

Model Cortical Association Fields Account for the Time Course and Dependence on Target Complexity of Human Contour Perception

Vadas Gintautas^{1,2*}, Michael I. Ham³, Benjamin Kunsberg⁴, Shawn Barr⁴, Steven P. Brumby⁵, Craig Rasmussen⁴, John S. George³, Ilya Nemenman⁶, Luís M. A. Bettencourt¹, Garret T. Kenyon^{3,4*}

1 Center for Nonlinear Studies and T-5, Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico, United States of America, **2** Physics Department, Chatham University, Pittsburgh, Pennsylvania, United States of America, **3** P-21 Applied Modern Physics (Biological and Quantum Physics), Los Alamos National Laboratory, Los Alamos, New Mexico, United States of America, **4** New Mexico Consortium, Los Alamos, New Mexico, United States of America, **5** Space and Remote Sensing Sciences, Los Alamos National Laboratory, Los Alamos, New Mexico, United States of America, **6** Departments of Physics and Biology and Computational and Life Sciences Initiative, Emory University, Atlanta, Georgia, United States of America

Abstract

Can lateral connectivity in the primary visual cortex account for the time dependence and intrinsic task difficulty of human contour detection? To answer this question, we created a synthetic image set that prevents sole reliance on either low-level visual features or high-level context for the detection of target objects. Rendered images consist of smoothly varying, globally aligned contour fragments (amoebas) distributed among groups of randomly rotated fragments (clutter). The time course and accuracy of amoeba detection by humans was measured using a two-alternative forced choice protocol with self-reported confidence and variable image presentation time (20–200 ms), followed by an image mask optimized so as to interrupt visual processing. Measured psychometric functions were well fit by sigmoidal functions with exponential time constants of 30–91 ms, depending on amoeba complexity. Key aspects of the psychophysical experiments were accounted for by a computational network model, in which simulated responses across retinotopic arrays of orientation-selective elements were modulated by cortical association fields, represented as multiplicative kernels computed from the differences in pairwise edge statistics between target and distractor images. Comparing the experimental and the computational results suggests that each iteration of the lateral interactions takes at least 37.5 ms of cortical processing time. Our results provide evidence that cortical association fields between orientation selective elements in early visual areas can account for important temporal and task-dependent aspects of the psychometric curves characterizing human contour perception, with the remaining discrepancies postulated to arise from the influence of higher cortical areas.

Citation: Gintautas V, Ham MI, Kunsberg B, Barr S, Brumby SP, et al. (2011) Model Cortical Association Fields Account for the Time Course and Dependence on Target Complexity of Human Contour Perception. *PLoS Comput Biol* 7(10): e1002162. doi:10.1371/journal.pcbi.1002162

Editor: Olaf Sporns, Indiana University, United States of America

Received: January 27, 2011; **Accepted:** June 29, 2011; **Published:** October 6, 2011

This is an open-access article, free of all copyright, and may be freely reproduced, distributed, transmitted, modified, built upon, or otherwise used by anyone for any lawful purpose. The work is made available under the Creative Commons CC0 public domain dedication.

Funding: This work was supported by Los Alamos National Laboratory LDRD program under project 20090006DR; the National Science Foundation, grant ID 0749348; and the DARPA NeoVision2 project. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. This publication qualified for unclassified release under DUSA BIOSCI with LA-UR 11-00499.

Competing Interests: The authors have declared that no competing interests exist.

* E-mail: vgintautas@chatham.edu (VG); gkenyon@lanl.gov (GTK)

Introduction

The perception of closed contours is fundamental to object recognition, as revealed by the fact that common object categories can be rapidly detected in black and white line drawings in which all shading and luminance cues have been removed [1]. Cortical association fields, hypothesized to capture spatial correlations between local image features via long-range lateral synaptic interactions, provide a natural substrate for rapid contour perception [2]. The link between cortical association fields and contour perception has been investigated through a variety of behavioral, experimental, and theoretical techniques [3–6]. Psychophysical measurements reveal that the detection of implicit contours, defined by sequences of Gabor-like elements presented against randomly oriented backgrounds, becomes more difficult as the local curvature increases and as the individual Gabor elements are spaced further apart or their alignment is randomly perturbed. This dependence on proximity and relative orientation implies that, in early visual areas,

cortical association fields are primarily local and aligned along smooth trajectories [2,7,8]. In related studies, collinear Gabor patches have been shown to both increase and decrease the contrast detection threshold of a central Gabor patch in a manner that depends on the relative timing, orientation and spatial separation of the flanking elements [9–11], providing further psychophysical evidence that lateral influences act at early cortical processing stages, although the contribution of collinear facilitation to contour integration remains controversial [12]. In primary visual cortex (V1), electrophysiological recordings indicate that the responses to optimally oriented and positioned stimuli can be facilitated by flanking stimuli placed outside the classical receptive field center [5,6,10,13], although these effects have also been ascribed to elongated central receptive fields [14,15] and facilitation has been attributed to increases in baseline activity [16]. Nonetheless, collinear facilitation is consistent with anatomical studies indicating that orientation columns are laterally connected to surrounding columns with similar orientation preference [17–19].

Author Summary

Current computer vision algorithms reproducing the feed-forward features of the primate visual pathway still fall far behind the capabilities of human subjects in detecting objects in cluttered backgrounds. Here we investigate the possibility that recurrent lateral interactions, long hypothesized to form cortical association fields, can account for the dependence of object detection accuracy on shape complexity and image exposure time. Cortical association fields are thought to aid object detection by reinforcing global image features that cannot easily be detected by single neurons in feed-forward models. Our implementation uses the spatial arrangement, relative orientation, and continuity of putative contour elements to compute the lateral contextual support. We designed synthetic images that allowed us to control object shape and background clutter while eliminating unintentional cues to the presence of an otherwise hidden target. In contrast, real objects can vary uncontrollably in shape, are camouflaged to different degrees by background clutter, and are often associated with non-shape cues, making results using natural image sets difficult to interpret. Our computational model of cortical association fields matches many aspects of the time course and object detection accuracy of human subjects on statistically identical synthetic image sets. This implies that lateral interactions may selectively reinforce smooth object global boundaries.

Because extensive association fields are present in the primary visual cortex [17–19], lateral interactions may be key to discriminating smooth object boundaries at very fast time scales (of the order of tens of ms), as observed in numerous speed of sight psychophysical experiments [1,20–23]. Correspondingly, theoretical models have proposed that V1 cortical association fields can be described mathematically on the basis of cocircularity, and that relaxation dynamics based on cocircular association fields can extract global contours by suppressing local variation [24]. Such models are qualitatively consistent with human judgments as to whether pairs of short line segments belong to the same or separate contours, with human judgments closely following the pairwise statistics of edge segments extracted from natural scenes [25]. Further, model cortical association fields, when used to detect implicit contours, can predict key aspects of human psychophysics, particularly the measured dependence on the density of foreground elements relative to background elements [8,26].

In this paper, we extend the above studies by investigating whether model cortical association fields can account not only for dependence of contour perception on intrinsic task difficulty, a relationship that has been previously explored [8,26], but also for the detailed time course of human contour detection, an aspect that has heretofore not been modeled explicitly, although the time-dependent influence of lateral interactions has been determined for several theoretical models [27,28]. In this work, we employ multiplicative relaxational dynamics to estimate the time course of contour detection from a computational model employing optimized kernels. Model results are then compared to speed-of-sight measurements from human subjects performing the same contour detection task. To obtain optimized cortical association fields, we design lateral connectivity patterns using a novel method that exploits the global statistical properties of salient contours relative to background clutter. Our procedure, which can be generalized beyond the present application, can be summarized as follows.

We begin by generating a large training corpus, divided into target and distractor images, from which we obtain estimates of

the pairwise co-occurrence probability of oriented edges conditioned on the presence or absence of globally salient contours. From the difference in these two probability distributions, we construct Object-Distractor Difference (ODD) kernels, which are then convolved with every edge feature to obtain the lateral contextual support at each location and orientation across the entire image. Edge features that receive substantial contextual support from the surrounding edges are preserved, indicating they are likely to belong to a globally salient contour, whereas edge features receiving minimal contextual support are suppressed, indicating they are more likely to be part of the background clutter. The lateral contextual support is applied in a multiplicative fashion, so as to prevent the appearance of illusory edges, and the process is iterated several times, mimicking the exchange of information along horizontal connections in the primary visual cortex. Our method is thus intended to capture the essential computational elements of cortical association fields that are hypothesized to mediate the pop-out of salient contours against cluttered backgrounds.

To obtain a large number of training images and to better isolate the role of cortical association fields linking low-level visual features, we employ abstract computer-generated shapes consisting of short, smooth contour segments that could either be globally aligned to form wiggly, nearly closed objects (*amoebas*), or else randomly rotated to provide a background of locally indistinguishable contour fragments (*clutter*). Amoeba targets lack specific semantic content, presumably reducing the influence of high level cortical areas, such as IT. However, our computer-generated images would not be expected to eliminate the contribution to contour perception from extrastriate areas [29–32]. Thus, our model of lateral interactions between orientation-selective neurons is designed to account for just one of several cortical mechanisms that likely contribute to contour perception.

Our amoeba/no-amoeba image set differs from stimuli used in previous psychophysical experiments that employed sequences of Gabor-like elements to represent salient contours against randomly oriented backgrounds [2,7,8]. An advantage of contours represented by random Gabor fields is that the target and distractor Gabor elements can be distributed at approximately equal densities, thereby precluding the use of local density operators as surrogates for global contour perception [2]. However, our amoeba/no-amoeba image set is more akin to the natural image sets used in previous speed-of-sight object detection tasks [33], particularly with respect to studies employing line drawings derived from natural scenes [1]. Humans can detect closed contours, whether defined by aligned Gabor elements or by continuous line fragments, in less than 200 ms [1,20], which is shorter than the mean interval between saccadic eye movements [34], thus mitigating the contribution from visual search. Like Gabor defined contours, our amoeba/no-amoeba image set implements a pop-out detection task involving readily perceived target shapes whose complexity can be controlled parametrically.

To benchmark the accuracy and the time course of the ODD kernel-based procedure applied to the amoeba/no-amoeba task, we compare our model results to the performance of human subjects on a 2AFC speed-of-sight task in which amoeba/no-amoeba images are presented very briefly side by side, followed by a mask designed to limit the time the visual system is able to process the sensory input [1,20–23]. Since it takes an estimated 100–300 ms for activation to spread through the ventral stream of the visual cortex [21], an effective mask presented within this time frame can potentially degrade object detection performance by interfering with the neural processing mechanisms underlying recognition [22,35]. By plotting task performance as a function of

the stimulus onset asynchrony (SOA)—the interval between image and mask presentation onsets—the resulting psychometric curves are hypothesized to estimate the neural processing time required to reach a given level of classification accuracy. Amoeba targets of low to moderate complexity were found to reliably pop-out against the background clutter, allowing subjects to achieve near perfect performance at SOAs less than 250 ms, even when followed by an optimized mask consisting of rotated versions of the target and distractor images [20]. Our model cortical association fields were able to account for the dependence of human performance on amoeba complexity as well as for aspects of the time course of contour perception as measured by the improvement in human performance with increasing SOA. Thus, we present the first network-level computational model to simultaneously account for spatial and temporal aspects of contour perception, as measured in human subjects performing the same contour detection task. Aspects of the experimental data for which our model fails to account, particularly data showing that human subjects require longer processing times to detect more complex targets, may indicate the possible involvement of extrastriate areas, which may be essential for the perception of more complex shapes.

Results

To investigate low-level cortical mechanisms for detecting smooth, closed contours presented against cluttered backgrounds with statistically similar low-level features, we designed an amoeba/no-amoeba detection task using a novel set of synthetic images (Figure 1). Amoebas are radial frequency patterns [36] constructed via superposition of periodic functions described by a discrete set of radial frequencies around a circle. In addition, we added clutter objects, or distractors, that were locally indistinguishable from targets. Both targets and distractors were composed of short contour fragments, thus eliminating unambiguous indicators of target presence or absence, such as total line length, the presence of line endpoints, and the existence of short gaps between opposed line segments. To keep the bounding contours

smooth, only the lowest K radial frequencies were included in the linear superposition used to construct amoeba targets. To span the maximum range of contour shapes and sizes, the amplitude and phase of each radial frequency component was chosen randomly, under the restriction that the minimum and maximum diameters could not exceed lower and upper limits. When only 2 radial frequencies were included in the superposition, the resulting amoebas were very smooth. As more radial frequencies were included, the contours became more complex. Thus, K , the number of radial frequencies included in the superposition, provided a control parameter for adjusting target complexity. Figure 1 shows target and distractor images generated using different values of K .

