

fMRI Activation in Response to Illusory Contours and Salient Regions in the Human Lateral Occipital Complex

Damian A. Stanley and Nava Rubin*

Center for Neural Science
New York University
4 Washington Place
New York, New York 10003

Summary

Regions in the human Lateral Occipital Complex (LOC) show fMRI responses to illusory surfaces. We show that the LOC activation is due to the globally completed region and occurs even when the region is not bounded by illusory contours (ICs). Kanizsa-type stimuli were modified by rounding the corners of the “pacmen” inducers and misaligning them slightly. The impression of an enclosed, salient region (SR) remained, although ICs were no longer perceived (psychophysical data). fMRI activity was elevated for both the IC and SR stimuli, compared to their control stimuli. The LOC response to salient regions may be the result of fast but crude region-based segmentation processes, which are useful for selecting parts of cluttered images for more detailed, computationally intensive processing.

Introduction

A central component of visual processing is segmenting the retinal image into surfaces corresponding to different objects. Visual segmentation involves computations which are more complex than mere edge detection, such as completing surfaces which are fragmented in the image due to occlusion, shadows, or changing illumination. Illusory contours (ICs) have been used extensively to study the mechanisms underlying scene segmentation. How IC completion is achieved by the brain is not well understood. There is evidence that cells in early visual cortical areas (V1/V2) in the macaque monkey respond to ICs (i.e., to stimuli which contain a perceptually completed edge but no luminance-defined edge, within their “classical receptive field”); Bakin et al., 2000; Grosf et al., 1993; Lee and Nguyen, 2001; Peterhans and von der Heydt, 1989; Ramsden et al., 2001; Sheth et al., 1996; Sugita, 1999; von der Heydt et al., 1984).

More recently, several brain imaging studies examined responses to ICs in the human brain. Hirsch et al. (1995) compared fMRI BOLD activation in response to Kanizsa-type ICs (similar to that in Figure 1A; Kanizsa, 1955, 1976) with activation in response to control stimuli which contained similar local features but lacked the globally completed surface (Figure 1B). They found that there was a cortical region that responded more strongly to the ICs than to the control stimuli and concluded that this excess activation reflected processing related to the IC surface (see also Ffytche and Zeki, 1996). A later study (Mendola et al., 1999) used similar IC stimuli but

went an important step further by localizing the IC-specific region with respect to known visual cortical areas. The strongest IC-related activation was found in the Lateral Occipital Complex (LOC), in a cortical strip anterior to the first three retinotopically organized visual areas, which included V3A, V4v, V7, and V8. These regions have been previously shown to respond preferentially to familiar objects and object fragments (Grill-Spector et al., 1998; Malach et al., 1995; responses to object stimuli also extend into more anterior portions of lateral occipito-temporal cortex; cf. Kanwisher et al., 1996). LOC therefore represents a later stage than V2 in the presumed hierarchy of visual cortex, both anatomically and functionally. In contrast, in the early visual cortical areas V1 and V2, Mendola et al. (1999) found very little fMRI activation in response to IC surfaces (but see Seghier et al., 2000, for a study eliciting fMRI responses in V1 using dynamic IC stimuli).

The use of the same term, “illusory contours,” to describe the stimuli that elicited responses in early visual cortex (V1/V2) and LOC overlooks an important difference between the neurons in these brain regions. V1/V2 neurons have small receptive fields, and they are known to respond to edges, and therefore the interpretation that their responses to Kanizsa-type stimuli represented portions of the (illusory) bounding contour followed naturally. In contrast, LOC neurons pool information from large portions of the visual field, often spanning both hemifields (Grill-Spector et al., 1998; Tootell et al., 1998; see also electrophysiological findings about receptive field sizes in monkey higher visual cortex, Gross et al., 1969; Seacord et al., 1979; Tanaka, 1996). Therefore, for LOC neurons it is difficult to talk about responses to the “illusory contour” parts of the stimulus separately from responses to the figure as a whole. This suggests that the increased fMRI activation in response to Figure 1A may have been elicited by the completed surface—the diamond—in a way which is independent of the existence of ICs in the image.

Is the distinction between a surface and its bounding contour a valid one? It is true that in general surfaces are delineated by bounding contours. But in real-world scenes, portions of the bounding contours are often hard to detect—because of surface similarity, lighting conditions, or blur. Identification of surfaces based solely on detection of bounding contours therefore requires contour completion computations, which are resource intensive and slow. This led computer vision scientists to try a different approach, of going *from the surface to its boundaries*, rather than the other way around. In this literature, the term “salient region” is used to refer to a set of contiguous image pixels that likely corresponds to a major surface in the scene. Algorithms for the detection of salient regions benefit from the possibility of propagating signals in all directions (rather than only along contours; cf. Pao et al., 1999; Sharon et al., 2000; Shi and Malik, 2000; see also Supplemental Figure S4 at <http://www.neuron.org/cgi/content/full/37/2/323/DC1>). This means that salient regions (SRs, our abbreviation) can be detected rapidly, even in rela-

*Correspondence: nava@cns.nyu.edu

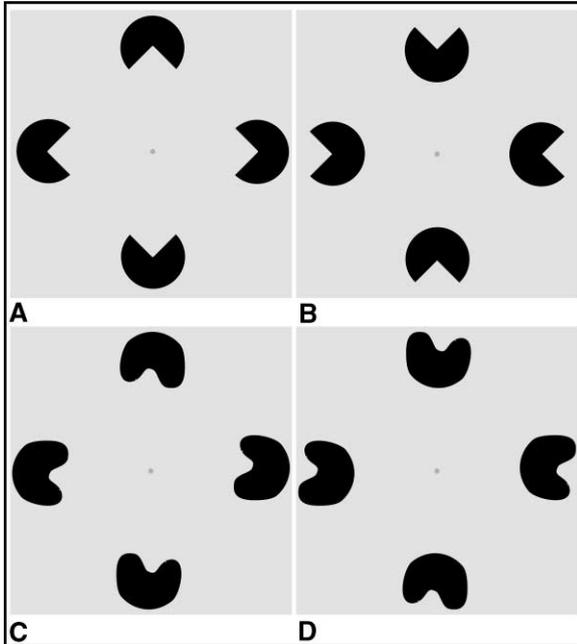


Figure 1. Illustration of the Stimuli Used in the fMRI Experiments (A) IC experimental stimulus; the perceptually completed diamond-shape surface appears to be bounded by a contour also in parts where there is no luminance gradient. (B) IC control stimulus; the rearrangement of the inducing elements eliminates the scene interpretation of an occluding diamond and, with it, the ICs. (C) SR experimental stimulus; the arrangement of the inducing elements creates an impression of an enclosed ("salient") region, but this region is not bounded by crisp illusory contours (note that bounding ICs cannot be seen even with mental effort). (D) SR control stimulus; the saliency of the central region is diminished by the rearrangement of the inducing elements.

tively cluttered scenes. Furthermore, since a region is always bound by contours all around (by necessity, topologically), the SR map gives a first approximation for the edges in the scene. At the same time, it is a crude map, since the bounding contours are not worked out in detail, and therefore more detailed processing is required. Detection of SRs has therefore emerged as a useful strategy to direct computational resources to select portions in the image. SRs can be thought of as "candidate surfaces," worthy of further processing, such as boundary completion.

