CHAPTER

79

Lateral Interactions between Targets and Flankers Require Attention

Elliot Freeman

ABSTRACT

What we see from moment to moment depends on the current sensory context, and also on the behavioural context of the task in which we are engaged. However, little is understood about how these different factors interact. Integration of visual context has often been studied using the phenomenon of lateral interactions, in which perceptual sensitivity and also neural activity in early visual cortex to a central oriented target is enhanced in the presence of collinear flanking patches (lined up to form a virtual contour). Our recent psychophysical studies suggest that task-directed attention to the flankers can directly modulate lateral interactions, affecting integration of the whole flanker-target configuration, rather than just the perceptibility of the individual parts. This attentional modulation is highly specific to collinear configurations and also dependent on tasks involving discrimination of global rather than local stimulus properties. Our results suggest that attention plays a critical role in the selective integration of task-relevant groups, directly modulating early visual mechanisms sensitive to collinear structures.

I. STIMULUS AND TASK-RELATED CONTEXT EFFECTS IN EARLY VISION

In everyday activity, the visual system routinely faces the problem of disentangling contours relating to objects of interest from those caused by reflections, shadows, or occlusions by other irrelevant objects. What we ultimately experience depends not only on stimulus-driven grouping processes, sensitive to the geometrical relationships between different parts of a scene, but also on task-driven attentional factors sensitive to which aspects of the scene are currently relevant. Grouping has traditionally been seen as a separate stage preceding grouping (Driver, Davis, Russell, Turatto, and Freeman, 2001). However, recent data suggest that grouping and selective processes can both operate at similarly early stages in visual processing (see Chapters 107 and 47).

Much has been learned about stimulus-determined context-integration, from the phenomenon of lateral interactions (Polat, 1999). For example, perceptibility of an oriented target element, and neural activity in early visual cortex (e.g., V1), can both be enhanced in the context of flanking elements (positioned outside the target cell's classical receptive field), when all are lined up to form a virtual contour (Kapadia, Ito, Gilbert, and Westheimer, 1995). Evidence suggests that this collinearity effect reflects spreading of excitatory impulses via horizontal connections, which tend to connect V1 neurones along lines parallel to their orientation preference (Callaway, 1998).

We devised a method to test whether this phenomenon of lateral interactions might be modulated as a function of changing attentional demands (Freeman, Sagi, and Driver, 2001). Our results suggest that the mechanisms by which local contextual components are integrated into global groups can indeed be selectively modulated by attention, depending on the relevance of specific perceptual groups to the task at hand. Given the neural underpinnings of lateral interactions, any such modulation should shed light on the mechanisms by which grouping and selective processes interface with each other.

II. BASIC DUAL-TASK DUAL-AXIS PROCEDURE

Our basic method used a novel dual-axis stimulus, with four suprathreshold flankers surrounding a central near-threshold target (see Fig. 79.1a). The orientation of the central target was collinear with the pair of flankers along just one of the two axes (though with small misalignments, see later), but was orthogonal to the orientation of each of the flankers in the pair along the other axis. The primary task was to indicate in which of the two successive intervals the near-threshold central target was presented. To direct covert attention to one or other flanker-pair, we imposed a secondary Vernier discrimination task to be performed concurrently with the central target detection task. Subjects were instructed to judge the relative position of the flankers on just one prespecified axis, indicating in which of two successive intervals they were displaced in a prespecified direction. The two flankers on the other axis were to be ignored. Two conditions may thus be compared: attending to the collinear flankers versus attending to the orthogonal flankers, while the stimuli and attention to the central target are kept constant.

This approach differs from most others, where the intensity or spatial distribution of attention is manipulated over the whole stimulus display, for example by varying the difficulty of a secondary task or by directing attention to one or several spatial positions (e.g., Gilbert, Ito, Kapadia, and Westheimer, 2000; see also Chapters 76 and 74). Such manipulations might alter local processing of display elements in a number of ways, affecting spatial uncertainty, spatial resolution, or internal noise levels within restricted regions (see Chapter 73). By controlling these factors it becomes possible, in principle, to measure specifically how attention influences lateral interactions between flankers and target, while eliminating the possible confounding effect of local attentionally driven changes in the processing of the target.