Human subjects are able to infer whether a two isolated line segments extracted from a natural scene are from the same or from separate contours using only distance, direction and relative orientation of the two segments as cues [25,37]. The performance of human subjects is well predicted by differences in the empirically calculated co-occurrence statistics of short line segments drawn from either the same or from different contours. To explore the ability of cortical association fields to account for the perception of smooth contours, we developed a network-level computational model of lateral interactions between orientation-selective elements governed by sigmoidal (piecewise linear) input/output synaptic transfer functions. To model lateral interactions, we constructed “Object-Distractor Difference (ODD) kernels” for the amoeba/no-amoeba task by computing coactivation statistics for the responses of pairs of orientation-selective filter elements, compiled separately for target and distractor images (Figure 2). Because the amoeba/no-amoeba image set was translationally invariant and isotropic, the central filter element may without loss of generality be shifted and rotated to a canonical position and orientation. Thus the canonical ODD kernel was defined relative to filter elements at the origin with orientation $\pi/16$ (to mitigate aliasing effects). Filter elements located away from the origin can be accounted for by a trivial translation. To account for filter

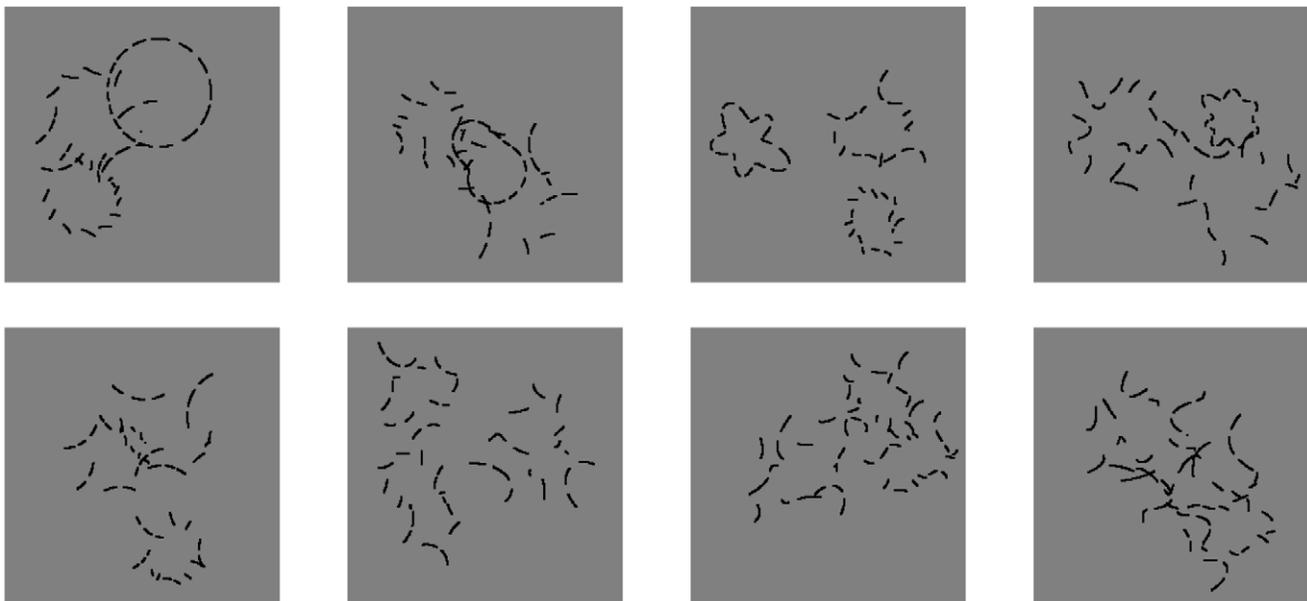


Figure 1. Examples of targets and distractors from the amoeba/no-amoeba image set for different K . From left to right: $K=2,4,6,8$. Top row: Targets; amoeba complexity increases with increasing numbers of radial frequencies. Clutter was constructed by randomly rotating groups of amoeba contour fragments. Bottom row: Distractors; only clutter fragments are present. doi:10.1371/journal.pcbi.1002162.g001

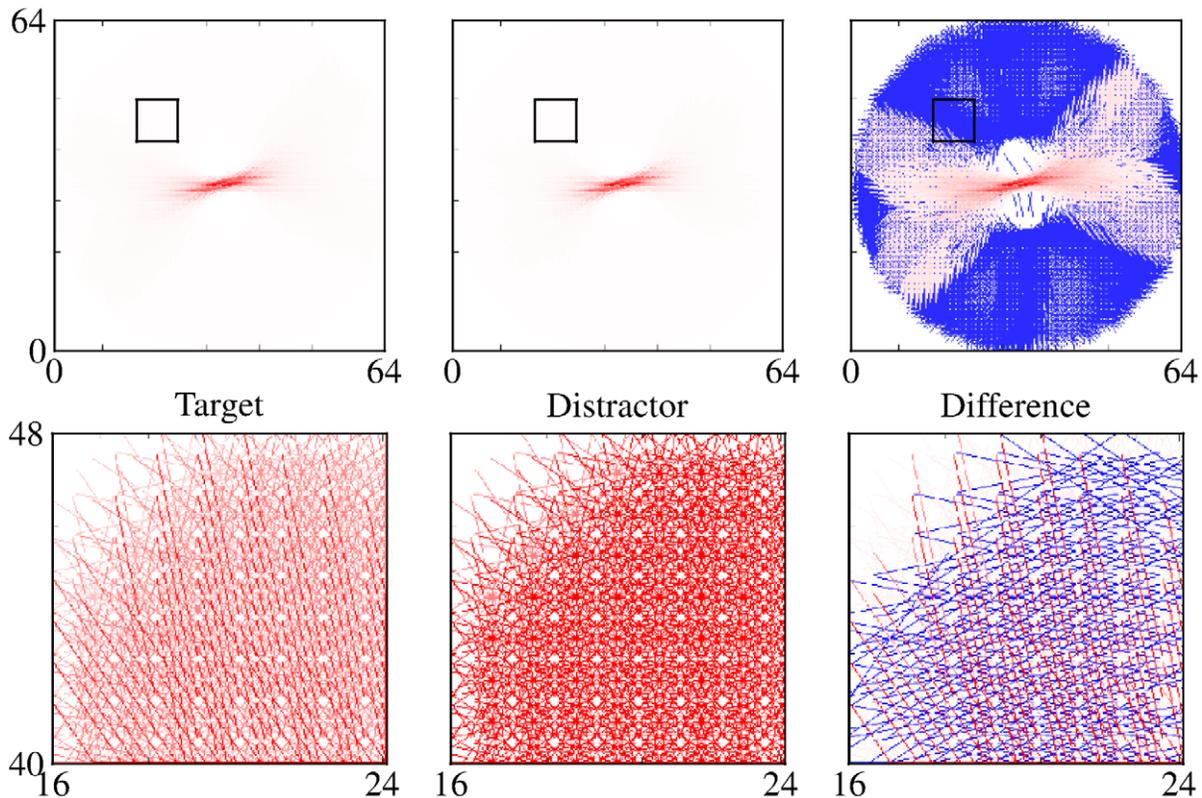


Figure 2. ODD kernels. Top Row: For a single short line segment oriented approximately horizontally at the center (not drawn), the co-occurrence-based support of other edges at different relative orientations and spatial locations is depicted. Axes were rotated by $(180^\circ/16)$ from vertical to mitigate aliasing effects. The color of each edge was set proportional to its co-occurrence-based support. The color scale ranges from blue (negative values) to white (zero) to red (positive values). Left panel: Co-occurrence statistics compiled from 40,000 target images. Center panel: Co-occurrence statistics compiled from 40,000 distractor images. Right panel: ODD kernel, given by the difference in co-occurrence statistics between target and distractor kernels. Bottom Row: Subfields extracted from the middle of the upper left quadrant (as indicated by black boxes in the top row figures), shown on an expanded scale to better visualize the difference in co-occurrence statistics between target and distractor images. Alignment of edges in target images is mostly cocircular whereas alignment is mostly random in distractor images, accounting for the fine structure in the corresponding section of the ODD kernel.

doi:10.1371/journal.pcbi.1002162.g002

elements with different orientations, separate ODD kernels were computed for 8 orientations then rotated to a common orientation and averaged to produce a canonical ODD kernel. The canonical kernel was then rotated in steps between 0 and π (offset by $\pi/16$) and then interpolated to Cartesian $x-y$ axes by rounding to the nearest integer coordinates.

The resulting ODD kernels were generally consistent with the predictions of cocircular constructions [24], except that support was mostly limited to line elements lying along low curvature contours, which follows naturally from the prevalence of low curvatures in our amoeba training set.

Curiously, the largest differences in the coactivation statistics occur close to the center of the kernel, where targets and distractors are presumably most similar. However, even at short distances, amoeba segments are still more likely to be aligned than clutter elements. Moreover, nearby pairs occur much more frequently than more distant pairs, amplifying their contribution to the difference map. Since, by design, the individual clutter fragments were locally indistinguishable from the target fragments, co-occurrence statistics of oriented fragments were necessary to solve the amoeba/no-amoeba task. The simplest solution, adopted here, was to focus on pairwise co-occurrences. Notably, in some neural preparations, pairwise interactions have been shown to be sufficient to account for a large fraction of all higher-order correlations [38,39].

At the retinal stage, target and distractor images were represented as 256×256 pixel monochromatic, binary line drawings. At the next stage, corresponding to an early cortical processing area such as V1, a set of filters was used to represent 8 orientations, uniformly-spaced and centered at each pixel, with the axes rotated slightly (by $\pi/16$) to mitigate aliasing artifacts. The bottom-up responses of each orientation-selective element were computed via linear convolution using filters composed of a central excitatory subunit flanked by two inhibitory subunits. Each subunit was an elliptical Gaussian with an aspect ratio of 7 : 1, consistent with the aspect ratios of V1 simple cell receptive fields measured experimentally [40] and similar to values employed in previously published models of V1 responses [41]. Likewise, we estimate that each image pixel subtended a visual angle of approximately 0.025° (see Methods), so that each orientation-selective element in the model subtended a visual angle of approximately 0.2° , consistent with physiological estimates of V1 receptive field sizes at small eccentricities [42]. All subunits had the same total integrated strength (to within a sign), whose magnitude was adjusted to yield relatively clean representations of the original image in terms of oriented edges. The synaptic transfer function was piecewise-linear with a minimum value of 0.0 and a maximum value of 1.0 and a fixed threshold of 0.5. A finite threshold and saturation level were essential in order to allow non-saturated

