Attentional Enhancement via Selection and Pooling of Early Sensory Responses in Human Visual Cortex

Franco Pestilli,[1,2,3,4,*] Marisa Carrasco,[2] David J. Heeger,[2] and Justin L. Gardner[1,2,*]

[1]Gardner Research Unit, RIKEN Brain Science Institute, 2-1 Hirosawa, Wako, Saitama 315-0198, Japan
[2]Department of Psychology and Center for Neural Science, New York University, New York, NY 10003, USA
[3]Mahoney-Keck Center for Brain and Behavior Research, Department of Neuroscience, Columbia University, 1051 Riverside Drive, Unit 87, New York, New York 10032, USA
[4]Present address: Department of Psychology and Stanford Center for Cognitive and Neurobiological Imaging, Stanford University, Stanford, CA 94305, USA
*Correspondence: fp2190@brain.riken.jp (F.P.), justin@brain.riken.jp (J.L.G.)
DOI 10.1016/j.neuron.2011.09.025

SUMMARY

The computational processes by which attention improves behavioral performance were characterized by measuring visual cortical activity with functional magnetic resonance imaging as humans performed a contrast-discrimination task with focal and distributed attention. Focal attention yielded robust improvements in behavioral performance accompanied by increases in cortical responses. Quantitative analysis revealed that if performance were limited only by the sensitivity of the measured sensory signals, the improvements in behavioral performance would have corresponded to an unrealistically large reduction in response variability. Instead, behavioral performance was well characterized by a pooling and selection process for which the largest sensory responses, those most strongly modulated by attention, dominated the perceptual decision. This characterization predicts that high-contrast distracters that evoke large responses should negatively impact behavioral performance. We tested and confirmed this prediction. We conclude that attention enhanced behavioral performance predominantly by enabling efficient selection of the behaviorally relevant sensory signals.
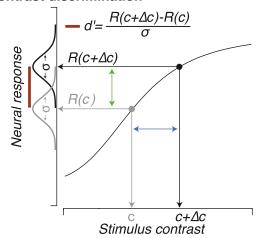
INTRODUCTION

Spatial attention allows us to see better by enhancing behavioral sensitivity and is associated with increased neural activity in early visual cortex. But what is the relation between these changes in behavioral performance measured psychophysically and the changes in neural activity measured physiologically? A satisfactory answer remains elusive because there is a paucity of empirical studies in which enhancement in behavioral sensitivity has been quantitatively linked to concurrently measured changes in neural activity (Cohen and Maunsell, 2009; Cook

and Maunsell, 2002; Sapir et al., 2005). Many psychophysical studies have documented changes in contrast sensitivity with attention but without measuring corresponding changes in neural activity (Carrasco et al., 2000; Lee et al., 1999; Lu and Dosher, 1998; Morrone et al., 2002; Pestilli et al., 2009). Single-unit monkey physiology (Martinez-Trujillo and Treue, 2002; McAdams and Maunsell, 1999; Mitchell et al., 2009; Reynolds and Heeger, 2009; Reynolds et al., 2000; Williford and Maunsell, 2006) and human neuroimaging (Buracas and Boynton, 2007; Li et al., 2008; Murray, 2008) studies have reported various effects of attention on neural response amplitudes and variability. These studies, however, have not quantitatively assessed whether measured neural changes could fully account for the improved behavioral performance with attention. Understanding how changes in cortical activity give rise to enhanced behavioral sensitivity requires concurrent measurements of behavioral sensitivity and cortical responses during tasks for which models can quantitatively link the two measurements.
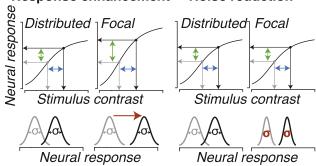
Contrast discrimination is a standard task for which plausible linkage hypotheses exist to relate amplitude and variability of neural responses in early sensory areas to behavioral sensitivity (Boynton et al., 1999; Geisler and Albrecht, 1997; Legge and Foley, 1980; Nachmias and Sansbury, 1974; Zenger-Landolt and Heeger, 2003). Neural responses in early visual cortex increase monotonically with contrast (see Figure 1A for an idealized example; Albrecht and Hamilton, 1982; Boynton et al., 1999; Zenger-Landolt and Heeger, 2003), suggesting that the brain can discriminate differences in contrast (Figure 1A, blue arrows) by monitoring differences in stimulus-evoked response amplitudes (Figure 1A, green arrows; Boynton et al., 1999; Legge and Foley, 1980; Nachmias and Sansbury, 1974; Zenger-Landolt and Heeger, 2003). According to this linkage hypothesis, attention may improve discrimination performance by increasing the slope of the contrast-response function: we refer to this as "response enhancement" (Figure 1B). Response enhancement would increase the difference in neural responses for the two corresponding contrasts and, therefore, improve discriminability (d'). Attention may also improve discrimination by reducing the noise in the sensory responses; we refer to this as "sensory noise reduction" (Figure 1C). Neural responses are inherently variable; slightly different responses are evoked on each presentation of

A Contrast discrimination

B Response enhancement C Noise reduction

D Efficient selection

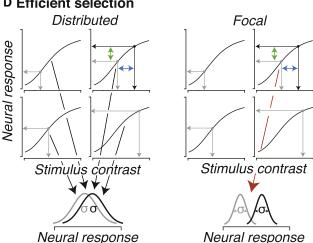**Figure 1. Signal Detection Theory, Sensitivity Enhancement, and Efficient Selection**

(A) Contrast-discrimination performance depends on the contrast-response function (black curve) and variability of response. Stimuli of contrast c and $c+\Delta c$ evoke neural responses with mean amplitude $R(c)$ and $R(c+\Delta c)$, respectively. Across trials, presentation of the same stimuli elicits slightly different responses, depicted by the distributions on the ordinate that have standard deviation σ. Behavioral sensitivity (d', see equation) is theorized to be equal to the difference between the responses to the two stimuli (green arrow) divided by σ.

the same stimulus, resulting in response distributions that can be characterized by their standard deviations (σ). Sensory noise reduction (smaller σ), which can be achieved either by reducing response variability in individual neurons and/or by reducing correlated noise across a population of neurons, would result in less overlap between two response distributions and would increase signal discriminability. Both these possibilities would increase contrast-discrimination performance with attention by improving the sensory representation—what we refer to as "sensitivity enhancement."

Attention may also improve behavioral performance by excluding irrelevant sensory signals from the decision process—what we refer to as "efficient selection." If attention were distributed across multiple stimuli (Figure 1D, *Distributed* condition), signals from relevant and irrelevant locations would be pooled together resulting in a large response variance, diluting the response differences between stimuli, and reducing stimulus discriminability. If, instead, attention were directed only to the target stimulus (Figure 1D, *Focal* condition), and if doing so selected only the relevant sensory signals (red arrow), then behavioral performance would be improved. Psychophysical experiments suggest that the effect of attention can be described by a class of pooling rules by which decisions are based on the neuronal subpopulations (or psychophysical channels) with the largest responses (Eckstein et al., 2000; Palmer et al., 2000; Pelli, 1985). Under such pooling rules, increasing responses to attended stimuli would improve performance by selecting those stimuli for decision and action.

Sensitivity enhancement and efficient selection are not mutually exclusive, and the degree to which each could, in principle, account for behavioral enhancement depends on what limits performance in any given task. We measured concurrently the psychophysical and physiological effects of spatial attention in a task that required high sensory discrimination and included multiple stimuli, thus potentially allowing attention to act via either, or both, sensitivity enhancement and efficient selection. By quantitatively linking the psychophysical and physiological measurements, using models of sensitivity enhancement (Figures 1B and 1C) and efficient selection (Figure 1D), we concluded that efficient selection plays the dominant role in improving visual sensitivity.