The dissociation between regions and bounding contours can be made in human vision, too. Figure 1C shows a modification of Figure 1A where the "pacman-shaped" inducers were altered so as to eliminate the sharp corners (discontinuities) in their outline and the colinearity of their edges. These manipulations were shown to reduce dramatically the strength of illusory contours, as measured by ratings of the perceived "clarity" of perceptually completed edges (Kellman and Shipley, 1991; Rubin, 2001; Shipley and Kellman, 1990; see also below, dot localization experiment). Nevertheless, at first glance, a clear impression of an enclosed region is formed. Borrowing the term from computer vision, we refer to it as a salient region (SR). Perceptually, we define a SR as an image region which creates a first impression

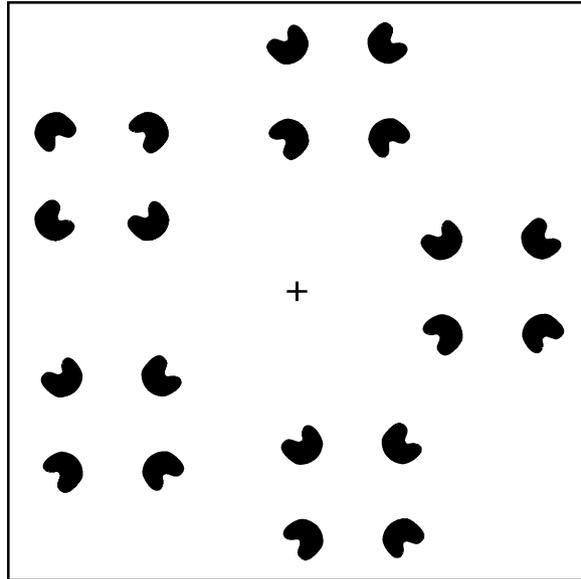


Figure 2. Illustration that the Salient Region "Pops Out" Even When It Is Not Bound by Illusory Contours

of a global surface but, upon closer scrutiny, may not correspond to an actual surface in the scene (and thus would not be supported by a bounding contour all around, either luminance defined or IC). Results of Gurnsey et al. (1996) offer a way to test experimentally that a region is "salient" (by the above definition). These authors followed up a study by Davis and Driver (1994), who reported that Kanizsa-type illusory surfaces "pop out" in a visual search task. Gurnsey et al. (1996) found that the rapid search performance was maintained even when the ICs bounding the Kanizsa squares were eliminated. They noted that, although ICs were no longer perceived, the targets all contained an enclosed region, and they hypothesized that a rapid detection of this region was the basis for the good performance. Thus, although they did not use the term "salient regions" to refer to their stimuli, their results support the idea that SRs exist and have functional significance in human vision. Gurnsey et al. (1996) used a different manipulation from ours to eliminate the ICs (interrupting them by short orthogonal bars). Therefore, to generalize their results and to make a direct connection to our case, we applied their search paradigm to SR stimuli resembling those in Figure 1C (target) and 1D (distractors). The rapid search performance was maintained (data not shown). An informal demonstration of this effect can be seen in Figure 2, where the SR "pops out" among the distractor elements.

The SR stimuli therefore offer a way to ask whether the previously observed LOC responses to Kanizsa-type illusory surfaces required the presence of bounding ICs. (Note that we restrict the term "illusory contours," or ICs, to refer to the [perceived] bounding contours and not to the entire Kanizsa-type stimulus.) We therefore decided to test the response of LOC to Figure 1C compared to Figure 1D (note the close parallelism with the design of Hirsch et al., 1995, and Mendola et al., 1999, illustrated in Figures 1A and 1B). If the response to Figure

1A found in previous studies required the existence of bounding ICs, we should not find fMRI activation in response to Figure 1C. If, on the other hand, the activation found previously was caused by the salient region, Figure 1C may lead to similar activation. In particular, if LOC participates in the detection of salient regions, then we should definitely expect a response to Figure 1C.

Results

In the scans of objects versus scrambled objects (LOC localization), all observers showed robust activation in lateral portions of the occipito-temporal cortex, responding more strongly to the objects stimuli. The amplitude modulation of the across-observer mean time course was 1.25%. Significant voxels appeared also in other brain regions: in the medial portions of the occipital lobe, with reversed phase indicating greater activation in response to the scrambled objects, and in the intraparietal sulcus. (Supplemental Figure S1 at <http://www.neuron.org/cgi/content/full/37/2/323/DC1> shows activation maps for one representative observer, 05).

Our critical fMRI experiment had two conditions, IC and SR. Each condition consisted of eight cycles of alternation between “experimental” and “control” blocks (cf. Figures 1A–1D). We conducted an ANOVA with the average activation level in each block as dependent variables (i.e., a set of 16 values per condition per observer). The ANOVA had three factors: condition, block, and observer (see Experimental Procedures). The analysis revealed a main effect of block ($F_{[1,224]} = 74.2, p < 10^{-14}$) with no interaction between block and condition. Thus, the LOC showed greater activation for the experimental than control stimuli in both the IC and SR conditions. The lack of interaction between block and condition ($F_{[1,224]} = 0.14, p = 0.71$) indicates that the main effect of block was present and comparable in both conditions. (Our statistical power would allow us to detect a difference between the IC and SR conditions if it were 50% or more of the modulation depths). This means that the LOC responds to salient regions even when they are not bounded by illusory contours.

In addition to the effect of block, there was an interaction between observer and block ($F_{[7,224]} = 6.2, p < 10^{-5}$). To examine this further, we analyzed individual observers’ data (two-tailed t test, $\alpha = 0.05$). All eight observers showed higher activation for the “experiment” than the “control” blocks in both the IC and SR conditions, but this difference reached significance only in four observers (01–04) for the IC condition and three of these for the SR condition. To verify that significant responses to the IC and SR stimuli can be obtained also in a population of observers who do not reach significance individually, we performed a second ANOVA including only those four observers (05–08). The results showed a main effect of block ($F_{[1,112]} = 8.6, p < 0.005$) and no interaction with observer or condition. Thus, although the difference between experimental and control stimuli may not reach significance individually for some observers, as a group, they show a highly significant response to both the IC and SR conditions.

