

Attention activates winner-take-all competition among visual filters

D.K. Lee, L. Itti, C. Koch and J. Braun

Computation and Neural Systems 139–74, California Institute of Technology, Pasadena, California 91125, USA

The first two authors contributed equally to this study.

Correspondence should be addressed to J.B. (achim@klab.caltech.edu)

Shifting attention away from a visual stimulus reduces, but does not abolish, visual discrimination performance. This residual vision with ‘poor’ attention can be compared to normal vision with ‘full’ attention to reveal how attention alters visual perception. We report large differences between residual and normal visual thresholds for discriminating the orientation or spatial frequency of simple patterns, and smaller differences for discriminating contrast. A computational model, in which attention activates a winner-take-all competition among overlapping visual filters, quantitatively accounts for all observations. Our model predicts that the effects of attention on visual cortical neurons include increased contrast gain as well as sharper tuning to orientation and spatial frequency.

Although attention profoundly alters visual perception¹, it is not equally important to all aspects of vision. For example, attention is of little or no help to many detection tasks (for example, detecting a luminance increment²), and the degree to which it benefits discrimination tasks varies widely with the discriminated attribute (for example, discriminating color, orientation, form^{3,4}). Here we report how attention alters thresholds for discriminating contrast, orientation and spatial frequency of simple patterns. Based on earlier work, we expected markedly different effects on different thresholds⁵. In addition, we describe how attention changes thresholds for detecting one pattern in the presence of another, superimposed pattern of different orientation or spatial frequency. Together, these measurements characterize the visual mechanisms that underlie basic pattern vision.

Perceptual thresholds for stimulus contrast, orientation and spatial frequency have been studied for several decades^{6–8}. Quantitative accounts of these thresholds have become increasingly refined and usually involve a population of ‘noisy filters’ tuned to different orientations and spatial frequencies. Although earlier models postulated filters that are independent of each other, there are serious shortcomings to this approach^{9,10}. More recent models postulate an interaction between filters with spatially overlapping receptive fields^{11–13}, specifically, the normalization of individual filter responses relative to the total response of the local filter population (‘divisive inhibition’¹⁴). This normalization accounts naturally for several otherwise puzzling observations, among them the initial decrease and later increase of contrast discrimination thresholds with increasing stimulus contrast^{12,13} (‘dipper function’^{6,8}) and the relative constancy of orientation and spatial frequency thresholds over a wide range of stimulus contrasts^{9,15}.

An intriguing parallel to these perceptual accounts can be found in certain models of visual cortical responses to stimulus contrast and orientation^{16,17}. Despite marked differences in detail, the models in question consider a population of neurons with overlapping receptive fields, broadly tuned to a range of different orientations, and normalize individual responses relative to

the population response. The normalization, which in some cases is implemented as a divisive inhibition, sharpens orientation tuning¹⁶ and renders it less dependent on stimulus contrast¹⁷. Thus, both perceptual and neuronal sensitivity to contrast and orientation seem to involve response normalization.

Here we report that attention modulates the response normalization that seems to underlie basic pattern vision. We reach this conclusion by comparing attentional changes in human thresholds to predictions of a computational model based on response normalization. Our model is similar to several others^{11–13,17} and comprises three stages: a population of overlapping filters responsive to different orientations and spatial frequencies at one visual location, non-linear interactions among this population to carry out the normalization and an ‘ideal observer’ decision that discriminates between stimulus alternatives on the basis of the maximum likelihood and is limited only by noise. Our observations are consistent with an attentional modulation of the second, but not the first or third, stage of the model.

RESULTS

Psychophysics

Although visual thresholds are usually measured when stimuli are fully attended, here we use a concurrent task to establish thresholds when stimuli are at best poorly attended^{4,18,19}. The concurrent task in question forces observers to withdraw attention from peripheral stimuli and to focus on stimuli near fixation (Fig. 1, Methods). This psychophysical manipulation is highly effective and causes substantial perceptual deficits in the periphery similar to the deficits obtained after a lesion in visual cortical area V4 of the monkey¹⁹. However, the perception of peripheral stimuli is not entirely abolished. Practiced observers enjoy a significant residual vision outside the focus of attention and render reliable threshold judgments about peripheral stimuli, especially when the display is uncluttered and contains only a few salient stimuli⁴.