RESULTS

Psychophysical Contrast-Discrimination Functions

Contrast-discrimination thresholds were measured concurrently with fMRI responses in early visual cortex. Each trial started with either a focal or distributed attention cue (Figure 2, interval 1). This was followed by two 0.6 s stimulus presentations (Figure 2, intervals 2 and 4) of four sinusoidal gratings with eight "pedestal"

(B–D) Top panels show possible effects of attention on the contrast-response functions. Bottom panels illustrate distributions of responses for two stimuli of different contrasts, showing how focal attention may affect stimulus discriminability. (B) Focal attention increases the slope of the contrast-response function. (C) Focal attention reduces the trial-to-trial variability in sensory neural responses. (D) Focal attention selects the relevant sensory signal (red arrow), so irrelevant information is ignored.
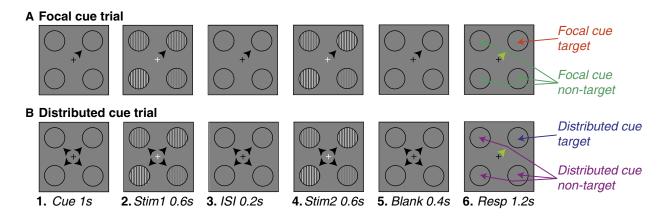
A Focal cue trial

B Distributed cue trial

1. *Cue 1s* **2.** *Stim1 0.6s* **3.** *ISI 0.2s* **4.** *Stim2 0.6s* **5.** *Blank 0.4s* **6.** *Resp 1.2s*

Figure 2. Behavioral Protocol

Observers performed a two-interval forced-choice contrast-discrimination task. On each trial, four gratings appeared in two temporal intervals (Stim1 and Stim2). Only one grating (target) had a slightly higher contrast in one of the intervals. A response cue (green arrow in Resp) appearing after stimulus offset indicated the location of the target. Observers reported the interval in which the target was higher in contrast. There were two kinds of attentional cues: (A) focal cue, a single arrow indicating the correct target location with 100% validity; and (B) distributed cue, four arrows indicating that the target was equally likely to appear in any quadrant. The protocol thus defined four stimulus cue combinations: (1) focal cue target, red arrow; (2) focal cue nontarget, green arrows; (3) distributed cue target, blue arrow; and (4) distributed cue nontarget, purple arrow.

contrasts (0%–84%). Different pedestal contrasts were selected for each of the four locations on each trial. During one of the two stimulus intervals, one of the four gratings (target, chosen at random) had a contrast increment, Δc, added to the pedestal contrast. After the second stimulus interval, a small green arrow at fixation pointed to the target location (Figure 2, interval 6), and the observers reported with a button press whether the first or second stimulus interval at that location had a higher contrast. An adaptive staircase procedure was used to find the *Δc* that resulted in 76% correct performance, i.e., the contrast-discrimination threshold (see Supplemental Experimental Procedures: Behavioral Protocol, available online). Contrast-discrimination thresholds were determined separately for each of the eight pedestal contrasts and two cue conditions by running independent and randomly interleaved staircases.

The contrast-discrimination functions (Figure 3, contrast-discrimination threshold as a function of pedestal contrast) had characteristics consistent with previous findings. First, as pedestal contrast increased from 1.75% to 28%, thresholds monotonically increased. This behavior is reminiscent of Weber's law, which predicts that discrimination thresholds maintain a constant ratio with the stimulus intensity (a slope of 1 plotted on a log-log axis). We found slopes <1 (blue curve, distributed cue, target stimulus, 0.73 ± 0.04; red curve, focal cue target stimulus, 0.78 ± 0.08; mean ± standard error of the mean [SEM] across observers), consistent with previous studies (Gorea and Sagi, 2001). Second, thresholds decreased for lower pedestal contrasts, resulting in a characteristic dipper shape of the contrast-discrimination function (Legge and Foley, 1980; Nachmias and Sansbury, 1974). Because we tested a large range of mid-to-high contrasts to reliably compare any slope changes in the fMRI measurements, we did not sample low enough contrast pedestals to fully characterize the dipper (compare blue and red curves). Third, thresholds decreased above 28%–56% with a slope on a log-log axis of −2.9 ± 0.18
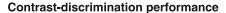
(blue curve, mean ± SEM across observers) and −3.22 ± 0.67 (red curve). This decrease in threshold at high contrast may be explained by the selection model presented below (see last section or Results).

The effect of focal attention on contrast-discrimination thresholds was characterized using spatial cues. On half of the trials, a focal cue (Figure 2A, small black arrow) was shown before the stimuli to be discriminated. This focal cue indicated the target location with 100% validity but did not provide information regarding the stimulus interval containing the higher contrast target. Observers were instructed to use this cue to direct spatial attention to the target. On the rest of the trials (randomly interleaved), a distributed cue was shown (Figure 2B, four small black arrows), which did not provide information about the target location; observers were instructed to distribute their spatial attention across the four stimuli. To minimize uncertainty about the target location, in both cases a response cue (green arrow) indicated the target location after stimuli offset. Contrast-discrimination thresholds were lower at all pedestal contrasts on focal than on distributed cue trials (Figure 3, red and blue curves, respectively). This vertical shift of the contrast-discrimination functions was reflected in smooth function fits to the data (Figure 3, solid curves; see Experimental Procedures: Psychophysical Contrast-Discrimination Functions) in which the parameter controlling vertical offset increased significantly for the distributed cue compared to the focal cue trials (g_r, p = 0.02, Student's t test across observers), but other parameters did not change (g_c, p = 0.58; s, p = 0.17; q, p = 0.4, Student's t test). Enhanced contrast discrimination could not be attributed to any change in eye position between focal and distributed cue trials (see Experimental Procedures: Eye Position Monitoring).

Testing Response Enhancement

Contrast-response functions were measured, in each of several visual cortical areas, for each of four stimulus cue combinations
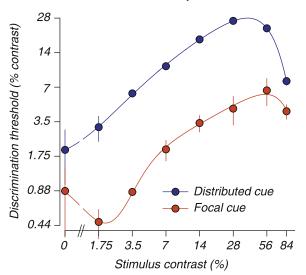
Figure 3. Contrast-Discrimination Performance
Contrast-discrimination thresholds for the focal (red) and distributed cue (blue) conditions (mean ± SEM across observers; some error bars are smaller than symbols) are plotted as a function of stimulus contrast. Thick curves are best fit of Equations 2–4.

(Figure 4): focal cue target, focal cue nontarget, distributed cue target, and distributed cue nontarget. fMRI responses increased monotonically with stimulus contrast (Figures 4A and 4B, representative observer), and depended on the stimulus cue combination (Figures 4B and 4C, representative observer and average, respectively). Response amplitudes were smallest for unattended stimuli (Figure 4C, green, focal cue nontarget), larger when attention was distributed (purple, distributed cue nontarget; blue, distributed cue target), and largest for attended targets (red, focal cue target).

There was no evidence for response enhancement in any of the visual cortical areas. We fit the data by adopting a parametric equation for the contrast-response functions (see Experimental Procedures: fMRI Contrast-Response Functions). Only one of the fitted parameter values differed significantly across the four stimulus cue combinations: the baseline response (b) that determined the vertical positions of the contrast-response functions. Allowing only this parameter to vary across cue conditions provided a fit that was statistically indistinguishable from the fit allowing all parameters (g_c, g_r, and b) to vary across cue conditions (V1, V2, V3, and hV4 each $F_{(14,8)}$, p = 0.3). Thus, we did not observe a change in gain or slope of the contrast-response functions, consistent with previous reports (Figure 1B; Buracas and Boynton, 2007; Murray, 2008). Instead, the cue effect was well described as a vertical additive shift of the contrast-response functions. The amount of additive offset increased across the hierarchy of visual cortical areas. Values for b increased from the focal cue nontarget curve to the distributed cue nontarget curve by 0.04, 0.08, 0.14, and 0.25 (percent [%] fMRI signal change) in visual areas V1, V2, V3, and hV4, respectively. The values increased from focal cue nontarget to distributed cue nontarget by 0.11, 0.18, 0.27 and 0.34, and they

increased from focal cue nontarget to focal cue target by 0.29, 0.39, 0.52, and 0.51. Therefore, the effect of attention was about 6×, 3×, or 1.75× larger in hV4 than in V1, depending on which pair of conditions was compared.

There was a small but reliable difference in responses between distributed cue target and nontarget stimuli (Figure 4C; blue and purple). Values for b in V1–hV4 differed by 0.07%, 0.10%, 0.13%, and 0.10% signal change, respectively. These response differences were evident even though these trials differed only after the stimuli had been removed from the display for 400 ms (Figure 2B), when the response cue was presented. This effect cannot be the result of differences in neural responses during the first interval because the response cue defined the target only after the second interval. Observers could have inferred the target location during the second interval, before the response cue, if they noticed where the change in contrast occurred between the two intervals. Consequently, they would have attended more to the identified target location during the second stimulus interval. However, we found no difference between correct and incorrect trials, either for the distributed cue target or for distributed cue nontarget responses (quantified by the b parameter; p > 0.1, paired Student's t test across subjects and visual areas). Thus, this small response difference likely originates from a poststimulus modulation during the response phase (Sergent et al., 2011).

Testing Sensory Noise Reduction

To test whether sensory noise reduction alone can account for enhanced behavioral performance with focal attention, fMRI and behavioral data were fit using the sensitivity model depicted in Figure 1 (see Experimental Procedures: Testing Sensory Noise Reduction). The sensitivity model fit the fMRI (contrast response) based on parameterized behavioral (contrast discrimination) data with two key parameters: the baseline response (b), and the sensory noise standard deviation (σ).

For the distributed cue condition (Figures 5A and 5B), the psychophysical contrast-discrimination data were again fit with a smooth function (Figure 5A, blue line), and then the σ and b parameters were optimized to find the best fit to the fMRI contrast-response function (Figure 5B, blue line). This procedure was repeated for each visual cortical area. The sensitivity model fit well the contrast-response measurements in each visual area (V1, r^2 = 0.95, Figure 5B; V2, r^2 = 0.97; V3, r^2 = 0.97; hV4, r^2 = 0.98; average across observers), and for each individual observer (observer 1, r^2 = 0.98; observer 2, r^2 = 0.94; observer 3, r^2 = 0.97; average across visual areas).

