

One “shape” fits all: The orientation bandwidth of contour integration

Bruce C. Hansen

Department of Psychology and Neuroscience Program,
Colgate University, Hamilton, NY, USA



Keith A. May

Centre for Applied Vision Research,
City University London, London, UK



Robert F. Hess

McGill Vision Research Unit,
Department of Ophthalmology,
McGill University, Montreal, QC, Canada



The ability of human participants to integrate fragmented stimulus elements into perceived coherent contours (amidst a field of distracter elements) has been intensively studied across a large number of contour element parameters, ranging from luminance contrast and chromaticity to motion and stereo. The evidence suggests that contour integration performance depends on the low-level Fourier properties of the stimuli. Thus, to understand contour integration, it would be advantageous to understand the properties of the low-level filters that the visual system uses to process contour stimuli. We addressed this issue by examining the role of stimulus element orientation bandwidth in contour integration, a previously unexplored area. We carried out three psychophysical experiments, and then simulated all of the experiments using a recently developed two-stage filter-overlap model whereby the contour grouping occurs by virtue of the overlap between the filter responses to different elements. The first stage of the model responds to the elements, while the second stage integrates the responses along the contour. We found that the first stage had to be fairly broadly tuned for orientation to account for our results. The model showed a very good fit to a large data set with relatively few free parameters, suggesting that this class of model may have an important role to play in helping us to better understand the mechanisms of contour integration.

Introduction

The early stages of vision involve parallel channels comprising neurons that can act like nonlinear filters.

Such filters are believed to decompose the retinal image into a set of images that each contain information specific to a particular band of spatial frequencies, orientations, and phases, to name a few germane to the current study (Carandini et al., 2005; De Valois, Albrecht, & Thorell, 1982; De Valois, Yund, & Hepler, 1982; Field & Tolhurst, 1986; Maffei & Fiorentini, 1973; Merigan & Maunsell, 1993; Ringach, Shapley, & Hawken, 2002; Shapley & Lennie, 1985). However, such an account on its own is an overly simplified description of the early visual representation—that is, it provides little informative value regarding how such a multidimensional representation is integrated to form coherent percepts. For instance, one classic problem facing the piecewise multidimensional visual representation is how such a representation can be used to solve the problem of establishing perceived continuity of background structure (e.g., a continuous edge or line following a contour defined by the shape of any given object) when occluded by content in the foreground. Thus, it seems crucial that theories pertaining to the early visual analysis of complex scenery incorporate an early multidimensional representation that emphasizes the integration of correlated information across local regions of the visual field (unencumbered by gaps or occluding foreground structure).

In an attempt to achieve the above, the ability of human participants to integrate separate, oriented, stimulus elements into perceived coherent contours (amidst a field of distracter elements) has been intensively studied across a large number of contour element parameters, with such parameters ranging from luminance contrast and chromaticity to motion and stereo (Field & Hayes, 2003; Hess, Hayes, & Field,

Citation: Hansen, B. C., May, K. A., & Hess, R. F. (2014). One “shape” fits all: The orientation bandwidth of contour integration. *Journal of Vision*, 14(13):17, 1–21, <http://www.journalofvision.org/content/14/13/17>, doi:10.1167/14.13.17.

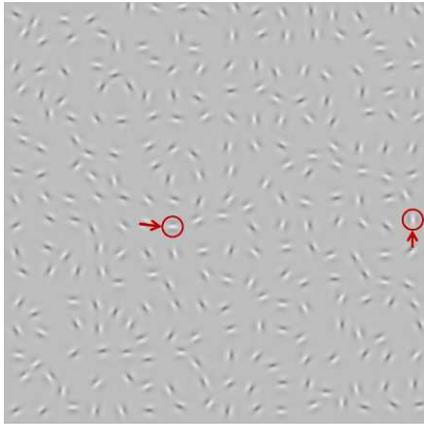


Figure 1. Example of a typical stimulus used in contour integration experiments. The stimulus is composed of a field of distractor Gabor elements, within which is embedded 12 Gabor elements that are aligned to a predefined path, thereby defining the contour. Here, the start and end points of the target contour have been circled to facilitate viewing (i.e., those elements were not highlighted during actual experiments).

2003). The stimuli are typically arrays of Gabor micropatches, each consisting of a sinusoidal function of luminance (i.e., light and dark bars of luminance) windowed with a Gaussian function to constrain the sinusoidal luminance modulation to a particular region in image space (see Figure 1 for an example of such a stimulus).

The evidence suggests that the contour integration process depends primarily on the low-level Fourier properties of the individual contour elements, rather than on their perceived appearance (Dakin & Hess, 1999). Since selectivity to the Fourier properties of the stimuli is mediated by the low-level filters in the visual system, Dakin and Hess' (1999) finding suggests that to understand contour integration, it is important to understand the properties the low-level visual filters used by the contour integration mechanisms. One way of characterizing the visual filters is to measure their spatial-frequency bandwidth. Dakin and Hess (1998) found that contour integration performance was fairly narrowly tuned to the spatial frequencies of the contour elements. In their study, half the stimulus elements were assigned one spatial frequency and the other half were assigned a different spatial frequency. Dakin and Hess (1998) fit Gaussian functions to the plots of performance against spatial frequency difference, and found bandwidths (full-width at half-height) of around 1.5 octaves for highly curved contours, and around 2–3 octaves for straight contours. Another way to characterize the filters is to measure their orientation bandwidth, and that is the subject of the current study. Cortical neurons have been reported to possess orientation tuning bandwidths (full-width at half-height) ranging from 6° to 360° (median $\sim 40^\circ$) in

primate cortex (e.g., De Valois, Yund, & Hepler, 1982; Ringach et al., 2002), but we currently have little knowledge of the orientation bandwidths of the neurons that are used in contour integration.

