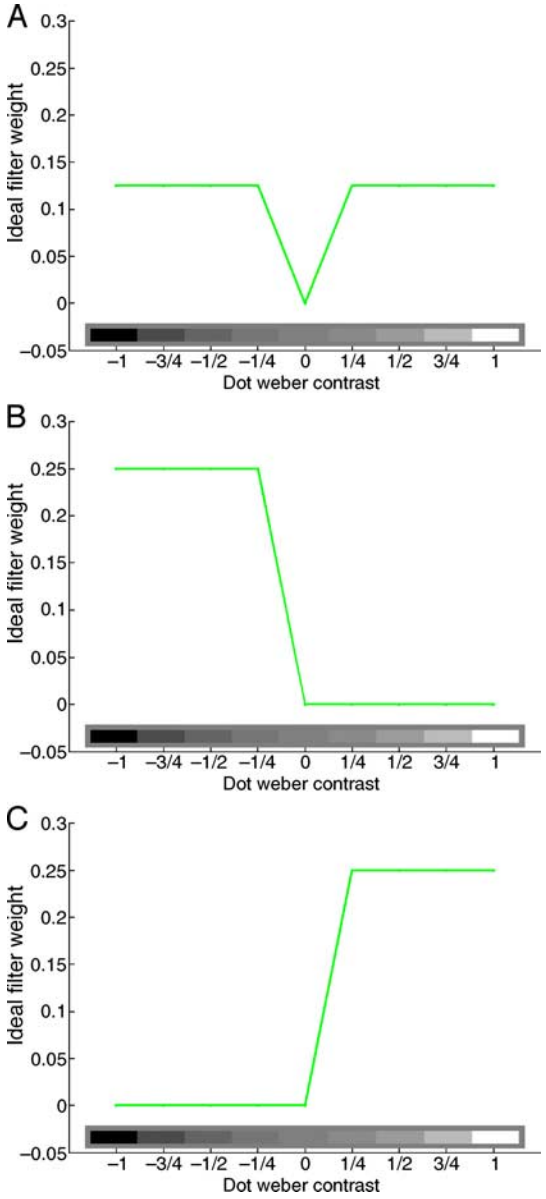

Figure 6. Target filters for (A) Attend-to-All, (B) Attend-to-Dark, and (C) Attend-to-Light instructions. A target filter gives equal weight to attended dot Weber contrasts and zero weight to unattended dot contrasts. The filter weights of all attended dots are normalized to sum to 1.

the x-location and y-location of z, respectively. The estimated standard deviations σ_x and σ_y of noise$_x$ and noise$_y$ did not differ systematically and so are subsumed in a single parameter σ, $\sigma = ((\sigma_x^2 + \sigma_y^2)/2)^{1/2}$, the standard deviation of the symmetrically distributed internal response noise.

In all conditions, we fix $F_{N,A}(0) = 0$ (i.e., implying that dots that match the background (invisible dots) exert zero weight in the centroid), and we constrain $F_{N,A}$ so that

$$\sum |F_{N,A}(c)| = 1, \tag{3}$$

where the summation ranges over the nine Weber contrasts $c = -1, -0.75, \ldots, 1$ used in each display. Under these constraints, for each condition (N, A), the filter $F_{N,A}$, and distance correction function $R_{N,A}$ are chosen to produce predicted response locations $\hat{X}_{N,A}(t)$ and $\hat{Y}_{N,A}(t)$ that minimize the sum of squared distances of the trial-by-trial observed responses. That is, we minimize

$$\sum (\hat{X}_{N,A}(t) - X_{N,A}(t))^2 + (\hat{Y}_{N,A}(t) - Y_{N,A}(t))^2, \tag{4}$$

where the sum is over all trials t in condition (N, A).

In the model, the number of degrees of freedom is 8. The attention filter consists of eight weights, those assigned to Weber contrasts $c = -1, -0.75, \ldots, 1$ (excluding 0 to which the attention filter assigns the weight 0). The absolute values of the attention filter weights are constrained to sum to 1.0 making the total number of degrees of freedom absorbed by the attention filter equal to 7. The eighth degree of freedom is contributed by the free parameter $\alpha_{N,A}$ controlling the form of the distance correction function $R_{N,A}$ (Equation A3 in Appendix A). The parameter σ giving the standard deviation of the internal response noise is derived from the deviation of the model fit from the observed responses; thus it does not constrain the fit of the model to the data.

Performance evaluation

Selectively attending to subsets of dots can adversely affect centroid judgments in two primary ways: (1) the attention filter may deviate from the target filter, and (2) the effort of filtering may increase internal response noise. Therefore, we gauge the cost of imposing a particular attention filter in two ways. First, we compute root-mean-square filter error $RMSFE_{N,A}$ (see Equation A7 in Appendix A), which reflects the *systematic deviation* of the attention filter[3] achieved by the observer in condition (N, A) from the target attention filter for condition (N, A). Second, we compute an unbiased estimate of the standard deviation σ of the *internal response noise* (see Equation 5).

Results

Preliminary measures

For each attention condition (N, A), Table 1 gives the mean distance of responses made by all observers in condition (N, A) from the centroids of (1) the dark dots, (2) the light dots, and (3) all dots in the display. These results provide preliminary evidence that observers can adjust their performance to match varying attentional demands. Note first that in the Attend-to-Dark condition

	Mean distance of response from centroid of dark dots	Mean distance of response from centroid of light dots	Mean distance of response from centroid of all dots
Attend-to-Dark (8)	24.2	47.1	32.1
Attend-to-Light (8)	44.4	23.9	33.6
Attend-to-All (8)	25.4	26.8	17.8
Attend-to-Dark (16)	23.8	37.2	26.6
Attend-to-Light (16)	38.0	23.7	26.6
Attend-to-All (16)	25.4	24.5	18.1

Table 1. Mean pixel distances of responses made in different attention conditions from the centroid of the dark dots, the centroid of the light dots, and the centroid of all dots in the display.

(for both 8- and 16-dot displays), responses are closer to the centroid of the Dark dots than they are to the centroid of the Light dots or to the centroid of All dots. Moreover, as might be expected, the mean distance of the response is greater from the centroid of the Light dots than it is from the centroid of All dots. An analogous pattern holds for responses made in the Attend-to-Light conditions. Interestingly, the most accurate responses (those that come closest to their actual target centroids) are those made in the Attend-to-All conditions (for both 8 and 16 dots).

