

Attending to orientation results in an inhibitory surround in orientation space

MICHAEL TOMBU

Vanderbilt University, Nashville, Tennessee

AND

JOHN K. TSOTSOS

York University, Toronto, Ontario, Canada

Subjects were required to attend to an orientation and make judgments about the stripes on briefly presented disks. Stripe orientation was varied so that they could be at, near, or far from the attended orientation. According to the selective-tuning model (Tsotsos, 1990; Tsotsos et al., 1995), attending to an orientation results in an inhibitory surround for nearby orientations, but not for orientations farther away. In line with this prediction, the results revealed an inhibitory surround. As in the spatial domain, attending to a point in orientation space results in an inhibitory surround for nearby orientations.

The visual system is an important tool that humans use to gather information about their surroundings. Vision is of high value for locomotion, threat detection, food acquisition, and a myriad of other tasks that, evolutionarily, have been (and continue to be) essential to survival. However, the computations required to fully analyze all of the information available to the human visual system exceed its capacity (Lennie, 2003; Tsotsos, 1990). Given this limitation, at any point in time, only some of the available information can be processed more fully. Although the system could select objects (or areas) of interest at random for further processing, this selection procedure would be inefficient and would lead to long search times to find objects of importance in complex visual scenes. From an evolutionary standpoint, when one applies this inefficient selection procedure to a threat detection situation, the likely end result is both tragic and fatal. Given the importance of being able to identify objects of interest in a visual scene quickly, it is critical that the visual system be equipped with an efficient mechanism that selects objects of interest for further processing on the basis of their likely importance. This mechanism is generally referred to as *visual attention*.

Linking Physiology and Theory

Any model of visual attention must abide by the constraints imposed by the biology of the visual system in order to be viable. Structurally, the visual system consists of multiple processing areas arranged hierarchically (Felleman & Van Essen, 1991). Early areas, such as V1, are retinotopically organized, with neurons that have relatively small receptive fields and respond optimally to sim-

ple visual stimuli such as oriented bars of light (Kastner & Ungerleider, 2000). At progressively higher visual areas, receptive field size increases, as does the complexity of the visual information being represented. This gradual increase in receptive field size and complexity of information necessitates a many-to-one mapping between neurons as information is fed through the visual system. Connections between areas are generally reciprocal, allowing for both feedforward and feedback processing (Felleman & Van Essen, 1991).

In addition to the orderly representation of spatial information within visual areas, other, nonspatial stimulus features are also organized in an orderly manner. Hubel, Wiesel, and colleagues have demonstrated in monkeys both ocular dominance (LeVay, Hubel, & Wiesel, 1975) and orientation (Hubel, Wiesel, & Stryker, 1977) columns in area V1. A similar neural organization may also exist for color (Dow, 2002). Orientation columns have also been observed in extrastriate area V2, and foveal representations in V4 (Ghose & Ts'o, 1997), as well as VP and V3 (Vanduffel, Tootell, Schoups, & Orban, 2002).

Orientation-selective neurons respond to a range of orientations near their preferred orientation. This range, as measured by median orientation bandwidth (bandwidth at half the maximal response), is less than 30° in areas V1 and V2 (Levitt, Kiper, & Movshon, 1994) and about 52° in area V4 (Desimone & Schein, 1987). Orientation bandwidth can be modulated by attention. Neurons in area V4 exhibit narrower orientation bandwidths when a particular orientation is attended to (Haenny & Schiller, 1988; Spitzer, Desimone, & Moran, 1988). Using 2-deoxyglucose autoradiography and retrograde labeling,

Gilbert and Wiesel (1989) have shown in the cat that cells tuned to vertical or horizontal orientations in area 18 receive inputs from multiple clusters of cells in area 17 with similar orientation tuning. As in the spatial domain, this implies a many-to-one mapping of orientation-selective neurons from visual area to subsequent visual area. In this case, 2-deoxyglucose autoradiography identifies cells as vertical or horizontal. Therefore, any cell that prefers an orientation closer to vertical than to horizontal will be identified as vertical by this technique. The fact that area 18 neurons are interconnected with similarly tuned cells in area 17 suggests that area 18 neurons selective for orientation are interconnected with cells in area 17 over a range of preferred orientations that are near their preferred orientation. This implies not only a many-to-one pattern of connectivity between areas 17 and 18, but also that neurons in area 18 receiving preferential input from a range of neurons with similar preferred orientations. Orientation-tuning curves support the hypothesis that orientation-selective neurons receive input from neurons over a range of preferred orientations.