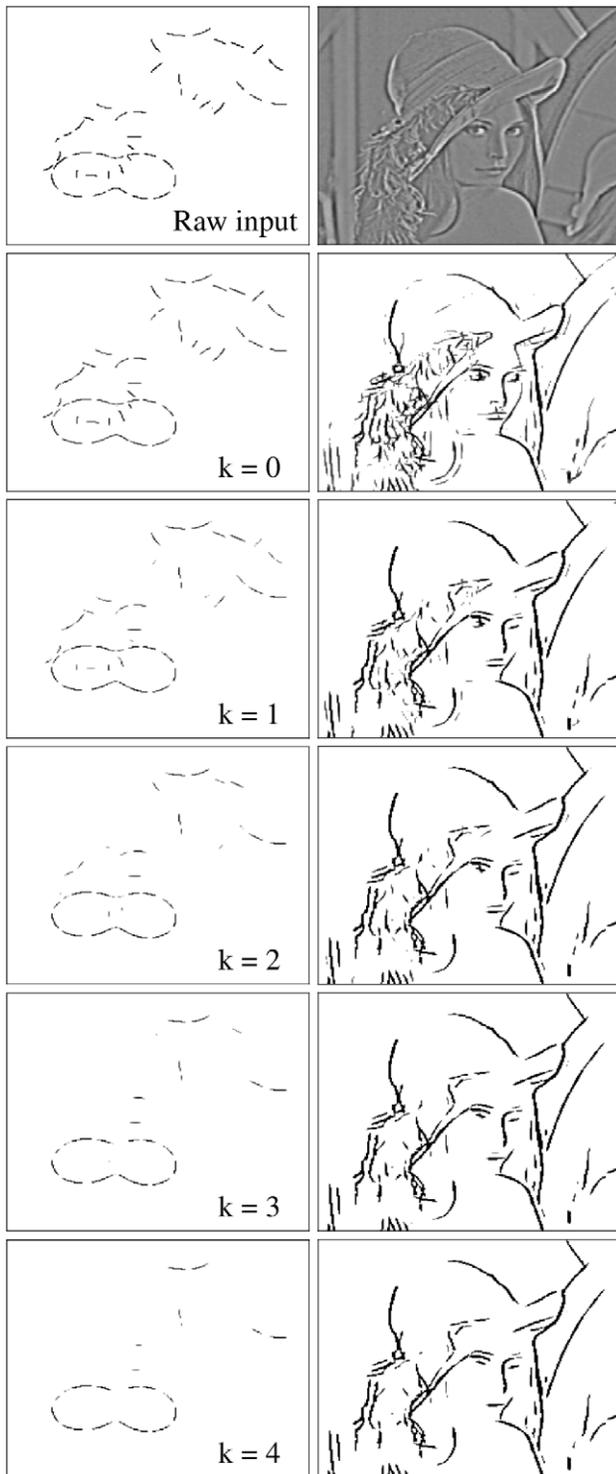


Figure 3. The effect of lateral interactions on example images. Left column: black and white amoeba-target image ($K=4$). Right column: Gray-scale natural image (the standard computer vision test image “Lena”) after applying a hard Difference of Gaussians (DoG) filter to enhance edges. Top row: Raw retinal input. Second row: Responses of orientation-selective elements before any lateral interactions ($k=0$). To aid visualization, the activity of the maximally responding orientation-selective element at each pixel location is depicted as a gray-scale intensity. Rows 3-6: Activity after $k=1,2,3,4$ iterations of the multiplicative ODD kernel, as labeled. For each iteration, activity was multiplied by the local support, computed via linear convolution of the

previous output activity with the ODD kernel. Lateral interactions tended to support smooth contours, particularly those arising from amoeba segments, while suppressing clutter or background detail. doi:10.1371/journal.pcbi.1002162.g003

contour fragments to be suppressed while preventing well-supported fragments from growing without bound. The precise values used for threshold and saturation were not critical, as responsiveness was controlled independently by adjusting the overall integrated strength of the bottom-up and lateral interaction kernels (see *Methods*).

Orientation-selective responses were modulated by 4 successive applications of the multiplicative ODD kernel. Lateral support was first computed via linear convolution of the ODD kernel with the surrounding orientation-selective elements, out to a radius of 32 pixels. Given that images were approximately $7^\circ \times 7^\circ$ in extent (see *Methods*), ODD kernels spanned a total visual angle of approximately 1.75 degrees, roughly in correspondence with the estimated visuotopic extent of horizontal projections in V1 [42]. The previous activity of each cell was multiplied by the current lateral support, passed through the piecewise-linear synaptic transfer function, and the process repeated for up to 4 iterations. Contour segments that received insufficient lateral support were thereby suppressed, whereas strongly supported elements were either enhanced or remained maximally activated. When applied to the amoeba/no-amoeba image set, the ODD kernels typically suppressed clutter relative to target segments (Figure 3, left column).

When applied in a similar manner to a natural gray-scale image to which a hard Difference-of-Gaussians (DoG) filter has been applied to maximally enhance local contrast (see Figure 3, right column), ODD-kernels tended to preserve long, smooth lines while suppressing local spatial detail. Although ODD kernels were trained on a narrow set of synthetic images, the results exhibit some generalization to natural images due to the overlap between the cocircularity statistics (see Figure 2) of the synthetic image set and those of natural images.

To quantify the ability of the model to discriminate between amoeba/no-amoeba target and distractor images, we used the total activation summed over all orientation-selective elements after k iterations of the ODD kernel. A set of 2,000 target and distractor images was used for testing; test images were generated independently from the training images. Histograms of the total activation show increasing separability between target and distractor images as a function of the number of iterations (Figure 4). To maximize the range of shapes and sizes spanned by our synthetic targets and distractors, we did not require that the number of ON retinal pixels be constant across images. Rather, the retinal representations of both target and distractor images encompassed a broad range of total activity levels, although the two distributions strongly overlapped and there was no evident bias favoring one or the other. At the next processing stage, prior to any lateral interactions, there was likewise little or no bias evident in the bottom-up responses of the orientation-selective elements. Each iteration of the multiplicative ODD kernel then caused the distributions of total activity for target and distractor images to become more separable, implying corresponding improvements in discrimination performance on the amoeba/no-amoeba task.

The general principles governing the operation of our model cortical association fields are conceptually straightforward. ODD kernels, which capture differences in the coactivation statistics of edge segments belonging to amoebas relative to edge segments belonging to the background clutter, are used to determine the

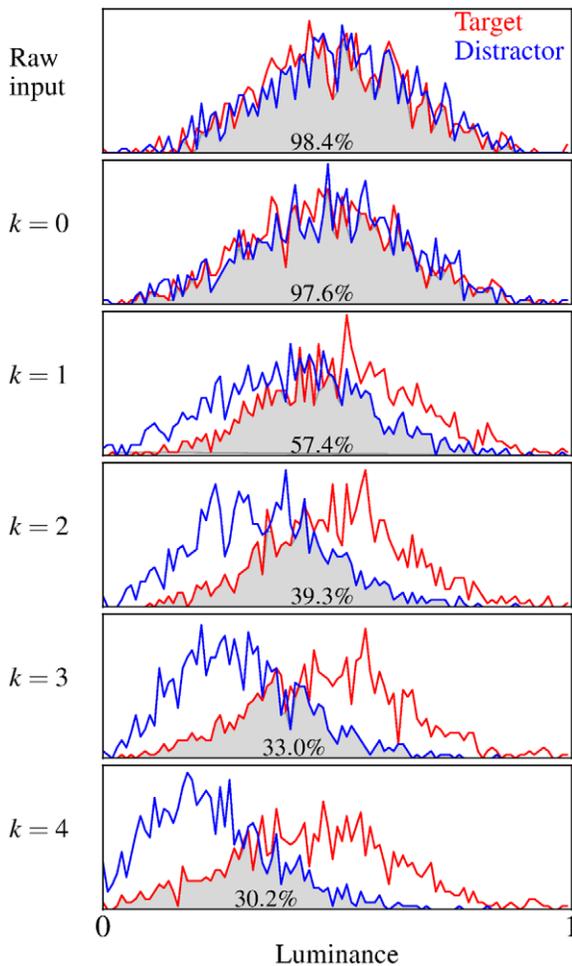


Figure 4. Histograms of total luminance in target and distractor images as a function of the number of iterations. Red bins: Total activity histograms for all 1,000 test target images. Blue bins: Total activity histograms for all 1,000 test distractor images. The degree that the two distributions overlap is shown as the gray shaded area, which provides a measure of whether total luminance can be used to distinguish targets from distractors. The percentage in each shaded area shows the approximate lower bound amount of overlap of the two histograms, for comparison. Top row: Total summed activity over all retinal pixels. Little, if any bias between target and distractor images was evident in the input black and white images as there is nearly complete overlap between the distributions. Subsequent rows: Total activity histograms summed over all orientation-selective elements. Second row: Bottom-up responses prior to any lateral interactions. Third - sixth rows: Total activity histograms after 1 - 4 iterations of the multiplicative ODD kernel, respectively. Total summed activity became progressively more separable with additional iterations, as evinced by a decrease in the overlapping areas.
doi:10.1371/journal.pcbi.1002162.g004

lateral contextual support for individual edge segments in an image. Edge segments receiving sufficiently strong support are preserved, indicating they are likely to be part of an amoeba, whereas edge segments receiving insufficient support are suppressed, indicating they are likely to belong to the background clutter.

To assess the ability of the model cortical association fields to account for the time course of human contour perception, we measured the stimulus presentation time required for human subjects to reach a given level of accuracy on an amoeba/no-amoeba task. The psychophysical experiment was implemented

using a speed-of-sight protocol employing a two-alternative forced choice (2AFC) design, with subjects using a slider bar to indicate which of two images, presented side-by-side, contained an amoeba (Figure 5). The distance the bar was displaced to the left or to the right was used to indicate confidence, see *Methods*. To effectively interrupt visual processing at a given SOA, both target and distractor images were replaced by an optimized mask, constructed by combining randomly rotated amoeba and clutter segments [20]. Our optimized masks were designed to render the amoeba targets virtually invisible in the fused target-mask composite.

As a measure of human performance on the amoeba/no-amoeba task, we constructed receiver operating characteristic (ROC) curves [43] (Figure 6), using each subject's reported confidence (slider bar location relative to the center position) as a noisy signal for estimating which side, either left or right, contained the target on a given trial. True positives corresponded to trials on which the subject reported the target was on the left

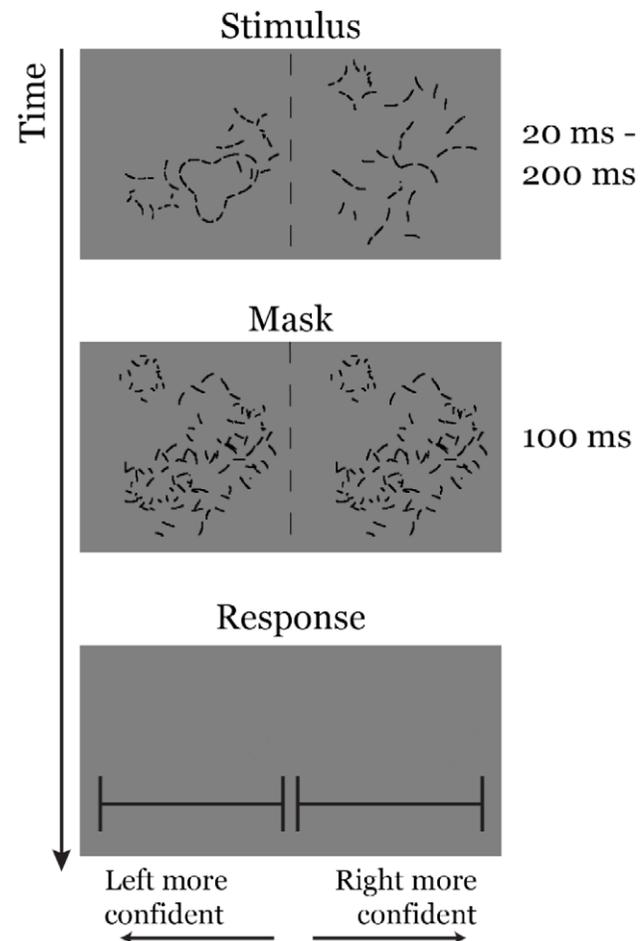


Figure 5. Psychophysical experiment schematic. The stimulus consisted of one target image and one distractor image (randomly positioned with equal probability on the left or right), presented simultaneously for an SOA between 20 ms and 200 ms, followed by an optimized 100 ms mask generated from randomly rotated groups of target and distractor segments. Subjects indicated which side contained the target object (amoeba) using a computer mouse to click along a horizontal slider bar. Clicking far to the left or right indicated strong confidence that the corresponding side contained the target; clicking close to the center indicated weak confidence. A narrow gap in the center forced subjects to choose between left and right.
doi:10.1371/journal.pcbi.1002162.g005