There are a few possible explanations why some of our observers did not reach significance individually,

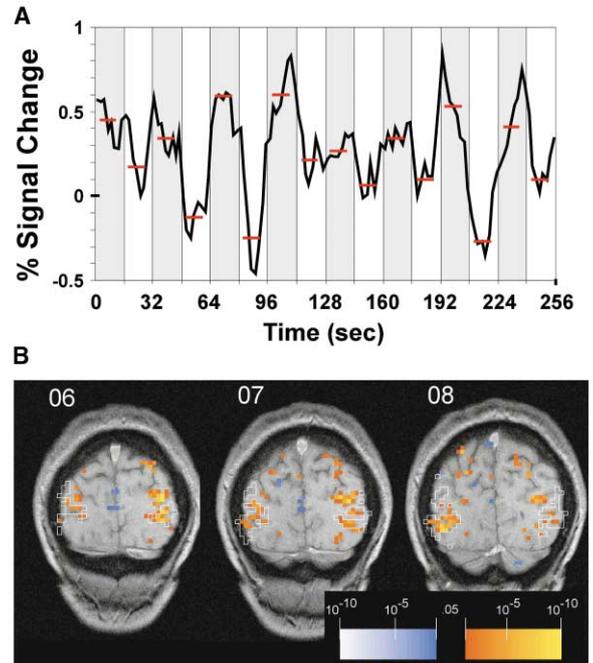


Figure 3. Responses in the LOC to Illusory Contours

(A) Across-observers average time course in LOC for the IC condition. The gray regions indicate “experimental” blocks, the white regions “control” blocks. Red bars indicate the average activation values computed within five-point windows centered at each epoch. (The same procedure was used with individual data to produce the 16 values used for statistical testing.)

(B) Activation maps for observer 04 in the IC condition (slices 06–08; for all slices, see Figure S2 in the Supplemental Data available at <http://www.neuron.org/cgi/content/full/37/2/323/DC1>). The red scale gives the significance level of voxels which were more active in the “experimental” block (Figure 1A); the blue scale is for the voxels which were more active in the “control” block (Figure 1B). The white outline shows the regions included in the functional definition of LOC.

while in another study (Mendola et al., 1999), most observers did. First, there might have been differences in signal strength related to fMRI methodology (e.g., coil sensitivity, pulse sequence used) and/or intrinsic scanner noise. Second, the analysis we employed was different and included more conservative statistical criteria (e.g., correction for multiple comparisons). Finally, some of the difference may lie in individual differences between observers. (Interestingly, the four observers who did not reach significance in our study were all novice fMRI observers, whereas Mendola et al. note that all of their observers were highly practiced fMRI participants. This issue may warrant further investigation.)

Figure 3A shows the across-observer mean LOC time course for the IC condition. Representative activation maps (observer 04, select slices) are shown in Figure 3B. Figure 4 shows the results for the SR condition in the same format. Both time courses clearly show a higher BOLD response to the “experimental” stimuli, i.e., the ICs and the SRs, compared with the “control” stimuli. The depths of modulation were comparable in the two experiments, averaging 0.45% (IC) and 0.41% (SR). Moreover, A linear regression between the IC and SR

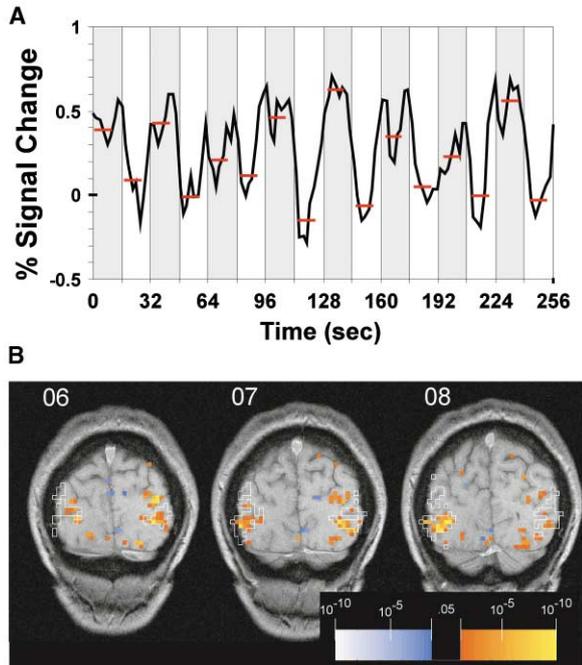


Figure 4. Responses in the LOC to Salient Regions
(A) Across-observers average time course in LOC for the SR condition.
(B) Activation maps for observer 04 in the SR condition (slices 06–08; for all slices, see Figure S3 in the Supplemental Data available at <http://www.neuron.org/cgi/content/full/37/2/323/DC1>).
Notations in (A) and (B) are the same as in Figure 3.

modulation depths in individual observers (which ranged from 0.25% to 1%) showed that the individual response strengths were highly correlated (slope = 1, $r = 0.88$). The pattern of activation is also very similar in the two cases. The activation maps looked similar also for the other observers who individually showed significant modulation of activity in both the IC and SR experiments. In these observers, there was also significant activation in parietal regions; since the present study focuses on LOC activation, those regions were not analyzed further.

Examining the individual observers' activation maps in the IC and SR experiments, we noticed that for some observers there were significant voxels in lateral occipito-temporal cortex which fell outside the functionally defined LOC (i.e., they were not activated in the objects versus scrambled objects runs). We therefore performed the following additional analysis. For each observer, we generated regions of interests (ROIs) based on anatomical markers. The ROIs consisted of the inferior and middle occipital gyri, the lateral and inferior occipital sulci, and the lunate sulcus, when present (cf. Moore and Engel, 2001). We then computed the mean time courses of these ROIs in the IC and SR experiments in each observer and conducted the same analysis as before. The eight-observer ANOVA showed a main effect of block ($F_{[1,224]} = 61.2$, $p < 10^{-12}$) and no interaction with condition ($F_{[1,224]} = 2.0$, $p = 0.16$) or observer ($F_{[7,224]} = 1.3$, $p = 0.23$). Similar results were obtained with the four-observer ANOVA (main effect of block, $F_{[1,112]} = 9.8$,

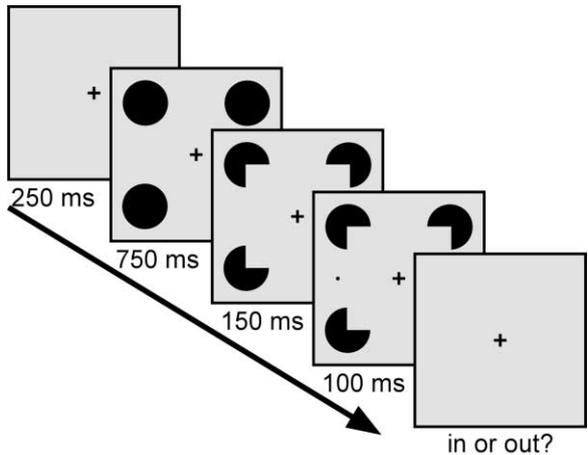


Figure 5. Experimental Paradigm, Dot Localization Task
After the pre-cue (750 ms), an IC or SR stimulus was presented for 150 ms, and then a probe dot appeared (shown here near the left illusory edge). Observers had to report whether the dot fell “in” or “out” of the enclosed region of the IC or SR.