This method was successful in demonstrating that lateral interactions depend on attention to the flankers, in the absence of any changes in the stimulus (Freeman et al., 2001). The graph in Fig. 79.2a (filled symbols) shows target contrast thresholds averaged across 16 subjects (all naïve except for EF). Central target detection was facilitated (i.e., target contrast thresholds were reduced) when the collinear flankers were attended, but not when the orthogonal flankers were attended instead, despite the continued presence of ignored collinear flankers. This result indicates that attention can modulate the lateral interactions phenomenon. On their own, however, these data would be consistent with a range of different accounts for how this attentional modulation might be achieved. Our new results help to narrow down these possibilities, characterizing the functional role played by attention, its constraints, and the likely mechanisms of modulation.

III. STRENGTH AND SPECIFICITY OF ATTENTIONAL MODULATION

A. Collinearity vs. Local Feature Similarity

The defining property of classic lateral-interaction phenomena is their sensitivity to *collinear* configurations in particular. If our attentional effect specifically modulates the underlying mechanisms of lateral interactions, it should likewise be configuration-specific, found only when stimulus elements have strictly collinear arrangements. An alternative account of the attentional effects, which would not be configuration-specific, can be derived from current influential accounts in terms of feature-based attention (e.g., Saenz, Buracas, and Boynton, 2002; Chapter 49). On such accounts, attention to a given feature at one location automatically enhances perception of all similar features throughout the visual field. This might explain our initial results, given that in collinear configurations the central target has the same local orientation as each of the collinear flankers. The central target might therefore in principle have been facilitated by attending just the local orientation of the relevant flankers, rather than by a global collinear arrangement *per se*.

A test of this feature-specific account was recently reported in an independent study by Saenz and Boynton (2002). This study used stimuli in which target and flankers on one axis were arranged parallel to each other (e.g., ••), so that they had the same local orientations but no global collinearity. In apparent support for the feature-specific hypothesis, target detection improved slightly when this axis was attended, compared to attention to the other orthogonal axis. However, some past studies have reported target facilitation even with this parallel target-flanker configuration, though in the absence of attentional manipulations and with only one pair of flankers present (see Polat, 1999). We have tested the feature-specific account using a configuration with which facilitation has never previously been found (Freeman, Driver and Sagi, 2004). We disrupted collinearity by rotating all patches by 45 degrees, so that local orientation similarity was preserved on one axis (see Fig. 79.1b, right) while the other axis remained orthogonal