Main result: Estimated model attention filters

Attention filters conform to attention instructions

The results of the experiment are the centroid judgments in each of the three attention conditions (Attend-to-Dark, Attend-to-Light, Attend-to-All) times two display sizes (8, 16 dots). However, these raw data are very difficult to interpret. Therefore, we present the results in terms of the model—the attention filter weights, the distance-distortion parameter α, and the standard deviation σ of internal response noise that best predict the observers' performances. An attention filter was computed for each observer in each condition. The attention filters, averaged over the 11 observers, for each of the 6 conditions, are shown in Figure 7. Figure 7A plots the best fitting attention filter (average of all observers) $F_{N,\text{Attend-to-All}}$ when observers attend to all dots for $N = 8$ and for $N = 16$. Also shown for reference is the target attention filter $F^{\text{Target}}_{\text{Attend-to-all}}$. Error bars give 95% confidence intervals. Figures 7B and 7C show the corresponding plots for the Attend-to-Dark and Attend-to-Light conditions.

The best fitting filters closely approximate the corresponding target filters; this shows that observers can modify their state of attention to conform to each of these three attention conditions. The mean $F_{N,\text{Attend-to-All}}$ filter (for each of $N = 8$ and 16) is strikingly flat across all non-zero contrasts, showing only a slight drop in sensitivity for the two Weber contrasts closest to 0. Evidently, observers are able to weight all dots nearly equally, despite the differences in Weber contrast between dots.

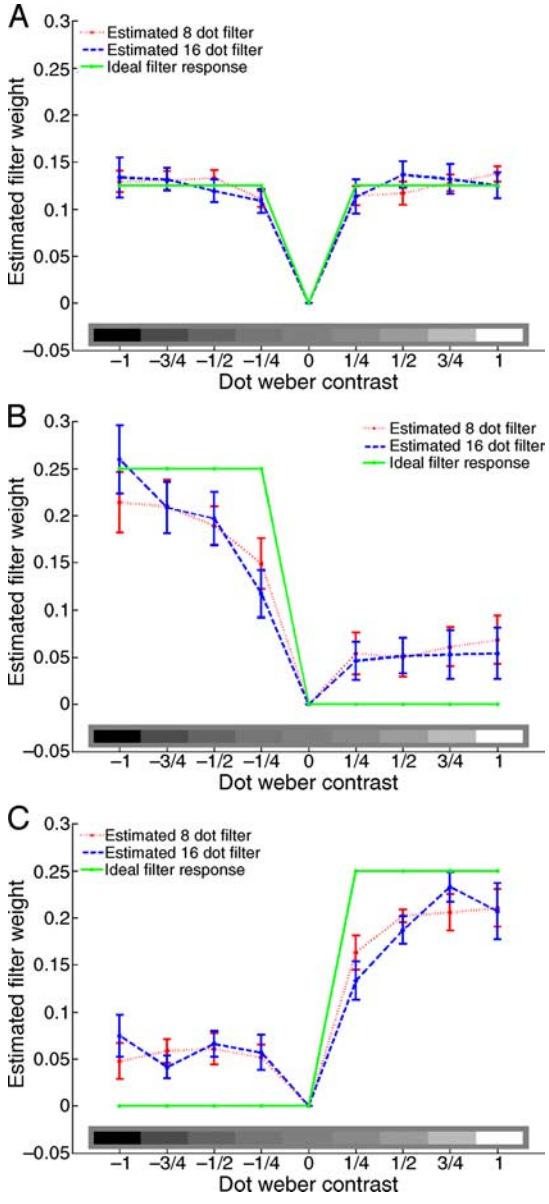

Figure 7. Estimated attention filters under (A) Attend-to-All, (B) Attend-to-Dark, and (C) Attend-to-Light task instructions. Filter weights were computed for each observer and averaged across 11 observers; the bars indicate 95% confidence intervals of the mean attention weight across observers.

	Mean RMSFE		Standard deviation of RMSFE	
Attention instructions	8-Dot clouds	16-Dot clouds	8-Dot clouds	16-Dot clouds
Attend-to-Dark	0.189	0.209	0.105	0.076
Attend-to-Light	0.165	0.205	0.053	0.053
Attend-to-All	0.058	0.075	0.020	0.025
Mean	0.137	0.163	0.059	0.051

Table 2. Deviation of attention filters from target filters: Root-mean-square filter errors (differences between target and achieved attention filters) for 8-dot and 16-dot displays averaged over the 11 observers and variation of these errors between observers.

Attention filter selectivity

The filters achieved in the Attend-to-Light conditions, $F_{8,\text{Attend-to-Light}}$ and $F_{16,\text{Attend-to-Light}}$, show a high degree of selectivity for positive vs. negative Weber contrasts; the Attend-to-Dark filters, $F_{8,\text{Attend-to-Dark}}$ and $F_{16,\text{Attend-to-Dark}}$, are equally selective. Nevertheless, the four selective filters, A = *Attend-to-Light*, *Attend-to-Dark* for N = 8 and 16, show systematic deviations from their target forms. The weights given to unattended contrasts—negative Weber contrasts by $F_{N,\text{Attend-to-Light}}$ and positive contrasts by $F_{N,\text{Attend-to-Dark}}$—are all significantly greater than 0; observers cannot completely ignore unattended dots. To characterize just how good the selective attention filters are, we consider just the three most extreme contrasts at each end of the range and omit the two contrasts closest to zero, which are less discriminable. The ratio of the average weights for these three attended to the three unattended contrasts ranges from 3.60:1 to 4.33:1. That is, among the 6 most discriminable of the 8 dot contrasts, the attention filter gives attended contrasts typically a 4:1 weight advantage over unattended contrasts. Not perfect selection, but surprisingly good.

Further results

Model attention filters depend on number of dots displayed

Mean $RMSFE_{N,A}$ values (across all observers) for N = 8, 16 and A = *Attend-to-All*, *Attend-to-Dark*, and *Attend-to-Light* are given in Table 2, and indeed the mean value of $RMSFE_{8,A}$ is lower than that for $RMSFE_{16,A}$ for all three attention instructions. A repeated measures ANOVA reveals a significant main effect for cloud size ($F(1, 10) = 9.601$, $p = 0.011$), indicating that observers are slightly but significantly better at matching their attention filter to the target filter when displays contain only eight versus sixteen dots.