Modeling Visual Attention

Most models of visual attention posit that objects of interest are selected on the basis of their saliency. Saliency is determined by both bottom-up and top-down factors. Bottom-up factors affecting saliency involve physical differences between stimuli, so that stimuli that are more different from their neighbors are more salient than stimuli that are less different from their neighbors. Top-down factors, such as foreknowledge of target identity, can also influence the selection process by biasing selection in favor of stimuli sharing physical attributes with the target or by emphasizing one physical attribute dimension over another.

The consequences of selection for attended and unattended stimuli vary from model to model. In the spatial domain, most models postulate that processing at the attended location is either protected (i.e., left alone) or enhanced, whereas processing at unattended locations is inhibited or ignored (left alone). Having said that, different models predict a different pattern of inhibition as a function of the spatial distance from the attended location. Models can be divided into three broad categories based on the predicted pattern of inhibition as a function of spatial distance from the attended location. The first class of models predicts that unattended locations will be equally inhibited (or not enhanced) and can be broadly termed *spotlight* models of attention (e.g., Koch & Ullman, 1985; Treisman & Gelade, 1980). The second class of models predicts that inhibition will gradually increase as the distance from the attended location increases (or that the enhancement caused by attending will slowly decrease as the distance from the attended location increases). These models are often referred to as *gradient* models of attention (e.g., LaBerge & Brown, 1989). Finally, a third class of models predicts an inhibitory surround around an attended stimulus, with inhibition gradually dissipating as the distance from the attended location increases (e.g., Desimone & Duncan, 1995; Tsotsos, 1990; Tsotsos et al., 1995). The explanation for this inhibitory surround

differs across models, with the *biased competition model* (Desimone & Duncan, 1995) ascribing it to competition for representation at the level of receptive fields, whereas the *selective-tuning model* (Tsotsos, 1990; Tsotsos et al., 1995) posits that it is a result of competition any time there is a many-to-one mapping of a stimulus dimension at successively higher levels of information processing.

Previous behavioral work (Caputo & Guerra, 1998; Cutzu & Tsotsos, 2003; Hopf et al., 2006; Mounts, 2000; Müller & Kleinschmidt, 2004; Müller, Mollenhauer, Rösler, & Kleinschmidt, 2005; Steinman, Steinman, & Lehmkuhle, 1995) has demonstrated that attending to a spatial location results in an inhibitory surround. Specifically, larger inhibitory effects were observed at spatial locations closer to an attended location than at spatial locations more distant from the attended location. This result is consistent with the third class of models described above. In this article, we intend to explore one such model, the selective-tuning model, and to test an additional prediction it makes.

The Selective-Tuning Model

The selective-tuning model (Tsotsos et al., 1995) hypothesizes that the process of selection begins at the highest level of the visual system and that, through feedback connections, it progressively works its way down to the earliest level, where it ultimately selects a focus of attention. Because of the progressive increase in both receptive field size and the complexity of information represented by neurons in higher visual areas, selecting a winning neuron (or group of neurons) at the highest level of the visual system acts to constrain the identity of the most salient stimulus in the visual field, but it does not identify it. As a result, the neurons that provide the input to the winning neuron must also compete to determine the most salient neuron. In order to eliminate interference at higher levels, the feedforward connections of losing neurons are inhibited, whereas the winner's inputs are spared. This process is repeated at successively lower levels of the visual system until an attentional "beam," consisting of the spared connections surrounded by the inhibited connections, forms. Visual information is then passed back through the inhibited visual system. This allows the representation of the selected stimulus to pass through the network without any interference being caused by other nearby representations that might have competed, by mutual suppression, with the selected stimulus. This attentional beam can be deployed in advance of stimulus presentation on the basis of top-down direction.