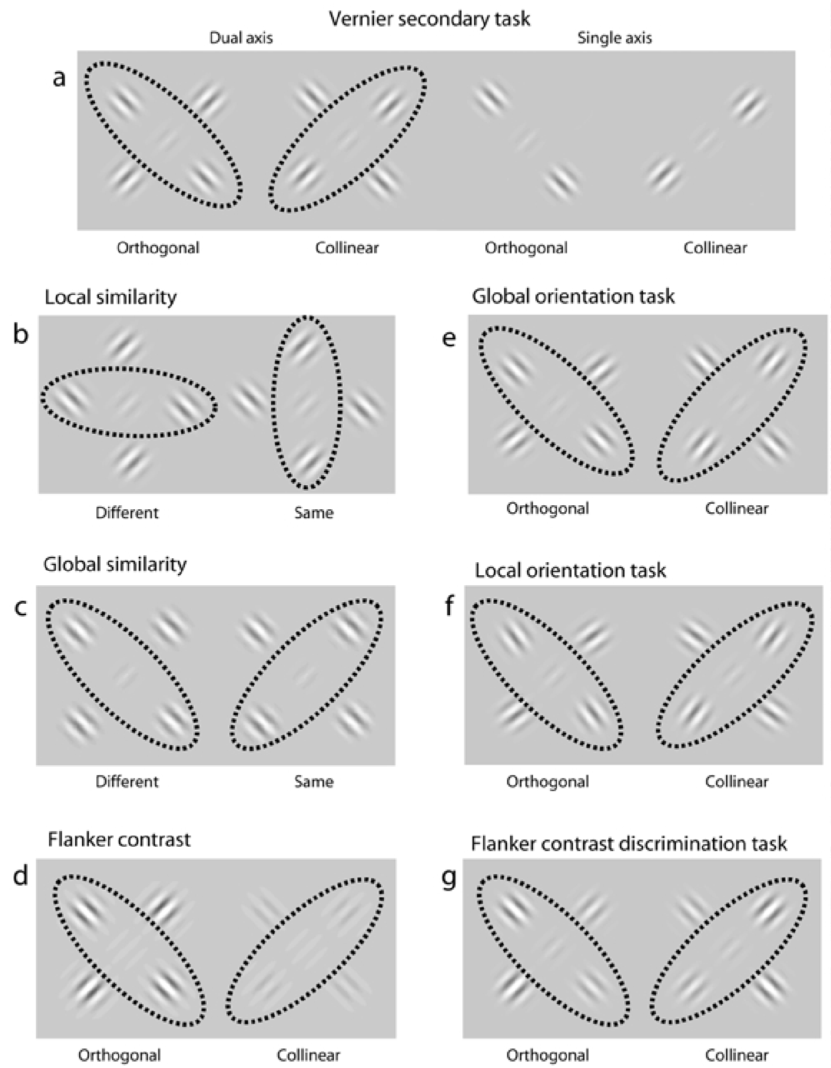

FIGURE 79.1 Sample stimuli, composed of a central low-contrast target Gabor patch (Gaussian envelope distribution and carrier wavelength both equal to 0.15 degrees), surrounded by one (single-axis) or two (dual-axis) pairs of flankers (flanker-target separation approx. four wavelengths). Dotted ellipses illustrate the two attentional conditions. Observers perform a task on one flanker-pair while ignoring the other, at the same time as detecting the central target (2×2-interval forced-choice, with 80 ms display intervals and 500 ms ISI). **A.** Stimuli for the Vernier task, with the two leftmost images depicting dual-axis stimuli in attend-orthogonal and attend-collinear conditions; the rightmost two images show single-axis orthogonal and collinear stimuli. **B.** Local similarity manipulation: target has either the same (right) or 90 degrees different (left) orientation to the flankers attended for the Vernier task. **C.** Global similarity: target has either the same or different orientation as the virtual contour connecting the task-relevant flankers. **D.** Flanker contrast manipulation: same as in **A** but with high or low contrast flankers. **E.** Global orientation task: observers discriminate between two rotations of a specified flanker-axis. **F.** Local orientation task: observers discriminate rotation of flankers around their own centre. **G.** Flanker-contrast discrimination task: observers decide which flanker of a relevant pair has the higher contrast.

to the target orientation (see Fig. 79.1b, left). We found that the attentional effect was completely eliminated for five out of six subjects (subject mean shown in Fig. 79.2b), compared to our standard stimuli (Figs. 79.1a and 79.2a) containing collinear elements. Feature-specific attention therefore cannot account for all of the attentional modulation effects found using the present lateral interactions paradigm.

B. Collinearity vs. Global Feature Similarity

We tested the specificity of the attentional modulation for collinearity further by manipulating similarity at the global level. Note that in a collinear configuration the central target shares not only the local orientation of the flankers, but also the global orientation of the virtual contour along which the target and collinear flankers are all aligned. A generalized feature-specific account might argue that attending to the collinear flankers improves detection of targets aligned with this global orientation. We tested this by rotating the flankers on one axis by 90 degrees, so that the target still had the same global orientation as the virtual contour drawn through this axis, but without having local orientation similarity (see Fig. 79.1c, right, and control stimuli on the left). This disruption of collinearity again resulted in a complete elimination of