Observers discriminated contrast, orientation or spatial frequency of a luminance-modulated pattern appearing at varying




Fig. 1. Measurement of visual thresholds with either full or poor attention. **(a)** Sequence of fixation, stimulus and mask displays (schematic). Observers fixate the center of all displays. The stimulus comprises a central and a peripheral component, which appear at varying locations of constant eccentricity. The central component consists of 5 Ts and/or Ls (central targets) and observers report “same” (that is, 5 Ts or 5 Ls) or “different” (that is, 4 Ts + 1 L or 4 Ls + 1 T). The peripheral component consists of the luminance-modulated patterns shown in Fig. 2a–e (peripheral target). For example, the peripheral component might be a grating pattern of vertical or tilted orientation, in which case observers would report “vertical” or “tilted.” The mask display limits visual persistence of central targets. **(b)** Single task (peripheral target ‘fully attended’), observers fixate the center but respond only to the peripheral task (see Fig. 2). **(c)** Double task (peripheral target ‘poorly attended’), observers fixate the center and respond first to the central task and second to the peripheral task.

locations of 4° eccentricity (peripheral target; Fig. 1a). To draw attention away from this pattern, we asked observers to discriminate whether five shapes near fixation (central targets) were the “same” or “different”. When observers carried out both tasks, they concentrated attention on the central task, which they were instructed to consider the primary task, and thus left the peripheral target poorly attended (double-task thresholds). In contrast, when observers viewed the same display but performed only the peripheral task, they fully attended to the peripheral target (single-task thresholds). The comparison of single- and double-task thresholds reveals if and how attention alters visual perception.

We compared five types of thresholds under single- and double-task conditions (Fig. 2a–e). When peripheral targets were fully attended, contrast detection thresholds (zero mask contrast) were about 20% lower, and contrast discrimination thresholds (mask contrast greater than zero) about 40–50% lower than when peripheral targets were poorly attended (Fig. 2a). In addition, the decrease of the discrimination threshold as mask contrast increases from zero (dipper) was evident only when targets were fully attended. Note that the target position varied from trial to trial (to forestall eye movements) and that positional uncertainty of this kind is known to reduce the dipper^{20–22}. Therefore, it is possible that our data underestimate the depth of the dipper.

The effects of attention on spatial frequency and orientation discrimination were even more pronounced (Fig. 2b and c). Spatial frequency thresholds were about 60% lower and orientation thresholds about 70% lower when peripheral targets were fully attended compared to when they were poorly attended. Note that both types of thresholds remained essentially constant for contrast values above 20%.

Interactions between superimposed stimuli of different orientation or spatial fre-

quency (target and mask; Fig. 2d and e) were also altered by attention. When target and mask had similar orientation or spatial frequency, attention lowered the maximal threshold by about 50% (consistent with Fig. 2a, mask contrast 0.5). As target and mask became progressively more different, fully and poorly attended thresholds decreased toward the same baseline level. The baseline




Fig. 2. Single- and double-task thresholds compared. Five types of thresholds were measured. In each case, observers discriminate between two alternative forms of the peripheral (4° eccentricity) target. Filled and open symbols represent fully attended (single-task) and poorly attended (double-task) thresholds, respectively (mean and standard error of two observers). Solid and dashed curves represent the corresponding model predictions. **(a)** Contrast detection and discrimination. Observers report the presence (arrows) or absence of a vertical target stripe from a circular masking pattern (contrast range 0.0–0.5). **(b, c)** Spatial frequency and orientation discrimination. Observers report whether a circular target grating (contrast 0.02–0.8) exhibits higher or lower spatial frequency **(b)** or whether its orientation is vertical or tilted clockwise **(c)**. (*) indicates further data points off scale. **(d, e)** Orientation and spatial frequency masking. Observers report the presence (arrows) or absence of a vertical target stripe from circular masking patterns (contrast 0.5) of different orientation (difference range 0–90; **d**) or different spatial frequency (difference range –1 to +1 oct; **e**). **(f)** Model parameters of plausible fits computed separately for single- and double-task data. Although all 10 parameters are permitted to differ, most parameters do not differ significantly (n.s.).




Fig. 3. Three-stage model of visual filters and their interactions (schematic). Each stimulus location is analyzed by linear filters sensitive to different orientations and spatial frequencies (Eq. 1, first-stage responses, $E_{\theta, \omega}$). Filter responses are subjected to excitatory and inhibitory interactions in the form of amplification and divisive normalization (Eqs. 3, 4, second-stage responses, $R_{\theta, \omega}$). The decision stage assumes that first-stage responses show a variance similar to that of cortical neurons (Eq. 2) and chooses between stimulus alternatives on the basis of maximal likelihood. See ref. 15 and Methods for details. Our results suggest that attention strengthens non-linear interactions between filters (gray box), but does not affect other parts of the model.

was comparable to thresholds without mask (Fig. 2a, mask contrast 0), indicating minimal interactions between targets and masks of very different orientation or spatial frequency.