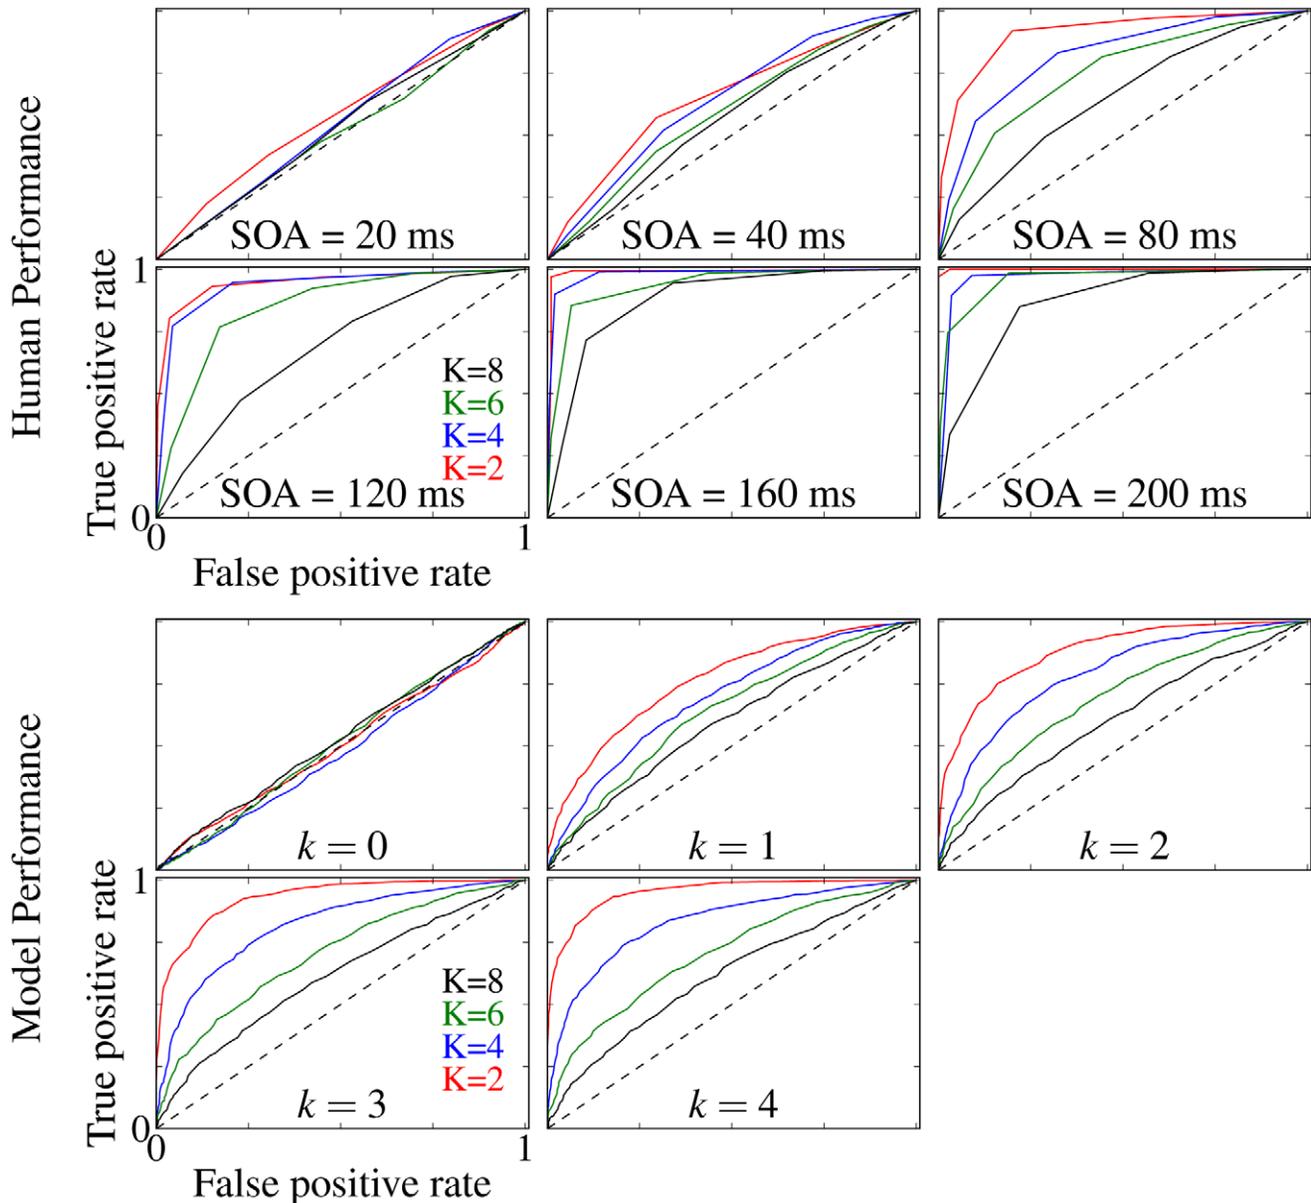


Figure 6. ROC curves comparing human and model performance on the amoeba/no amoeba task. Top two rows: ROC curves averaged over four different human test subjects using reported confidence (points). The dashed diagonal line in each plot indicates the curve corresponding to chance. Red, blue, green, black correspond to $K=2,4,6,8$, respectively. Bottom two rows: ROC curves for model cortical association fields computed from total activity histograms.
doi:10.1371/journal.pcbi.1002162.g006

(relative to threshold) and the target was actually on the left (relative to threshold). False positives corresponded to trials on which the subject reported the target was on the left whereas the target was actually on the right (relative to threshold). To construct each ROC curve, the confidence scale along the slider bar was divided into 6 discrete threshold values. For each threshold value, a cumulative proportional true positive rate was calculated by considering only those trials as true positives in which the confidence value was above threshold. The cumulative proportional false positive rate for each threshold value was calculated similarly. Each threshold value thus contributed one point on the ROC curve, with true positive rate plotted as the ordinate and the false positive rate as the abscissa. The complete set of points was connected by straight lines to guide the eye (Figure 6), with a

separate ROC curve computed for each combination of SOA and target complexity.

ROC curves for quantifying the performance of the model on the amoeba/no-amoeba task were computed similarly, using the difference in total luminance between the left and right images as the raw signal for estimating which side contained the target on a given trial. If the total luminance of the left image was higher than that of the right (relative to threshold), the response of the model would be reported as target on the left. Ideally, after several iterations of the ODD kernel, no segments would remain in the distractor image and only amoeba segments would remain in the target image; in practice, the total luminance served as a measure of confidence. Given the much larger number of trials (1000) available for assessing model performance, 100 equally spaced

threshold values were used to calculate the corresponding ROC curves. As with the ROC curves constructed from the confidence values reported by the human subjects, the ROC curves computed from the confidence values reported by the model give the cumulative proportional true positive rate as a function of cumulative proportional false positive rate, with the confidence threshold varied from zero to maximum. Graphically, the area under the ROC curves is given by the amount of overlap between the total luminance histograms (see figure 4) for the target and distractor images [44].

ROC curves for human subjects show performance increasingly above chance, indicated by a diagonal line of slope 1, as a function of both increasing SOA and decreasing target complexity. For amoeba targets of low to moderate complexity, ROC curves obtained from human subjects were well matched to those generated by the model cortical association fields, consistent with the hypothesis that lateral interactions between orientation-selective neurons contribute to human contour perception, at least for simple targets.

The area under the ROC curve (AUC) gives the probability that a randomly chosen target image will be correctly classified relative to a randomly chosen distractor image, and thus provides a threshold-independent assessment of performance on the 2AFC task. Both the average over human subjects and the model cortical association fields exhibited qualitatively similar performance on the 2AFC amoeba/no-amoeba task (Figure 7). Performance declined as a function of increasing target complexity, both for human subjects, measured at a fixed SOA, and for the model, measured at a fixed number of iterations, implying that K was an effective control parameter for adjusting task difficulty. At 20 ms SOA, the performance of human subjects was indistinguishable from chance, suggesting that our optimized masks effectively prevented the development of bottom-up cortical responses, even for the simplest targets ($K=2$). Although some studies report that line drawings are processed more rapidly than natural images, with above chance performance being observed at short SOA values [1,26], the fact that performance on the amoeba/no-amoeba task was no better than chance at a 20 ms SOA implies that our optimized masks effectively interrupted visual processing of the amoeba targets. Since the model used here did not include any account for the time course of bottom-up retinocortical dynamics, we assumed that the performance of human subjects at 20 ms SOA should be equated to model performance at 0

iterations (prior to any lateral interactions), a time frame consistent with the distribution of the shortest measured response latencies recorded in primary visual cortex [45].

Overall, average human performance improved as a function of increasing SOA in a manner analogous to the improvement in model performance as a function of the number of iterations of the ODD kernel. This correspondence was especially evident for amoebas of low to moderate complexity ($K \leq 4$). For more complex targets, model performance lagged well behind that of human subjects. Studies suggest that low and high radial frequencies are processed by different cortical channels [46]. Model performance might have been improved by training a new set of ODD kernels specifically for targets containing $K \geq 6$ radial frequencies, thereby utilizing a hypothetical sub-population of orientation-selective neurons optimized for detecting high-curvature contours. Here, our model was limited to a single multiplicative kernel for detecting all predominately smooth contours.

To quantify how average human performance on the 2AFC amoeba/no-amoeba task varied with SOA, and to compare with the dependence of model performance on the number of iterations of the ODD kernel, areas under both sets of ROC curves were fit to a monotonically increasing function of the following sigmoidal form:

$$f(t) = \frac{F_{\infty}}{1 - (1 - 2F_{\infty})e^{-\lambda(t-t_0)}}. \quad (1)$$

For human experiments, the parameter t corresponds to the SOA in ms. Since we expect humans to perform close to 100% accuracy for very long SOA, we set $F_{\infty} = 1$. Since humans perform essentially at chance (50%) for 20 ms SOA, we set $t_0 = 20$ ms. Thus λ was the only free parameter; fits to the average human data were denoted by λ_H ; λ_H has units of 1/ms. Likewise, model performance was fit to a curve with the same functional form, with $t = k$ measuring the number of iterations; λ_M was used to denote curve fits to the model data. However, visual inspection of the model data suggests that its performance saturates at less than 100% accuracy even after an infinite number of iterations, thus we forced the sigmoidal curve fit to the model results to asymptote at the final measured value of AUC: $F_{\infty} = AUC_{k=4}(K)$. Since the model performs better than chance after only 1 iteration, we set

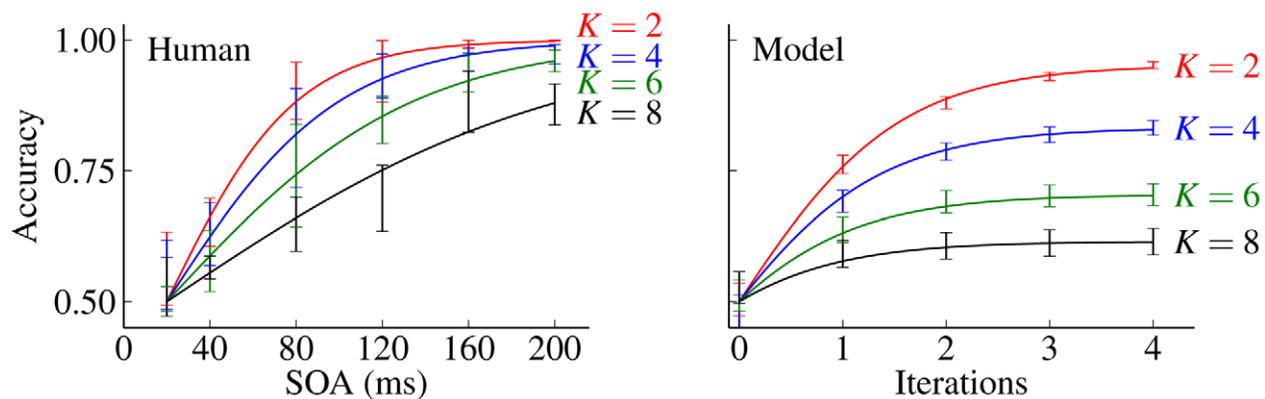


Figure 7. A comparison of human and model performance on the 2AFC amoeba/no amoeba task. Left: Average human performance for different SOA in milliseconds. Right: Performance of model cortical association fields for increasing numbers of iterations. Both panels: Accuracy, which is equivalent to area under the ROC curve, (error bars) fitted to single sigmoidal functions (solid lines). The four curves from top to bottom correspond to $K=2,4,6,8$ radial frequencies. doi:10.1371/journal.pcbi.1002162.g007

$t_0=0$. For both the human experiments and the model performance, the functional form of $f(t)$ ensures that $f(t_0)=\frac{1}{2}$, corresponding to a minimal performance equal to chance.

We find that $1/\lambda_H$ and $1/\lambda_M$ behave quite differently as a function of K , the number of radial frequencies used in amoeba generation (Figure 7). As anticipated for a relaxational process governed by a single kernel, the model data was well described by a single value of λ_M (in units of 1/iterations), equal to 1.26. For the human subjects data, values of λ_H increased from 0.034 to 0.011 as a function of amoeba complexity, corresponding to lateral processing times of 29.8 to 90.6 ms, respectively. If human performance depended on only a single set of lateral connections, then, at least in the linear approximation case, we might expect human performance to be well described by a single dominant time constant, representing the dominant eigenmode of the horizontal interactions [47,48]. Multiple time scales in the human performance case may emerge from any number of physiological mechanisms not included in the present model, including additional non-linearities in the action of the horizontal connections and/or contributions to contour perception from extrastriate areas. Our data do not allow us to make a firm distinction between these possibilities.