Having fit the sensitivity model parameters to the data in the distributed cue condition, we asked whether these parameters could account for the data in the focal cue condition. Had the slope of the contrast-response function changed in a way that could account for the behavioral data (Figure 5C), then fixing the σ and b parameters to what had been estimated in the distributed cue condition would have provided a good fit in the focal cue condition. It did not. To show this, we used the σ and b that were fitted to the distributed cue condition (Figure 5A) to predict the contrast-response function from the focal cue condition. The resulting contrast-response function had a much steeper slope than that measured in the focal
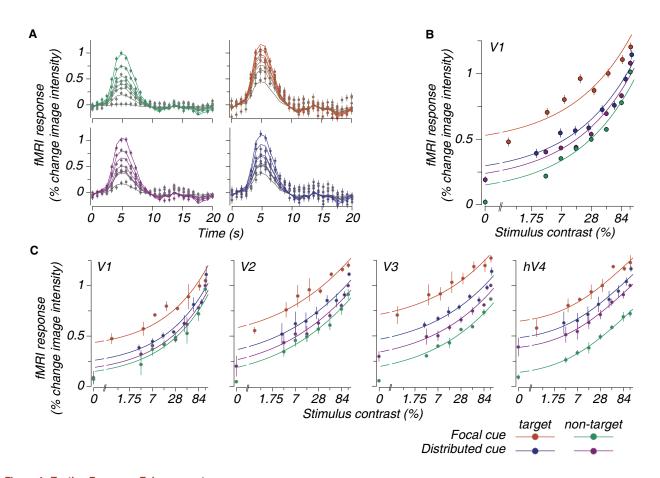
Figure 4. Testing Response Enhancement

(A) Example fMRI response time courses from one observer's V1. Colors indicate different stimulus cue combinations (see legend at bottom right). Shading indicates stimulus contrast; darker colors correspond to lower contrasts. Error bars, SEM over repeated trials.

(B) Example contrast-response functions, corresponding to the response time courses in (A). Continuous curves, best fit of Equation 3.

(C) Contrast-response functions for each visual area and stimulus cue combination averaged across observers (mean ± 1 SEM) for each visual area.

See also Figure S1.

cue condition, and did not fit well the measured contrast-response functions in any of the visual areas (V1, $r^2 = 0.58$, Figure 5D, blue curve; V2, $r^2 = 0.63$; V3, $r^2 = 0.63$; hV4, $r^2 = 0.64$; average across observers), nor for any observer (observer 1, $r^2 = 0.63$; observers 2, $r^2 = 0.59$; observer 3, $r^2 = 0.64$; average across visual areas).

Allowing the standard deviation (σ) and the baseline response (b) to be adjusted for the focal cue condition (Figure 5E) resulted in good fits to the contrast-response functions for each visual area (Figure 5F; V1, $r^2 = 0.89$; V2, $r^2 = 0.85$; V3, $r^2 = 0.89$; hV4, $r^2 = 0.83$; average across observer) and for each individual observer (observer 1, $r^2 = 0.90$; observer 2, $r^2 = 0.77$; observer 3, $r^2 = 0.91$; average across visual areas). For V1, the best-fit value of the sensory noise standard deviation (σ) was 0.085% signal change for the distributed cue and 0.016% signal change for the focal cue condition. The best-fit value of the baseline response (b) was 0.34% signal change for the distributed cue and 0.55% signal change for the focal cue condition. Thus, there was no evidence for a change in the response gain of the fMRI responses, only for a change in the sensory noise standard deviation and baseline response parameters.

A similar result was found for each visual area and observer; the ostensible effect of the focal cue was to decrease sensory noise and increase the baseline response. These two model parameters were fit separately for the distributed cue and focal cue conditions for each visual area and each observer. The average σ value for the distributed cue (σ_d) condition was 0.064% ± 0.02% and 0.016% ± 0.01% for the focal cue (σ_f) condition. The ratio of σ_d to σ_f was significantly greater than 1 in all observers and visual areas (p < 0.01, bootstrap test; see Supplemental Experimental Procedures: Statistical Tests in Individual Observers) and implied approximately a 4-fold reduction in sensory noise (Figure 6A). The average b value increased from 0.58% ± 0.02% for the distributed cue condition to 0.74% ± 0.04% for the focal cue condition (Figure 6B, b_d and b_f, respectively). The difference between b_d and b_f was significantly different from zero in all observers and visual areas (p < 0.05, except for hV4 in one observer, p = 0.38, bootstrap test).

Testing Efficient Selection

The approximately 400% reduction in σ between the distributed and focal cue conditions could be due to a decrease in early
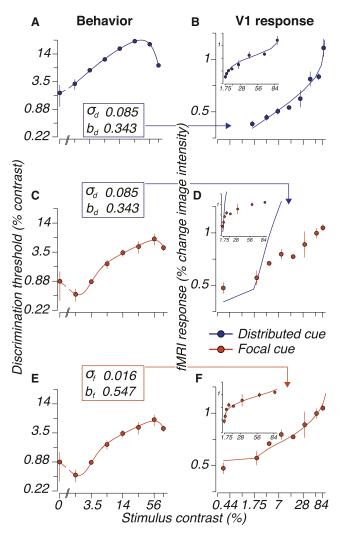
Figure 5. Testing Sensory Noise Reduction

Left column shows contrast-discrimination functions (mean ± SEM across observers) for distributed (blue, A) and focal (red, C and E) cues. Right column illustrates V1 contrast-response functions (mean ± 1 SEM across observers) for distributed (blue, B) and focal (red, D and F) cues. Contrast-response functions are plotted on log-linear axis in the large panels and on linear-linear axis in the insets. The smooth curves are fits of the model, with parameter values listed in each box (σ, sensory noise; b, baseline response).

signal-to-noise ratio, to greater inefficiencies in "reading out" the sensory signals, or to a combination of the two. Monkey electrophysiology experiments have shown that attention can reduce sensory noise, but not by such a large amount. For visual cortical area hV4, attention decreases both trial-to-trial variability in firing rates in individual neurons and trial-to-trial correlations in firing rates across neurons, such that an overall reduction in sensory noise of approximately 50% is achieved when averaged across a pool of neurons (Cohen and Maunsell, 2009; Mitchell et al., 2009). This suggests that inefficiencies in sensory pooling and decision making play a large role in explaining the difference in performance accuracy for focal and distributed cue trials.

We propose a particular example of a model that exhibits such inefficiencies, which we call the "selection model." The selection

model pools sensory responses across the four stimulus locations according to a max-pooling rule (it weighs the largest response the most). This ensures that decisions on focal cue trials are based primarily on responses to the target stimuli (which are larger than responses to nontargets because the baseline responses are larger for attended stimuli), leading to good behavioral performance. On distributed cue trials, one of the nontarget stimuli evokes the largest responses (noting that in our experimental protocol one of the nontargets typically had a higher contrast than the target). Max pooling thereby causes decisions to be based primarily on irrelevant sensory signals corresponding to incorrect locations, leading to correspondingly poor behavioral performance.
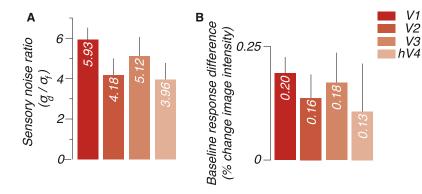
We begin by considering attentional selection via max pooling in a focal cue trial. Figure 7A shows simulation results, idealized sensory response distributions for the two intervals in the task at each of the four target locations. Each location elicited some response as measured by the contrast-response function for target and nontarget stimuli. Only the target location had an actual difference in mean response between the two intervals (because there was a contrast increment added only at this location). For these simulations, the means of the sensory response distributions in Figure 7 were set to be the mean fMRI response amplitudes (from V1) for the target and nontarget locations, and the standard deviation of the sensory response distributions was set to the best-fit value from the sensitivity model fit (see above, Testing Sensory Noise Reduction) for the focal cue condition. To readout the responses, the max-pooling operation weighted responses differently depending on their relative amplitude (Figure 7B):

$$R_p = \frac{1}{4}\sqrt[k]{\sum_{i=1}^{4} r_i^k}, \qquad (1)$$

where r_i was the response at each of the four stimulus locations, R_p was the pooled readout of the responses, and k was a model parameter that changed the pooling operation from averaging ($k = 1$) to maximizing ($k = \infty$). With a large k, the largest amplitude response dominated the readout distribution from which the decision was made. For focal cue trials, attention served to boost the target response above the nontarget responses, and therefore, the readout distribution was dominated by the response to the target (i.e., the readout distributions in Figure 7C are virtually identical to the sensory response distributions at the target location in Figure 7A).