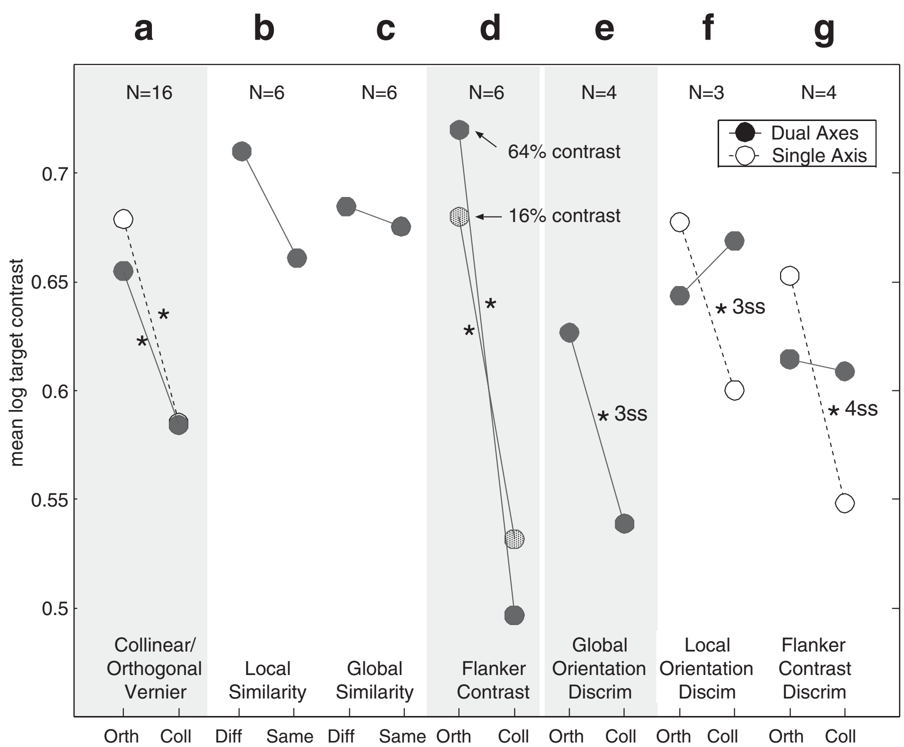

FIGURE 79.2 Target contrast thresholds averaged across observers, in log units, estimated using method of constant stimuli (except for **D**; see text). Refer to Figs. 79.1a–g for example dual-axis stimuli corresponding respectively to conditions labelled A–G here. Shaded regions indicate conditions where a significant attentional effect was found for dual-axis stimuli. **A.** Standard attend-collinear versus attend-orthogonal condition. All 16 observers experienced this plus one or more of the other conditions, but with the exception of EF were naïve to the purpose of the studies. **B, C.** Attentional effects were abolished following disruption of target-flanker collinearity by manipulating stimulus geometry at global or local levels. **D.** Similar stimuli and task to **A**, showing that the significant attentional effect is not diminished by increased flanker contrast. Differences in threshold-estimation method might account for the increased apparent effect-size for just this condition; see main text. Asterisks indicate significant difference between conditions ($p < .01$) in group analyses. **E–G.** Dual-axis attentional effect depends critically on the specific task performed on the flankers, while the single-axis stimulus effect remained unaffected. Here, statistical reliability is assessed by individual: numbers by asterisks indicate how many observers in each sample showed significant differences between conditions ($p < .01$).

attentional modulation effects (means for six subjects in Fig. 79.2c).

Neither local nor global feature similarity can explain away these critical results. We therefore may conclude that attention can specifically modulate sensitivity to collinear configurations.

C. Strength of Modulation

Although attention may modulate sensitivity to collinearity, this does not necessarily imply that attention directly modulates the integrative lateral interactions mechanisms themselves that underlie this sensitivity. Attention might instead just affect how information relating to collinear structures is subsequently used after the initial analysis of collinearity is complete. For example, one possibility is that the target reciprocally influences the contrast response to collinear flankers (Solomon, Watson, and Morgan, 1999). This might provide an extra cue in the flankers for detecting the target, which might be detectable only when the flankers are attended. Another possibility is that the secondary Vernier task is easier with collinear flankers, leading to a redistribution of attentional resources in favour of the primary target detection task (though note there was no evidence of consistent configuration-dependent variations in Vernier accuracy in any of our studies). Both of these accounts predict attentional modulation that is specifically dependent on collinearity. However, they assume no direct attentional interaction with the mechanisms underlying flanker-target facilitation. Such indirect attentional effects therefore should ride on top of stimulus-driven effects of collinearity. If this were the whole explanation of the observed modulation, it should always be possible to measure some residual facilitation from collinear flankers even when they are unattended, because the mechanisms underlying this facilitation would still be functioning as normal.