Because only the inputs to a winning neuron compete in the selection process at the next lower level, only inputs to a winner can be inhibited. Therefore, the model predicts that stimuli that do not contribute to a winner will not be inhibited. For example, in the spatial domain, receptive field size increases at successively higher visual areas, indicating that many neurons with neighboring receptive fields contribute to neurons at higher levels. As a result, a winner at one level is fed by a group of neurons with nearby receptive fields at the next lower level. One of these neurons will then be selected as the most salient,

whereas the remainder will be inhibited. The selective-tuning model therefore predicts that spatial locations near a selected stimulus will be inhibited, whereas stimulus locations further away will not, which is exactly the pattern of results that has been observed behaviorally (e.g., Caputo & Guerra, 1998; Hopf et al., 2006; Mounts, 2000; Müller & Kleinschmidt, 2004; Müller et al., 2005; Steinman et al., 1995; Cutzu & Tsotsos, 2003).

There is also neurological evidence supporting this prediction. Results from single-cell recording studies in monkeys indicate that when two stimuli are simultaneously present in the receptive field of a neuron, the resulting firing rate is a value intermediate to the firing rates observed when either stimulus is presented alone (Luck, Chelazzi, Hillyard, & Desimone, 1997; Moran & Desimone, 1985). However, when one of the two stimuli is attended, the firing rate of the neuron closely resembles the firing rate observed when only the attended stimulus is presented. Importantly, this suppression of unattended stimuli is restricted to stimuli sharing a receptive field with the attended stimulus. When the attended stimulus is moved outside of the receptive field of the neuron under investigation, the firing rate to the unattended stimulus is affected to a much lesser degree, or not at all (Moran & Desimone, 1985; Motter, 1993). Directing attention to a spatial location in the absence of visual stimulation also modulates firing rates. Increases in baseline firing rates are observed for neurons at attended locations prior to the presentation of the to-be-attended stimulus (Luck et al., 1997).

It is important to emphasize that the selective-tuning model also predicts inhibitory surrounds for other feature dimensions where there is an ordered representation of the stimulus dimension and a many-to-one mapping of neurons from lower to higher visual areas. As was reviewed earlier, a many-to-one mapping from one visual area to the next also exists for orientation selectivity. Neurons with preferred orientations are expected to have preferential input from neurons with similar preferred orientations. Therefore, the selective-tuning model predicts that directing attention to a point in orientation space will result in the inhibition of stimuli at nearby orientations, but not of stimuli at more distant orientations. It is worth noting that the selective-tuning model makes predictions in the spatial domain similar to those of the biased competition model (Desimone & Duncan, 1995). However, the prediction that inhibitory surrounds should be observed any time there are many-to-one mappings is explicitly made only by the selective-tuning model. Our aim in the present investigation was to test this prediction in the orientation domain.

The Present Experiment

The purpose of the present experiment was to assess the effects of directing attention in orientation space. Specifically, when subjects direct attention to a point in orientation space (i.e., an orientation), does an inhibitory surround form for orientations near the attended orientation, as is predicted by the selective-tuning model?

In this experiment, subjects attended to a particular orientation and judged whether stripes on a disk were straight or jagged (see Figure 1). The orientation of the




Figure 1. Striped disks, with stripe orientations (clockwise from vertical) of 22°, 67°, 112°, or 157° (22° orientation disk shown), were employed as stimuli. The subjects were instructed to indicate whether the stripes on the disk were straight (A) or jagged (B). Stimulus presentation times varied between subjects from 67 to 150 msec, as well as from block to block for each subject, in order to keep performance levels at approximately 80% correct (C). A 166-msec mask preceded and followed stimulus presentation. At the beginning of each block, the subjects were instructed to attend to one orientation. For that block, 70% of the trials were at the attended orientation, with each unattended orientation presented 10% of the time. Within each condition, the stripes on the disk were straight 50% of the time and jagged 50% of the time. Each orientation was the attended orientation for two blocks per session.

stripes, although irrelevant to the task, was varied, with stripe orientations of 22°, 67°, 112°, and 157° clockwise from vertical being employed. During each block, 70 trials were presented at the attended orientation (which varied from block to block), as well as 10 trials at each of the unattended orientations. If inhibitory surrounds exist in orientation space, performance should be highest for stripe orientations at the attended orientation, impaired for stripe orientations near the attended orientation, and show sparing for orientations far from the attended orientation. An important strength of this experimental design is that each stimulus (straight or jagged stripes; 22°, 67°, 112°, or 157° oriented stripes) contributes equally to each trial type, thereby eliminating potential stimulus artifact confounds.

METHOD

Subjects

Eight subjects (1 female), including the first author, between 21 and 30 years of age, participated in this experiment. All the subjects had normal or corrected-to-normal vision.