The max-pooling rule with exactly the same k value predicted a larger threshold contrast (Δc, i.e., worse-performance) on distributed cue trials. On distributed cue trials, a much larger Δc evoked a larger sensory response difference at the target location (Figure 7A, compare sensory response distributions corresponding to the target location, top left, for focal and distributed cues). In spite of the much larger target contrast difference on distributed versus focal cue trials, and the correspondingly larger separation between the sensory response distributions at the target location, the readout distributions were virtually identical (Figure 7C, compare response distributions for focal versus distributed cues). Because of the max-pooling rule, the readout distributions were dominated by the

Figure 6. Noise Estimates and Baseline Responses Inferred from Model Fits

(A) Ratio of sensory noise on distributed cue trials (σ_d) to focal cue trials (σ_f).

(B) Difference between baseline on focal cue trials (b_f) and distributed cue trials (b_d).

Error bars, ± 1 SEM across observers.

stimulus location evoking the highest response. For focal cue trials, this was nearly always the target location. For distributed cue trials, none of the sensory responses were preferentially increased by attention, so the max-pooling rule biased the readout distributions to correspond to the stimulus with the highest contrast, which was not usually the target. A larger Δc was consequently needed in the distributed cue trials compared to focal cue trials, to get the same separation between the readout distributions and correspondingly the same performance accuracy.

Unlike the sensitivity model described above, this selection model quantitatively predicted behavioral enhancement based on the measured differences in cortical response amplitudes without any sensory noise reduction. We adjusted the k and σ parameters to fit the contrast-discrimination functions (see Experimental Procedures: Testing Efficient Selection). We used a single σ value across both focal cue and distributed cue conditions, and found that the selection model provided excellent fits (e.g., Figure 8A plots behavioral data and V1 contrast-response functions averaged across observers). Fitting the k and σ parameters across individual observers and visual areas, we found k values with a mean near the maximizing end of the spectrum ($k = 68.08$). We used an information criteria (AIC) and cross-validated r^2 to compare the quality of the model fits (see Experimental Procedures: Model Comparisons). Across all visual areas, the fits to the data averaged across observers were better (AIC difference = -23.94, -10.90, -59.88, -21.09 V1–hV4, respectively) for the selection model using a single σ value for both focal cue and distributed cue conditions (cross-validated $r^2 = 0.84$, 0.88, 0.89, 0.89, V1–hV4, respectively) compared to the sensitivity model (fit without allowing σ to vary; cross-validated $r^2 = 0.06$, 0.20, 0.13, 0.16). The selection model also provided better fits than the sensitivity model for the data from individual observers (selection model cross-validated $r^2 = 0.82$, 0.83, 0.84, 0.83, V1–hV4, respectively, computed separately for each individual observer and then averaged across observers; sensitivity model cross-validated $r^2 = 0.10$, 0.41, 0.34, 0.40; AIC difference = -23.21, -33.20, -40.26, -41.03).

We confirmed the result that the max-pooling selection rule accounted for contrast-discrimination performance, by adopting a single k value (the mean across V1–hV4) for each observer and applying it to all visual areas. It is not necessarily the case that each visual area should have exactly the same balance of maximization versus averaging as implied by a single k value.

Nonetheless, this analysis was used to test how well a fixed pooling rule could account for the behavioral data. With a fixed k for each observer, σ for the distributed and for the focal cue trials was allowed to vary separately to fit the contrast-discrimination functions. The ratio of σ_d to σ_f (1.04 ± 0.05, mean and SEM across visual areas and observers) was statistically indistinguishable from 1 ($p = 0.56$; bootstrap test), demonstrating that the selection model could account for the difference in behavioral performance between the two conditions without requiring any sensory noise reduction. This result is contrasted with the sensitivity model in which σ_d was on average about four times larger than σ_f (4.12 ± 0.23, mean and SEM across visual areas and observers; Figure 8C, a recapitulation of the result in Figure 6).

A combination of sensory noise reduction with our selection model also fit the data well. As noted above, the largest sensory noise reduction reported in the literature is about 50% (Cohen and Maunsell, 2009), but our contrast-discrimination functions were not adequately fit with a 50% sensory noise reduction, disregarding pooling of the sensory responses (Figure 8B; cross-validated $r^2 = 0.46$, 0.55, 0.52, 0.53, V1–hV4, respectively). However, this amount of sensory noise reduction when coupled with the selection model provided good fits to our data (cross-validated $r^2 = 0.89$, 0.92, 0.92, 0.92), resulting in slightly smaller k values (61.03, averaged over areas) than the selection model alone. This combined model also provided a good fit to the data from individual observers (cross-validated $r^2 = 0.86$, 0.84, 0.86, 0.85, V1–hV4, respectively, computed separately for each individual observer and then averaged across observers). This fit was virtually indistinguishable from the selection model alone (compare to selection model r^2, two paragraphs above), but it was better than the sensitivity model with noise reduction (cross-validated $r^2 = 0.40$, 0.63, 0.60, 0.63; AIC difference = -19.05, -18.06, -17.35, -23.78).

Robustness of Results

We confirmed the robustness of our conclusions via the following analyses:

(1) We removed anticipatory hemodynamic effects separately for focal cue and distributed cue trials (see Figure S2A).

(2) We varied the sizes of the regions of interest in each visual area, corresponding to each stimulus location, using different statistical thresholds (see Figure S2B).

(3) We used different functional forms (polynomial, a skewed Gaussian, and the form fit by the sensitivity model) to parameterize contrast-discrimination functions. These
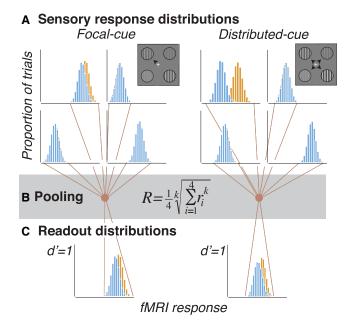
A Sensory response distributions

B Pooling

$$R = \frac{1}{4} \sqrt[k]{\sum_{i=1}^{4} r_i^k}$$

C Readout distributions

$d'=1$ $d'=1$

fMRI response

Figure 7. Selection Model
(A) Sensory response distributions. Each panel plots simulated response distributions, the proportion of trials on which a given response was evoked, for each stimulus location. Blue indicates first stimulus interval. Orange illustrates second stimulus interval (that included the contrast increment at the target location). Left panel shows focal cue trials. Right panel indicates distributed cue trials. Insets are schematics of stimulus displays and attention cues (see Figure 2).
(B) Max-pooling operation. Sensory responses to each stimulus interval were pooled across locations according to the max-pooling rule (best-fit k = 68.08).
(C) Readout distributions after pooling across locations.

changes made quantitative but not qualitative changes in the parameter estimates of the various models and did not affect the overall conclusions (see Figure S2C).

Testing High-Contrast Distracter Prediction

One important prediction of our selection model is that behavioral performance is not determined by the properties of the target alone—the configuration of distracters also affects behavior. For example, a distracter with high contrast that evokes a large response will preferentially pass through the selection mechanism and, therefore, be expected to disrupt behavioral performance more than a distracter that evokes a smaller response. We confirmed this prediction in the following two ways.

First, we found that our selection model, given the configuration of distracter contrasts in the main experiment, predicted the prominent dip at high contrast of the measured contrast discrimination functions (Figure 3). Distracter contrasts were always randomized around the target contrast. However, for the highest contrast pedestal, physical constraints (a maximum of 100% contrast is achievable) necessitated presenting lower contrast distracters. Thus, these high-contrast pedestals were paired with distracters that evoked comparatively smaller responses and, therefore, were excluded to a great extent by our selection rule. This resulted in a prediction of better perfor-

mance at high than at lower pedestal contrasts. This effect was even more pronounced given that contrast-response functions saturated at higher contrast, resulting in comparatively weaker distracter responses. Thus, our selection model predicted a prominent dip at high contrast for the distributed cue condition (Figure 8, blue curve), despite the fact that the form of the contrast-response functions used in the model fits did not include any accelerating nonlinearity at high contrast. The dip in the modeled distributed cue discrimination function was due solely to the selection mechanism excluding the smaller response of the distracters at high contrast from the readout distributions. Our selection model also predicted that the focal cue condition would be less susceptible to these distracter effects due to the enhanced response at the focal cue target (Figure 8, red curve). While our selection model overpredicts the ability of focal attention to overcome the effect of distracters (i.e., predicts no, rather than a small, dip), there was indeed a much smaller dip in the contrast-discrimination performance at high contrast for the focal cue condition (Figure 3, red curve).

As a second, more direct confirmation of the prediction of our selection model, we conducted behavioral experiments similar to the ones described above but added a second set of conditions in which we replaced the lowest contrast distracter in each condition with a distracter of 84% contrast (see Supplemental Experimental Procedures: Behavioral Protocol for details). As before, thresholds were lower for the focal cue condition than the distributed cue condition (Figure 9A); indeed, there was an ~4.2-fold difference (Figure 9B; p < 0.001, two-way nested ANOVA main effect of cue), thus replicating the behavioral effect of focal attention. As predicted, behavioral performance for both distributed cue and focal cue conditions was worse when there was a high-contrast distracter (Figure 9A; p < 0.001, two-way nested ANOVA main effect of distracter condition), with no evidence of any individual differences between subjects (p = 0.49, two-way nested ANOVA). The effect of distracter contrast was greater for the distributed cue condition than the focal cue condition (Figure 9B) as expected by our selection model, given that in the focal cue condition the target location was predicted (by the model) to have an enhanced response that could better compete with the high-contrast distracter.