$p < 0.005$). Thus, the exclusion of a few voxels that showed significant modulation in the IC and/or SR experiments from the functionally defined LOC did not change the main result.

One possible concern about the LOC response to the SR stimuli is that the manipulation we performed to eliminate the bounding ICs was not effective enough. Although it has been shown that rounding the corners (L junctions) and misaligning the edges greatly reduce IC perception, this has only been done with subjective rating procedures (Kellman and Shipley, 1991; Rubin, 2001; Shipley and Kellman, 1990). There was the risk that there were some weak or residual ICs that observers did not report but that affected LOC activity. To test this, we conducted a psychophysical experiment designed to measure the strength of the bounding contours of the IC and SR stimuli (as well as two control stimuli, see below). Figure 5 illustrates the experimental design. The task involved the localization of a probe dot which was presented briefly near the (putative) boundary of the enclosed region. Observers had to report whether the dot fell inside or outside the region. The results for the IC and SR conditions are shown in Figure 6. The psychometric curves in the top panel show the across-observer average fraction of “out” responses as a function of dot position. The slopes of the curves provide an estimate of the sharpness of the perceived boundary: responses change more quickly from “in” to “out” in the vicinity of sharp boundaries. There was a marked difference between the slopes of the curves in the IC and SR conditions. Defining threshold as the dot displacement needed to shift responses from 50% to 82% “out,” the thresholds were 9.2 and 19.9 min-arc for ICs and SRs, respectively. The bottom panel shows the thresholds for the individual observers. The two control conditions assessed performance when the boundary was a luminance-defined contour and when no salient region was present at all (RC and XH, respectively; see Experimental Procedures). The XH threshold was 16.1 min-arc (95% confidence interval: 14.3–17.9), indicating that dot local-

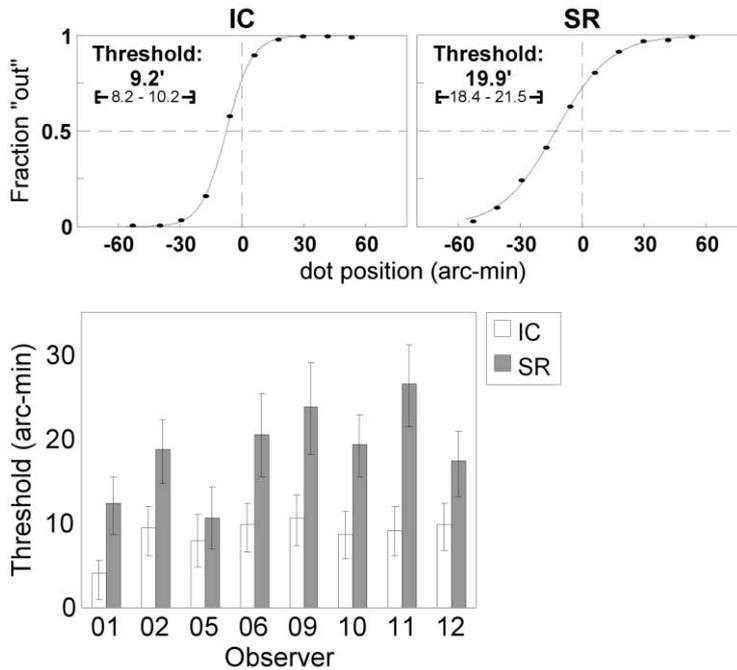


Figure 6. Results, Dot Localization Task (Top panel) Psychometric curves, across-observer means. The fraction of “out” responses is plotted against dot position, for the IC (left) and SR (right) conditions. The greater slope in the IC condition indicates a markedly sharper perceived boundary than in the SR condition. Thresholds (see text) are indicated on the top left of each graph, with the 95% confidence intervals below. (Bottom panel) Thresholds for individual observers in the IC (white) and SR (gray) conditions (error bars show the 95% confidence intervals).

ization in the SR condition did not benefit at all from the presence of the salient region. In other words, there were no residual IC boundaries. (The slightly better performance for XHs than SRs is probably due to the crosses, which provided some alignment information.) The RC threshold was 6.0 min-arc (confidence interval: 4.3–7.6), indicating that perceptually completed ICs can support boundary localization almost as well as luminance-defined contours (although the difference between the RC and IC conditions is statistically significant).

Discussion

We compared fMRI BOLD responses in the Lateral Occipital Complex (LOC) to illusory contour (IC) and salient region (SR) stimuli. It was previously shown that the LOC is activated in response to Kanizsa-type surfaces (Hirsch et al., 1995; Mendola et al., 1999; cf. Figures 1A and 1B). We asked whether this activation required the presence of the perceptually completed (“illusory”) bounding contours of the Kanizsa surfaces or whether it was related to the presence of the global surface itself. To test this, we used stimuli which did not contain bounding ICs but nevertheless created an impression of an enclosed, salient region (Figures 1C and 1D). We found that the SR stimuli led to a similar level of activation in the LOC, compared with IC stimuli.

Perceptual differences between Kanizsa-type illusory surfaces and modified versions resembling our SR stimuli (rounded corners, misaligned edges) have been documented, but so far they have been restricted to subjective rating procedures (Kellman and Shipley, 1991; Rubin, 2001; Shipley and Kellman, 1990). Given the similar LOC responses we obtained, it became more critical to insure that there was a robust perceptual difference between the two types of stimuli. We therefore designed

a dot localization task to quantify the sharpness of the (putative) boundaries of ICs and SRs. Test dots were briefly positioned in the vicinity of the boundaries, and observers had to judge whether they were “inside” or “outside” the completed region. The slopes of the psychometric functions indicated that the (perceived) boundary of the IC stimuli was indeed sharp and well localized, while that of the SR stimuli was not.

The psychophysical results also ruled out the possibility that the BOLD activation observed in the LOC with the SR stimuli occurred because observers could perceive bounding ICs in those stimuli by exerting “mental effort,” or “imagery.” Surely, if that were possible, they would have done so in the behavioral task, where perception of bounding ICs improved performance. The poor dot localization performance in the SR condition indicates that those stimuli cannot support the perception of ICs, with or without mental effort. This still leaves room for the possibility that our observers were exerting (intentionally or not) some mental effort to see the enclosed region. This possibility, while worth mentioning, does not pose a problem for the present work, since our goal was to test whether the bounding ICs were critical for LOC activation. The possible interaction between SR perception and attention remains a question for future research.