Stimuli

The stimuli employed were striped circular disks oriented at one of four (22°, 67°, 112°, or 157° clockwise from vertical) orientations (see Figure 1A for an example). The disks had a radius of 24 mm, and there were 36 stripes per disk (stripe width of 1.33 mm). Stripes were of alternating light (mean luminance = 6.960 cd/m²) and dark gray (mean luminance = 2.785 cd/m²) on a medium gray background (mean luminance = 4.710 cd/m²) at a viewing distance of 57 cm. The disks and backgrounds were roughly equiluminant. The stripes on the disks could be either straight (see Figure 1A) or jagged (see Figure 1B).

Procedure

Each trial began with the presentation of a 4.8 × 4.8 cm fixation cross at screen center, which indicated the position where the stimulus sequence was to be presented. Each trial was initiated with a keypress, and 500 msec later, the stimulus sequence was presented. Each target stimulus (the striped disk; see, e.g., Figures 1A and 1B) was preceded and followed by an identical mask stimulus (see Figure 1C) for 166 msec. The target stimulus presentation time varied from subject to subject over the course of the experiment. Overall, presentation times varied from 67 to 150 msec. Following stimulus presentation, the subject indicated, with a manual buttonpress, whether the stripes on the target stimulus were straight or jagged.

The session began with a brief description of the experiment. A scoring system was implemented to encourage optimal performance. Before each block, feedback about the score on the previous block, total score this session, and the high score for a single block (by any subject) was provided. The subjects were then shown the to-be-attended orientation for that block. Each block consisted of 100 trials, of which 70 were at the to-be-attended orientation and 10 were at each of the unattended orientations, randomly intermixed. In all conditions, half of the trials in each block were jagged and half were straight. If accuracy on a given block was worse than 65% or exceeded 90%, the presentation time for the target stimulus was adjusted up or down by 16.7 msec. Each session consisted of eight blocks, and each orientation was selected as the to-be-attended orientation twice. Four orientations, four attentional conditions (attended, unattended near clockwise, unattended near counterclockwise, and unattended far), and two target types (jagged and straight) were employed in this experiment, for a total of 32 conditions. For each of the 24 unattended conditions, the subjects performed 40 trials, and for each of the 8 attended conditions, the subjects performed 280 trials (per session).

The subjects performed between four and eight sessions, of which only the final four were analyzed. Some subjects performed more sessions than others because of variability in determining the optimal presentation time. The presentation time had to be sufficiently fast to engage attention but sufficiently slow to be able to perform the task.

RESULTS

A 2 × 4 ANOVA was performed on accuracy results with target type (jagged or straight) and attention condition (attended, unattended near clockwise, unattended near counterclockwise, or unattended far) as variables. The results are displayed in Figure 2.

The results indicate a significant interaction between attention condition and target type [$F(3,21) = 17.1$, $MS_e = 0.00177$, $p < .0001$]. There were no differences in performance across attention conditions for the straight target type condition. However, when the stripes on the disks were jagged, the subjects performed most accurately when the stripes were in the attended orientation (proportion of .796 correct), followed by the unattended far condition (.707 correct), and finally, were least accurate in the unattended near (clockwise and counterclockwise) conditions (.637 and .625, respectively). The main effects of target type and attention condition were also significant [$F(1,7) = 6.8$, $MS_e = 0.06578$, $p < .04$, and $F(3,21) = 22.6$, $MS_e = 0.00090$, $p < .0001$, respectively]. The sub-