Model

The visual thresholds measured here are thought to reflect the activity of a population of 'noisy filters' selective for stimuli of different orientations and spatial frequencies^{7,8}. We define a filter tuned to orientation θ and spatial frequency ω by

$$E_{\theta, \omega} = A c_s e^{-\frac{(\theta_s - \theta)^2}{2\sigma_\theta^2}} e^{-\frac{(\omega_s - \omega)^2}{2\sigma_\omega^2}} + B \quad (1)$$

where $E_{\theta, \omega}$ is the linear response, c_s , θ_s and ω_s are the contrast, orientation and spatial frequency of the stimulus, A is the contrast gain, B is the background activity, and σ_θ and σ_ω are the sharpness of tuning (for a sinusoidal grating stimulus). When the properties of such filters are inferred from behavioral threshold measurements, they tend to match the response properties of neurons in visual cortical areas V1 and V2 (refs. 23–25). Accordingly, each visual filter is thought to correspond to a population of visual cortical neurons tuned to a particular orientation and spatial frequency.

Can the observed effects of attention be understood simply as a change in the properties of individual visual filters? To answer this question, we first examined the case in which filters are independent, so that the output of each filter, $R_{\theta, \omega}$, is a monotonic (and perhaps non-linear) function of its linear response, $E_{\theta, \omega}$. We also assume that the variance of the filter output, $V_{\theta, \omega}^2$, is given by

$$V_{\theta, \omega}^2 = \beta (R_{\theta, \omega} + \epsilon) \quad (2)$$

where β is the 'light noise' and ϵ is the 'dark noise'. This approximates the response variance of visual cortical neurons²⁵. Given these assumptions, the observed 20% difference between contrast detection thresholds with full and poor attention implies (see Methods) that either the gain A decreases or the light noise β increases by about 20%. The 60% to 70% difference in spatial frequency and orientation thresholds, on the other hand, cannot be explained by a 20% change in A and/or β . To account for all observations, we must therefore assume that attention alters not only the contrast gain or noise level of visual filters, but also their tuning widths, σ_θ and σ_ω , for orientation and spatial frequency.

Next we analyzed a more complex model, which also proves to be consistent with attentional changes of both gain and tuning (Fig. 3). The first stage of this model consists of a population of overlapping linear filters responsive to different orientations and spatial frequencies at one visual location (Eq. 1). The second stage assumes that filters are not independent but interact so as to nor-

malize individual responses relative to the filter population. Specifically, second-stage responses, $R_{\theta, \omega}$, are obtained by subjecting first-stage responses, $E_{\theta, \omega}$, to a power law followed by divisive inhibition:

$$R_{\theta, \omega} = \frac{(E_{\theta, \omega})^\gamma}{S^\delta + \sum_{\theta', \omega'} W_{\theta\theta', \omega\omega'} (E_{\theta', \omega'})^\delta} \quad (3)$$

The exponents of the power law, δ and γ , are of particular consequence: their absolute values govern the strength of the interaction between filters, and their difference determines the saturation of responses at high contrast. The semi-saturation constant, S , determines the response at low stimulus contrast. The distribution of weight factors, $W_{\theta\theta', \omega\omega'}$,

$$W_{\theta\theta', \omega\omega'} = e^{-\frac{(\theta - \theta')^2}{2\Sigma_\theta^2}} e^{-\frac{(\omega - \omega')^2}{2\Sigma_\omega^2}} \quad (4)$$

whose Gaussian widths are given by Σ_θ and Σ_ω , determines whether the 'inhibitory pool' includes the entire filter population or only filters tuned to similar orientations and spatial frequencies.

The third stage of the model discriminates between stimulus alternatives based on the maximum likelihood of second-stage responses. This corresponds to an ideal observer whose performance is limited only by the variance (noise) of second-stage responses (Eq. 2). Further details about the decision are given elsewhere¹⁵ (also see Methods).

When we fit this model (10 free parameters: γ , δ , σ_θ , σ_ω , Σ_θ , Σ_ω , S , B , β , ϵ) separately to single- or double-task data, we obtained good agreement between predicted and observed thresholds with physiologically plausible parameter values (solid curves in Fig. 2). Note in particular the realistic widths of filter tuning, with half-widths at half-maximum between 12° and 15° for orientation and 0.42 octaves (oct) and 0.52 oct for spatial frequency, compared to $20 \pm 9^\circ$ and 0.76 ± 0.30 oct for neurons in monkey visual cortex²⁵. Note also that orientation and spatial frequency thresholds remain constant for contrast values above 20% (Fig. 2b and c) and that the curves for full and poor attention appear displaced vertically rather than horizontally. This shows clearly that attention changes more than contrast gain, because a difference in gain of the linear filter stage would merely produce a horizontal displacement. The main discrepancy between model and data is that the model predicts a more pronounced dipper for contrast discrimination thresholds than is actually observed (Fig. 2a). Because our data may underestimate the dipper (see above), this prediction may be correct.