Internal response noise

Total response error is the difference between the observer's response and the true centroid. The model assumes that every stimulus dot is represented internally (no dots are forgotten) and that response error occurs because of (1) an imperfect attention filter, (2) spatial distortion in the centroid computation, and (3) internal response noise. In a given condition (N, A), the internal response noise standard deviation, which we assume is common to each of the horizontal and vertical components of the internal response noise, is estimated by

$$\hat{\sigma}_{N,A} = \sqrt{\frac{\sum(\hat{X}_{N,A}(t) - X_{N,A}(t))^2 + (\hat{Y}_{N,A}(t) - Y_{N,A}(t))^2}{200 - df \text{ in model}}}, \quad (5)$$

where the sum is over all trials t, $1 \le t \le 200$ in condition (N, A). In this equation, the number of degrees of freedom in the model is 8, $X_{N,A}(t)$ and $Y_{N,A}(t)$ give the x- and y-coordinates of the observer's response on trial t in condition (N, A), and $\hat{X}_{N,A}(t)$ and $\hat{Y}_{N,A}(t)$ give the x- and y-locations predicted by the model.

As noted above, an unknown proportion of the deviation of predicted responses from observed responses is due to model failure. There is, however, no reason to think that the contribution of model failure will differentially influence performance in the different conditions. Thus, any significant differences in $\hat{\sigma}_{N,A}$ across N or A can plausibly be assumed to reflect differences in the cost incurred in adapting to these task variations.

Table 3 gives the means and standard deviations (across 11 observers) of the unbiased estimates $\hat{\sigma}_{N,A}$ of the standard deviation of random response error given by Equation 5 for all conditions (N, A). There are several trends to note in Table 3. First, $\hat{\sigma}_{N,A}$ is lower for N = 8 versus N = 16 dots for each of the three attention instructions. Second, for each of N = 8 and N = 16, $\hat{\sigma}_{N,\text{Attend-to-All}}$ is lower than $\hat{\sigma}_{N,\text{Attend-to-Dark}}$, and similarly $\hat{\sigma}_{N,\text{Attend-to-Dark}}$ is lower than $\hat{\sigma}_{N,\text{Attend-to-Light}}$. A repeated measures ANOVA confirms that all of these differences are significant: for the effect due to dot number, $F(1, 10) = 10.61$, $p = 0.009$. In addition, within-subject contrasts reveal that the difference in noise between the *Attend-to-All* vs. the *Attend-to-Dark* conditions is significant ($F(1, 10) = 5.831$, $p = 0.036$), as is the difference in noise between the *Attend-to-Dark* vs. the *Attend-to-Light* conditions ($F(1, 10) = 5.380$, $p = 0.043$). As one might expect given these two results, the difference

Attention instructions	Unbiased estimates $\hat{\sigma}_{N,A}$ of the standard deviation of internal response noise		Standard deviation of $\hat{\sigma}_{N,A}$	
	$N = 8$	$N = 16$	$N = 8$	$N = 16$
Attend-to-Dark	15.89	17.31	4.43	3.61
Attend-to-Light	17.38	19.13	4.84	4.13
Attend-to-All	14.36	15.19	2.71	3.56
Mean	15.87	17.21		

Table 3. Means and standard deviations (in pixels) of the standard deviation σ of internal response noise in centroid judgments across 11 observers.

in noise between the *Attend-to-All* vs. the *Attend-to-Light* conditions is also highly significant ($F(1, 10) = 20.820$, $p = 0.001$).

Note that if the noise that corrupted performance entered, for example, in the form of random mislocalization of the individual dots in the internal representation used by the observer to compute the centroid, and if the standard deviation of these random dot mislocalizations were equal across the $N = 8$ and $N = 16$ conditions, then the averaging performed in the 16-dot conditions would be expected to yield response errors substantially smaller than those observed in the 8-dot conditions. Specifically, we should expect the values of $\hat{\sigma}_{8,A}$ to be approximately equal to $\sqrt{2} \times \hat{\sigma}_{16,A}$. As observed above, we find on the contrary that $\hat{\sigma}_{8,A}$ is slightly but significantly *smaller* than $\hat{\sigma}_{16,A}$. The simplest account of this finding is that the dominant source of noise compromising performance is late noise, e.g., motor noise or random, perceptual mislocalization of the (noiselessly computed) centroid prior to response production. Assuming this is true, the current results imply that increasing the number of dots from 8 to 16 produces a slight but significant increase in this late noise.

Response errors for equal numbers of dots

Note that in the *Attend-to-Light* and *Attend-to-Dark* conditions using 16 dots, only 8 dots are supposed to contribute to the target centroid. The same is true of the *Attend-to-All* condition using 8 dots. Thus the difference between $\hat{\sigma}_{16,\text{Attend-to-Light}}$ and $\hat{\sigma}_{16,\text{Attend-to-Dark}}$ provides a measure of the cost in additional internal response noise incurred in attempting to ignore the irrelevant dots in the *Attend-to-Light* and *Attend-to-Dark* conditions. As might be expected, given the results noted above, the paired comparison *t*-test results in Table 4 reveal that $\hat{\sigma}_{16,\text{Attend-to-Light}}$ and $\hat{\sigma}_{16,\text{Attend-to-Dark}}$, suggesting that the presence of the to-be-ignored dots and the task of ignoring them in the *Attend-to-Light* and *Attend-to-Dark* conditions introduces increased levels of internal response noise into the centroid calculation process.

Distance-distortion effects

The median is considered a "robust" measure of central tendency because it is not influenced by the extreme values in a data set. We anticipated that observers might use a robust centroid computation—a statistic that is not greatly influenced by extreme values. Alternatively, observers might draw a virtual polygon around the display (including only the convex hull of the dot cloud) and use the center of the virtual polygon as the judged centroid. The polygon algorithm exclusively weights extreme values and is anti-robust. Our aim was not to make an accurate model of centroid judgments (which would be quite complex) but to reduce the unexplained variance in the data so as to provide more accurate estimations of attention filters. For this purpose, it was sufficient to include one additional parameter, $\alpha_{N,A}$, in the model to characterize distance distortions. This elaborated model enables us to determine whether our observers use computations that are robust or anti-robust, i.e., to test the hypothesis that observers differentially weighted peripheral versus central elements in the cloud (for computational details, see Appendix A).

To estimate the distance-distortion parameter α, we use a two-pass fitting procedure: First, we fit an attention filter under the assumption that the observer is using a center of gravity computation free of distance distortion. This yields an estimated centroid for each display from which we derive a distance-from-centroid $d(z)$ for every dot z in each

Conditions compared	*t*	*df*	Significance (two-tailed)
Attend-to-All (8)–Attend-to-Dark (16)	4.011	10	$p = 0.0012$
Attend-to-All (8)–Attend-to-Light (16)	4.669	10	$p = 0.0004$

Table 4. Comparison of the estimated magnitude of internal response noise in inclusive versus selective attention conditions: Attend to 8 of 8 versus Attend to 8 of 16.