DISCUSSION

The behavioral and cortical effects of attention were concurrently measured using psychophysics and fMRI, and a computational analysis was used to quantitatively link these measurements. Cortical responses in early visual areas increased when spatial attention was focused on a single location as compared to when attention was distributed across all stimuli, consistent with previous studies (Buracas and Boynton, 2007; Li et al., 2008; Liu et al., 2005; Murray, 2008). Concurrent behavioral performance also improved (contrast-discrimination thresholds decreased) when observers were cued to the target location, also consistent with previous studies (Foley and Schwarz, 1998; Lee et al., 1999; Lu and Dosher, 1998; Morrone et al., 2002; Pestilli et al., 2009). We considered whether sensitivity enhancement, in the form of response enhancement or noise
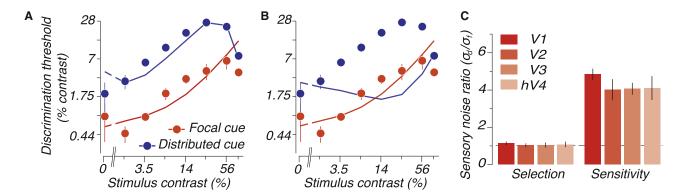
Figure 8. Comparison of Sensory Noise Reduction and Efficient Selection
(A) Contrast-discrimination functions (data replotted from Figure 3) fit with the selection model (continuous curves), constrained by the average V1 contrast-response functions (not shown, but see Figure 4). The two free parameters, k (max-pooling parameter, Equation 1) and σ (sensory noise standard deviation), were constrained to be the same for the focal and distributed cue conditions.
(B) Contrast-discrimination functions fit with the sensitivity model, again constrained by the V1 contrast-response functions. Noise standard deviation (σ) was constrained to be 50% smaller for focal than distributed cue trials.
(C) Changes in noise standard deviation (σ) for focal and distributed cue trials estimated by the selection (left) and the sensitivity (right) models. Error bars, ± 1 SEM across observers.
See also Figure S2.

reduction, and efficient selection, in the form of a max-pooling selection rule, could quantitatively link the two measurements. We concluded that efficient selection played the dominant role in accounting for the behavioral enhancement observed in the contrast discrimination task. Finally, we confirmed one prediction of our selection model, that high-contrast distracters disrupt behavioral performance.

fMRI and Neural Activity
In describing our effort to quantitatively link fMRI responses and behavioral enhancement with attention, an underlying assumption of our analysis is that the fMRI responses were approximately proportional to a measure of local average neuronal activity (Boynton et al., 1996; Heeger and Ress, 2002). It has been claimed that fMRI responses are most closely related to synaptic input and intracortical processing within a cortical area, not the spiking output (Logothetis and Wandell, 2004). Cortical circuits are, however, dominated by massive local connectivity in which most synaptic inputs originate from nearby neurons (Douglas and Martin, 2007). Thus, synaptic "inputs" in cerebral cortex are mostly produced by local spiking of neighboring neurons, leading typically to a tight coupling between synaptic and spiking activity, as well as vascular responses. It is not surprising, therefore, that fMRI responses have been found to be highly correlated with neural spiking (Heeger et al., 2000; Mukamel et al., 2005). Even suppression of neuronal activity, which probably involves an increase in synaptic inhibition, has been found to be correlated with smaller fMRI responses (Shmuel et al., 2006; Zenger-Landolt and Heeger, 2003).

There are clear demonstrations that vascular responses can be dissociated from spiking activity. A striking example of such dissociation is a spatially global anticipatory hemodynamic modulation during regularly paced trials that is not reflected in spiking activity (Sirotin and Das, 2009). Our methodology removed such anticipatory hemodynamic modulation by randomizing the inter-

trial intervals and subtracting a spatially homogenous component of the responses (see Figure S1A). After subtracting this spatially global component, the residual vascular responses are tightly linked with spiking activity, such that the magnitude of the vascular responses evoked by different stimulus contrasts is linearly proportional to the magnitude of spiking activity as assumed by our analysis (A. Das, personal communication).

Attentional Modulation of Activity in Visual Cortex: fMRI versus Electrophysiology
We considered whether potential conflicts between fMRI and single-unit measurements of the effect of attention on contrast-response suggest another possible dissociation of vascular and spiking activity. Attention has been reported to have a wide variety of effects on the contrast-response functions of neurons in visual cortex. Contrast-gain changes (Martinez-Trujillo and Treue, 2002; Reynolds et al., 2000; Williford and Maunsell, 2006), response-gain changes (Lee and Maunsell, 2010; Williford and Maunsell, 2006), activity-gain changes (Williford and Maunsell, 2006), additive offsets dependent on visibility (Pooresmaeili et al., 2010; Thiele et al., 2009), and baseline shifts in the absence of a stimulus (Reynolds et al., 2000; Williford and Maunsell, 2006) have all been observed, even different changes in different neurons during the same experiment (Williford and Maunsell, 2006).

Some of these inconsistent results from single-unit studies may be due to uncontrolled task parameters. For example, the normalization model of attention predicts different effects (response-gain changes, contrast-gain changes, or a combination of the two that can mimic a baseline shift) in different neurons depending on stimulus size and attention field size (i.e., the spatial and featural extent of attention), with respect to receptive field size (Reynolds and Heeger, 2009). To date, stimulus size and attention field size have only been manipulated systematically in one behavioral and neuroimaging study (Herrmann et al., 2010), and have not been systematically manipulated in electrophysiology
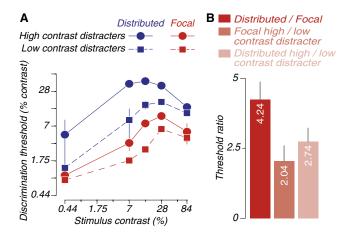
Figure 9. Testing Prediction of Selection Model that High-Contrast Distracters Disrupt Behavioral Performance

(A) Contrast discrimination thresholds (mean ± SEM across observers; some error bars are smaller than symbols) for focal (red) and distributed (blue) cue conditions as a function of distracter condition (see legend).

(B) Ratio of thresholds averaged across pedestal contrast. Error bars, SEM across observers.

experiments. In addition, task difficulty is known to modulate neuronal responses (Boudreau et al., 2006; Chen et al., 2006), and task difficulty varies with contrast (e.g., orientation discrimination is typically harder at low contrast than at high contrast; Lu and Dosher, 1998; Pestilli et al., 2009). In our experiment, separate staircases were run for each contrast, thus ensuring the same threshold level of discrimination difficulty at each contrast.

Whereas attentional modulation of single-unit firing rates has shown inconsistent effects across and within experiments, the results may be more uniform when averaged across a large population of neurons. Our results are consistent with those from previous fMRI experiments (Buracas and Boynton, 2007; Murray, 2008) reporting additive offsets with attention as well as a voltage-sensitive dye experiment that reached a similar conclusion about selection (E. Seidemann, personal communication). Equal increases in responses at all contrasts may result when responses are averaged across populations of neurons for at least two reasons. First, if some neurons show enhancement primarily at low and intermediate contrasts (contrast-gain like changes) and other neurons show enhancement primarily at high contrasts (response-gain like changes), then the overall sum of activity (and, consequently, any population readout that depends on this sum) would be expected to show enhancement at all contrasts (i.e., an additive offset). Indeed, an electrophysiological study has reported that some neurons exhibit contrast-gain, others response-gain, and yet others exhibit additive changes in the same experiment (Williford and Maunsell, 2006). Moreover, the normalization model of attention (Reynolds and Heeger, 2009) can yield contrast-gain or response-gain like changes in different neurons dependent on the locations and sizes of their receptive fields. These effects in individual neurons can appear as an additive offset change when averaged across neurons (unpublished simulations). Second, the majority of single-unit electrophysiology experiments used stimulus param-

eters that were matched to the tuning properties of the individual units being recorded. But in fact, any stimulus that is the target of attention will give rise to activity in many neurons whose receptive fields and tuning properties may only partially match with the stimulus. Small baseline shifts with attention (Luck et al., 1997; Reynolds et al., 2000; Williford and Maunsell, 2006) in each of many neurons may sum to a large effect in the overall population output, evident in the fMRI responses. The behavioral performance improvements with attention may, for some stimuli and tasks, depend primarily on this component of the population responses that is correlated across neurons (not the response- and/or contrast-gain changes evident in each individual neuron's responses). Our max-pooling selection rule exemplifies how such a baseline shift can lead to improved behavioral performance. Hence, it is possible to reconcile the attentional modulation effects that have been measured with fMRI with those measured electrophysiologically.

Response Enhancement

Is it possible that behavioral performance improvements with attention are due to response enhancement (Figure 1B) but that we were simply unable to measure the response enhancement with fMRI? In particular, is it possible that increases in fMRI responses that we and others (Buracas and Boynton, 2007; Murray, 2008) have measured with attention were entirely due to neuromodulatory input and, therefore, did not reflect signals used by the brain for contrast discrimination? We considered three specific possibilities to address these questions.