That a single set of parameter values accounts for all thresholds observed with full attention is not a matter of course. One might have expected that attending to stimulus orientation would




Fig. 4. Predicted thresholds when attention changes some model parameters but not others. Format and experimental data are identical to Fig. 2. The thick curves represent a simultaneous fit to both single- and double-task data (solid and dashed, respectively), in which only the exponents δ and γ take different values depending on attention (12 free parameters). Observed and predicted thresholds agree reasonably well, and parameter values are physiologically plausible (two leftmost columns in f). The thin curves represent the optimal joint fit to single- and double-task data (solid and dashed, respectively) when all parameters except the exponents take different values depending on attention (18 free parameters). Neither the dipper (a) nor the maximal extent of contrast masking (d, e) are predicted, and parameter values are biologically implausible (two rightmost columns in f).

affect visual processing differently than, say, attending to spatial frequency. Instead, our results are consistent with the possibility that attention alters visual processing in the same way for all examined tasks. However, a strict test of task independence would require that model parameters be determined independently for different tasks. Unfortunately, such a test is not feasible because the data from any one task do not constrain all model parameters.

Although there are several differences between the parameters obtained with full and poor attention, the change in the exponents of the power law, γ and δ , is particularly significant (Fig. 2f). We assess the significance of a difference in the values of a given parameter by determining how rapidly the quality of fit deteriorates when this value is changed (see Methods). To study the role of γ and δ , we fit the model simultaneously to both single- and double-task data, while allowing only γ and δ to take different values depending on attention. In other words, γ and δ take two values, whereas all other parameters take a single value (12 free parameters total). Once again, we obtained acceptable fits with physiologically plausible parameter values ('12-dimensional joint fits', solid curves in Fig. 4a–e, left columns in Fig. 4f).

In contrast, when we allow all parameters except γ and δ to take different values depending on attention (18 free parameters, total), there are no acceptable fits with plausible parameter values. The optimal fit under these assumptions predicts neither the dipper in the contrast discrimination thresholds (Fig. 4a) nor the maximal extent of contrast masking (Fig. 4d and e). To obtain this poor fit, the tuning widths for orientation and spatial frequency and the size of the inhibitory pool have to change dramatically (σ_θ from 17° to 5° , σ_ω from 0.7 oct to 0.3 oct, and Σ_θ from $0.6 \sigma_\theta$ to $5 \sigma_\omega$). It seems unlikely that attention would alter cortical interactions so profoundly.

DISCUSSION

We measured thresholds for discriminating the contrast, orientation and spatial frequency of simple patterns that are either fully or poorly attended. We observed differences of 20% in contrast detec-

tion thresholds, 40–50% in contrast discrimination thresholds (and appearance of the dipper), 60–70% in orientation and spatial frequency discrimination thresholds, and up to 50% in contrast masking thresholds. These observations tightly constrain any effect attention may have on the visual filters that are thought to underlie basic pattern vision. Comparison with a computational model shows that the observed effects of attention are consistent with stronger interactions among filters, but not with a change in noise parameters without change in interactions, as is sometimes thought^{26,27}. Essentially, the effects of attention on different thresholds are too disparate to be accommodated by a single change in noise parameters.




Fig. 5. Effect of attention on early visual processing. Predictions based on 12-dimensional joint fit in Fig. 4f. Attention increases the contrast gain (3.3-fold, a), causes the contrast response to assume sigmoidal shape at low contrast (b) and sharpens orientation tuning by 40% (c) and spatial-frequency tuning by 30% (d). To the extent that the visual filters of our model reflect individual neurons in visual cortex, this predicts that attention both increases the gain and sharpens the tuning of such neurons.




Fig. 6. Attentional change in the response distribution. Predictions based on 12-dimensional joint fit in Fig. 4f. Responses $R_{\theta,\omega}$ of filters tuned to orientations between -20° to $+20^\circ$, to a grating stimulus of orientation 0° and contrasts between 0 and 5% (threshold regime). Responses to fully and poorly attended stimuli are represented by the red and blue surfaces, respectively (shown interleaved for clarity). By strengthening a winner-take-all competition among visual filters, attention restricts responses to the filters tuned best to the stimulus at hand.