Evidence from two sources indicates that there can be no such residual facilitation. First, in our original studies, we compared dual-axis stimuli (filled symbols in Fig. 79.2a) with single-axis stimuli (see Fig. 79.1a, right) in which only one collinear or one orthogonal pair of flankers was displayed (open symbols). The physical presence or absence of task-irrelevant collinear flankers made no difference to target detection. The effect of ignoring collinear flankers (i.e., reduction of their impact on the target) was therefore equivalent in size to physically removing them altogether. Second, in the previous experiment manipulating global congruity (Freeman et al., 2004), we could compare the influence of *unattended* flankers that are either collinear or noncollinear with the target (compare leftmost points of Figs. 79.2a and 79.2c and corresponding stimuli left of Figs. 79.1a and 79.1c). This comparison showed there was no significant residual target facilitation from ignored collinear flankers, indicating that attentional exclusion of their influence is complete. This strong modulation appears consistent with attention directly modulating the processing of collinear configurations, rather than merely filtering the output of this processing at some later stage.

IV. MECHANISM OF ATTENTIONAL MODULATION

Our psychophysical evidence for attentional modulation of lateral interactions appears consistent with recent influential single-cell studies, which showed not only that the response of neurons in primary visual cortex to an oriented target falling inside the classical receptive field (CRF) can be influenced by a collinear flanker outside the CRF, but also that this influence can be modulated by manipulations of spatial attention to the target (Gilbert et al., 2000). This neurophysiological result was taken to imply that attention can modulate flanker-target interactions by weighting the horizontal lateral connections between cells responding to target and flankers.

Such a *connection-weighting* account is not the only possible explanation of these attentional modulation effects, however. A mere shift in the distribution of spatial attention may be sufficient to modulate perceptual sensitivity to a stimulus, by amplifying the response of cells whose CRF includes the attended stimulus (Treue, 2001). Attended elements would thus have a greater effective contrast. Unlike the Gilbert et al. (2000) study, our experiments by design kept spatial attention to the target constant, ensuring that any observed attentional effects most likely reflect modulation of flanker effects on the target. However, our attentional manipulation did involve directing spatial attention toward or away from the collinear flankers, which on the previous account might have modulated the local contrast response to unattended versus attended flankers. Thus, the response to unattended flankers might have been suppressed to such an extent that they no longer provided any facilitating input to the central target. This *flanker-modulation* account would imply that attention does not change the lateral-interaction mechanisms themselves, but just controls the input to them.

We devised a psychophysical method for distinguishing between these connection-weighting and flanker-modulation accounts (Freeman, Driver, Sagi, and Zhaoping, 2003). We derived divergent predictions from a known nonlinearity relating flanker contrast to psychophysical target detection threshold. Facilitation of central target detection by collinear flankers has been shown to increase to a maximum as flanker contrast crosses contrast threshold; thereafter, further flanker contrast increments produce no further improvements in target visibility (Polat, 1999). If the flanker-modulation account of our attentional effects were true, then ignored collinear flankers would have to be suppressed to their contrast threshold before their facilitation of the target is appreciably reduced. This account therefore predicts elimination of the attentional effect when physical flanker contrast is too high to be suppressed to threshold levels by being ignored (see Fig. 79.1d for examples of stimuli with high and low flanker contrast). Figure 79.2d plots mean thresholds for six subjects, showing significant attentional effects at low flanker contrasts (16%), but with no significant reduction in the attentional effect at higher flanker contrast and (64%). Increasing flanker contrast further to 80% did nothing to reduce the attentional effect. Note that for this study only, thresholds were estimated from d' measured at fixed target contrast across a range of flanker contrasts, whereas all other studies employed the method of constant stimuli. This methodological difference might explain the increased effect-size apparent in the graph.

To explain this result, the flanker modulation hypothesis would have to assume that even at these high contrasts, flankers can still be suppressed to below their detection threshold when ignored. Such strong modulation seems highly implausible in the light of existing data on attention and contrast (Carrasco et al., 2000; Lee, Itti, Koch, and Braun, 1999; Zenger, Braun, and Koch, 2000). On the other hand, the connection-weighting account described earlier predicts exactly the independence of the attentional effect from suprathreshold changes in flanker contrast that we observed. According to this account, attention should amplify or attenuate the impact of flanker on the target by the same amount regardless of the actual contrast of the flankers. Our results thus allow us to conclude that attention modulates integration of the whole flanker-target configuration, rather than just the processing of the individual parts, thus providing the first clear psychophysical evidence that attention can directly influence contextual integration in early vision.