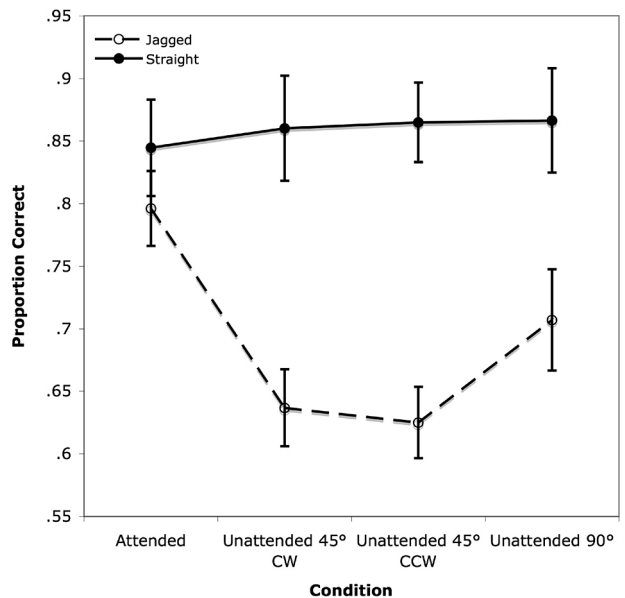


Figure 2. Results for 8 subjects collapsing across stripe orientation. When the stripes were straight, no effect of attending to a particular orientation was observed. However, when the stripes were jagged, the subjects performed best when the stimulus was at the attended orientation, at an intermediate level when the stimulus was at an orientation far from the attended orientation, and worst when the stripe orientation was near the attended orientation. Performance when the stripes were near the attended orientation was significantly worse than when the stripe orientation was far from the attended orientation. This pattern of results is consistent with the selective-tuning model. CW, clockwise; CCW, counterclockwise.

jects responded more accurately when the stripes on the disks were straight than when they were jagged (.859 and .691 correct, respectively), indicating a response bias for straight stripes. Performance was best when the presented orientation matched the to-be-attended orientation (.820 correct), at an intermediate level when the presented orientation was far from the attended orientation (.787 correct) and worst when the presented orientation was near the attended orientation (.748 and .745 correct for the 45° clockwise from attended and 45° counterclockwise from attended, respectively).

It was important to determine whether performance at unattended orientations near the attended orientation was inhibited and whether performance at the attended orientation was facilitated, relative to a baseline. The unattended orientation far from the attended orientation condition (unattended 90° from the attended orientation) was used as a baseline to test for facilitation and inhibition.

Test of Inhibition

A separate 2 (jagged or straight) \times 2 (unattended near or unattended far) ANOVA was conducted to test for inhibition of unattended orientations near the attended orientation. For the purpose of this set of analyses, the unattended near clockwise and counterclockwise conditions were collapsed and compared with unattended orientations far from the attended orientation.

Most important, the subjects responded more accurately when the stripes were presented at an orientation far from the attended orientation than when the stripes were presented at an orientation near to the attended orientation [$F(1,7) = 60.9$, $MS_e = 0.00021$, $p < .0001$]. The subjects also responded more accurately when the stripes on the disks were straight (.864 correct) than when they were jagged (.669 correct) [$F(1,7) = 8.3$, $MS_e = 0.03689$, $p < .03$]. The interaction between target type and attention condition approached significance [$F(1,7) = 5.3$, $MS_e = 0.00198$, $p < .06$]. A test of simple effects indicates that the inhibitory effect (worse performance for orientations near the attended orientation than for orientations far from the attended orientation) was reliable only when the stripes on the disks were jagged [.076 disadvantage in proportion correct for the unattended near condition relative to the unattended far condition; $F(1,7) = 18.4$, $MS_e = 0.00126$, $p < .004$]. When the stripes on the disks were straight, no reliable inhibitory effect (.004 disadvantage in proportion correct for the unattended near condition relative to the unattended far condition) was observed ($F < 1$). In line with the selective-tuning model, significant inhibition was observed for orientations near the attended orientation relative to orientations farther away from the attended orientation. However this was the case only when the stripes on the disks were jagged. When the stripes were straight, no inhibition was observed.

Test of Facilitation

Another separate 2 (jagged or straight) \times 2 (attended or unattended far) ANOVA was conducted on accuracy results.

The results indicated a significant interaction between attention condition and target type [$F(1,7) = 9.5$, $MS_e = 0.00260$, $p < .02$]. Facilitation was observed when the target was jagged (.089 benefit for the attended condition in proportion correct), but not when the target was straight (.022 disadvantage for the attended condition in proportion correct).

Relative to unattended orientations far from the attended orientation, significant facilitation was observed for orientations presented at the attended orientation; however, this was the case only when the stripes on the disks were jagged. When the stripes on the disks were straight, no facilitation was observed.

DISCUSSION

In this experiment, the subjects attended to an expected orientation and were presented with a striped disk to which they made a judgment as to whether the stripes were jagged or straight. The stripes could be in the orientation to which the subjects were attending, at an orientation near to the attended orientation, or at an orientation far from the attended orientation. According to the selective-tuning model, directing attention to a point in orientation space should result in an inhibitory surround for orientations near the attended orientation, but not for orientations farther away.