However, one possible interpretation of the present results is that the perception of simple contours is dominated by relatively fast lateral interactions placed early in the visual processing pathway, thereby accounting for the good fit between the model and experimental results for targets of low to moderate complexity. Building on this interpretation, we postulate that the perception of more complex contours requires more extensive, and therefore slower, processing mechanisms involving higher cortical areas, thus explaining the discrepancy between model and experimental performance as target complexity increases. Under the assumption that human perception of simple amoeba targets ($K \leq 4$) depends primarily on recurrent lateral interactions between orientation-selective neurons, we can estimate the time required for each iteration of the multiplicative ODD kernel. This rate is estimated using the $K=2$ time constants from the fits: $\lambda_{M,K=2}/\lambda_H = 37.5$ ms

per iteration, a value consistent with estimates of lateral conduction delays within the same cortical area [13].

Having shown that the lateral interactions based on multiplicative ODD kernels can account for both spatial and temporal aspects human contour perception, we seek to identify model details that are essential to the performance reported here. First, we demonstrate that the proposed model is robust and does not require that the magnitude of the ODD kernel be carefully titrated to a precise value. Model performance on the 2AFC amoeba/no-amoeba task, measured by the area under the ROC curve (AUC) for increasing numbers of iterations ($k=0,1,2,3,4$), was plotted for different values of the strength of the ODD kernel, given by the total integrated strengths of the equal and opposite target and distractor contributions (Figure 8). The number of radial frequencies was fixed at $K=4$. Qualitatively similar performance was obtained for ODD kernel strengths ranging from 300 to 400. The ODD kernel used in the present study, whose strength was set to 325, produced near optimal performance and also exhibited monotonic improvement with increasing numbers of iterations. That performance was generally insensitive to the value of the main free parameter in the model provides strong evidence for the robustness of the proposed contour detection mechanism based on multiplicative lateral interactions.

A second aspect of the model that merits scrutiny is the detailed structure of the ODD kernels, which were trained using computer-generated images in which the pairwise edge statistics uniquely identifying globally salient contours could be calculated directly. Previous models of contour perception typically employed much simpler patterns of lateral connectivity, in which excitatory interactions were either collinear or cocircular, and inhibitory interactions were approximately independent of relative orientation [8,24,27,47–49]. To determine if the detailed structure of the ODD kernel was critical to the observed performance, we repeated the amoeba/no-amoeba experiment using a much simpler kernel whose basic form was consistent with a number of previously published models (see Figure 8). Specifically, we used a “Bowtie” kernel in which excitatory connections fanned out with

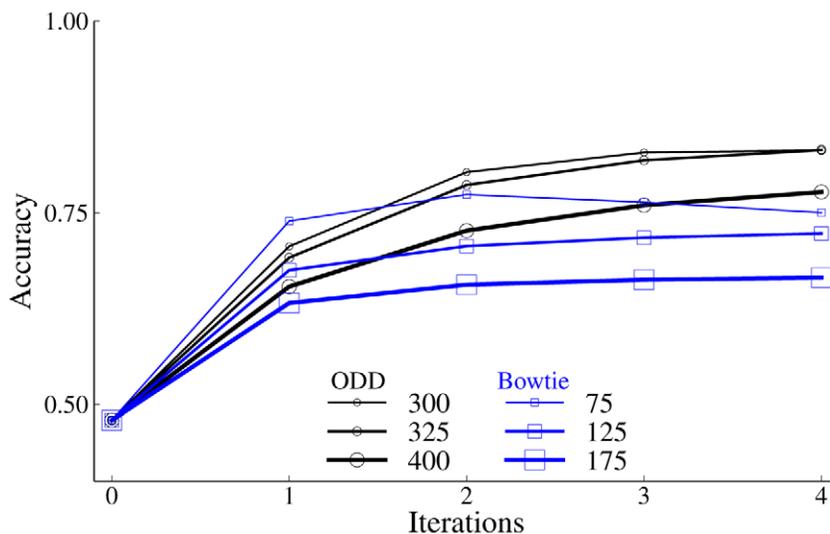


Figure 8. A comparison of ODD and simpler “Bowtie” kernel performance on the on the 2AFC, $K=4$ amoeba/no amoeba task plotted as a function of the number of iterations for a range of different kernel strengths. Line width and marker size denote values on kernel strength, which was the main free parameter in the model. Kernel strength is a dimensionless constant. Black lines: ODD kernel performance. Blue lines: “Bowtie” kernel performance. Qualitative behavior was similar for both kernels, demonstrating that multiplicative lateral interactions act robustly to reinforce smooth closed contours.

doi:10.1371/journal.pcbi.1002162.g008

an opening angle of $\pi/6$ and the difference in the preferred orientations of the pre- and post-synaptic elements differed by no more than $\pm\pi/6$. Both excitatory and inhibitory connection strengths fell off in a Gaussian manner, with inhibition strength being insensitive to orientation. Although the overall accuracy of the Bowtie kernels was lower than that achieved by the ODD kernels, performance on the amoeba/no-amoeba tasks was qualitatively similar, particularly regarding the general monotonic improvement with the number of iterations and the absence of a sensitive dependence on kernel strength. Thus, we conclude that multiplicative lateral interactions are able to preserve smooth closed contours while suppressing clutter in a manner that is robust to broad changes in model details.

Discussion

We have shown that simple models of neural activity in primary visual cortex, enriched with lateral association kernels, reproduce some of the behavioral features regarding the human perception of broken closed contours. Our results agree not only with the measured dependence on contour complexity but also with the temporal dependence of human perception as a function of SOA, suggesting that horizontal connections in V1 may play a non-trivial and global computational role in the perception of closed contours on very fast timescales.

A number of studies relate to the potential contribution of cortical association fields to human contour perception; these encompass a range of anatomical, physiological, psychophysical, and theoretical techniques [2–5,7–10,11,13,16–19,50]. In particular, a number of theoretical models have sought to account for human contour perception at the level of biologically-plausible neural circuits [8,27,28,49,51–54], with most studies incorporating some form of cortical association field configured to reinforce smoothness [24]. Although biologically plausible models of cortical association fields have been used to account for the dependence of contour visibility on key parameters controlling task difficulty, such as smoothness, closure, and density of background clutter [8], model cortical association fields have not been directly compared to the time course of human contour perception as a function of contour complexity. Here, we used cortical association fields based on ODD kernels, which were computed from differences in the pairwise coactivation statistics of orientation-selective elements arising from target as opposed to distractor images. While we designed the kernels specifically for the amoeba-clutter disambiguation, we emphasize that the algorithm for the ODD kernel construction is completely general and can be used to improve detection of salient image features in any situation where generative models of targets and distractors are known, or there exists data sets of sufficient size to characterize the contour co-occurrence statistics empirically for both targets and distractors. In our experiments, ODD kernels were able to account for the experimentally observed variations in the saliency of closed contours as a function of parametric complexity and for the time course with which smooth contours are processed by cortical circuits. Crucial for these results was our use of a synthetic target/distractor data set with controllable complexity and the absence of top-down contextual features or local cues that might give away target presence.

Here, we used a semi-supervised training scheme to learn lateral connectivity patterns optimized for performing the amoeba/no-amoeba task. Necessarily, we sought to model only a subset of the lateral interactions between orientation-selective neurons, namely, those horizontal connections configured to reinforce smooth, closed contours. We did not attempt to capture the full range of

spatial relationships between features extracted at early cortical processing stages [24,55]. Presently, databases containing sufficient numbers of fully *annotated* and *segmented* natural images needed to reproduce the weeks (or months) of visual experience required to train the full complement of horizontal connections in the primary visual cortex do not exist. Moreover, the computational resources to exploit such databases, even if they did exist, are highly non-trivial to assemble. Thus, we focused here on a subset of horizontal connections for which it was possible to construct synthetic surrogate images. At most, the proposed model represents a subset—and only a subset—of the lateral connections between orientation-selective cortical neurons. Moreover, even a complete set of such horizontal connections would, at most, represent but a subset of the cortical mechanisms that contribute to the time course and shape-dependence of contour perception.

The supervised training scheme employed here might be related to perceptual learning phenomena, which take place over time scales much shorter than those typically associated with developmental processes [56–58]. It is possible that known physiological mechanisms, such as spike-timing-dependent plasticity (STDP), especially with accounts for realistic conduction delays [59], could mediate a rapid refinement of lateral connections so as to facilitate the perception of amoeba targets. Moreover, physiological plasticity mechanisms might produce different patterns of connectivity for orientation-selective elements representing points of low as opposed to high local curvature, thereby optimizing lateral interactions for contours of varying complexity. Here, we made no attempt to customize distinct ODD kernels for detecting contours of varying complexity. Instead, a single ODD kernel was trained using a complete set of images in which different numbers of radial frequency components were equally represented. Although we did not investigate whether, or to what extent, the performance of human subjects improved over the course of the amoeba/no-amoeba experiment, such investigations might shed insight into the role of perceptual learning in the detection of closed contours.

The question of how lateral connectivity based on ODD kernels might be acquired during development was not addressed explicitly. In principle, coactivation statistics between pairs of orientation-selective neurons could be accumulated over time in an unsupervised manner by a Hebbian-like learning rule [60]. Under natural viewing conditions, we expect that contour fragments consistent with smooth, closed boundaries would tend to occur simultaneously, whereas contour fragments inconsistent with object boundaries would tend to occur at random temporal delays. Thus, a Hebbian-like learning rule sensitive to temporal correlations, such as certain mathematical forms of STDP-like learning rules [61], might under normal developmental conditions lead to connectivity patterns that reinforce smooth contours.

Of course, human contour perception may have nothing to do with cortical association fields, or lateral interactions may play a subordinate role. Early models showed how spatial filtering could enhance texture-defined contours in the absence of orientation-specific interactions [4] and short-range lateral interactions can accentuate texture-defined boundaries [31,62]. However, psychophysical studies employing implicit contours [2,7,8], in which foreground and background elements are present at equal density and which lack explicit texture cues, appear to rule out explanations that omit long-range, orientation-specific interactions. An influential class of biologically-inspired computer vision models achieves a degree of viewpoint-invariant object recognition by constructing feed-forward hierarchies to extract progressively more complex and viewpoint invariant features [33,63]. By analogy with such models, scale- and position-independent

representations for detecting long, smooth contours could in principle be constructed hierarchically, starting with simple edge detectors and building up progressively longer, more complex curves using a “bag-of-features” approach. Presently, there appear to be insufficient data to decide whether human contour perception involves primarily lateral, feed-forward, or even top-down connections [30,32,64]. Hypothetically, the cortical association fields used in the present study could have been implemented as a feed-forward architecture, using a hierarchy of orientation-selective neurons to link progressively more widely separated contour fragments. Functionally, there may not exist a clean distinction between lateral, feed-forward and feed-back topologies, with the possibility that all three types of connectivity contribute to human contour perception.

To quantify the temporal dynamics underlying visual processing, we performed speed-of-sight psychophysical experiments that required subjects to detect closed contours (amoebas) spanning a range of shapes, sizes and positions, whose smoothness could be adjusted parametrically by varying the number of radial frequencies (with randomly chosen amplitudes). To better approximate natural viewing conditions, in which target objects usually appear against noisy backgrounds and both foreground and background objects consist of similar low-level visual features, our amoeba/no-amoeba task required amoeba targets to be distinguished from locally indistinguishable open contour fragments (clutter). For amoeba targets consisting of only a few radial frequencies ($K \leq 4$), human subjects were able to perform at close to 100% accuracy after seeing target/distractor image pairs for less than 200 ms, consistent with a number of studies showing that the recognition of unambiguous targets typically requires 150–250 ms to reach asymptotic performance [22,23,35], here likely aided by the high intrinsic saliency of closed shapes relative to open shapes [7]. Because mean inter-saccade intervals are also in the range of 250 ms [34], speed-of-sight studies indicate that unambiguous targets in most natural images can be recognized in a single glance. Similarly, we found that closed contours of low to moderate complexity readily “pop out” against background clutter, implying that such radial frequency patterns are processed in parallel, presumably by intrinsic cortical circuitry optimized for automatically extracting smooth, closed contours. As saccadic eye movements were unlikely to play a significant role for such brief presentations, it is unclear to what extent attentional mechanisms are relevant to the speed-of-sight amoeba/no-amoeba task.