First, it could be that the fMRI measurements were dominated by attention-related synaptic input that was constant for all stimulus contrasts and, hence, looked like an additive offset. Such would be the case if the fMRI measurements reflected only the neuromodulatory input that specified the attention field (i.e., the changes in synaptic gain corresponding to the spatial extent of attention), which would be only indirectly evident in extracellular electrophysiological measurements of spiking activity. However, we measured a monotonically increasing contrast-response function in all task conditions (Figure 4) that indicated that at least part of the fMRI responses was driven by the stimulus. Moreover, the gain changes that would have been needed to account for the behavioral enhancement with attention were approximately 4-fold (Figure 5) and should have been easily measurable as they would have been much larger than contrast-gain changes with adaptation measured with fMRI using similar procedures (Gardner et al., 2005).

Second, could it be that the contrast-response functions we measured reflected only bottom-up input? Had this been the case, gain changes within a cortical area would not have been evident in the fMRI responses from that area, but rather, those gain changes would have been displaced to a later visual area. For example, even if one area, say V1, were dominated by bottom-up inputs, e.g., from the LGN, we would have expected to see gain changes in the areas to which V1 projects. However, no gain changes were observed in V2, V3, and hV4.

Third, could it be that signals used to perform the contrast-discrimination task were encoded at a spatial scale below the resolution afforded by hemodynamic measurements? Whereas we cannot fully rule out this possibility, it is unlikely because

single-unit studies (Martinez-Trujillo and Treue, 2002; McAdams and Maunsell, 1999; Mitchell et al., 2009; Reynolds et al., 2000; Williford and Maunsell, 2006) have uniformly measured gain changes that are too modest to explain the large (∼4-fold) response-gain changes needed to account for the observed behavioral effect. Indeed, population sensitivity measures from single-unit data agree with our conclusion that gain changes can account for only a very small fraction of behavioral enhancement (Cohen and Maunsell, 2009).

Sensory Noise Reduction

Sensory noise reduction (Figure 1C) is another possible mode of sensitivity enhancement, which could have been missed by fMRI measurements (Cohen and Maunsell, 2009; Mitchell et al., 2009). Direct measurements of the variability of neural responses with fMRI are difficult if not impossible as fMRI is corrupted by various other sources of noise (thermal, physiological, movement artifacts, hemodynamic, etc.). Indeed, the sensitivity model estimated that the trial-to-trial fluctuations of the fMRI signal needed to account for behavior (i.e., the estimate of the noise variance due to neural sources that influence perception) were less than 0.1% change in fMRI image intensity—an order of magnitude smaller than the overall trial-to-trial variability in the fMRI responses we measured (approximately 1%).

However, although fMRI may not be able to measure changes in neural variability directly, the sensitivity model estimated that an unrealistically high 400% reduction in noise was needed to account for behavioral enhancement. This amount of noise reduction was an order of magnitude larger than the reduction in the response variance inferred from monkey electrophysiology (Cohen and Maunsell, 2009; Mitchell et al., 2009). We note, however, that there are still few studies that have examined changes in response variation and correlation between neurons with attention and that there is considerable uncertainty about how much reduction in variability at the level of populations of neurons can be inferred from the existing data. Nonetheless, our analysis suggested that response enhancement, coupled with a realistic amount of noise reduction, would not suffice to account for the behavioral performance improvements that we observed.

We assumed additive noise when estimating neural variability to link the contrast-discrimination and contrast-response functions, but single-unit studies have found that firing rate response variances scale with the mean firing rates, similar to a Poisson process (Softky and Koch, 1993). Therefore, it might seem that contrast discrimination should be modeled with multiplicative noise, which scales with response. However, because perceptual decisions are likely based on populations of neural activity, behavioral performance is not necessarily limited by Poisson-like noise evident in single neurons. If the neural noise that scales with the response amplitudes is independent across neurons, then the Poisson-like noise will be averaged out, and only correlated components of the noise will remain. This remaining correlated noise component might be additive. Indeed, the standard deviation of the population response measured with voltage-sensitive dyes does not change with contrast in V1 (Chen et al., 2006). Moreover, psychophysical data suggest that perceptual performance is limited by an additive noise component (Gorea and Sagi, 2001).

Efficient Selection

Not being able to account for the behavioral enhancement with the forms of sensory enhancement discussed above, we considered the possibility that attention improved behavioral performance by efficiently selecting relevant sensory signals (Eckstein et al., 2000; Palmer et al., 2000; Pelli, 1985). And we found that a simple max-pooling selection mechanism could fully and realistically account for the behavioral enhancement. The selection model we implemented did not change the pooling rule across focal and distributed cue trials (a fixed k and σ across both trial types could account adequately for the behavioral data), and therefore does not suggest that the pooling mechanism itself becomes more efficient for focal cue trials. Instead, we hypothesize that larger responses (in the form of an additive offset) aid in propagating the relevant visual information through a static pooling rule leading to more efficient selection of relevant signals.

Whereas the particular form of max-pooling selection rule used was not essential, we used it because it has a plausible neural implementation. We implemented a continuum of selection rules from averaging to max pooling by taking the sum of the exponent of input signals. Other selection rules such as a soft-max operator (Kouh and Poggio, 2008) could have been used to achieve the same function. However, an exponential relation to inputs has been observed for visual neurons in sensory areas; these neurons are well modeled as linear operators with a static output nonlinearity in the range of two to four (Albrecht and Hamilton, 1982). Higher exponent values might be achieved as sensory signals pass from one area to the next, each area contributing a part of the full exponent value. Our selection rule also includes a root operator, the purpose of which was simply to keep the output of the selection rule in the same range as the input, and could also be achieved by other computations such as divisive normalization (Heeger, 1992).

A prediction of our selection model is that distracters that evoke large responses (for example, those presented with higher contrast) will be better able to pass through the selection mechanism and, thus, disrupt performance. We tested and confirmed this prediction. These results parallel other reports (Palmer and Moore, 2009; Yigit-Elliott et al., 2011) that show that high-contrast distracters (foils) can be incorrectly selected, leading to errors in behavioral performance. Similarly, searching for a high-contrast target among low-contrast distracters is less impaired relative to searching for a low-contrast target among high-contrast distracters when attention is allocated elsewhere (Braun, 1994) or V4 is lesioned (Schiller and Lee, 1991). These results all suggest that high-contrast stimuli preferentially access perception (but see Jonides and Yantis, 1988). Efficient selection with winner-take-all like selection mechanisms as described here and elsewhere (e.g., Koch and Ullman, 1985; Lee et al., 1999) provides a unified framework that can explain both these types of bottom-up effects as well as top-down effects of focal attention.

Working Memory

Our conceptualization of the processes involved in the contrast discrimination task did not consider the limits of working memory in changing behavioral performance. To perform a two-interval discrimination task, observers must hold the contrast perceived in the first interval in working memory to compare with the

perceived contrast in the second interval. Our task was designed to minimize demands put on working memory by using a very brief ISI (200 ms; Figure 2) and a small set size (four) that has been shown to minimize performance decrements due to working memory (Luck and Vogel, 1997). Nonetheless, performance may have improved for the focal cue because working memory was needed only to hold the relevant one item instead of all four items with the distributed cue. If this account holds, it raises the question of what process acts to exclude irrelevant information from working memory in the focal cue condition. One possibility is that efficient selection in a matter akin to what we have formulated here acts as a gatekeeper that excludes irrelevant information from working memory. Indeed, exclusion of irrelevant items in working memory is a key factor improving performance in working memory tasks (Vogel et al., 2005), thus suggesting that attentional enhancement in the form of efficient selection may be a key process in determining the efficacy of working memory.

Selection and Sensitivity

Whether attention improves performance through sensory enhancement or efficient selection may critically depend on the types of tasks used to probe attentional effects. Sensory enhancement and efficient selection are not mutually exclusive, rather they are both likely to contribute to the computational processes by which attention improves performance (Eckstein et al., 2000; Lu and Dosher, 1998; Palmer et al., 2000). On the one hand, many experiments have limited the number of behaviorally relevant stimuli; for example by presenting one or two stimuli on a blank background, thus limiting demand on the neural processes that govern the efficiency of selection (Carrasco et al., 2000; Lu and Dosher, 1998; Morrone et al., 2002; Pestilli et al., 2009). For these types of tasks, the bottleneck in performance may therefore be in the fidelity of the stimulus representation. Correspondingly, single-unit studies using such tasks have reported signal enhancement in the form of gain changes (Martinez-Trujillo and Treue, 2002; McAdams and Maunsell, 1999; Reynolds et al., 2000; Williford and Maunsell, 2006), and reductions of correlated noise (Cohen and Maunsell, 2009; Mitchell et al., 2009). On the other hand, tasks in which the relevant signals must be selected out of many possible alternatives place higher demands on selection efficiency (Eckstein et al., 2000; Palmer et al., 2000; Pelli, 1985). For these tasks, the bottleneck in performance may not be the fidelity of the stimulus representation, but the efficiency of selection. Moreover, tasks in which relevant and irrelevant stimuli are placed in near proximity to each other may result in selection of relevant signals and suppression of irrelevant signals at stages of the visual system in which both stimuli are within the same receptive field (cf. "biased-competition"; Desimone and Duncan, 1995). Real-world situations usually involve complex and cluttered visual environments in which efficient selection mechanisms may be particularly crucial for optimal behavior.