The results clearly indicate that directing attention to a point in orientation space, even in the absence of competing stimuli, facilitates performance for the attended orientation and inhibits performance for orientations near the attended orientation. However, this pattern of results was observed only when the stripes on the disk were jagged. When they were straight, no effect of attention was observed. One explanation for this pattern of results is that attending to an orientation sharpens perception of that orientation, allowing the jaggedness of the stripes to be better perceived, which would explain why an inhibitory surround was observed when the stripes were jagged. However, this sharpening effect would aid perception only for jagged stimuli. When the stripes were straight, both sharpened (attended) and blurred (unattended) stimuli would appear to be straight, resulting in equivalent performance for all the attention conditions. The results support this interpretation; performance was better when the stripes on the disk were straight (proportion of .859 correct) than when they were jagged (.691 correct), and an inhibitory surround was observed only when the stripes were jagged. In line with this hypothesis, Carrasco and colleagues (Carrasco, Loula, & Ho, 2006) have shown that covert spatial attention results in a sensitivity shift toward high spatial frequencies. Taken as a whole, these results support the specific prediction of the selective-tuning model of an inhibitory surround in orientation space when subjects direct attention to a point in orientation space, as well as the class of models positing inhibitory attentional surrounds. Other models in this class, such as the biased competition model (Desimone & Duncan, 1995), would need to be extended to account for the observed inhibitory surround in orientation space that we observed in this experiment.

A result of the design of this experiment is that trials in the attended orientation outnumbered trials in each of the unattended conditions seven to one. Stimulus and condition repetitions may, therefore, account for some or all of the observed facilitation for the attended condition. However, the inhibition observed for unattended orientations near the attended orientation was also significant, relative to unattended orientations far from the attended orientation. Because all of the unattended conditions were presented infrequently and the same number of times, it is unlikely that stimulus or condition repetitions were responsible for the inhibition observed for unattended orientations near to the attended orientation.

Previous examinations of attention in nonspatial domains often directed subjects to attend to one of two (or more) nonspatial feature domains (e.g., attend to color, ignore orientation; see, e.g., Corbetta, Miezin, Dobmeyer, Shulman, & Petersen, 1990). In this sense, attention is manipulated *between* nonspatial domains. In most spatial attention studies, attention is directed to one of two (or more) spatial locations. In this sense, attention is manipulated *within* the spatial domain (e.g., Cutzu & Tsotsos, 2003). In order to make comparisons between spatial and nonspatial attention, it seems important that attention be manipulated in the same manner for both cases. An important feature of this experiment is that nonspatial attention is manipulated *within* the orientation domain as is done in the spatial domain. When this is done, the results in the orientation domain appear to agree with those from studies examining spatial attention, suggesting the parsimonious conclusion that directing attention spatially and directing attention to orientation operate in a similar manner. In both cases, an inhibitory surround is found for stimulus features near the attended stimulus feature, whether it be in the spatial or the orientation domain—a conclusion in line with the predictions made by the selective-tuning model.

AUTHOR NOTE

This work was supported by an NSERC postdoctoral fellowship to the first author. Correspondence concerning this article should be addressed to M. Tombu, Department of Psychology, Vanderbilt University, Wilson Hall, 111 21st Avenue South, Nashville, TN 37203 (e-mail: mike.tombu@vanderbilt.edu).

REFERENCES

CAPUTO, G., & GUERRA, S. (1998). Attentional selection by distractor suppression. *Vision Research*, **38**, 669-689.

CARRASCO, M., LOULA, F., & HO, Y.-X. (2006). How attention enhances spatial resolution: Evidence from selective adaptation to spatial frequency. *Perception & Psychophysics*, **68**, 1004-1012.

CORBETTA, M., MIEZIN, F. M., DOBMEYER, S., SHULMAN, G. L., & PETERSEN, S. E. (1990). Attentional modulation of neural processing of shape, color, and velocity in humans. *Science*, **248**, 1556-1559.

CUTZU, F., & TSOLOS, J. K. (2003). The selective tuning model of attention: Psychophysical evidence for a suppressive annulus around an attended item. *Vision Research*, **43**, 205-219.

DESIMONE, R., & DUNCAN, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, **18**, 193-222.

DESIMONE, R., & SCHEIN, S. J. (1987). Visual properties of neurons in area V4 of the macaque: Sensitivity to stimulus form. *Journal of Neurophysiology*, **57**, 835-868.