In the framework of our model, the strength of interactions among filters is controlled by the exponents of a power law, γ and δ . The immediate reasons why larger exponents account for the observed effects of attention are as follows: for small stimulus contrasts, higher exponents reduce background firing and the sigmoidal shape of the contrast response (Fig. 5a and b), which explains the improved contrast detection thresholds and the enhanced dipper of the contrast discrimination curve. For larger stimulus contrasts, higher exponents entail a 3.3-fold increase in contrast gain (Fig. 5b), which accounts for lower contrast discrimination and contrast-masking thresholds. Additionally, higher exponents sharpen the tuning for orientation by 40%, and for spatial frequency by 30% (Fig. 5c and d), lowering thresholds for discriminating orientation and spatial frequency still further. To the extent that visual filters can be identified with individual neurons in visual cortex, our model thus predicts that attention changes both the gain and tuning of such neurons.

The more fundamental reason, however, is that larger exponents activate what is best described as a winner-take-all competition among visual filters. Attention (larger exponents) shifts the distribution of responses across the population of filters (Fig. 6). Attention accentuates existing differences between filter responses, boosting filters that respond relatively well to a given stimulus, while suppressing filters that respond relatively poorly. This explains the perceptual advantage conferred by attention: attention enhances the sensory representation by restricting responses to the filters that are tuned best to the stimulus at hand.

Previous studies of attentional changes in visual thresholds are broadly consistent with our results, even though our effects are larger. This includes reports that attention reduces contrast thresholds by 17% (ref. 27), orientation acuity by 15% (for an individual target without distractors)²⁸, and size acuity by 20% (ref. 29). However, these studies manipulated attention with a spatial cue rather than with a concurrent task, which complicates quantitative comparison. We believe that a concurrent task detains attention more consistently than spatial cueing; certainly concurrent tasks induce substantially larger changes in thresholds. An effect we have not considered here is that attention is able to improve perceptual deci-

sions in the face of positional uncertainty^{20–22}. We estimate that eliminating spatial uncertainty about the target location would improve contrast detection thresholds by 19%, but would have no appreciable effect on any thresholds at stimulus contrasts higher than 10% (see Methods). Thus, reduced uncertainty cannot explain the pattern of attentional effects we have observed.

Our model is also consistent with recent findings in the visual cortex of humans and monkeys. Attentional changes in neuronal activity have been reported in several early visual cortical areas, including areas V1, V2, V4 and MT/MST^{30–35} (see also S.P. Gandhi, D.J. Heeger & G.M. Boynton, *Inv. Ophthalm. Vis. Sci. (Suppl.)* 39, 5194, 1998). Furthermore, the notion that attention modulates a local competition in visual cortex has been proposed independently based on theoretical^{36,37} and single-neuron studies³⁸. In the macaque, attentional modulation of responses in the visual cortex is weak or absent if only a single stimulus is present in the receptive field, suggesting that attention modulates interactions between neurons with overlapping receptive fields^{30,33}. Finally, our model is consistent with reports that attention increases contrast gain in areas V2 and V4 of the macaque³⁹ (see also J. Reynolds, T. Pasternak & R. Desimone, *Inv. Ophthalm. Vis. Sci. (Suppl.)* 38, 3206, 1997). Whether attention sharpens the orientation tuning of visual cortical neurons remains controversial^{39,40}. Our model predicts that fully focused attention sharpens orientation tuning in the parts of visual cortex that mediate basic pattern vision (presumably areas V1 and/or V2). In area V4, increased competition would presumably sharpen tuning along other, more complex, stimulus dimensions.

Finally, we do not wish to claim that attention is restricted to local interactions at one particular level of visual cortex. More than likely, attention has additional effects on long-range interactions at the same level and, indeed, at all levels of visual cortex. Nevertheless, our results show that the activation of a winner-take-all competition among overlapping visual filters explains many basic perceptual consequences of attention.

METHODS

Psychophysics. Stimuli were displayed on an SGI Indigo (1024 × 1286 pixels RGB). Viewing was binocular at 120 cm distance (1° corresponds to 80 pixels). Mean luminance was 30 cd/m², with linear increments of 0.07 cd/m² (obtained by gamma correction and 'color bit stealing'⁴¹), and room luminance was 3 cd/m². Central targets appeared at 0–0.8° eccentricity and measured 0.4° across. Peripheral targets appeared at 4° eccentricity, in a circular aperture of 1.5° (timing as shown in Fig. 1). They were either sinusoidal gratings (Fig. 2b and c) or vertical stripes whose luminance profile was given by the 6th derivative of a Gaussian (Fig. 2a, d and e). Mask patterns were generated by superimposing 100 Gabor filters, positioned randomly within the circular aperture (Fig. 2a, d and e). The spatial frequency was 4 cpd (vertical stripes in Fig. 2a, d and e; sinusoidal gratings in Fig. 2b; superimposed Gabors in Fig. 2a and d), and the mask contrast was 0.5 (Fig. 2d and e). All thresholds were established with standard adaptive staircase methods (80 trials per block). The values given are averages from between 12 and 20 blocks of trials and 2 naive observers. Standard deviations were computed separately for each observer, and error bars represent the average value. In the double-task situation, observers were required to match or exceed a certain level of central performance. Approximately 15% of double-task blocks were discarded because of poor central performance. In both single- and double-task situations, observers fixated the display center, ensuring identical visual stimulation. The brief presentation effectively precludes saccades toward the peripheral target.