Our results further indicate that subjects perform no better than chance at SOAs shorter than approximately 20 ms. Other studies, however, report above chance performance on unambiguous target detection tasks at similarly short SOA values [1,23,26,33]. The discrepancy may be attributed to the different masks employed. Whereas the above cited studies used masks consisting of either spatially filtered (e.g. $1/f$) noise, distractor images, or scrambled versions of the target image set, we constructed rotation masks that were optimized for each target/distractor image pair [20]. Our working hypothesis was that an optimized mask should completely obscure the target object in the target-mask composite image; also referred to as pattern masking. The requirement that the mask completely hide the target follows from the assumption that at very short SOA, the target and mask images are likely to be effectively fused due to the finite response time of neurons and receptors in the early visual system [65]. For the amoeba/no-amoeba task, we created optimized masks by rotating the amoeba and clutter fragments with the goal of producing the maximum amount of interference in the responses of orientation-selective cells. Presumably, maximum interference occurs when orientation-selective neurons are presented with randomly rotated contour

fragments in rapid succession. Although backward masks can have heterogeneous effects, with performance in some cases showing a *U*-shaped dependence on SOA [66], for the masks used here performance always increased monotonically with SOA. Empirically, the fact that performance was no better than chance at 20 ms SOA suggests that our optimized masks were able to effectively interrupt the processing of smooth, closed contours at early cortical processing stages. Indeed, the ability to drive overall performance down to chance at SOA values shorter than 20 ms could provide an operational criteria for assessing the degree to which a given backward pattern mask is able to effectively interrupt visual processing.

The amoeba/no-amoeba task required the integration of information over length scales spanning viewing angles of approximately $3-4^\circ$, larger than the classical excitatory receptive field size of parafoveal V1 neurons. The amoeba/no-amoeba image set (see Figure 1) was configured so that purely local information, such as a few adjoining contour fragments, would not be sufficient to solve the target detection problem. Rather, distinguishing amoebas from clutter required integrating global information across multiple contour fragments. Our results suggest that such global integration can be accomplished via lateral interactions between local, orientation-selective filters. Although the density of target and clutter segments was not precisely equilibrated in our amoeba/no-amoeba image set, the wide range of target sizes and shapes spanned by our image generation algorithm makes it unlikely that the near perfect performance of human subjects at long SOA could have been attained using density cues alone [4]. Here, lateral inputs were used to modulate the bottom-up responses in a multiplicative fashion, so that our cortical association fields acted primarily as gates that suppressed contour fragments that did not receive sufficiently strong contextual support. By preventing lateral inputs from producing activity unless there was already a strong bottom-up input, a multiplicative non-linearity prevented the activation of contour fragments not present in the original image.

The phenomenon of illusory contours suggests that in some cases contextual effects can produce activity even in the absence of a direct bottom-up response [30]. The precise form of the multiplicative interaction used here was adopted for algorithmic simplicity rather than for biological realism. We observed that including a small additive contribution from the lateral interactions did not fundamentally affect our conclusions. This suggests that ODD kernels, if implemented more generally, might account for the perception of illusory contours as well. However, a more realistic description of the underlying cellular and synaptic dynamics would likely be necessary to model a relaxation process that includes both additive and multiplicative elements.

Both the model and the psychophysical experiments employed a 2AFC design (see Figure 5) in which the goal was to correctly identify which of a pair of images contained an amoeba target. Since each trial involved a forced choice between two images, the model used a simple classifier that labeled the image with greater total activity as the target. For both human subjects and the model, the number of radial frequencies K proved to be a good control parameter for adjusting task difficulty (see Figure 7). For targets of low to moderate complexity, both model performance (as a function of number of iterations) and human performance (as a function of increasing SOA) monotonically approached nearly perfect asymptotic performance as described by a single sigmoidal function with a characteristic scale, representing either time or number of iterations, that increased with K (see Figure 7). Based on comparison with human performance at different SOA values, each iteration of the ODD kernels was estimated to require

approximately 37.6 ms of cortical processing time, consistent with measured conduction delays between laterally connected cortical neurons [13].

Prior to any lateral interactions, the stimulus was projected onto a retinotopic array of orientation-selective filter elements, providing a convenient representation for learning cortical association fields by computing differences in pairwise coactivation statistics between target and distractor images. We found that each iteration of the ODD kernel increased the activity of contour fragments that were part of amoebas compared to the activity of clutter fragments, so that after several iterations the mean overall activity, summed across all orientation-selective filter elements, was higher on average for target images than distractor images (see Figure 4). Even in trials that were incorrectly classified, contour fragments belonging to amoebas were typically still favored relative to background clutter. Because the total number of contour fragments varied from trial to trial, with only the average number of fragments being fixed across the entire image set, our relatively crude criterion for discriminating between target and distractor images sometimes led to classification errors even when amoeba fragments had been partially segmented from the background clutter, simply because the distractor image initially contained more fragments. A more sophisticated classifier might have led to a closer correspondence between model and human performance. Although performance of the present multiplicative model appeared to saturate after only a few iterations of the ODD kernel (e.g. $k \leq 4$), it is possible that a different implementation might have continued to show improvements after additional iterations. However, the longer processing time implied by additional iterations suggests that other physiological mechanisms, particularly visual search, would likely come into play. Granted, there is an apparent mismatch between the fading of clutter elements in the model and the persistence of such elements perceptually in human subjects. To reconcile this apparent mismatch, it has been suggested that the initial perception of brightness might be driven by the initial bottom-up response of the individual orientation-selective feature detectors, whereas persistent responses across these same feature detectors might drive salience [28].

The amoeba/no-amoeba image set was designed to allow for parameterized complexity (in terms of the amount of clutter, number of radial frequencies, etc.) while avoiding reference to exogenous world knowledge. Since the amoeba/no-amoeba image set was machine generated, it was possible to produce a very large number of training images; 40,000 target and 40,000 distractor images at 256×256 pixel resolution were used to train ODD kernels in the present study. Many computer vision systems employ standard image classification datasets such as the Caltech 101 [67], which allows for uniform benchmarking and thus facilitates direct comparison between models. Datasets based on natural images, however, suffer from several shortcomings. First, the resolution and number of images are fixed when the set is created. While some man-made datasets, such as MNIST [68], consist of tens of thousands of handwritten characters, annotated sets of natural photographs ideal for speed-of-sight experiments are typically limited to a few hundred images. In contrast, humans are exposed to millions of natural scenes during visual development. Biologically motivated models that attempt to replicate human performance might require similar numbers of examples. A second shortcoming of natural image datasets is prevalence of high-level contextual information that utilizes exogenous world knowledge, such as the increased a priori likelihood of finding a car on a road, or an animal in a forest. Exploiting such exogenous world knowledge poses a formidable challenge for existing

computational models and, on tasks that employ natural images, may obscure the ability of such models to extract behaviorally meaningful information from low-level visual cues. Third, natural image datasets typically provide limited capability for adjusting intrinsic task difficulty. For example, one widely used dataset [33] includes photographs of animals at different distances, but only a few discrete distances are annotated and the relationship of target distance to task difficulty is not easily quantified. Here, we illustrated how a synthetic set of images could be used to compare model and human performance in a task with parametric difficulty, potentially validating the use of artificial as opposed to natural images.

The present study addressed the role of cortical association fields in the perception of closed contours, which are presumably important for detecting visual targets based on shape or outline. Although studies show that human subjects can rapidly distinguish between images containing target and non-target object categories using only the line drawings obtained by filtering natural scenes [1], normal experience involves a number of complementary visual cues, such as texture, color, motion and stereopsis. Presumably, cortical association fields also act to reinforce features representing these complementary visual cues as well. Human subjects, for example, can distinguish whether pairs of texture patches were drawn from the same natural object or two different natural objects in a manner that exhibits a similar dependence on pairwise co-occurrence statistics as was found for orientated edges [55]. We may speculate that an analysis of coactivation statistics for features selective to a combination of cues such as local orientation, texture, color, motion, and disparity may lead to a more general and more powerful set of kernels capable of fast and effective determination of global object properties, which in turn can play an important role in complex object identification.

Methods

Synthetic amoeba/no-amoeba image set

An *amoeba* is a type of radial frequency pattern [36] consisting of a deformed circle in which the radius varies as a function of the polar angle. By choosing the number and relative amplitudes of the different frequency components, the radius can describe an arbitrarily complex shape, exactly analogous to how a Fourier basis can be used to construct an arbitrary waveform on a finite interval. Each radial frequency component was represented by a sinusoidal function defined at $C=1024$ discrete polar angles, spaced uniformly on the interval $[0, 2\pi)$. The cutoff radial frequency used in constructing the closed contour provided a control parameter for regulating the complexity of the resulting figure, which ranged from nearly circular, when only the 2 lowest radial frequencies had non-zero amplitudes, to highly sinusoidal and irregular, when the first 8 radial frequencies had non-zero amplitudes. All amoeba shapes generated here may be considered smooth, in that local curvature was always bounded.

In detail, the radius of an amoeba at each polar angle was:

$$r(\phi_c) = A_0 + \sum_{n=1}^K A_n \cos\left(\frac{2\pi n}{C} + \alpha_n\right). \quad (2)$$

All amplitudes A_n were initially drawn from normal distributions with 0 mean and unit variance. All phases α_n were drawn from uniform distributions over the interval 0 and $\pi/2$. The resulting radial frequency pattern was then linearly rescaled so that the maximum radius, r_{\max} , was equal to a random number drawn

from a uniform distribution such that $L/4 \leq r_{\max} \leq L/2$, where L is the linear size of the square image ($L=256$ pixels), and the minimum radius was given by a second randomly chosen value so that $r_{\max}/4 \leq r_{\min} \leq r_{\max}/2$. Uniform pseudo-random numbers were generated by the intrinsic MATLAB 7.0 function RAND, or its Octave 3.2.3 equivalent.

To facilitate the construction of locally indistinguishable clutter and model contour occlusion in natural images, amoeba contours were divided into 16 periodically-spaced fragments by removing short sections whose lengths varied within a specified range. Specifically, the gaps between amoeba fragments varied from 16 to 32 in units of discrete polar angle ($2\pi/C$). Amoeba contours were then broken into fragments by periodically inserting 16 gaps of variable width ranging from 16 to 32, spaced $C/16$ segments apart. Gaps were deleted from the underlying contour, so that the polar angle subtended by each fragment varied in accordance with the changes in preceding gap width. The starting point of the first gap was chosen randomly on the interval $1, C/G$, so that over the entire image set the inserted gaps were distributed uniformly around the circle.

To create clutter fragments, an amoeba was first generated using the above procedure. Consecutive amoeba fragments were then grouped, with the number of fragments in each group determined by a Poisson process with a mean value of 2 and an upper cutoff of 3. Each group of amoeba fragments was then rotated about its center of mass through random angles on the interval $\pi/8$ to $7\pi/8$. The resulting clutter consisted of the same fragments as the original amoeba but rotated so that collectively the rotated fragments no longer supported the perception of a closed object. Clutter fragments constructed in this manner were thus locally indistinguishable from amoeba fragments. To create clutter in both target and distractor images, several amoebas were first superimposed at random positions and then groups of fragments rotated following the procedure described above. All amoebas contained the same total number of contour fragments (and therefore the same number of gaps) but varied in both maximum diameter and total contour length.

The center of each amoeba was chosen randomly under the restriction that no contour be allowed to cross an image boundary. Specifically, the x -coordinate of the amoeba center, x_0 , was chosen randomly on a restricted interval, $r_{\max} \leq x_0 \leq L - r_{\max}$, and likewise for the y -coordinate, y_0 . When groups of amoeba fragments were randomly rotated to make clutter, portions of a contour belonging to a clutter fragment would occasionally cross an image boundary. In such cases, any out-of-bounds portions of a contour were reflected back into the image region using mirror boundary conditions.

Target images always consisted of 1 set of amoeba fragments and 2 sets of clutter fragments. Distractor images consisted of 3 sets of clutter fragments and thus, averaged over the entire image set, had the same mean luminance and the same variance as the target images. Mask images were constructed following a procedure nearly identical to that used for constructing distractor images, except that mask images consisted of 6 sets of clutter fragments, obtained by randomly rotating the 6 original amoeba objects used in constructing the corresponding target and distractor images. All contour fragments were initially represented as a set of points in polar coordinates, corresponding to the radius at each discrete polar angle. Points along the contour were then transformed back to Cartesian coordinates and rounded to the nearest discrete pixel value. MATLAB scripts for generating the image set used in this study are publicly available at: <http://petavision.sourceforge.net>.