Given the behavioral evidence for a weak bounding contour in the SR stimuli, the increased BOLD activation they produced in the LOC may appear surprising. But turning our attention to the perceptual similarities between IC and SR stimuli (rather than their differences) offers a possible clue. The perceptual similarities between IC-bound Kanizsa-type stimuli and their modified counterparts, salient regions not bounded by ICs, have been established with the technique of visual search. As mentioned before, Gurnsey et al. (1996) showed that stimuli which created an impression of an enclosed re-

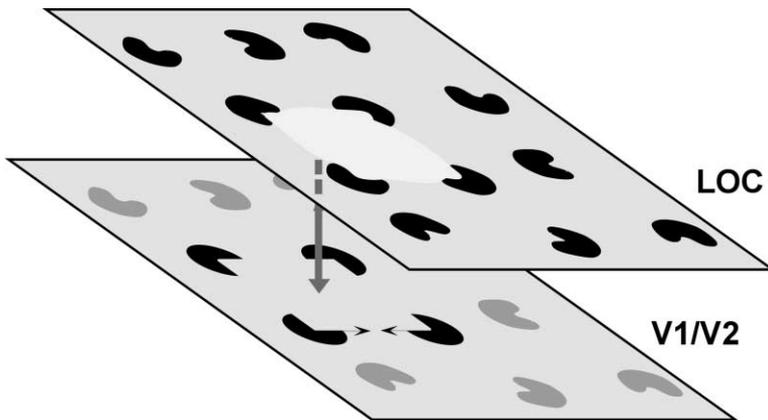


Figure 7. A Possible Role for Salient Region Detection in the LOC

Feedback from the LOC about the location of salient regions may direct resource-intensive contour completion processes in early visual cortex (V1/V2) to those regions.

gion “popped out” in a visual search task even if, upon further inspection, the enclosed region was not perceptually bound by ICs. Their modified Kanizsa squares (which conform to our definition of SRs) yielded search slopes of 8.5 ms/item, virtually identical to those found by Davis and Driver for (intact) Kanizsa-type illusory surfaces (10 ms/item). In contrast, Gurnsey et al. (1996) found that detection of an IC-bound surface among SR distractors was extremely slow (98 ms/item), indicating that the two types of stimuli are hard to distinguish by preattentive mechanisms.

The rapid detection of SRs brings up the possibility that human vision may be employing a similar strategy to that used in computer vision for segmentation. Recall that in that field researchers process real-world, cluttered images by first detecting salient regions using fast-but-crude algorithms and then combining them with more computationally intensive contour-based algorithms applied to restricted parts of the image—the salient regions. It is true that occasionally a salient region may turn out not to correspond to a global surface, but quite often salient regions do correspond to the main objects in the scene, and therefore detecting them first is a useful strategy to reduce computational load. The idea that this strategy is used by the brain offers a simple interpretation for our fMRI results. They suggest that the LOC plays a role in the rapid detection of salient regions but that it plays a lesser role (or not at all) in the contour-based computations that distinguish SRs bound by ICs from those which are not.

If the LOC is not sensitive to the presence of the sharp bounding illusory contours, what might be the neural basis for their perception? A reasonable conjecture is that contour-based processes that support the perception of ICs are performed in early, retinotopically organized visual areas (V1, V2, or both). This, in turn, suggests that early visual cortex relies on feedback from LOC to restrict computationally intensive processes to salient regions. Figure 7 illustrates this idea schematically. There are many observations that support the idea that contour completion relies on early cortex. First, there are the well known electrophysiological results showing that activity of early visual cortical cells can be highly correlated with the perception of illusory and occluded contours (Bakin et al., 2000; Lee and Nguyen, 2001; Peterhans and von der Heydt, 1989; Sugita, 1999;

von der Heydt et al., 1984). There is also behavioral evidence that early cortical areas are critical for the perception of ICs (Pillow and Rubin, 2002). The observation that image details at high spatial resolution can affect whether IC surfaces are perceived or not (e.g., the rounded corners in Figure 1C; Minguzzi, 1987; Rubin, 2001; Shipley and Kellman, 1990) provides a computational motivation why early visual cortical areas should be involved in contour completion. These areas have immediate access to high-resolution information in the image, and this could necessitate their constant involvement in the processing (Lee, 2002; Lee et al., 1998). Finally, it was recently reported that LOC responses to IC stimuli occur extremely fast, within 90–100 ms, making the idea of feedback from LOC to early cortex plausible (Murray et al., 2002). There is also the possibility that communication between early cortex and LOC takes place iteratively, refining information about global regions as well as contours (Grossberg, 1988; Ringach and Shapley, 1996).

If the processing of completed contours indeed involves early cortex, then it should be possible to observe differential response to our IC versus SR stimuli in those areas. In particular, early cortical neurons whose receptive fields lie on the perimeter of SRs should not show the same kind of responses as when bounding ICs fall within their receptive fields. This question could be pursued with electrophysiological experiments; it may be more difficult to approach with fMRI. The existing electrophysiological results suggest that only a fraction of the orientation-selective V1/V2 neurons respond to completed contours and that their responses are often weaker than to luminance-defined contours (Bakin et al., 2000; Peterhans and von der Heydt, 1989; Sugita, 1999; von der Heydt et al., 1984). Nevertheless, precise localization of the early visual cortical sites responsive to a portion of the (completed) boundary is technically possible with fMRI, and with enough averaging it may be possible to observe responses to bounding ICs in those sites. This would allow testing the prediction that when the ICs are eliminated early cortical responses should disappear, too. However, it is unlikely that fMRI methods will allow, in the near future, to test the hypothesis that feedback from LOC is important for contour completion processes in early cortex. Methods with much higher temporal resolution would need to be employed for this purpose.

Another important direction for further research is how salient regions are detected in the brain. Here we showed that they activate the LOC and summarized the psychophysical evidence that SRs are rapidly detected by human observers and the computational motivation for why SR detection may be advantageous. But we did not go into the question of how they are detected—in other words, what are the image cues that are used by the brain to identify SRs. Here, too, some insight might be gained from work in computer vision, which suggests that a large variety of image properties—brightness, color, texture, size, etc.—can be useful in the initial, crude parsing of an image into regions. Identifying the role that those image cues play in SR detection in the brain would require further work, both behavioral and physiological.

Experimental Procedures

Observers

Eight observers (three women) participated in the fMRI experiments. Two of the observers were the authors (01 and 02, respectively). Eight observers participated in the psychophysical experiment. Four of them also participated in the fMRI experiments (01, 02, 05, and 06). The observers were not told the specific purpose of the experiments. Four of the fMRI participants were very experienced at being scanned in MRI machines (01–04). Five observers (six for the psychophysics experiment) received payment for their participation. All observers were between the ages of 18 and 40 and had normal or corrected-to-normal vision. Informed consent was obtained from all observers (procedures approved by the New York University Committee on Activities Involving Human Subjects and the Princeton University Institutional Review Panel).