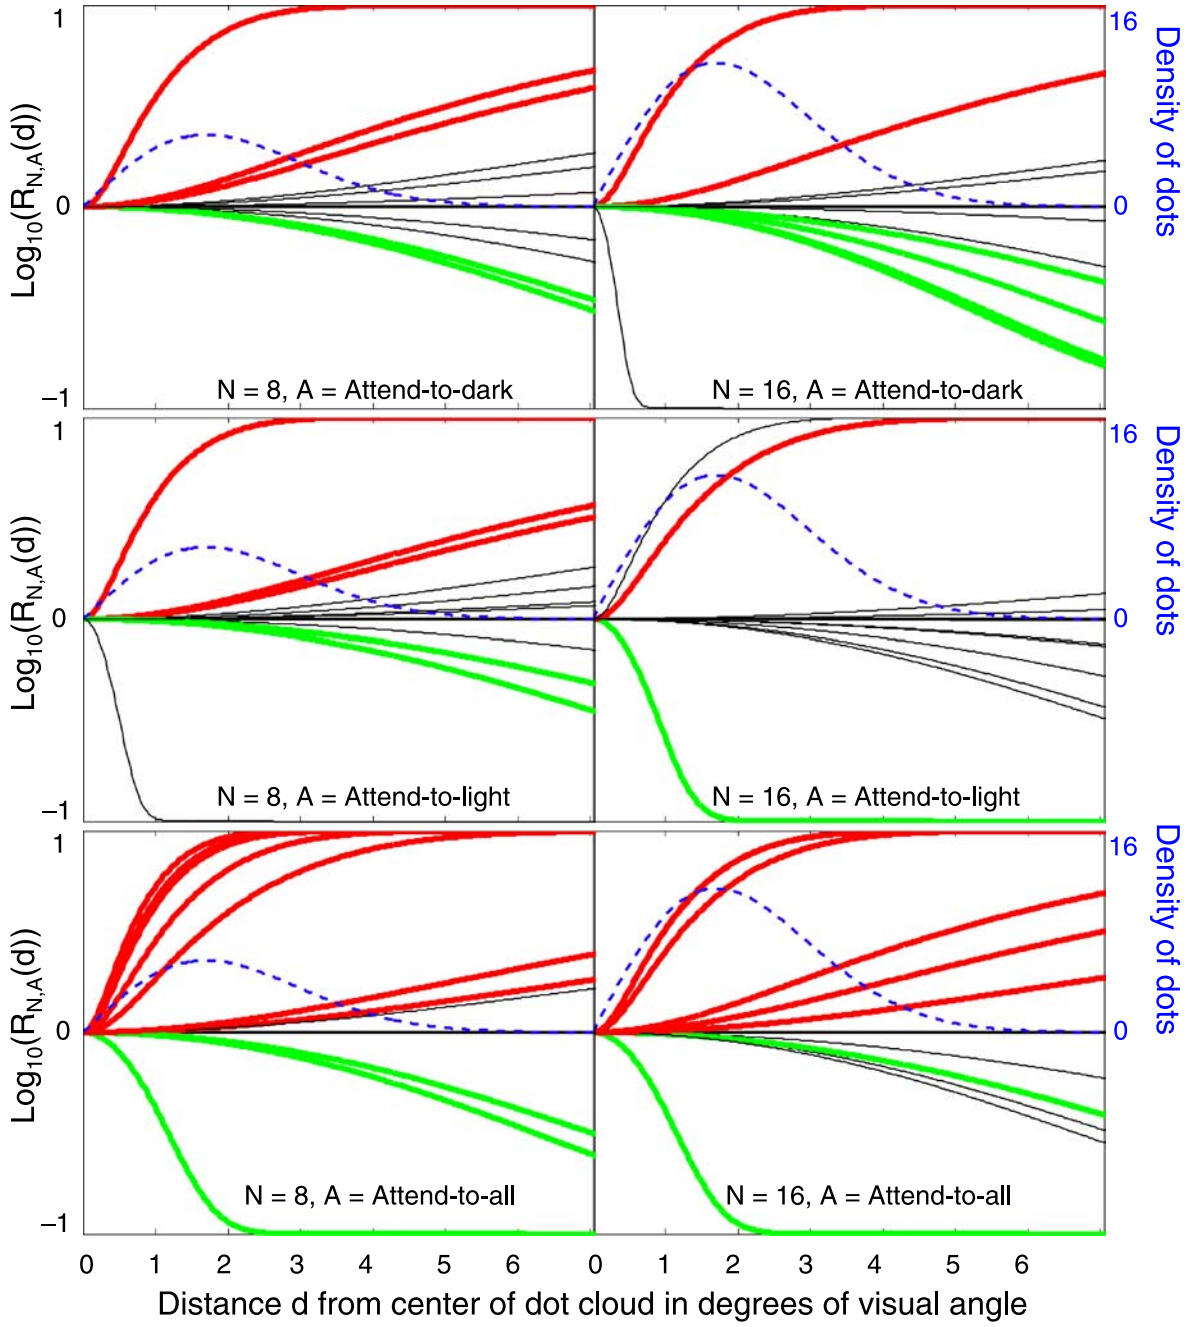

Figure 8. The log of the distance distortion function $R_{N,A}$ as a function of dot peripherality for 11 observers in each of 6 conditions (N, A). Green lines indicate statistically significant levels of robustness. Red lines indicate significant levels of anti-robustness. The blue curve gives the expected number of dots at each distance in the condition shown in the given panel. The vertical scale for the blue line is given on the right side of the panel.

display. In Pass 2, we fit a refined model to the same data in which the weight exerted by each dot becomes a product of two elements: (1) the weight attributed by the attention filter (based on the dot's contrast—as in Pass 1) and (2) a function $R_{N,A}(d(z))$ of the dot z's distance-from-centroid $d(z)$ acquired in Pass 1. $R_{N,A}$ can be either a decreasing function of distance, making the weighting a "robust" computation, or an increasing function of distance, yielding an "anti-robust" computation (see Figure 8 and Appendix A for details).