Ethics statement

The Los Alamos National Laboratory (LANL) Human Subjects Research Review Board (HSRRB) has reviewed the following experimental protocol and determined that it provides adequate safeguards for protecting the rights and welfare of human subjects involved in the protocol. The protocol was reviewed and approved in compliance with the U.S. Department of Health and Human Services (DHHS) regulations for the Protection of Human Subjects, 45 CFR 46, and in accordance with the LANL Federal Wide Assurance (FWA#00000362) with the National Institutes of Health/Office for Human Research Protections (NIH/OHRP). The identification number is LANL 08-03 X.

Human psychophysics

Human performance was evaluated using two-alternative forced choice (2AFC) psychophysical experiments. There were 5 subjects, all with normal or corrected-to-normal vision. One subject only contributed data for a portion of the tested SOAs. Each subject was seated in a dark room, at an approximate distance of 65 cm from a 19-inch nominal (36.2×27 cm actual size) Hitachi 751 CRT monitor. Images spanned a viewing angle of approximately $7^\circ \times 7^\circ$. The monitor resolution was 1024×768 pixels and the refresh rate was 100 Hz. The display was driven by a dual-core 3.0 GHz Mac Pro, with MATLAB 7.6 running Psychtoolbox [69].

After a short training period to familiarize the subject with the task, one target image and one distractor image were shown side by side, followed by a mask intended to interrupt cognitive processing of the target and distractor images. Two separate sets of experiments were conducted for each subject. In one set, the SOA was chosen randomly from the values 40, 80, 120 ms. For the second set of experiments, the SOA was chosen randomly from the values 20, 160, 200 ms. The duration of the stimulus was always the same as the SOA, and thus both the target and distractor images remained visible until mask onset. The duration of the mask was always 100 ms. Each subject was shown 1200 images divided into 10 blocks of 120 images, with rest breaks in between blocks (rest break duration was at the discretion of each subject). The pace of the experiment was under the control of the subject, who initiated each trial using the space bar. A small temporal jitter, chosen uniformly between 0 to 250 ms, was added to the interval preceding each trial, to prevent entrainment. Task conditions, consisting of variations in both the SOA and the number of radial frequencies K , were randomly interleaved such that each condition occurred the same number of times over the course of the entire experiment.

On each trial, subjects indicated which side contained the target, using a mouse-driven slider bar to report confidence (see Figure 5). The reported confidence values were used to construct receiver operating characteristic (ROC) curves, which plot the percentage of true positives (or hits) against the percentage of false positives (or false alarms), with each true/false positive pair obtained by setting a confidence threshold at a different location along the slider bar. A correct response was not necessarily considered a true positive: to generate one point on the ROC curve, the reported confidence on each trial was measured relative to the current threshold position, which could be to either the left or to the right of center. Thus, a trial might be labeled as incorrect, even though the subject moved the slider bar in the correct direction, as long as the threshold level was not exceeded. Specifically, whenever the reported confidence fell to the left of threshold, the corresponding trial was treated as though the subject reported the target as being to the left, even if the threshold location had been set to the right of center and the confidence bar had actually been slid to the right. Likewise, when the reported

confidence fell to the right of the current threshold position, the trial was always treated as if the subject had reported the target to the right, again regardless of how the subject moved the slider bar relative to the center position. By choosing a range of threshold positions, spanning the full range of reported confidence values, a complete ROC curve was obtained. Note that as the threshold was moved closer to the left edge of the slider bar, the percentage of true and false positives both approached minimum values, since only trials with very high reported confidence could contribute to either the true positive or false positive rate (most trials were rejected as either true or false negatives). As the threshold position moved closer to the center of the confidence slider bar, the percentage of true positives increased. Finally, as the threshold was moved closer to the right edge of the slider bar, both the true positive rate and the percentage of false positives approached maximum values. The true positive rate averaged over all false positive rates, or the area under the ROC curve (AUC), was used as an overall measure of subject performance. The AUC is equivalent to the probability that a randomly chosen target image will be correctly classified relative to a randomly chosen distractor image, and thus directly predicts performance on the 2AFC task. Results for each SOA and for each value of K were averaged over 5 subjects. Error bars denote the standard deviation over the 5 subjects.

Model

Model cortical association fields were based on differences in the coactivation statistics of orientation-selective filter elements drawn from target and distractor images. Geisler and Perry measured co-occurrence statistics for oriented edges in human segmented natural images [25], and found a close correspondence to human judgments as to whether pairs of short line fragments were drawn from the same or different contours. Thus, we refer to the difference in coactivation statistics between target object and distractor images as Object-Distractor Difference (ODD) kernels. ODD kernels were trained using 40,000 target and 40,000 distractor images, each divided into 4 sets of 10,000 images each, with each set associated with a different value of $K=2,4,6,8$. The order in which the images were presented had no bearing on the final form of the ODD kernel; that is, there was no temporal component to the training. Training with more images did not substantively improve performance, although small differences were observed in the ODD kernels trained using a smaller number of images (10,000 target and 10,000 distractor images).

Each 256×256 pixel training image activated a regular array of 256×256 retinal elements whose outputs were either 0 or 1, depending on whether the corresponding image pixel was ON or OFF, respectively. Each retinal unit activated a local neighborhood of orientation-selective filters, which spanned 8 angles spaced uniformly between 0 and π . To mitigate aliasing effects, the orientation-selective filters were rotated by a small, fixed offset, equal to $\pi/16$, relative to the axis of the training images. All orientation-selective filters were 7×7 pixels in extent and consisted of a central excitatory subunit, represented by an elliptical Gaussian with a standard deviation of 7.0 in the longest direction and an aspect ratio of 7.0, flanked by two inhibitory subunits whose shapes were identical to the central excitatory subunit but were offset by ± 1.4 pixels in the direction orthogonal to the preferred axis.

The weight $W_\theta(x_1 - x_2, y_1 - y_2)$, from a retinal element at (x_2, y_2) to a filter element at (x_1, y_1) with dominant orientation θ , was given by a sum over excitatory and inhibitory subunits:

$$\begin{aligned} W_\theta(x_1 - x_2, y_1 - y_2) &= W_\theta(\mathbf{r}_1 - \mathbf{r}_2) = \\ &A \left\{ \exp\left[\frac{1}{2}(\mathbf{r}_1 - \mathbf{r}_2) \cdot \mathbf{R}_\theta^{-1} \sigma^{-1} \mathbf{R}_\theta \cdot (\mathbf{r}_1 - \mathbf{r}_2)^T\right] \right. \\ &- \exp\left[\frac{1}{2}(\mathbf{r}_1 + \mathbf{f} - \mathbf{r}_2) \cdot \mathbf{R}_\theta^{-1} \sigma^{-1} \mathbf{R}_\theta \cdot (\mathbf{r}_1 + \mathbf{f} - \mathbf{r}_2)^T\right] \\ &\left. - \exp\left[\frac{1}{2}(\mathbf{r}_1 - \mathbf{f} - \mathbf{r}_2) \cdot \mathbf{R}_\theta^{-1} \sigma^{-1} \mathbf{R}_\theta \cdot (\mathbf{r}_1 - \mathbf{f} - \mathbf{r}_2)^T\right] \right\}, \end{aligned} \quad (3)$$

where the position vector is given by $\mathbf{r}_i = [x_i, y_i]$ and the matrix $\sigma = \begin{bmatrix} 1 & 0 \\ 0 & 7 \end{bmatrix}$ describes the shape of the elliptical Gaussian subunits for $\theta=0$. In Eq. 3, \mathbf{R}_θ is a unitary rotation matrix,

$$\mathbf{R}_\theta = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix}, \quad (4)$$

and $\mathbf{f} = [0.0, 1.4]$ is a translation vector in the direction orthogonal to the dominant orientation when $\theta=0$. The amplitude A was determined empirically so that the total integrated strength of all excitatory connections made by each retinal unit equaled 20.0 (and thus the total strength of all inhibitory connections made by each retinal unit equaled -40.0). Mirror boundary conditions were used to mitigate edge effects. The retinal input to each orientation-selective filter element $\mathbf{s}(x_1, y_1, \theta)$ was then given by

$$\mathbf{s}(x_1, y_1, \theta) = \sum_{x_2, y_2} \mathbf{W}_\theta(x_1 - x_2, y_1 - y_2) \mathbf{I}_{(x_1, y_1)}(x_2, y_2), \quad (5)$$

where $\mathbf{I}_{(x_1, y_1)}$ is the 7×7 binary input image patch centered on (x_1, y_1) . The sum is over all pixels (x_2, y_2) that are part of this image patch. The initial output of each orientation-selective filter element $z_0(x_1, y_1, \theta)$ was obtained by comparing the sum of its excitatory and inhibitory retinal input to a fixed threshold of 0.5. Values below threshold were set to 0 whereas values above unity were set to 1.0. Thus

$$z(x_1, y_1, \theta) = g(\mathbf{s}(x_1, y_1, \theta)), \quad (6)$$

where the function,

$$g(s) = \begin{cases} 0 & s < 0.5 \\ s & 0.5 \leq s \leq 1.0, \\ 1 & s > 1.0 \end{cases}, \quad (7)$$

is an element-wise implementation of these thresholds. The responses of all suprathreshold orientation-selective filters contributed to the coactivation statistics, with only the relative distance, direction, and orientation of filter pairs recorded. Because of the threshold condition, only the most active orientation-selective filters contributed to the coactivation statistics.

For every suprathreshold filter element extracted from the i -th target image, coactivation statistics were accumulated relative to all surrounding suprathreshold filter elements extracted from the same image. Thus the ODD kernel G is given by

$$G_{\theta-\theta_0}^{ii}(\rho(x-x_0, y-y_0), \phi_{\theta_0}(x-x_0, y-y_0)) = \sum_{x, y} z_0(x, y, \theta), \quad (8)$$

where the radial distance ρ is a function of the (x, y) coordinates of

the two filter elements, the direction ϕ is the angle measured relative to θ_0 , the sum is over all suprathreshold elements within a cutoff radius of 32, the superscript t_i denotes the i -th target image, and the difference in the orientations of the two filter elements $\theta - \theta_0$ is taken modulo π . Because the amoeba/no-amoeba image set was translationally invariant and isotropic, the central filter element may without loss of generality be shifted and rotated to a canonical position and orientation, so that the dependence on ρ_0, ϕ_0, θ_0 may be omitted. The coactivation statistics for the i -th target image can then be written simply as $G_{\theta}^{t_i}(\rho, \phi)$, where (ρ, ϕ) gives the distance and direction from the origin to the filter element with orientation θ , given that the filter element at the origin has orientation $\pi/16$. An analogous expression gives the coactivation statistics for the j -th distractor image $G_{\theta}^{d_j}(\rho, \phi)$. The ODD kernel $G_{\theta}(\rho, \phi)$ is given by the difference

$$G_{\theta}(\rho, \phi) = W_+ \sum_i G_{\theta}^{t_i}(\rho, \phi) - W_- \sum_j G_{\theta}^{d_j}(\rho, \phi), \quad (9)$$

where the sums are taken over all target and distractor images and the normalization factors W_+ and W_- are determined empirically so as to yield a total ODD strength of 325 (see Figure 8 and Results), defined as the sum over all ODD kernel elements arising from either the target or distractor components. By construction, the sum over all ODD kernel elements equals zero, so that the average lateral support for randomly distributed edge fragments would be neutral. Our results did not depend critically on the RMS magnitude of the ODD kernel (see Figure 8). To minimize storage requirements individual connection strengths were stored as unsigned 8-bit integers, so that the results of the present study did not depend on computation of high precision kernels.

As described above, the canonical ODD kernel is defined relative to filter elements at the origin with orientation $\pi/16$. Filter elements located away from the origin can be accounted for by a trivial translation. To account for filter elements with different orientations, separate ODD kernels were computed for all 8 orientations then rotated to a common orientation and averaged to produce a canonical ODD kernel. The canonical kernel was then rotated in steps between 0 and π (offset by $\pi/16$) and then interpolated to Cartesian $x-y$ axes by rounding to the nearest integer coordinates. Although it has been demonstrated that global contour saliency is enhanced for orientations along the cardinal axes [58], this bias is by construction absent from this model.