MRI Acquisition

Observers were scanned on a 3 Tesla head-only Siemens Allegra MRI machine located at Princeton University. A Siemens head coil (transmitter-receiver) was used for all scans. Two types of high-resolution T1-weighted scans were obtained for each observer. A set of 170 slices were collected in the sagittal plane with the MPRAGE sequence (TR 11 ms, TE 4 ms). Slices were 1 mm apart (no gap) and had 256×256 voxels (1×1 mm) each. In addition, a set of 16 slices were acquired in a near-coronal plane (parallel to the brainstem), using a T1-weighted EPI pulse sequence (TR 200 ms, TE 5 ms). Interslice distance was 3 mm (no gap), resolution 256×256 , FOV 192 mm, resulting in $3 \times 0.75 \times 0.75$ mm voxels. The slices were taken at the same position and orientation as the functional scans (see below), covering the occipital lobe and posterior portions of the temporal and parietal lobes. The T1-weighted EPI images were used, in conjunction with the MPRAGE images, to determine the precise anatomical location of functional EPI data, as well as to display those data on high-resolution anatomical background. Functional (T2*-weighted) EPI images (TR = 2 s, TE = 30 ms, flip angle 90°) were acquired using the same slice prescription as the T1-weighted EPIs, except that the in-plane resolution was 64×64 , resulting in $3 \times 3 \times 3$ mm voxels. Each functional scan consisted of 136 consecutive acquisitions, thus lasting 272 s. The visual stimulation started with the fifth acquisition and was preceded by a visual “countdown” stimulus, to minimize artifacts due to the onset of the loud noise and prepare the observer for the onset of the stimulus. To offset for this delay and the hemodynamic delay, data from the first eight acquisitions of each functional scan were subsequently discarded, leaving data series of 128 time points for each voxel. (For four observers, there were 134 acquisitions, stimulus onset was on the third acquisition, and the first six were discarded.) Observers remained in the scanner for approximately 90 min. Foam was used to minimize head motion.

Experimental Design and Visual Stimuli: fMRI

All functional scans used a two-condition blocked design, alternating “experimental” and “control” blocks. Each block lasted 16 s,

and the pair of blocks was repeated eight times for a total of 256 s (128 fMRI acquisitions). Each experiment was repeated twice, and data were averaged (after motion correction, see below; for four observers, two experiments were repeated three times; see Table S1 in the Supplemental Data available at <http://www.neuron.org/cgi/content/full/37/2/323/DC1>). Visual stimuli were generated by a Macintosh Powerbook G3 computer. The graphical output (640×480 pixels) was fed into an Epson PowerLite 7250 LCD projector equipped with an extra focusing lens. The projected image passed through a focusing lens, was deflected 90° by a mirror, and appeared on a plastic rear-projection screen. Observers viewed the stimuli in a mirror mounted on the head coil. Observers were required to maintain fixation during the scans.

The LO Complex was defined functionally by its response to pictures of everyday objects. (This method is different from that used by Mendola et al., who localized the LOC with respect to early retinotopic cortex, but has been widely used by other researchers; Grill-Spector et al., 1998; Kanwisher et al., 1996; Malach et al., 1995). Observers viewed alternating blocks which contained grayscale pictures of either objects or scrambled versions of the same objects. The objects were shown in isolation on a homogeneous background and subtended, on average, 14° visual angle. The scrambled versions were generated by breaking down each image into 20×20 tiles and randomizing their position while maintaining the distribution of the luminance in the image approximately equal to that in the intact object pictures. This ensured that the low spatial frequency content of the two types of stimuli was comparable. To make the high spatial frequency content of the two types of stimuli comparable, a grid of thin (two-pixel wide) dark lines was superimposed on both sets of images. The lines were positioned right on the borders between the tiles in the scrambled-object images (and in the same positions in the intact-object images), resulting in identical sharp-edge content in the two sets. Tiles were small enough so that object parts could not be identified in the scrambled pictures. Each block consisted of 16 pictures of objects or their scrambled counterparts, presented for 1 s each. There were 32 objects in all, grouped into two sets. The 16 objects in each set (and in the scrambled sets) were presented four times in pseudorandomized order, and the sets were presented in alternating order.

There were two types of experimental stimuli: illusory contours (IC) and salient regions (SR). IC stimuli were Kanizsa-type diamond surfaces like that shown in Figure 1A, and the control stimuli consisted of four outward-facing “pacmen” like those shown in Figure 1B. The “pacmen” radius and the sides of the illusory diamonds were 1° and 5° visual angle, respectively. Therefore, the support ratio (the ratio between the luminance-defined portion and the entire bounding contour) was 0.4, and the stimuli subtended 9° visual angle (distance between furthest edges of inducers; two observers were shown stimuli which were twice as large; see Table S1 with the Supplemental Data available at <http://www.neuron.org/cgi/content/full/37/2/323/DC1>). Stimuli were contrast-reversed at a rate of 0.5 Hz. The experimental and control SR stimuli looked like those illustrated in Figures 1C and 1D, respectively. They were generated by manually tracing the outlines of the IC inducers while eliminating their sharp corners and then rotating each inducer by 10° clockwise to eliminate the alignment of the straight parts of the edges. Other than that, all other details of the SR stimuli were identical to those of the IC stimuli. Perceptual data were not collected from the observers, but in debriefings they confirmed that the IC stimuli seemed to have a bounding contour all around, while the SR stimuli did not (see also Shipley and Kellman, 1990).

fMRI Data Analysis

Data from the functional runs were corrected for head motion (in three dimensions) by means of the Automated Image Registration (AIR) package (Woods et al., 1998). Because the motion correction procedure creates artifacts at the boundaries of the volume, the first and last slices, as well as a two-voxel border around all of the remaining slices, were discarded. Data from multiple scans of the same experiment were subsequently averaged. All subsequent analysis procedures were written in-house in Matlab (The Mathworks, Inc., Natick, MA, version 6.1, 2001). Time course data were preprocessed to remove linear trends (using linear regression) from