DOW, B. M. (2002). Orientation and color columns in monkey visual cortex. *Cerebral Cortex*, **12**, 1005-1015.

FELLEMAN, D. J., & VAN ESSEN, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, **1**, 1-47.

GHOSE, G., & TS'O, D. (1997). Form processing modules in primate area V4. *Journal of Neurophysiology*, **77**, 2191-2196.

GILBERT, C. D., & WIESEL, T. N. (1989). Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *Journal of Neuroscience*, **9**, 2432-2442.

HAENNY, P. E., & SCHILLER, P. H. (1988). State dependent activity in monkey visual cortex: I. Single cell activity in V1 and V4 on visual tasks. *Experimental Brain Research*, **69**, 225-244.

HOPF, J.-M., BOEHLER, C. N., LUCK, S. J., TSOLOS, J. K., HEINZE, H.-J., & SCHOENFELD, M. A. (2006). Direct neurophysiological evidence for spatial suppression surrounding the focus of attention in vision. *Proceedings of the National Academy of Sciences*, **103**, 1053-1058.

HUBEL, D. H., WIESEL, T. N., & STRYKER, M. P. (1977). Orientation columns in macaque monkey visual cortex demonstrated by 2-deoxyglucose autoradiographic technique. *Nature*, **269**, 328-330.

KASTNER, S., & UNGERLEIDER, L. G. (2000). Mechanisms of visual attention in the human cortex. *Annual Review of Neuroscience*, **23**, 315-341.

KOCH, C., & ULLMAN, S. (1985). Shifts in selective visual attention: Towards the underlying neural circuitry. *Human Neurobiology*, **4**, 219-227.

LABERGE, D., & BROWN, V. (1989). Theory of attentional operations in shape identification. *Psychological Review*, **96**, 101-124.

LENNIE, P. (2003). The cost of cortical computation. *Current Biology*, **13**, 493-497.

LEVAY, S., HUBEL, D. H., & WIESEL, T. N. (1975). The pattern of ocular dominance columns in macaque visual cortex revealed by a reduced silver stain. *Journal of Comparative Neurology*, **159**, 559-576.

LEVITT, J. B., KIPER, D. C., & MOVSHON, J. A. (1994). Receptive fields and functional architecture of macaque V2. *Journal of Neurophysiology*, **71**, 2517-2542.

LUCK, S. J., CHELAZZI, L., HILLYARD, S. A., & DESIMONE, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, **77**, 24-42.

MORAN, J., & DESIMONE, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, **229**, 782-784.

MOTTER, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *Journal of Neurophysiology*, **70**, 909-919.

MOUNTS, J. R. W. (2000). Evidence for suppressive mechanisms in attentional selection: Feature singletons produce inhibitor surrounds. *Perception & Psychophysics*, **62**, 969-983.

MÜLLER, N. G., & KLEINSCHMIDT, A. (2004). The attentional "spotlight's" penumbra: Center-surround modulation in striate cortex. *NeuroReport*, **15**, 977-980.

MÜLLER, N. G., MOLLENHAUER, M., RÖSLER, A., & KLEINSCHMIDT, A. (2005). The attentional field has a Mexican hat distribution. *Vision Research*, **45**, 1129-1137.

SPITZER, H., DESIMONE, R., & MORAN, J. (1988). Increased attention enhances both behavioral and neuronal performance. *Science*, **240**, 338-340.

STEINMAN, B. A., STEINMAN, S. B., & LEHMKUHL, S. (1995). Visual attention mechanisms show a center-surround organization. *Vision Research*, **35**, 1859-1869.

TREISMAN, A. M., & GELADE, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, **12**, 97-136.

TSOLOS, J. K. (1990). Analyzing vision at the complexity level. *Behavioral & Brain Sciences*, **13**, 423-469.

TSOLOS, J. K., CULHANE, S. M., WAI, W. Y. K., LAI, Y., DAVIS, N., & NUFLO, F. (1995). Modeling visual attention via selective tuning. *Artificial Intelligence*, **78**, 507-545.

VANDUFFEL, W., TOOTELL, R., SCHOUPS, A., & ORBAN, G. (2002). The organization of orientation selectivity throughout macaque visual cortex. *Cerebral Cortex*, **12**, 647-662.