References

- Velisavljević L, Elder JH (2009) Cue dynamics underlying rapid detection of animals in natural scenes. *J Vision* 9.
- Field DJ, Hayes A, Hess RF (1993) Contour integration by the human visual system: Evidence for a local "association field". *Vision Res* 33: 173–193.
- Loffler G (2008) Perception of contours and shapes: Low and intermediate stage mechanisms. *Vision Res* 48: 2106–2127.
- Hess R, Field D (1999) Integration of contours: new insights. *Trends Cogn Sci* 3: 480–486.
- Fitzpatrick D (2000) Seeing beyond the receptive field in primary visual cortex. *Curr Opin in Neurobiol* 10: 438–443.
- Serićs P, Lorenceau J, Frégnac Y (2003) The "silent" surround of v1 receptive fields: theory and experiments. *J Physiol Paris* 97: 453–474.
- Kovács I, Julesz B (1993) A closed curve is much more than an incomplete one: effect of closure in figure-ground segmentation. *Proc Natl Acad Sci USA* 90: 7495–7497.
- Pettet MW, McKee SP, Grzywacz NM (1998) Constraints on long range interactions mediating contour detection. *Vision Res* 38: 865–879.
- Polat U, Sagi D (1993) Lateral interactions between spatial channels: Suppression and facilitation revealed by lateral masking experiments. *Vision Res* 33: 993–999.
- Kapadia MK, Ito M, Gilbert CD, 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.
- Polat U, Terkin A, Yehzekel O (2008) Spatio-temporal low-level neural networks account for visual masking. *Adv Cogn Psych* 3: 153.
- Huang PC, Hess RF (2007) Collinear facilitation: Effect of additive and multiplicative external noise. *Vision Res* 47: 3108.
- Bringuier V, Chavane F, Glaeser L, Frégnac Y (1999) Horizontal Propagation of Visual Activity in the Synaptic Integration Field of Area 17 Neurons. *Science* 283: 695–699.
- Cavanaugh JR, Bair W, Movshon JA (2002) Nature and interaction of signals from the receptive field center and surround in macaque v1 neurons. *J Neurophys* 88: 2530–2546.
- Cavanaugh JR, Bair W, Movshon JA (2002) Selectivity and spatial distribution of signals from the receptive field surround in macaque v1 neurons. *J Neurophys* 88: 2547–2556.
- Pooresmaeili A, Herrero JL, Self MW, Roelfsema PR, Thiele A (2010) Suppressive Lateral Interactions at Parafoveal Representations in Primary Visual Cortex. *J Neurosci* 30: 12745–12758.
- Bosking WH, Zhang Y, Schofield B, Fitzpatrick D (1997) Orientation Selectivity and the Arrangement of Horizontal Connections in Tree Shrew Striate Cortex. *J Neurosci* 17: 2112–2127.
- Gilbert C, Wiesel T (1989) Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex. *J Neurosci* 9: 2432–2442.
- Malach R, Amir Y, Harel M, Grinvald A (1993) Relationship between intrinsic connections and functional architecture revealed by optical imaging and in vivo

ODD kernels were used to compute lateral support for each orientation-selective filter element, via linear convolution. The output of each filter element was then modulated in a multiplicative fashion by the computed lateral support. The procedure was iterated by calculating new values for the lateral support \mathfrak{s} , which were again used to modulate filter outputs in a multiplicative fashion:

$$s_k(x, y, \theta) = z_{k-1}(x, y, \theta) \sum_{x', y', \theta'} G_{\theta}(\rho(x-x', y-y'), \phi_{\theta}(x-x', y-y')) z_{k-1}(x', y', \theta'), \quad (10)$$

where the subscript k denotes the k -th iteration. The same kernel was used for all iterations. All source code used to train and apply cortical association fields is publicly available at

<http://sourceforge.net/projects/pevision/>.

To measure model performance, in each trial 1 target image and 1 distractor image were tested as a pair, so as to emulate the 2AFC format of the human experiments. The orientation-selective filter responses to both test images were evaluated after $k=0, 1, 2, 3, 4$ iterations of the ODD kernel. The total activation across all filter elements, $T = \sum_{x, y, \theta} z_k(x, y, \theta)$, was used to compare the two test images. Since the model cortical association fields tended to support contour fragments belonging to amoebas while inhibiting clutter fragments, the image with higher total activation T was assumed to be the target image. Error bars for the model performance (as shown in Figure 7) were estimated using the standard deviation of a binomial distribution with probability p equal to percent correct and N equal to the number of trials.

Acknowledgments

The authors wish to thank Steven Zucker for stimulating discussions that helped initiate this project.

Author Contributions

Conceived and designed the experiments: VG JSG IN LMAB GTK. Performed the experiments: MIH SB. Analyzed the data: VG MIH BK SB GTK. Contributed reagents/materials/analysis tools: SPB CR GTK. Wrote the paper: VG IN LMAB GTK. Conceived the research program: JSG IN LMAB GTK. Designed and implemented model cortical association fields: VG GTK.

- targeted biocytin injections in primate striate cortex. *Proc Natl Acad Sci USA* 90: 10469–10473.
20. Hess RF, Beaudot WHA, Mullen KT (2001) Dynamics of contour integration. *Vision Res* 41: 1023–1037.
 21. Keyser C, Xiao DK, Földiák P, Perrett DI (2001) The speed of sight. *J Cognitive Neurosci* 13: 90–101.
 22. Keyser C, Perrett DI (2002) Visual masking and rsvp reveal neural competition. *Trends Cogn Sci* 6: 120–125.
 23. Bacon-Macé N, Macé MJM, Fabre-Thorpe M, Thorpe SJ (2005) The time course of visual processing: Backward masking and natural scene categorisation. *Vision Res* 45: 1459–1469.
 24. Ben-Shahar O, Zucker S (2004) Geometrical computations explain projection patterns of long-range horizontal connections in visual cortex. *Neural Comput* 16: 445–476.
 25. Geisler WS, Perry JS (2009) Contour statistics in natural images: Grouping across occlusions. *Visual Neurosci* 26: 109–121.
 26. Mandon S, Kreiter AK (2005) Rapid contour integration in macaque monkeys. *Vision Res* 45: 291–300.
 27. Ursino M, Cara GEL (2004) A model of contextual interactions and contour detection in primary visual cortex. *Neural Networks* 17: 719.
 28. Sterkin A, Sterkin A, Polat U (2008) Response similarity as a basis for perceptual binding. *J Vis* 8: 1.
 29. Bair W, Cavanaugh JR, Movshon JA (2003) Time course and time-distance relationships for surround suppression in macaque v1 neurons. *J Neurosci* 23: 7690.
 30. Zhang NR, von der Heydt R (2010) Analysis of the Context Integration Mechanisms Underlying Figure-Ground Organization in the Visual Cortex. *J Neurosci* 30: 6482–6496.
 31. Schwabe L, Obermayer K, Angelucci A, Bressloff PC (2006) The Role of Feedback in Shaping the Extra-Classical Receptive Field of Cortical Neurons: A Recurrent Network Model. *J Neurosci* 26: 9117–9129.
 32. Angelucci A, Levitt JB, Walton EJS, Hupe JM, Bullier J, et al. (2002) Circuits for Local and Global Signal Integration in Primary Visual Cortex. *J Neurosci* 22: 8633–8646.
 33. Serre T, Oliva A, Poggio T (2007) A feedforward architecture accounts for rapid categorization. *Proc Natl Acad Sci USA* 104: 6424.
 34. Martínez-Conde S, Macknik SL, Troncoso XG, Hubel DH (2009) Microsaccades: a neurophysiological analysis. *Trends Neurosci* 32: 463–475.
 35. Rolls ET, Tovee MJ (1994) Processing Speed in the Cerebral Cortex and the Neurophysiology of Visual Masking. *P Roy Soc Lond B Bio* 257: 9–15.
 36. Wilkinson F, Wilson HR, Habak C (1998) Detection and recognition of radial frequency patterns. *Vision Res* 38: 3555–3568.
 37. Geisler WS, Perry JS, Super BJ, Gallogly DP (2001) Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res* 41: 711–724.
 38. Schneidman E, Berry II MJ, Segev R, Bialek W (2006) Weak pairwise correlations imply strongly correlated network states in neural population. *Nature* 440: 1007–1012.
 39. Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, et al. (2006) The Structure of Multi-Neuron Firing Patterns in Primate Retina. *J Neurosci* 26: 8254–8266.
 40. Jones JP, Palmer LA (1987) An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J Neurophys* 58: 1233–1258.
 41. Troyer TW, Krukowski AE, Priebe NJ, Miller KD (1998) Contrast-invariant orientation tuning in visual cortex: feedforward tuning and correlation-based intracortical connectivity. *J Neurosci* 18: 5908–5927.
 42. Angelucci A, Levitt JB, Walton EJS, Hupe JM, Bullier J, et al. (2002) Circuits for local and global signal integration in primary visual cortex. *J Neurosci* 22: 8633–8646.
 43. Azzopardi P, Cowey A (1997) Is blindsight like normal, near-threshold vision? *Proc Natl Acad Sci USA* 94: 14190–14194.
 44. Macmillan NA, Creelman CD (1991) *Detection theory: a user's guide*. Cambridge: CUP Archive.
 45. Maunsell JHR, Gibson JR (1992) Visual response latencies in striate cortex of the macaque monkey. *J Neurophys* 68: 1332.
 46. Bell J, Badcock DR, Wilson H, Wilkinson F (2007) Detection of shape in radial frequency contours: Independence of local and global form information. *Vision Res* 47: 1518–1522.
 47. Li Z (2001) Computational design and nonlinear dynamics of a recurrent network model of the primary visual cortex. *Neural Comput* 13: 1749–1780.
 48. Li Z (1998) A neural model of contour integration in the primary visual cortex. *Neural Comput* 10: 903–940.
 49. Mundhenk TN, Itti L (2005) Computational modeling and exploration of contour integration for visual saliency. *Biol Cybern* 93: 188.
 50. Li W, Piéch V, Gilbert CD (2006) Contour saliency in primary visual cortex. *Neuron* 50: 951.
 51. Grossberg S, Mingolla E (1985) Neural dynamics of perceptual grouping: textures, boundaries, and emergent segmentations. *Percept Psychophys* 38: 141.
 52. Ullman S, Gregory RL, Atkinson J (1992) Low-Level Aspects of Segmentation and Recognition [and Discussion]. *Philos T R Soc Lon B* 337: 371–379.
 53. Yen SC, Finkel LH (1998) Extraction of perceptually salient contours by striate cortical networks. *Vision Res* 38: 719–741.
 54. Garrigues PJ, Olshausen BA (2007) Learning horizontal connections in a sparse coding model of natural images. In: *Adv Neur In*.
 55. Ing AD, Wilson AJ, Geisler WS (2010) Region grouping in natural foliage images: Image statistics and human performance. *J Vision* 10: 1–19.
 56. Yao H, Shi L, Han F, Gao H, Dan Y (2007) Rapid learning in cortical coding of visual scenes. *Nat Neurosci* 10: 772–778.
 57. Hua T, Bao P, Huang CB, Wang Z, Xu J, et al. (2010) Perceptual learning improves contrast sensitivity of V1 neurons in cats. *Curr Biol* 20: 887–894.
 58. Li W, Gilbert CD (2002) Global contour saliency and local colinear interactions. *J Neurophysiol* 88: 28462856.
 59. Knoblauch A, Sommer FT (2004) Spike-timing-dependent synaptic plasticity can form “zero lag links” for cortical oscillations. *Neurocomputing* 58-60: 185.
 60. Hoyer PO, Hyvärinen A (2002) A multi-layer sparse coding network learns contour coding from natural images. *Vision Res* 42: 1593–1605.
 61. Song S, Miller KE, Abbott LF (2000) Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci* 3: 919.
 62. Li Z (2002) A saliency map in primary visual cortex. *Trends Cogn Sci* 6: 9–16.
 63. Fukushima K (1980) Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol Cybern* 36: 193.
 64. Gilbert CD, Sigman M (2007) Brain states: Top-down influences in sensory processing. *Neuron* 54: 667.
 65. Schneeweis D, Schnapf J (1995) Photovoltage of rods and cones in the macaque retina. *Science* 268: 1053–1056.
 66. Enns JT, Lollo VD (2000) What's new in visual masking? *Trends Cogn Sci* 4: 345–352.
 67. Fei-Fei L, Fergus R, Perona P (2004) Learning generative visual models from few training examples: an incremental bayesian approach tested on 101 object categories. In: *CVPR 2004, Workshop on Generative-Model Based Vision*.
 68. LeCun Y, Bottou L, Bengio Y, Haffner P (1998) Gradient-based learning applied to document recognition. In: *P IEEE* volume 86: 2278.
 69. Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10: 433.