EXPERIMENTAL PROCEDURES

Observers

Three healthy males (age 33–36) with normal or corrected-to-normal vision who provided written informed consent participated in the study. Experimental

procedures were in compliance with the safety guidelines for MRI research and were approved by the University Committee on Activities Involving Human Subjects at New York University. Each observer participated in multiple fMRI experiments: one 1.5-hr-long session of retinotopic mapping, and five 2-hr-long sessions of the contrast discrimination experiment. To test the effect of high-contrast distracters, we conducted behavioral experiments on six observers (ages 23–39, one female), including two from the main experiment, all with normal or corrected-to-normal vision. Experimental procedures were conducted with the written consent of each observer and were approved by the RIKEN Brain Science Institute Functional MRI Safety and Ethics Committee.

Behavioral Protocol

The behavioral protocol is described in the Results and in detail in the Supplemental Experimental Procedures.

Stimulus Presentation

Visual stimuli were generated using MATLAB (The MathWorks, Inc., Natick, MA, USA) and MGL (http://justingardner.net/mgl) and presented via an LCD projector. See Supplemental Experimental Procedures.

MRI Acquisition and Preprocessing

MRI data were acquired on a 3 Tesla Allegra head-only scanner (Siemens, Erlangen, Germany) using standard procedures. See Supplemental Experimental Procedures.

Psychophysical Contrast-Discrimination Functions

Contrast-discrimination thresholds were computed separately for each pedestal contrast and each cue condition, and the resulting contrast-discrimination functions were then fit, following previous research (Boynton et al., 1999; Legge and Foley, 1980; Nachmias and Sansbury, 1974; Zenger-Landolt and Heeger, 2003), by assuming that behavioral performance is limited by the fixed difference in response amplitude (ΔR) divided by the standard deviation of sensory noise (σ). Then the contrast-discrimination threshold for a pedestal contrast, $\Delta c(c)$, satisfies:

$$d' = \frac{R(c + \Delta c(c)) - R(c)}{\sigma}, \qquad (2)$$

where R is the underlying contrast-response function. The contrast-response functions were parameterized as:

$$R(c) = b + g_r\left(\frac{c^{s+q}}{c^q + g_c^q}\right), \qquad (3)$$

where b is the baseline response, g_r is the response gain that determines the maximum response, g_c is the contrast gain that determines the horizontal position of the function along the contrast axis, and s and q are exponents that control how quickly the function rises and saturates. For the sensitivity and selection model fits, g_r (the response gain of the contrast-response function, Equation 3) and ΔR (the response difference at threshold, Equation 2) were constrained by measurements of the contrast-response functions. However, ΔR, σ, and g_r were codependent variables when fitting the contrast-discrimination functions on their own. We therefore set σ and ΔR to 1 and fit (nonlinear least-squares) the other parameters of the contrast-response function to the measured contrast thresholds. Increasing or decreasing either ΔR, σ or g_r resulted in a vertical shift upward or downward of the fit to the contrast-discrimination function, and b (baseline response; Equation 3) had no effect on the contrast-discrimination function and was set to zero.

We show that the decrease in contrast-discrimination thresholds at high contrast is explained by the selection model (see Results), but we also fit the data without ascribing it to any particular mechanism, by multiplying the thresholds, $\Delta c(c)$, from the aforementioned model (Equations 2 and 3) with a scaling factor:

$$\Delta c'(c) = \Delta c(c)e^{-\left(c/\gamma\right)^{\rho}}, \qquad (4)$$

where γ is the contrast at which threshold has decreased by 37%, and ρ is the slope of the decrease on a log-log axis.

In summary, the contrast-discrimination functions were fit (nonlinear least-squares) using a combination of Equations 2–4 (see Figure 3). There were a total of eight data points for each of two cue conditions (focal and distributed). These data were fit with six free parameters for each cue condition: g_r (response-gain), s, q (exponents), g_c (contrast-gain shift), γ, and ρ (center and slope of threshold dip at high contrast, Equation 4).

fMRI Contrast-Response Functions

While observers performed the contrast-discrimination task, cortical responses to the stimuli were measured in visual areas V1, V2, V3, and hV4. In a separate scanning session, we identified the four subregions of each visual area corresponding to each of the four stimulus apertures (see Supplemental Experimental Procedures, Retinotopic Mapping and Visual Field Quadrant Localizer). Responses corresponding to each stimulus contrast, for each stimulus cue combination (i.e., focal cue target, focal cue nontarget, distributed cue target, and distributed cue nontarget; see Figure 1), were then averaged across these four subregions of each visual area.

The mean fMRI response time courses were estimated using deconvolution (i.e., linear regression), baseline normalized to the nontarget focal cue condition, and the amplitude of response was estimated. These amplitudes were then fit using Equation 3. See Supplemental Experimental Procedures.

Testing Sensory Noise Reduction

The sensitivity model (Figure 1) was fit (nonlinear least-squares) to the contrast-response functions (see Figures 5A–5F), using Equations 2–4. The particular parameterization of the contrast discrimination functions was not essential for our results in that simplified forms (with fewer parameters; see Supplemental Experimental Procedures: Alternate Functional Forms Used to Fit Contrast-Response) did not qualitatively change the results (see Figure S2C).

To perform the fit, the contrast-discrimination functions were numerically integrated, using the following procedure, to predict the contrast-response functions. Given values for the noise, σ, and baseline response, b, a contrast-discrimination function uniquely specified a contrast-response function. The first point on the contrast-response function at 0% contrast was simply the baseline response:

$$R(0) = b, \tag{5}$$

We assumed that behavioral sensitivity (d′) was equal to the neural response difference, $R(c + \Delta c) - R(c)$, divided by σ (see Equation 2). The interpolated contrast discrimination functions gave the threshold contrast, Δc, for any contrast, c, for a behavioral sensitivity of d′ = 1 (see above, Psychophysical Contrast-Discrimination Functions), thus:

$$R(c + \Delta c) = R(c) + \sigma. \tag{6}$$

To compute the next point on the contrast-response function, we thus applied Equation 6, for c = 0 and Δc as estimated from the interpolated contrast discrimination function, i.e., $R(\Delta c) = b + \sigma$. Subsequent values of R were computed by repeated application of Equation 6 in which each new c was set to c+Δc from the previous iteration and Δc for that new contrast c, was retrieved from the interpolated contrast discrimination function (see Supplemental Experimental Procedures, for more details on the fitting procedure). σ and b were adjusted to produce the best fit of the contrast-response functions in the least-squares sense.

Testing Efficient Selection

The contrast-discrimination functions were fit (nonlinear least-squares) by the selection model, using Equations 1 and 3. To perform the fit, the contrast discrimination performance of the selection model (percent [%] correct) was computed by simulating synthetic trials based on responses computed from the measured contrast-response functions. Contrast-response functions were interpolated with a simplified version of Equation 3 (a Naka-Rushton type equation), which lacked the exponent s. The exact form of the interpolation function was not essential (see Supplemental Experimental Procedures). For any fixed value of k (Equation 1) and value of the sensory noise (σ), the selection model performance (percent [%] correct) was computed as follows.

For each pedestal contrast, Gaussian response distributions were computed for each stimulus location and each interval of the task (Figure 7A). The mean of each response distribution was determined according to the interpolated contrast-response functions. The standard deviations of the Gaussian response distributions were set to the σ parameter. Responses were then combined into "readout" distributions using the max-pooling rule (Equation 1) and the parameter k (Figure 7B). On each of 10,000 simulated trials, a response was taken from the readout distribution for each interval. If the larger of these two responses was in the same interval as the increment in contrast, the trial was marked as correct. The Δc that produced 76% correct values using this procedure was taken as the discrimination threshold. Values of k and σ were adjusted to produce the best fit of the contrast discrimination functions in the least-squares sense.

We also computed two variations of the aforementioned model (see Figure 8). One variation included two σ values (σ_f and σ_d), one for the focal cue and one for the distributed cue trials. The second variation of the model used a pooling rule in which the readout distribution was taken solely from the correct target location on both focal cue and distributed cue trials. This version had both σ_d and σ_f parameters, but no k parameter.

Model Comparisons

Model fits were compared using two different measures that account for differences in number of model parameters: cross-validated r^2 and AIC. See Supplemental Experimental Procedures.

Eye Position Monitoring

Eye position was monitored during the experiments, and analysis of the data did not reveal any potential artifacts. See Supplemental Experimental Procedures.

SUPPLEMENTAL INFORMATION

Supplemental Information includes five figures and Supplemental Experimental Procedures and can be found with this article online at doi:10.1016/j.neuron.2011.09.025.