Table 5 shows the results of F-tests (df numerator = 1, df denominator = 192) assessing whether the parameter $\alpha_{N,A}$ significantly improved the model fit (the model's estimate of the centroid location for each individual stimulus) to the data (the observer's estimate of the centroid) for each observer in each condition. For each

	Attend-to-Dark		Attend-to-Light		Attend-to-All		
Observer	8-Dot Cloud	16-Dot Cloud	8-Dot Cloud	16-Dot Cloud	8-Dot Cloud	16-Dot Cloud	Measure
1	−0.568[a]	0.912	*0.958*	−0.956	*0.920*	−0.979	α
	0.000[b]	2.740	*22.200*	0.660	*9.080*	3.110	*F*
	0.991[c]	0.100	*0.000*	0.418	*0.003*	0.080	*P*
2	0.903	0.896	0.919	**−0.997**	**−0.997**	−0.968	α
	1.310	1.210	1.740	**4.520**	**7.610**	1.050	*F*
	0.254	0.272	0.189	**0.035**	**0.006**	0.308	*P*
3	**−0.980**	**−0.985**	0.841	−0.971	**−0.979**	**−0.977**	α
	7.620	**8.570**	0.340	1.200	**6.960**	**4.650**	*F*
	0.006	**0.004**	0.563	0.274	**0.009**	**0.032**	*P*
4	*0.994*	−0.972	0.890	−0.799	*0.995*	*0.993*	α
	15.280	1.050	0.600	0.000	*11.940*	*12.660*	*F*
	0.000	0.308	0.441	0.985	*0.001*	*0.001*	*P*
5	0.921	*0.965*	*0.953*	*0.992*	*0.989*	*0.967*	α
	1.300	*18.520*	*13.160*	*14.140*	*95.560*	*28.590*	*F*
	0.256	*0.000*	*0.000*	*0.000*	*0.000*	*0.000*	*P*
6	−0.971	**−0.981**	**−0.978**	−0.977	**−0.982**	**−0.997**	α
	1.320	**5.560**	**4.200**	2.080	**10.190**	**10.030**	*F*
	0.251	**0.019**	**0.042**	0.151	**0.002**	**0.002**	*P*
7	*0.966*	*0.994*	*0.995*	0.994	*0.993*	*0.992*	α
	14.290	*11.070*	*9.890*	3.870	*26.370*	*22.550*	*F*
	0.000	*0.001*	*0.002*	0.051	*0.000*	*0.000*	*P*
8	**−0.978**	**−0.984**	−0.960	−0.979	*0.996*	*0.993*	α
	10.670	**16.210**	0.960	2.690	*5.450*	*12.620*	*F*
	0.001	**0.000**	0.329	0.103	*0.021*	*0.001*	*P*
9	−0.962	**−0.975**	**−0.973**	0.779	*0.995*	*0.922*	α
	3.390	**6.580**	**6.110**	0.090	*6.620*	*5.230*	*F*
	0.067	**0.011**	**0.014**	0.771	*0.011*	*0.023*	*P*
10	*0.960*	−0.942	−0.999	0.873	0.909	−0.980	α
	16.930	0.050	2.570	0.180	1.370	3.610	*F*
	0.000	0.816	0.110	0.669	0.244	0.059	*P*
11	0.824	−0.999	0.818	−0.958	*0.940*	*0.953*	α
	0.470	1.370	0.280	0.520	*20.060*	*43.640*	*F*
	0.496	0.244	0.596	0.472	*0.000*	*0.000*	*P*

Table 5. *F*-tests assessing the improvement of model fits to data attributable to the distance-correction parameter α in every condition. *Notes*: [a]Negative values of α indicate robustness (overvaluing central versus peripheral dots); positive values indicate anti-robustness (overvaluing peripheral dots in the centroid computation). Statistically significant ($p < 0.05$) negative α values (robust) are shown and appear in bold; statistically significant positive α values (anti-robust) appear in italic; α values not statistically different from 0.0 are black. [b]Degrees of freedom for *F*: *df* numerator = 1, *df* denominator = 192. [c]Probability that α has not improved the model, i.e., level of statistical significance.

observer, the top row gives the estimated value of $\alpha_{N,A}$, which can take values between −1 and 1; negative values indicate centroid estimators that tend toward robustness; positive values indicate estimators that tend toward anti-robustness. The second row gives the F-value, and the third row gives the corresponding p-value. Cells that appear in bold give conditions in which the centroid computation was robust at the $p < 0.05$ level of significance, and cells that appear in italic give conditions in which the centroid computation was anti-robust. (It should be noted that whether or not the parameter $\alpha_{N,A}$ improves the model fit significantly in a given instance depends on other factors in addition to value of $\alpha_{N,A}$ itself; thus, for example, it might happen for some observer that $0 < \alpha_{A_1,N_1} < \alpha_{A_2,N_2}$, yet α_{A_1,N_1} improves the fit significantly in condition (N_1, A_1), whereas α_{A_2,N_2} does not produce a significant improvement in the fit in condition (N_2, A_2).)

Interestingly, out of 66 possible cases (11 observers × 2 dot sizes × 3 attention conditions), 22 showed anti-robust computations with $p < 0.05$, whereas only 14 showed robust computations with $p < 0.05$. Note in addition that individual observers tended to be consistent in direction of their distance-distortion bias: each row of Table 5 tends to contain exclusively bold cells or italic cells. Only two observers show a mixture both. Figure 8 shows the best fitting distance correction functions for each of the observers in each condition (N, A). See Appendix A for the definition of $R_{N,A}$. The dashed blue curve is proportional in height to the expected number of dots at each distance in the condition corresponding the given panel; it indicates the distance values where α most influences the centroid computation.

Discussion

Although centroid extraction has been shown to occur automatically given brief displays (Zhou, Chu, Li, & Zhan, 2005), the current results demonstrate that these computations are susceptible to top-down influences, i.e., to attention instructions given at the beginning of a session. Observers are able to base their centroid computations on (a) all the dots in the display, (b) only those dots brighter than the background, or (c) only those dots darker than the background.

Minimal models

Only one default attention filter

Suppose our observers had only one default attention filter that they applied in all three attention conditions. The most cursory inspection of the data excludes that possibility.