V. FUNCTIONAL ROLE OF LATERAL INTERACTIONS AND ATTENTIONAL MODULATION THEREOF

A. Task Specificity vs. Object-based Attention

Lateral-interaction mechanisms commonly are assumed to perform the function of grouping local elements into a global pattern. If attention directly modulates lateral interactions, then this modulation might be strongest whenever observers are required to judge the properties of the global pattern formed by the flankers, compared to when only local properties of the individual flankers are task-relevant. This task-specific function of attention would contrast with the less specific but influential notion of object-based attention (Driver et al., 2001). On such object-based accounts, all the visual properties of an attended object are extracted and integrated into global descriptions, whether these properties are individually task-relevant or not.

We have found that our attentional effect successfully generalized from the original Vernier flanker task to a *global orientation* discrimination task. In this task, flankers on each axis are positioned along a virtual contour that can have two global orientations, switching unpredictably between display intervals. Flankers therefore rotate around a single central origin, and observers decide in which interval they have a pre-specified rotation (see Fig. 79.1e). Mean thresholds are plotted in Fig. 79.2e, with significant attentional effects for three out of four observers tested. However, the attentional effect was completely abolished (i.e., no difference between attend-collinear and attend-orthogonal conditions; see means for three observers in Fig. 79.2f) given a task requiring purely *local orientation* discrimination of the relevant flankers, where each flanker rotates around its own centre (see Fig. 79.1f). A similar null effect of attention was observed for discrimination of relative flanker contrast (see Fig. 79.1g and filled symbols in Fig. 79.2g, for four observers) and also flanker colour (Freeman and Driver, submitted), despite care being taken to ensure that both flankers had to be simultaneously attended. This task-dependency runs counter to the predictions of object-based attention models, in which attention to collinear flankers should be sufficient to encode all their local and global properties. Instead, it implies that the global properties of the whole target-flanker configuration may be encoded (via lateral interactions) only when these properties are specifically task-relevant. Furthermore, these results confirm the popular but hitherto untested assumption that lateral interactions play a critical role in perceptual grouping, challenging accounts in which this phenomenon is merely a by-product of local perceptual processing, resulting, for example, from signal averaging within spatial analysers (e.g., Solomon et al., 1999).

B. Dependence of Task-Specificity on Rival Flanker Axes

Interestingly, the preceding task-specific results were found only when two pairs of flankers are present, with one pair attended for the flanker task, and one pair ignored. The nature of the flanker task had no effect on target facilitation when only a *single* pair of flankers was present. Significant differences in threshold were found between single-axis collinear versus orthogonal stimuli effects for all observers tested (see open symbols in Fig. 79.2f and 79.2g). This reveals an important limitation of what top-down modulation by the imposed task can achieve: it cannot arbitrarily switch on or off lateral interactions whenever collinear stimuli are present. Rather, attentional modulation arises only when observers have to select between multiple sets of flankers. This result concurs with physiological data showing that weak attentional modulation of cell responses to an isolated stimulus can be greatly enhanced when other irrelevant stimuli are simultaneously presented, as if selective attention is biasing an active competition between them (Desimone and Duncan, 1995; see also Chapter 50). In the present case, attentional modulation may reflect top-down bias of the competition between alternative perceptual groupings (formed by each of the two flanking axes), in favor of one that is relevant.

VI. PHYSIOLOGICAL CORRELATES OF LATERAL INTERACTIONS AND ATTENTIONAL MODULATION

Given the known neurobiology of lateral interactions, the most likely mechanism for attentional modulation would involve direct modulation of the integrative horizontal connections in early visual cortex. Evidence from two new EEG studies supports such an early anatomical locus. First, we found early event-related potentials (ERPs) over the occipital midline (180–230 ms) are correlated with target facilitation from collinear versus orthogonal flankers in single-axis displays (Khoe, Freeman, Woldorff, and Mangun, in press). Second we found a later component (180–230 ms) that is dependent on attention to collinear (versus orthogonal) context in dual-axis displays (Khoe, Freeman, Woldorff, and Mangun in preparation). The relatively late timing of these attention-related ERP effects is consistent with a role for top-down feedback in modulating the stimulus-driven response to collinear configurations.

VII. CONCLUSIONS

Our findings to date reveal both the power and the limitations of attention in modulating processing of collinear structures, which taken altogether allow us to be more precise about how and for what task-driven attentional factors might interface with mechanisms of context integration in early vision.