Concurrent-task method. An important concern in concurrent-task experiments is the level of processing at which the two tasks interfere. In general, interference can reflect limitations of attention, memory

and/or response generation^{42–44}. In the present experiments, observers have ample time to respond to each task in turn, so that limitations of response generation are not likely. Because interference disappears when central and peripheral targets are presented successively (for example, with an onset asynchrony of 200 ms or more)^{4,45,46}, a limitation of memory is also unlikely. Further evidence that the critical factor is attention is that interference does not depend on the nature of the central task (as long as it poses a sufficient demand on attention). For example, central tasks based on form, color or motion discrimination⁴⁷, as well as on an 'attentional blink'⁴⁸, produce comparable interference.

Independent-filter model. If the 'transducer' function $R_{\theta,\omega} = t(E_{\theta,\omega})$ is linear over small ranges of contrast, and if the decision between stimulus alternatives is based on maximum likelihood, one can derive simple proportional relationships for the contrast detection threshold, Δc_{det} , and the orientation and spatial frequency discrimination thresholds, $\Delta\theta_{\text{dis}}$ and $\Delta\omega_{\text{dis}}$:

$$\Delta c_{\text{det}} \propto \frac{\beta}{A} \quad (5)$$

$$\Delta\theta_{\text{dis}} \propto \sigma_{\theta} \sqrt{\frac{\beta}{A c_s}} \quad \Delta\omega_{\text{dis}} \propto \sigma_{\omega} \sqrt{\frac{\beta}{A c_s}} \quad (6,7)$$

These relations indicate how thresholds depend on filter parameters such as β/A , σ_{θ} and σ_{ω} . Note that not all thresholds depend on all parameters.

Interacting-filter model. The first stage of the model comprises 150 filters with Fourier representations in the shape of a 2-dimensional Gaussian centered on 30 orientations ($0 \leq \theta_i < \pi$) and 5 spatial frequencies ($2 \text{ cpd} \leq \omega_i < 8 \text{ cpd}$; Eq. 1). Increasing the number of filters does not alter model predictions substantially. From each Fourier representation, a filter pair in quadrature phase is reconstructed, and linear responses $E_{\theta,\omega}$ are computed numerically. The effective tuning widths are σ_{θ} and σ_{ω} for sinusoidal gratings, 1.1 σ_{θ} and 1.3 σ_{ω} for vertical stripes, and 1.7 σ_{θ} and 1.3 σ_{ω} for mask patterns. The second stage of the model is given by Eqs. 3 and 4. For $\gamma, \delta \gg 1$, only the largest first-stage responses produce significant second-stage responses. For $\gamma, \delta = 1$, first- and second-stage responses are proportional. At high stimulus contrast, second-stage responses follow a power law with exponent $\gamma - \delta$. The transducer function becomes increasingly sigmoidal with higher exponents γ, δ . The third stage of the model assumes that second-stage filter responses exhibit Gaussian-distributed noise (Eq. 2), and uses maximum-likelihood principles to predict ideal observer thresholds from the means and variances of these responses. Specifically, the Fisher information for each filter and for the entire population provides a lower limit for the variance of any unbiased estimate of a stimulus parameter such as contrast, orientation or spatial frequency^{15,48,49}.

Model fits. Fits involve downhill simplex error minimization⁵⁰, with simulated annealing overhead. For 'separate fits', the 10 free model parameters were fit either to single- or double-task data (each data set comprised 32 values from 5 experiments). Fig. 2 shows a plausible fit, whose fit error is 8% larger than the optimal fit. For the '12-dimensional joint fit' in Fig. 4, two parameters take different values for single- and double-task data. Thus, 12 free parameters are fit to 64 data values. Fig. 4 shows a plausible fit, whose fit error is 12% larger than the optimal fit. For the '18-dimensional joint fit' in Fig. 4, eight parameters assume different values for single- and double-task data. Thus, 18 free parameters are fit to 64 data values. Fig. 4 shows the optimal fit, in which attention has a number of physiologically implausible effects, among them a 6-fold increase in contrast gain, 70% sharper orientation tuning and 60% sharper spatial frequency tuning.

To assess the significance of different parameter values, a 'tolerance range' was computed for each parameter, within which the error of fit increases by no more than 10% (a small but noticeable degradation of the quality of fit). The difference between single- and double-task values of a given parameter was considered significant if it fell outside this tolerance range (Figs. 2f and 4f).