Two attention filters

As a working hypothesis, suppose that human vision is equipped with two, hard-wired, preattentive filters, F^+ and F^-, differentially sensitive to dots of different Weber contrasts. Under this hypothesis, each of these two filters operates in a strictly bottom-up fashion, i.e., the sensitivity of each of these two filters is immutable and completely beyond the reach of attention control. One filter F^+ is strongly activated (approximately equally well) by dots of positive Weber contrast but only slightly activated by dots of negative Weber contrast. This is the filter used under the "Attend-to-Light" instruction. The second filter F^- is strongly activated (approximately equally well) by dots of negative Weber contrast but only slightly by dots of positive Weber contrast and used under "Attend-to-Dark" condition. The observer synthesizes the filter to be used in "Attend-to-All" by assigning maximal gain to both filters. The resulting attention filter would have approximately equal sensitivity to dots of all contrasts, positive and negative (Figure 7a).

To investigate the two-attention-filters hypothesis, we submitted the data to a further analysis. This analysis focused on the data from all three attention conditions (N, *Attend-to-All*), (N, *Attend-to-Dark*), and (N, *Attend-to-Light*) separately for each observer and each of $N = 8$ and $N = 16$. The analysis compared four nested models that we will refer to as Model_0, Model_1, Model_2, and Model_3. Each model assumed a fixed distance correction function $R_{N,A}$ across all three attention conditions (i.e., $\alpha_{N,\text{Attend-to-Dark}} = \alpha_{N,\text{Attend-to-Light}} = \alpha_{N,\text{Attend-to-All}_\iota} = \alpha$).

Model_0 has no free parameters. This model assumes that the subject computes a true centroid (no distance distortion) giving equal weights to all dots (this would be the optimal single filter to use if you had to pick one filter to deal with all three attention conditions).

Model_1 still assumes that the subject has to use the same filter for all three attention conditions but allows this filter to vary freely and also allows distance distortion (α is free to roam). Thus, Model_1 has 8 free parameters.

Model_2 allows the three attention filters to be drawn from a 2-dimensional (instead of a 1-dimensional) subspace of functions. This allows for two arbitrary linearly independent filters; any third filter must be a combination of the first two. In addition to letting α to vary, this model has 7 free parameters for $F_{N,\text{Attend-to-Dark}}$, 7 for $F_{N,\text{Attend-to-Light}}$, and one additional parameter θ used to generate $F_{N,\text{Attend-to-All}} = \cos(\theta) \times F_{N,\text{Attend-to-Dark}} + \sin(\theta) \times F_{N,\text{Attend-to-Light}}$. This gives a total of 16 free parameters.

Model_3 places no constraints on the three attention filters. This model has 22 free parameters: α and 7 degrees of freedom for each of the three attention filters.

The results of the analysis of the four models are summarized in Table 6. The main thing to note about Table 6 is that percent of variance accounted for by allowing three instead of only two attention filter dimensions (column three of Table 6) tends to be quite small in comparison the

Observer	*N*	Percent of variance above $Model_0$ accounted for by $Model_1$ (first filter dimension)	Percent of additional variance accounted for by $Model_2$ (second filter dimension)	Percent of additional variance accounted for by $Model_3$ (third filter dimension)	*p*-value for *F*-test[a] assessing significance of $Model_3$ vs. $Model_2$ (third vs. second filter dimension)
1	8	1.9%	49.6%	1.3%	**0.018**
	16	7.5%	24.4%	3.3%	**0.000**
2	8	4.7%	27.3%	0.7%	0.382
	16	1.6%	15.6%	1.2%	0.213
3	8	10.4%	25.1%	0.7%	0.406
	16	7.0%	20.2%	1.3%	0.106
4	8	9.0%	3.1%	1.6%	0.094
	16	5.8%	2.6%	1.1%	0.330
5	8	5.5%	20.4%	1.8%	**0.024**
	16	6.1%	12.1%	2.7%	**0.003**
6	8	3.9%	21.8%	0.9%	0.289
	16	6.6%	12.1%	1.6%	0.079
7	8	9.8%	11.1%	1.6%	0.070
	16	6.7%	10.3%	1.0%	0.314
8	8	2.7%	44.8%	1.2%	**0.039**
	16	2.5%	29.8%	1.1%	0.101
9	8	2.4%	51.3%	1.1%	**0.026**
	16	1.1%	33.8%	0.5%	0.566
10	8	3.7%	21.0%	0.4%	0.771
	16	3.8%	15.4%	0.7%	0.528
11	8	2.6%	50.5%	0.8%	0.126
	16	2.7%	23.1%	2.9%	**0.001**
Mean	8	5.1%	29.6%	1.1%	0.204
	16	4.7%	18.1%	1.6%	0.204

Table 6. Percent of variance accounted for by assuming different numbers of attention filter dimensions available to the observer. *Notes*: [a]*F*-test degrees of freedom: 6 (numerator) and 578 (denominator).

percent of variance accounted for by allowing two attention filter dimensions instead of only one (column two of Table 6). Note, however, that including three instead of only two attention filter dimensions in the model does significantly improve the fit of the model in some cases (7 out of 22).

Another interesting finding is that the effect of including a second filter dimension in the model is much more dramatic for stimuli comprising only 8 vs. 16 dots. We do not yet understand the basis of this difference.

Model comparison summary

Although it is quite possible that humans might have more than two attention filter dimensions for Weber contrast, two attention filters are sufficient to account for most of the accounted-for variance in the three attention conditions of the present data.

Imperfect centroid computation

Table 5 showed that many observers do not have a perfect centroid computation in all conditions, and that adding a parameter α produces a statistically significant improvement in the fit of the model to the data. So, although the model may seem complicated, all the elements are essential and statistically justified.

Internal noise

Even with a full description of the imperfect attention filters and an imperfect centroid computation, the model does not perfectly predict the observers' responses. This residual prediction error is conceptualized here as due to internal response noise although internal noise undoubtedly includes a component of model failure.

Model insights

Limitations of a two-filter model

If F^- and F^+ were the only human preattentive filters sensitive to dots of different Weber contrasts, then it would follow that the only attention filters people could achieve would be combinations of F^- and F^+. Of course, if human vision had other filters selectively sensitive to

dots of different Weber contrasts, then observers should be able to achieve filters outside the plane of filters spanned by F^- and F^+. Research is under way to address this question.

Costs of selective attention

The current results show that costs are incurred when observers attempt to attend only to the dots of a single contrast polarity. First (as shown in Figures 7B and 7C), the to-be-ignored dots are not fully discounted; in all cases, they enter positively (although with much less weight than the to-be-attended dots) into the centroid computation. Second, the random noise injected into the centroid computation is significantly lower in the 8-dot, *Attend-to-All* condition than it is in either of the 16-dot, *Attend-to-Dark* conditions even though eight dots are counted in all of these conditions.