each voxel independently. For each experiment, the Fourier spectra of the time course data of each voxel were computed using the FFT procedure. Since the visual stimulation had periodicity of 16 time points (8 experimental, 8 control), voxels that are affected by the experimental manipulation have a large 8th Fourier coefficient (128/16; dc coefficient indexed 0), whereas for voxels unaffected by the manipulation the magnitude of this coefficient is determined by the noise. For the affected voxels, the phase of the 8th Fourier coefficient reveals whether they were more activated in the “experimental” or “control” condition. To identify the voxels affected by the manipulation, we computed the function $f(i) = \tilde{N} \times ||C_8(i)||^2 / \sum_{j \in S} ||C_j(i)||^2$, where $C_j(i)$ is the j th Fourier coefficient of the time course data of voxel i , S is the set of all coefficients up to 64, excluding those at three-point windows around all harmonics $\{[n \times 8] + / - 1, n = 1, 2, \dots\}$ and coefficients 0, 1, 2 (to discard the dc and slow nonlinear drift), and $\tilde{N} = 39$ (the size of S). This function yields high values for voxels affected by the experimental manipulation (independent of the shape of the periodic time course) and an F statistic distribution over the remaining voxels (degrees of freedom $[2, 2\tilde{N}]$). Next, statistical significance values $p(i)$ were computed from $f(i)$ for each voxel. Those values were then adjusted for multiple comparisons by converting them to values $Q^{BH}_{adj}(p(i)) \equiv Q(i)$ using the procedures introduced by Benjamini and Hochberg, 1995, and Yekutieli and Benjamini, 1999 (see also Genovese et al., 2002). These procedures control for the False Discovery Rate, which is the expected proportion of erroneously rejected hypotheses among the rejected ones. Setting our threshold level at $\alpha = 0.05$, this means that in the set of voxels deemed “significant” by our analysis ($i | Q(i) < \alpha$) for a given experiment, there were (on average) no more than 5% voxels which were in fact just noise. FDR provides a less strict and therefore more powerful adjustment for multiple comparisons than procedures that control the family-wise error (such as Bonferroni) while at the same time offering a meaningful scale of significance values. Next, voxels with $Q(i) < \alpha$ which did not have at least one contiguous significant voxel (in the 3D, 26-neighbor volume) were eliminated from the set of voxels deemed significant. This procedure eliminated more voxels which appeared outside than inside the brain and, overall, no more than 8% of the voxels, on average. Finally, the phase of each significant voxel was examined to determine if it was more active in the “experimental” or “control” block. To distinguish the two types of voxels, their significance levels were illustrated with two different color ranges in the displayed activation maps.

The Lateral Occipital Complex of each observer was defined as the set of voxels in the lateral portions of occipito-temporal cortex which came out significant (by the procedure above) in the scans of objects versus scrambled objects. For all observers, phase information revealed that all significant voxels in the LOC showed more activation in the “experimental” (objects) condition. (For two observers, significant voxels extended continuously between lateral and medial cortex in some slices, and in those cases, phase information was used to aid in identifying the medial/lateral border.) For each condition (IC and SR), an average time course was computed for the LOC of each observer. The average values of fMRI activation in five-point windows centered at the “experimental” and “control” blocks were computed for each observer’s time course (the remaining time points were excluded to avoid across-block boundary artifacts). The mean was subtracted from the resulting values to remove differences in DC offset. These data (16 values, 8 experimental blocks and 8 control blocks) were then used as the dependent variable in an ANOVA with categorical predictors of block (experimental versus control; fixed), condition (IC versus SR; fixed), and observer (random). Finally, for each condition, the grand-average time course was computed across observers. The average LOC time course data were converted to percent signal change relative to the mean activation in the “control” conditions.

Psychophysical Experiment

Observers performed a dot localization task. An IC or SR stimulus was briefly presented (in separate blocks), and then a probe dot appeared (the duration of the probe dot ensured that observers could not make eye movements to it; for timing, see Figure 5). Observers indicated whether the probe dot was located inside or outside of the region enclosed by the inducers by pressing buttons

on a computer mouse. The probe dot appeared at one of ten fixed locations in the vicinity of the (presumed) boundary, on either side of the region (side randomized across trials; vertically, all dots were at the center of the stimulus). The fraction of “out” responses was plotted against dot position. These data were fit with a psychometric function $0.5 * [1 + \tanh(0.745(x - \beta)/\alpha)]$, where α is threshold performance (82% “out”) and β is bias. 95% confidence intervals for threshold were calculated using a parametric bootstrap method with $n = 2000$ samples for each psychometric function.

Stimuli were presented at two retinal sizes, equal to those used in the fMRI experiment (oriented as squares). The two sizes were presented within the same block in randomized order. The diameter of the probe dot subtended 9.4 arc-min. The probe dot randomly appeared at one of ten equidistant locations within the range of -53 to 53 arc-mins (-27 to 27 arc-mins for the small stimuli), centered at the location of the interpolated contours of the IC stimulus. Each condition thus consisted of 40 different trials (ten dot locations \times two sides \times two sizes), repeated ten times in randomized order. Data were subsequently collapsed across the two sides and sizes of stimuli (analyzed separately, there were no significant differences between the thresholds and biases of those subsets of data, when the x axis was scaled for stimulus size). The order of different condition blocks was randomized across observers.

In addition to the IC and SR conditions, we ran two other conditions. In the “real contour” (RC) condition, luminance-defined contours were added to the IC stimulus where the illusory contours would normally appear. The purpose was to measure the best possible performance. In the “crosshair” (XH) condition, the inducers were replaced by crosses enclosed within circular outlines (same diameter as the IC inducers). The purpose was to evaluate baseline performance (when no enclosed region is present). Other aspects of the design were identical to those of the IC and SR conditions, except that there were only two repetitions per side per size in the RC condition. Blocks of the RC and XH conditions were interleaved with the IC and SR conditions.

The Weber contrast of the inducers and crosshairs was 67%; RC contrast was -27% . Stimuli were generated on a SONY VAIO PC computer at a refresh rate of 75 Hz, using in-house graphics software. Viewing distance was 57 cm, and observers’ heads were stabilized with a chin rest. Observers were instructed to fixate on a fixation point at the center of the screen at all times.

Acknowledgments

We thank Jonathan Cohen, Tie-Qiang Li, and the colleagues and staff at the Princeton University Center for the Study of Brain, Mind and Behavior for their help and cooperation; Josh Fishman, Valerio Luccio, and Jessica Ackert for help in programming and data analysis; Yoav Benjamini, Daniel Yekutieli, and Larry Maloney for statistics advice; Nancy Kanwisher for the object stimuli; Anders Dale, Jean-Michel Hupe, Jonathan Pillow, and Robert Shapley for helpful discussions. Supported by National Science Foundation program for Learning in Intelligent Systems and the Sloan Foundation.

Received: April 15, 2002

Revised: November 25, 2002

References

- Bakin, J.S., Nakayama, K., and Gilbert, C.D. (2000). Visual responses in monkey areas V1 and V2 to three-dimensional surface configurations. *J. Neurosci.* 20, 8188–8198.
- Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B* 57, 289–300.
- Davis, G., and Driver, J. (1994). Parallel detection of Kanizsa subjective figures in the human visual system. *Nature* 371, 791–793.
- Flytche, D.H., and Zeki, S. (1996). Brain activity related to the perception of illusory contours. *Neuroimage* 3, 104–108.
- Genovese, C.R., Lazar, N.A., and Nichols, T. (2002). Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage* 15, 870–878.