Positional uncertainty. As an alternative to the present model, we considered the possibility that attention lowers thresholds by reducing posi-

tional uncertainty. We used the formalism of ref. 22 and chose a noise level that reproduced 'fully attended' thresholds without positional uncertainty. Then we introduced positional uncertainty among the eight possible target locations to predict the 'poorly attended' thresholds. Positional uncertainty increased contrast detection thresholds by 19%, but left all other thresholds almost unchanged. For example, contrast discrimination thresholds at 10% pedestal contrast increased only by 1.5%. Thus, positional uncertainty cannot explain the pattern of attentional effects we have observed.

ACKNOWLEDGEMENTS

Supported by NSF, NIMH, ONR and the NSF-ERC at Caltech. We thank T.

Albright and T. Sejnowski for access to facilities and J. Gallant, A. Manwani, S. Shimojo, K. Watanabe and B. Zenger for comments and discussions.

RECEIVED 24 NOVEMBER 1998; ACCEPTED 24 FEBRUARY 1999

- Pashler, H. *The Psychology of Attention* (MIT Press, Cambridge, Massachusetts, 1997).
- Mueller, H. J. & Findlay, J. M. Sensitivity and criterion effects in the spatial cueing of visual attention. *Percept. Psychophys.* **42**, 383–399 (1987).
- Treisman, A. & Gormican, S. Feature analysis in early vision: evidence from search asymmetries. *Psychol. Rev.* **12**, 97–136 (1988).
- Braun, J. & Julesz, B. Withdrawing attention at little or no cost: Detection and discrimination tasks. *Percept. Psychophys.* **60**, 1–23 (1998).
- Lee, D. K., Koch, C. & Braun, J. Spatial vision thresholds in the near absence of attention. *Vision Res.* **37**, 2409–2418 (1997).
- Nachmias, J. & Sansbury, R. V. Letter: Grating contrast: discrimination may be better than detection. *Vision Res.* **14**, 1039–1042 (1974).
- Wilson, H. R. A transducer function for threshold and suprathreshold human vision. *Biol. Cybern.* **38**, 171–178 (1980).
- Legge, G. E. & Foley, J. M. Contrast masking in human vision. *J. Opt. Soc. Am.* **70**, 1458–1471 (1980).
- Bowne, S. F. Contrast discrimination cannot explain spatial frequency, orientation or temporal frequency discrimination. *Vision Res.* **30**, 449–461 (1990).
- Wilson, H. R. & Wilkinson, F. Evolving concepts of spatial channels in vision: from independence to nonlinear interactions. *Perception* **26**, 939–960 (1997).
- Wilson, H. R. & Humanski, R. Spatial frequency adaptation and contrast gain control. *Vision Res.* **33**, 1133–1149 (1993).
- Foley, J. M. Human luminance pattern-vision mechanisms: masking experiments require a new model. *J. Opt. Soc. Am.* **11**, 1710–1719 (1994).
- Zenger, B. & Sagi, D. Isolating excitatory and inhibitory nonlinear spatial interactions involved in contrast detection. *Vision Res.* **36**, 2497–2513 (1996).
- Carandini, M. & Heeger, D. J. Summation and division by neurons in primate visual cortex. *Science* **264**, 1333–1336 (1994).
- Itti, L., Braun, J., Lee, D. K. & Koch, C. in *Advances in Neural Information Processing Systems*, Vol. 9 (eds. Mozer, M. C., Jordan, M. I. & Petsche, T.) 173–179 (MIT Press, Cambridge, Massachusetts, 1997).
- Somers, D. C., Nelson, S. B. & Sur, M. An emergent model of orientation selectivity in cat visual cortical simple cells. *J. Neurosci.* **15**, 5448–5465 (1995).
- Carandini, M., Heeger, D. J. & Movshon, J. A. Linearity and normalization in simple cells of the macaque primary visual cortex. *J. Neurosci.* **17**, 8621–8644 (1997).
- Braun, J. & Sagi, D. Vision outside the focus of attention. *Percept. Psychophys.* **48**, 45–58 (1990).
- Braun, J. Visual search among items of different salience: removal of visual attention mimics a lesion in extrastriate area V4. *J. Neurosci.* **14**, 554–567 (1994).
- Palmer, J. Attention in visual search: distinguishing four causes of set-size effects. *Curr. Dir. Psychol. Sci.* **4**, 118–123 (1995).
- Solomon, J. A., Lavie, N. & Morgan, M. J. Contrast discrimination function: spatial cueing effects. *J. Opt. Soc. Am.* **14**, 2443–2448 (1997).
- Foley, J. M. & Schwarz, W. Spatial attention: effect of position uncertainty and number of distractor patterns on the threshold-versus-contrast function for contrast discrimination. *J. Opt. Soc. Am.* **15**, 1036–1046 (1998).
- DeValois, R. L. & DeValois, K. K. *Spatial Vision* (Oxford Univ. Press, New York, 1988).
- Wilson, H. R., Levi, D., Maffei, L., Rovamo, J. & DeValois, R. in *Visual Perception: The Neurophysiological Foundations* (eds. Spillmann, L. & Werner, J. S.) 231–272 (Academic, San Diego, California, 1990).
- Geisler, W. S. & Albrecht, D. G. Visual cortex neurons in monkeys and cats: detection, discrimination, and identification. *Vis. Neurosci.* **14**, 897–919 (1997).
- Bonnel, A.-M. & Miller, J. Attentional effects on concurrent psychophysical discriminations: investigations of a sample size model. *Percept. Psychophys.* **55**, 162–179 (1994).
- Lu, Z. L. & Doshier, B. A. External noise distinguishes attention mechanisms. *Vision Res.* **38**, 1183–1198 (1998).