How should we think about these costs in terms of the "gain-adjustment" hypothesis? First, filter error costs (i.e., deviations of the filters achieved by observers from the target filters used to provide them feedback) are easy to understand. Under the "gain-adjustment" hypothesis, observers can only achieve attention filters that lie within the subspace spanned by the sensitivity functions of the preattentive filters hard-wired in their visual systems. When the target filter required by an attention instruction A projects only weakly into this space of achievable filters, then we expect the filter combination used by the observer to correlate poorly with the target filter, leading to a high value of filter error $RMSFE_{N,A}$.

A more perplexing question is: how should we understand the variations in internal response noise observed across attention instructions. Specifically, why should there be more internal response noise in the 16-dot, *Attend-to-Light* and *Attend-to-Dark* conditions than in the 8-dot *Attend-to-All* condition? A model in which some fraction of stimulus dots was lost, i.e., not represented perceptually, and in which the fraction of lost dots increased with the number of stimulus dots, would seem to have more internal noise for 16 than 8 presented dots. In other respects, the class of dot-loss model is very difficult to distinguish from the class of dot-weighting models that we consider here. Additional studies are required to resolve these complex issues.

Summary and conclusions

Eleven observers judged the centroids of briefly flashed clouds of dots varying in contrast from dark black to bright white on a neutral gray background. Observers attended either to only dots darker than the background, only to dots lighter than the background, or to all the dots. Observers succeeded quite well in these tasks: attended dots typically had 4× the weight of unattended dots in the centroid computation. The main results were accounted for by a model that assumed observers had only two attention filters, one preferentially sensitive to contrasts darker than the background, another to contrasts lighter than the background; the combination of both attention filters was used when attending to all the dots.

Other results: In 7 of 22 number of dots × observer conditions, there was statistically significant evidence of a weak, third attention filter. Internal noise that produces random position errors in the judged centroids is better modeled as "late" noise (after a preliminary centroid has been computed) versus early noise (independent errors in the representation of dot locations). There was significantly more internal noise and a less precise attention filter when observers attended to 16- versus 8-dot clouds and when observers selectively attended to 8 of 16 dots versus attending to all dots in an 8-dot cloud. In slightly more than half the 66 conditions, there was evidence of statistically significant distance distortion in the centroid computations: in 22 cases, peripheral dots were overweighted relative to central dots; in 14 cases, central dots were overweighted relative to peripheral dots.

Appendix A

Model

Data from each observer were obtained for dot-cloud sizes $N = 8$ and $N = 16$ in each of the three attention instructions (A = Attend-to-Dark, A = Attend-to-Light, and A = Attend-to-All). We will refer to the condition in which N dots were displayed and the attention instruction was A as condition (N, A). The data from condition (N, A) consist of the x- and y-locations, $X_{N,A}(t)$ and $Y_{N,A}(t)$, of the observer's response across all trials t. It is assumed that on a given trial the observer's response (the location he/she clicks on) is an estimate of the centroid of the dots in the display following the application of an attention filter differentially sensitive to different dot Weber contrasts.

We describe the model more precisely using the following notation. Let $F_{N,A}$ be the attention filter used by the observer in condition (N, A). That is, for any Weber contrast c, $F_{N,A}(c)$ gives the weight assigned to a dot z of contrast c in the observer's centroid estimates in condition (N, A). In this case, on trial t of condition (N, A) the x- and y-locations of the actual centroid of the attention-filtered display are

$$\bar{x}_{N,A}(t) = \frac{\sum F_{N,A}(c(z))x(z)}{\sum F_{N,A}(c(z))} \quad \text{and} \quad \bar{y}_{N,A}(t) = \frac{\sum F_{N,A}(c(z))y(z)}{\sum F_{N,A}(c(z))}, \tag{A1}$$

where each summation ranges over all dots z presented on trial t, and $x(z)$, $y(z)$, and $c(z)$ are the x-location, y-location, and Weber contrast of z, respectively.

We expect the x- and y-locations, $X_{N,A}(t)$ and $Y_{N,A}(t)$, of the observer's centroid estimates to deviate from $\bar{x}_{N,A}(t)$ and $\bar{y}_{N,A}(t)$ for two reasons; first, the responses will be degraded by noise. Additionally, we expect that observers may not combine information about disparate dots precisely according to Equation A1. We anticipate that the weight given a particular dot z may deviate from Equation A1 depending on the distance of z from the true centroid. Specifically, we assume that the x- and y-locations of the observer's response on trial t in condition (N, A) are given by Equation 2 where the weight $w(z)$ assigned to a given dot z on trial t is

$$w(z) = F_{N,A}(c(z)) \times R_{N,A}\left(\sqrt{(x(z)-\bar{x}_{N,A}(t))^2 + (y(z)-\bar{y}_{N,A}(t))^2}\right), \tag{A2}$$

for $R_{N,A}$ as a function intended to capture biases in the weighting of central versus peripheral points in the dot cloud. Specifically, we estimate a parameter $\alpha_{N,A}$ to determine the distance correction factor $R_{N,A}$ as follows:

$$R_{N,A}(\text{distance}) = \begin{cases} 0.1 + 0.9 \times \exp\left[-\left(\dfrac{\text{distance}}{10000 \times (1+\alpha_{N,A})}\right)^2\right] & -1 < \alpha_{N,A} < 0 \\ 10.0 - 9.0 \times \exp\left[-\left(\dfrac{\text{distance}}{10000 \times (1-\alpha_{N,A})}\right)^2\right] & 0 \le \alpha_{N,A} < 1 \end{cases}. \tag{A3}$$

$R_{N,A}$ is defined for $-1 < \alpha_{N,A} < 1$. For $\alpha_{N,A} < 0$, $R_{N,A}$ is a decreasing function of the distance of a dot from $(\bar{x}_{N,A}(t), \bar{y}_{N,A}(t))$; thus, in this case, the bias introduced by $R_{N,A}$ yields a "robust" estimator of the centroid—i.e., an estimator that gives most weight to dots near the middle of the cloud and is relatively immune to dots located in the periphery. The reverse is true if $\alpha_{N,A} > 0$; in this case, $R_{N,A}$ is an increasing function of distance, yielding an "anti-robust" estimator of the centroid that gives more weight to peripheral than to central dots in the cloud. In actuality, for $\alpha_{N,A}$ in a large neighborhood around 0, $R_{N,A}$ is essentially constant across all dot distances in any of our displays; specifically, for $\alpha_{N,A}$ greater than around 0.025 and less than around 0.975, dots contribute to the centroid computation with weights that are roughly independent of their distance (in pixels) from $(\bar{x}_{N,A}(t), \bar{y}_{N,A}(t))$; that is, the computation is approximately a standard centroid (as given by Equation A1).