- Grill-Spector, K., Kushnir, T., Hendler, T., Edelman, S., Itzhack, Y., and Malach, R. (1998). A sequence of object-processing stages revealed by fMRI in the human occipital lobe. *Hum. Brain Mapp.* 6, 316–28.
- Grosf, D.H., Shapley, R.M., and Hawken, M.J. (1993). Macaque V1 neurons can signal 'illusory' contours. *Nature* 365, 550–552.
- Gross, C.G., Bender, D.B., and Rocha-Miranda, C.E. (1969). Visual receptive fields of neurons in inferotemporal cortex of the monkey. *Science* 166, 1303–1306.
- Grossberg, S. (1988). Cortical dynamics of three-dimensional form, color, and brightness perception. I. Monocular theory. In *Neural Networks and Natural Intelligence*, S. Grossberg, ed. (Cambridge, MA: MIT Press), pp. 1–54.
- Gumsey, R., Poirier, F.J., and Gascon, E. (1996). There is no evidence that Kanizsa-type subjective contours can be detected in parallel. *Perception* 25, 861–874.
- Hirsch, J., DeLaPaz, R.L., Relkin, N.R., Victor, J., Kim, K., Li, T., Borden, P., Rubin, N., and Shapley, R. (1995). Illusory contours activate specific regions in human visual cortex: evidence from functional magnetic resonance imaging. *Proc. Natl. Acad. Sci. USA* 92, 6469–6473.
- Kanizsa, G. (1955). Margini quasi-percettivi in campi con stimolazione omogenea. *Rivista di Psicologia* 49, 7–30. Translated in S. Petry and G. Meyer, Eds. (1987). *The Perception of Illusory Contours* (New York: Springer).
- Kanizsa, G. (1976). Subjective contours. *Sci. Am.* 234, 48–52.
- Kanwisher, N., Chun, M.M., McDermott, J., and Ledden, P.J. (1996). Functional imaging of human visual recognition. *Brain Res. Cogn. Brain Res.* 5, 55–67.
- Kellman, P.J., and Shipley, T.F. (1991). A theory of visual interpolation in object perception. *Cognit. Psychol.* 23, 141–221.
- Lee, T.S. (2002). The nature of illusory contour computation. *Neuron* 33, 667–668.
- Lee, T.S., and Nguyen, M. (2001). Dynamics of subjective contour formation in the early visual cortex. *Proc. Natl. Acad. Sci. USA* 98, 1907–1911.
- Lee, T.S., Mumford, D., Romero, R., and Lamme, V.A. (1998). The role of the primary visual cortex in higher level vision. *Vision Res.* 38, 2429–2454.
- Malach, R., Reppas, J.B., Benson, R.R., Kwong, K.K., Jiang, H., Kennedy, W.A., Ledden, P.J., Brady, T.J., Rosen, B.R., and Tootell, R.B. (1995). Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex. *Proc. Natl. Acad. Sci. USA* 92, 8135–8139.
- Mendola, J.D., Dale, A.M., Fischl, B., Liu, A.K., and Tootell, R.B. (1999). The representation of illusory and real contours in human cortical visual areas revealed by functional magnetic resonance imaging. *J. Neurosci.* 19, 8560–8572.
- Minguzzi, G.F. (1987). Anomalous figures and the tendency to continuation. In *The Perception of Illusory Contours*, S. Petry and G. Meyer, eds. (New York: Springer-Verlag), pp. 71–75.
- Moore, C., and Engel, S.A. (2001). Neural response to perception of volume in the lateral occipital complex. *Neuron* 29, 277–286.
- Murray, M.M., Wylie, G.R., Higgins, B.A., Javitt, D.C., Schroeder, C.E., and Foxe, J.J. (2002). The spatiotemporal dynamics of illusory contour processing: combined high-density electrical mapping, source analysis, and functional magnetic resonance imaging. *J. Neurosci.* 22, 5055–5073.
- Pao, H., Geiger, D., and Rubin, N. (1999). Measuring convexity for Figure/Ground separation. *Proc. 7th IEEE Intl. Conf. Computer Vision (Corfu, Greece)*, 948–955.
- Peterhans, E., and von der Heydt, R. (1989). Mechanisms of contour perception in monkey visual cortex. II. Contours bridging gaps. *J. Neurosci.* 9, 1749–1763.
- Pillow, J., and Rubin, N. (2002). Perceptual completion across the vertical meridian and the role of early visual cortex. *Neuron* 33, 805–813.
- Ramsden, B.M., Hung, C.P., and Roe, A.W. (2001). Real and illusory contour processing in area V1 of the primate: a cortical balancing act. *Cereb. Cortex* 11, 648–665.
- Ringach, D.L., and Shapley, R. (1996). Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Res.* 36, 3037–3050.
- Rubin, N. (2001). The role of junctions in surface completion and contour matching. *Perception* 30, 339–366.
- Seacord, L., Gross, C.G., and Mishkin, M. (1979). Role of inferior temporal cortex in interhemispheric transfer. *Brain Res.* 167, 259–272.
- Seghier, M., Dojat, M., Delon-Martin, C., Rubin, C., Warnking, J., Segebarth, C., and Bullier, J. (2000). Moving illusory contours activate primary visual cortex: an fMRI study. *Cereb. Cortex* 10, 663–670.
- Sharon, E., Brandt, A., and Basri, R. (2000). Fast Multiscale Image Segmentation. *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (South Carolina)*, 70–77.
- Sheth, B.R., Sharma, J., Rao, S.C., and Sur, M. (1996). Orientation maps of subjective contours in visual cortex. *Science* 274, 2110–2115.
- Shi, J., and Malik, J. (2000). Normalized Cuts and Image Segmentation. *IEEE Trans. Patt. Anal. Mach. Intel.* 22, 888–905.
- Shipley, T.F., and Kellman, P.J. (1990). The role of discontinuities in the perception of subjective figures. *Percept. Psychophys.* 48, 259–270.
- Sugita, Y. (1999). Grouping of image fragments in primary visual cortex. *Nature* 401, 269–272.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annu. Rev. Neurosci.* 19, 109–139.
- Tootell, R.B., Mendola, J.D., Hadjikhani, N.K., Liu, A.K., and Dale, A.M. (1998). The representation of the ipsilateral visual field in human cerebral cortex. *Proc. Natl. Acad. Sci. USA* 95, 818–824.
- von der Heydt, R., Peterhans, E., and Baumgartner, G. (1984). Illusory contours and cortical neuron responses. *Science* 224, 1260–1262.
- Woods, R.P., Grafton, S.T., Holmes, C.J., Cherry, S.R., and Mazziotta, J.C. (1998). Automated image registration: I. General methods and intrasubject, intramodality validation. *J. Comput. Assist. Tomogr.* 22, 139–152.
- Yekutieli, D., and Benjamini, Y. (1999). Resampling-based false discovery rate controlling multiple test procedures for correlated test statistics. *J. Stat. Plann. Inference* 82, 171–196.