28. Morgan, M. J., Ward, R. M. & Castet, E. Visual search for a tilted target: tests of spatial uncertainty models. *Q. J. Exp. Psychol.* **51**, 347–370 (1998).
29. Yeshurun, Y. & Carrasco, M. Spatial attention improves performance in spatial resolution task. *Vision Res.* **39**, 293–305 (1999).
30. Moran, J. & Desimone, R. Selective attention gates visual processing in the extrastriate cortex. *Science* **229**, 782–784 (1985).
31. Motter, B. C. Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J. Neurophysiol.* **70**, 909–919 (1993).
32. Treue, S. & Maunsell, J. H. Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* **382**, 539–541 (1996).
33. Luck, S. J., Chelazzi, L., Hillyard, S. A. & Desimone, R. Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J. Neurophysiol.* **77**, 24–42 (1997).
34. Roelfsema, P. R., Lamme, V. A. & Spekreijse, H. Object-based attention in the primary visual cortex of the macaque monkey. *Nature* **395**, 376–381 (1998).
35. Brefczynski, J. A. & DeYoe, E. A. A physiological correlate of the 'spotlight' of visual attention. *Nat. Neurosci.* **2**, 370–374 (1999).
36. Niebur, E. & Koch, C. A model for the neuronal implementation of selective visual attention based on temporal correlation among neurons. *J. Comput. Neurosci.* **1**, 141–158 (1994).
37. Tsotsos, J. K. *et al.* Modeling visual attention via selective tuning. *Artif. Intell.* **78**, 507–547 (1995).
38. Desimone, R. Visual attention mediated by biased competition in extrastriate visual cortex. *Phil. Trans. R. Soc. Lond. B* **353**, 1245–1255 (1998).
39. McAdams, C. J. & Maunsell, J. H. R. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* **19**, 431–441 (1999).
40. Spitzer, H., Desimone, R. & Moran, J. Increased attention enhances both behavioral and neuronal performance. *Science* **240**, 338–340 (1988).
41. Tyler, C. W. Colour bit-stealing to enhance the luminance resolution of digital displays on a single pixel basis. *Spatial Vis.* **10**, 369–377 (1997).
42. Duncan, J. Selective attention and the organization of visual information. *J. Exp. Psychol. Gen.* **113**, 501–517 (1984).
43. Pashler, H. Dual-task interference in simple tasks - Data and theory. *Psychology B* **116**, 220–244 (1994).
44. Doshier, B. A. & Sperling, G. in *Handbook of Perception and Cognition, Perception and Cognition at Century's End: History, Philosophy, Theory* (ed. Hochberg, J.) 201–254 (Academic, New York, 1998).
45. Duncan, J., Ward, R. & Shapiro, K. Direct measurement of attentional dwell time in human vision. *Nature* **369**, 313–315 (1994).
46. Braun, J. Vision and attention: the role of training. *Nature* **393**, 424–425 (1998).
47. Lee, D. K., Koch, C. & Braun, J. Visual attention is undifferentiated: concurrent discrimination of form, color, and motion. *Percept. Psychophys.* **61** (in press).
48. Seung, H. S. & Sompolinsky, H. Simple models for reading neuronal population codes. *Proc. Natl. Acad. Sci. USA* **90**, 10749–10753 (1993).
49. Pouget, A., Zhang, K., Deneve, S. & Latham, P. E. Statistically efficient estimation using population coding. *Neural Comput.* **10**, 373–401 (1998).
50. Press, W. H., Teukolsky, S. A., Vetterling, W. T. & Flannery, B. P. *Numerical Recipes in C*, 2nd ed. (Cambridge Univ. Press, Cambridge, 1992).