In all conditions, we fix $F_{N,A}(0) = 0$ and constrain $F_{N,A}$ to satisfy Equation 3. Under these constraints, for each condition (N, A), the filter $F_{N,A}$, and the value of $\alpha_{N,A}$ are chosen to minimize the sum of squared distances of the responses estimated by the model ($\hat{X}_{N,A}(t)$ and $\hat{Y}_{N,A}(t)$) from the observed responses. That is, we minimize Equation 4.

Note that if the observer were able to perfectly achieve all of the attention goals in this experiment, then he/she should have $\alpha_{N,A} = 0$ for each condition (N, A), and in addition, for any dot contrast c,

$$F_{N,Attend-to-Dark}(c) = F^{Target}_{Attend-to-Dark}(c) = \begin{cases} 1/4 & \text{if } c = -1, -0.75, -0.5, -0.25 \\ 0 & \text{otherwise} \end{cases}, \tag{A4}$$

$$F_{N,Attend-to-Light}(c) = F^{Target}_{Attend-to-Light}(c) = \begin{cases} 1/4 & \text{if } c = 0.25, 0.5, 0.75, 1 \\ 0 & \text{otherwise} \end{cases}, \tag{A5}$$

$$F_{N,Attend-to-All}(c) = F^{Target}_{Attend-to-All}(c) = \begin{cases} 1/8 & \text{if } c = \pm 0.25, \pm 0.5, \pm 0.75, \pm 1 \\ 0 & \text{if } c = 0 \end{cases}. \tag{A6}$$

The cost of imposing a particular attention filter is reflected in part by deviations of the parameter values achieved by the observer from the target values (the $F_{N,A}$'s given in Equations A4, A5, and A6). We use the following statistic to estimate the root-mean-square filter error that compromises performance in a given condition:

$$RMSFE_{N,A} = \sqrt{\sum (F_{N,A}(c) - F^{Target}_{A}(c))^2}, \tag{A7}$$

where the sum ranges over $c = -1, -0.75, \ldots, 1$.

Note, however, that even if an observer had $\alpha_{N,A} = 0$ and $RMSE_{N,A} = 0$, he/she might still deviate significantly trial by trial from the target centroid if his/her responses were contaminated by high levels of noise. In assessing the cost of imposing an attention filter, one must take into account not only systematic error reflected by $\alpha_{N,A}$ and $RMSFE_{N,A}$ but also such random response error. We use the statistic $\hat{\sigma}_{N,A}$ given by Equation 5 to estimate the contribution of random noise to observers' responses in a

given condition (N, A). $\hat{\sigma}_{N,A}$ is an unbiased estimate of the standard deviation of the noise component of the observer's responses.

Acknowledgments

This work is supported by NSF Award # BCS-0843897 and by NIH Award 2R01MH068004-04A2.

Commercial relationships: none.
Corresponding authors: Stefanie A. Drew, Charles F. Chubb, George Sperling.
Email: stefanie.drew@uci.edu, cfchubb@uci.edu, sperling@uci.edu.
Address: Department of Cognitive Sciences, University of California at Irvine, 3151 Social Sciences Plaza, Irvine, CA 92697, USA.

Footnotes

[1]The process the brain uses to extract centroids is undoubtedly much more complex than our one-parameter approximation; nevertheless, this approximation accounts for a significant fraction of data variance and characterizes individual differences.

[2]The "deviations of predicted from observed centroids," which are here modeled as internal response noise, consist of two components: "true internal noise" and "model prediction error." These cannot be separated in the present experiments. Experiments using a double-pass procedure (Drew, Chubb, & Sperling, 2009) are underway to determine the relative contributions of these two factors to deviations of predicted from observed centroids.

[3]Alternatively, we could have used the sum of absolute differences between the attention filter achieved by the observer and the target attention filter. The correlation between this measure and the measure we used ($RMSFE_{N,A}$) is 0.99.

References

Baud-Bovy, G., & Soechting, J. (2001). Visual localization of the centre of mass of compact, asymmetric, two-dimensional shapes. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 692–706.

Bergen, J. R., & Adelson, E. H. (1988). Early vision and texture perception. *Nature, 333,* 363–364.

Broadbent, D. E. (1954). The role of auditory localization in attention and memory span. *Journal of Experimental Psychology, 47,* 191–196.

Cherry, E. C. (1953). Some experiments on the recognition of speech with one and two ears. *Journal of the Acoustical Society of America, 25,* 975–979.

Cohen, E. H., Schnitzer, B. S., Gersh, T. M., Singh, M., & Kowler, E. (2007). The relationship between spatial pooling and attention in saccadic and perceptual tasks. *Vision Research, 47,* 1907–1923.

Drew, S., Chubb, C., & Sperling, G. (2009). Quantifying attention: Attention filtering in centroid estimations [Abstract]. *Journal of Vision, 9*(8):229, 229a, http://www.journalofvision.org/content/9/8/229, doi:10.1167/9.8.229.

Driver, J. (2001). A selective review of selective attention research from the past century. *British Journal of Psychology, 92,* 53–78.

Foster, D. H., & Ward, P. A. (1991). Asymmetries in oriented-line detection indicate two orthogonal filters in early vision. *Proceedings of the Royal Society of London B, 243,* 75–81.

Friedenberg, J., & Liby, B. (2002). Perception of two-body center of mass. *Perception & Psychophysics, 64,* 530–539.

James, W. (1918). Attention. In *The principles of psychology* (p. 402). New York: Henry Holt.

Pashler, H. E. (1997). *The psychology of attention*. Cambridge, MA: The MIT Press.

Rubin, T. N., Chubb, C. F., Wong, S. A., & Sperling, G. (2008). Spatiotemporal dynamics of the perception of dot displays [Abstract]. *Journal of Vision, 8*(6):282, 282a, http://www.journalofvision.org/content/8/6/282, doi:10.1167/8.6.282.

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs, 74,* 1–30.

Wolfe, J. M. (1998). Visual search. In H. Pashler (Ed.), *Attention*. London: University College London Press.

Yantis, S. (2000). Goal-directed and stimulus driven determinants of attentional control. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes* (pp. 73–103). Cambridge, MA: The MIT Press.

Zhou, X., Chu, H., Li, X., & Zhan, Y. (2005). Center of mass attracts attention. *Neuroreport, 17,* 85–88.