REFERENCES

Albrecht, D.G., and Hamilton, D.B. (1982). Striate cortex of monkey and cat: contrast response function. J. Neurophysiol. 48, 217–237.

Boudreau, C.E., Williford, T.H., and Maunsell, J.H. (2006). Effects of task difficulty and target likelihood in area V4 of macaque monkeys. J. Neurophysiol. 96, 2377–2387.

Boynton, G.M., Engel, S.A., Glover, G.H., and Heeger, D.J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. J. Neurosci. 16, 4207–4221.

Boynton, G.M., Demb, J.B., Glover, G.H., and Heeger, D.J. (1999). Neuronal basis of contrast discrimination. Vision Res. 39, 257–269.

Braun, J. (1994). Visual search among items of different salience: removal of visual attention mimics a lesion in extrastriate area V4. J. Neurosci. *14*, 554–567.

Buracas, G.T., and Boynton, G.M. (2007). The effect of spatial attention on contrast response functions in human visual cortex. J. Neurosci. *27*, 93–97.

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.

Chen, Y., Geisler, W.S., and Seidemann, E. (2006). Optimal decoding of correlated neural population responses in the primate visual cortex. Nat. Neurosci. *9*, 1412–1420.

Cohen, M.R., and Maunsell, J.H. (2009). Attention improves performance primarily by reducing interneuronal correlations. Nat. Neurosci. *12*, 1594–1600.

Cook, E.P., and Maunsell, J.H. (2002). Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. J. Neurosci. *22*, 1994–2004.

Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. Annu. Rev. Neurosci. *18*, 193–222.

Douglas, R.J., and Martin, K.A. (2007). Mapping the matrix: the ways of neocortex. Neuron *56*, 226–238.

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.

Foley, J.M., and Schwarz, W. (1998). Spatial attention: effect of position uncertainty and number of distractor patterns on the threshold-versus-contrast function for contrast discrimination. JOSA A *15*, 1036–1047.

Gardner, J.L., Sun, P., Waggoner, R.A., Ueno, K., Tanaka, K., and Cheng, K. (2005). Contrast adaptation and representation in human early visual cortex. Neuron *47*, 607–620.

Geisler, W.S., and Albrecht, D.G. (1997). Visual cortex neurons in monkeys and cats: detection, discrimination, and identification. Vis. Neurosci. *14*, 897–919.

Gorea, A., and Sagi, D. (2001). Disentangling signal from noise in visual contrast discrimination. Nat. Neurosci. *4*, 1146–1150.

Heeger, D.J. (1992). Normalization of cell responses in cat striate cortex. Vis. Neurosci. *9*, 181–197.

Heeger, D.J., and Ress, D. (2002). What does fMRI tell us about neuronal activity? Nat. Rev. Neurosci. *3*, 142–151.

Heeger, D.J., Huk, A.C., Geisler, W.S., and Albrecht, D.G. (2000). Spikes versus BOLD: what does neuroimaging tell us about neuronal activity? Nat. Neurosci. *3*, 631–633.

Herrmann, K., Montaser-Kouhsari, L., Carrasco, M., and Heeger, D.J. (2010). When size matters: attention affects performance by contrast or response gain. Nat. Neurosci. *13*, 1554–1559.

Jonides, J., and Yantis, S. (1988). Uniqueness of abrupt visual onset in capturing attention. Percept. Psychophys. *43*, 346–354.

Koch, C., and Ullman, S. (1985). Shifts in selective visual attention: towards the underlying neural circuitry. Hum. Neurobiol. *4*, 219–227.

Kouh, M., and Poggio, T. (2008). A canonical neural circuit for cortical nonlinear operations. Neural Comput. *20*, 1427–1451.

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.

Lee, J., and Maunsell, J.H. (2010). The effect of attention on neuronal responses to high and low contrast stimuli. J. Neurophysiol. *104*, 960–971.

Legge, G.E., and Foley, J.M. (1980). Contrast masking in human vision. J. Opt. Soc. Am. *70*, 1458–1471.

Li, X., Lu, Z.L., Tjan, B.S., Dosher, B.A., and Chu, W. (2008). Blood oxygenation level-dependent contrast response functions identify mechanisms of covert attention in early visual areas. Proc. Natl. Acad. Sci. USA *105*, 6202–6207.

Liu, T., Pestilli, F., and Carrasco, M. (2005). Transient attention enhances perceptual performance and FMRI response in human visual cortex. Neuron *45*, 469–477.

Logothetis, N.K., and Wandell, B.A. (2004). Interpreting the BOLD signal. Annu. Rev. Physiol. *66*, 735–769.

Lu, Z.L., and Dosher, B.A. (1998). External noise distinguishes attention mechanisms. Vision Res. *38*, 1183–1198.

Luck, S.J., and Vogel, E.K. (1997). The capacity of visual working memory for features and conjunctions. Nature *390*, 279–281.

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-Trujillo, J., and Treue, S. (2002). Attentional modulation strength in cortical area MT depends on stimulus contrast. Neuron *35*, 365–370.

McAdams, C.J., and Maunsell, J.H. (1999). Effects of attention on the reliability of individual neurons in monkey visual cortex. Neuron *23*, 765–773.

Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. Neuron *63*, 879–888.

Morrone, M.C., Denti, V., and Spinelli, D. (2002). Color and luminance contrasts attract independent attention. Curr. Biol. *12*, 1134–1137.

Mukamel, R., Gelbard, H., Arieli, A., Hasson, U., Fried, I., and Malach, R. (2005). Coupling between neuronal firing, field potentials, and FMRI in human auditory cortex. Science *309*, 951–954.

Murray, S.O. (2008). The effects of spatial attention in early human visual cortex are stimulus independent. J. Vis. *8*, 2.1–2.11.

Nachmias, J., and Sansbury, R.V. (1974). Letter: Grating contrast: discrimination may be better than detection. Vision Res. *14*, 1039–1042.

Palmer, J., and Moore, C.M. (2009). Using a filtering task to measure the spatial extent of selective attention. Vision Res. *49*, 1045–1064.

Palmer, J., Verghese, P., and Pavel, M. (2000). The psychophysics of visual search. Vision Res. *40*, 1227–1268.

Pelli, D.G. (1985). Uncertainty explains many aspects of visual contrast detection and discrimination. JOSA A *2*, 1508–1531.

Pestilli, F., Ling, S., and Carrasco, M. (2009). A population-coding model of attention's influence on contrast response: estimating neural effects from psychophysical data. Vision Res. *49*, 1144–1153.

Pooresmaeili, A., Poort, J., Thiele, A., and Roelfsema, P.R. (2010). Separable codes for attention and luminance contrast in the primary visual cortex. J. Neurosci. *30*, 12701–12711.

Reynolds, J.H., and Heeger, D.J. (2009). The normalization model of attention. Neuron *61*, 168–185.

Reynolds, J.H., Pasternak, T., and Desimone, R. (2000). Attention increases sensitivity of V4 neurons. Neuron *26*, 703–714.

Sapir, A., d'Avossa, G., McAvoy, M., Shulman, G.L., and Corbetta, M. (2005). Brain signals for spatial attention predict performance in a motion discrimination task. Proc. Natl. Acad. Sci. USA *102*, 17810–17815.

Schiller, P.H., and Lee, K. (1991). The role of the primate extrastriate area V4 in vision. Science *251*, 1251–1253.

Sergent, C., Ruff, C.C., Barbot, A., Driver, J., and Rees, G. (2011). Top-down modulation of human early visual cortex after stimulus offset supports successful postcued report. J. Cogn. Neurosci. *23*, 1921–1934.
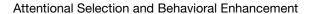
Shmuel, A., Augath, M., Oeltermann, A., and Logothetis, N.K. (2006). Negative functional MRI response correlates with decreases in neuronal activity in monkey visual area V1. Nat. Neurosci. *9*, 569–577.

Sirotin, Y.B., and Das, A. (2009). Anticipatory haemodynamic signals in sensory cortex not predicted by local neuronal activity. Nature *457*, 475–479.

Softky, W.R., and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. J. Neurosci. *13*, 334–350.

Thiele, A., Pooresmaeili, A., Delicato, L.S., Herrero, J.L., and Roelfsema, P.R. (2009). Additive effects of attention and stimulus contrast in primary visual cortex. Cereb. Cortex *19*, 2970–2981.

Vogel, E.K., McCollough, A.W., and Machizawa, M.G. (2005). Neural measures reveal individual differences in controlling access to working memory. Nature *438*, 500–503.

Williford, T., and Maunsell, J.H. (2006). Effects of spatial attention on contrast response functions in macaque area V4. J. Neurophysiol. *96*, 40–54.

Yigit-Elliott, S., Palmer, J., and Moore, C.M. (2011). Distinguishing blocking from attenuation in visual selective attention. Psychol. Sci. *22*, 771–780.

Zenger-Landolt, B., and Heeger, D.J. (2003). Response suppression in v1 agrees with psychophysics of surround masking. J. Neurosci. *23*, 6884–6893.