First, the strong and specific modulation of collinear configurations, and the independence of attentional effects from flanker contrast, are all consistent with attention directly modulating the integration of target and flankers into a contour, rather than just affecting the visibility of the individual stimulus elements. Second, although it seems that task-directed attention might have the power to exert direct control over integrative circuitry in early vision, the scope of its influence is strongly constrained by the precise sensory and behavioral context in which the stimulus is encountered. The modulation we observed does not depend simply on where the subject is attending or to what particular features, but is highly specific to stimuli containing collinear elements, and then again only when performing a global task that arguably requires perception of these elements as a continuous contour with supervening global properties. A further important criterion may be the presence of competing stimuli or stimulus groupings to select between; in our case, two intersecting virtual contours.

These constraints hint at the likely functional role played by this top-down control of lateral interactions, in *selectively integrating* local image elements with their behaviorally relevant context. Thus, given a typical scene containing multiple overlapping contours, the task of selecting one object (e.g., looking at one's reflection in the window while ignoring the background) may be achieved by adjusting the neural circuitry so that only the relevant contour segments are integrated with each other, and other competing contours are excluded from the perceived group.

Acknowledgments

This research was supported by project grants to EF from the BBSRC (UK).
Thanks to Jon Driver for helpful criticism.

References

Callaway, E. M. (1998). Local circuits in primary visual cortex of the macaque monkey. *Annual Review of Neuroscience* **21**, 47–74.

Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience* **18**, 193–222.

Driver, J., Davis, G., Russell, C., Turatto, M., and Freeman, E. (2001). Segmentation, attention and phenomenal visual objects. *Cognition* **80**, 61–95.

Freeman, E., Sagi, D., and Driver, J. (2001). Lateral interactions between targets and flankers in low-level vision depend on attention to the flankers. *Nature Neuroscience* **4**, 1032–1036.

Freeman, E., Driver, J., Sagi, D., and Zhaoping, L. (2003) Top-down modulation of lateral interactions in early vision. Does attention affect integration of the whole or just perception of the parts? *Curr Biol*. **13**, 985–989.

Freeman, E., Sagi, D., and Driver, J. (2004) Configuration-specific attentional modulation of flanker-target lateral interactions. *Perception* **33**, 181–194.

Freeman, E., Driver, J. (Submitted). Task-dependent modulation of target-flanker lateral interactions: effects of manipulating attention to spatial and non-spatial flanker attributes.

Gilbert, C., Ito, M., Kapadia, M., and Westheimer, G. (2000). Interactions between attention, context and learning in primary visual cortex. *Vision Research* **40**, 1217–1226.

Kapadia, M. K., Ito, M., Gilbert, C. D., and Westheimer, G. (1995). Improvement in visual sensitivity by changes in local context: Parallel studies in human observers and in V1 of alert monkeys. *Neuron* **15**, 843–856.

Khoe, W., Freeman, E., Woldorff, M. G., and Mangun, G. R. (In press). Probing lateral interactions in early visual areas: an event-related potential study of collinear and orthogonal stimulus configurations. *Vision Research*.

Khoe, W., Freeman, E., Woldorff, M. G., and Mangun, G. R. (In prep). Longer Latency Effects of Attention on Integrative Processes in Early Visual Areas.

Lee, D. K., Itti, L., Koch, C., and Braun, J. (1999). Attention activates winner-takes-all competition among visual filters. *Nature Neuroscience* **2**, 375–381.

Polat, U. (1999). Functional architecture of long-range perceptual interactions. *Spatial Vision* **12**, 143–162.

Saenz, M., Buracas, G. T., and Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience* **5**, 631–632.

Saenz, M., and Boynton, G. M. (2002). Feature-based attention modulates lateral interactions in contrast detection. Society for Neuroscience, 32nd annual meeting.

Solomon, J. A., Watson, A. B., and Morgan, M. J. (1999). Transducer model produces facilitation from opposite-sign flanks. *Vision Research* **39**, 987–992.

Treue, S. (2001). Neural correlates of attention in primate visual cortex. *Trends in Neurosciences* **24**, 295–300.

Zenger, B., Braun, J., and Koch, C. (2000). Attentional effects on contrast detection in the presence of surround masks. *Vision Research* **40**, 3717–3724.