To address this issue, we measured contour integration performance as a function of various parameters related to the orientation bandwidth of the contour elements. While it is usual to use narrowband spatial frequency elements (e.g., Gabors) in tasks designed to measure contour integration (Field, Hayes, & Hess, 1993), any manipulation of the orientation bandwidth of such stimulus elements results in concurrent changes in element coverage, and hence their second-order content (i.e., contrast envelope size). To avoid these related changes, we used elements composed of filtered noise (filtered to possess a narrow bandwidth of spatial frequency, and variable orientation bandwidth) so that we could manipulate the orientation bandwidth of the elements without any consequent changes to the global layout of the stimulus (i.e., maintain a constant envelope). Thus, the elements themselves do not possess phase alignment along any particular angle, but instead possess globally distributed (within the elements themselves) contrast along the range of orientations passed by the filter. Since contour integration has been shown to be minimally influenced by the phase of the contour elements (Dakin & Hess, 1999; Field, Hayes, & Hess, 2000; Hansen & Hess, 2006), we expected participants to have little difficulty integrating contours made up of elements as described above.

Lastly, there is a critical distinction between the orientation tuning of a psychophysical effect and the orientation tuning of the underlying contour integration mechanism. The contour integration process is exceedingly complicated, such that is not possible to infer the orientation bandwidth of the contour integration mechanisms directly from the psychophysical effects. We therefore sought to provide a quantitative model of the underlying mechanisms. Given that the current study is motivated to assess the effect of orientation bandwidth on contour integration, it is crucial to employ a functional model that is selective for such low-level Fourier properties, and largely unconstrained regarding the linking strength. The latter point is important, as any preconceived assumptions regarding contour element linking rules could seriously influence estimates of element orientation bandwidths leading to successful integration. We therefore simulated all experiments in the current study using the two-stage filter-overlap contour integration model introduced by May and Hess (2008). Briefly, the key feature of the May and Hess (2008) model is that it represents the stimulus in a space that has a dimension representing orientation, in addition to the two dimensions of the image. The spread of the filter's response across image space and orientation has the

effect of blurring the representation along all three dimensions so that the responses due to elements that are close in space and orientation join up to form a contour. The third dimension of the space allows contours to overlap in space while forming separate contours if they are sufficiently separated along the orientation dimension at the point of overlap (May & Hess, 2008, figure 9 and movie 1). All previous attempts to model contour integration experiments quantitatively have used models based on Field and colleagues' (1993) notion of an association field, which explicitly links stimulus elements according to a set of rules. May and Hess' (2008) model works differently in that there is no explicit linking between stimulus elements, and the linking occurs simply by virtue of the overlap between the model filter responses to different stimulus elements, with no assumptions regarding the strength of linkage between differently oriented contour elements. The critical role played by the filters in the model suggested to us that the parameters of the filters would be strongly constrained by our data on the effect of orientation bandwidth of the stimulus elements. The two-stage filter-overlap model is therefore ideally suited to provide a quantitative depiction of the orientation tuning bandwidth of the front-end of the contour integration mechanism. So far, the performance of this kind of model has been demonstrated by examining its output to a handful of example stimuli (May & Hess, 2008; Rosenholtz, Twarog, Schinkel-Bielefeld, & Watenberg, 2009). Thus, in addition to employing this type of model to estimate the front-end orientation tuning bandwidth of contour integration mechanisms, the current study also provides the first thorough quantitative comparison of the performance of this kind of model with that of human participants.

The current study consists of two parts. Part I was an empirical study (Experiments 1 through 3) of contour detection while manipulating orientation bandwidth of the stimulus elements. Experiment 1 was designed to allow for an indirect observation regarding the tolerance of the contour integration mechanism to increases in the orientation bandwidth of the stimulus elements. Specifically, we were interested in estimating the element orientation bandwidth of contour integration. We systematically varied the stimulus element orientation bandwidth for different levels of contour curvature and several different degrees of contour element-to-contour path orientation alignment, and measured contour detection performance for each combination of stimulus parameters. Experiments 2 and 3 served as follow-ups to Experiment 1 (i.e., same task, but with different stimulus manipulations) and consisted of two different paradigms designed to provide further data that would allow us to draw meaningful conclusions regarding the orientation bandwidth of contour integration. In Experiment 2 we

held stimulus element orientation bandwidth at a fixed amount (25° full-width at half-height) and systematically removed increasing amounts of orientation contrast from the center of the orientation passband (i.e., notch filtered along the central element orientation). Experiment 3 employed contour elements that alternated between narrow and broad orientation bandwidths or alternated between different notch bandwidths. Part II of the current study involved simulating Experiments 1 through 3 with the May and Hess (2008) two-stage filter-overlap contour integration model.

Methods and materials

Apparatus

All stimuli were presented with an Intel Pentium IV (2.84 GHz) processor equipped with 1GB RAM via a VSG 2/5 graphics card (Cambridge Research Systems, Rochester, UK) on a Sony G520 CRT monitor with a maximum luminance of 80 cd/m^2 , frame rate of 120 Hz, and resolution of 1024×768 pixels. The stimuli were generated using C routines called from MATLAB (The MathWorks, Inc., Natick, MA) version 7.0 and were linearly scaled to fit the range 0–255 and stored in an 8-bit frame store on the VSG card and subsequently scaled to the correct contrast and gamma corrected by mapping the 8-bit values onto 15-bit values. An analogue input to the monitor was generated from these 15-bit values using two 8-bit digital-to-analogue converters in the VSG card. Participants viewed the display binocularly from a distance of 60 cm.

Participants

Two of the authors, BCH and RFH (both experienced psychophysical observers), participated in all experiments.¹ Both participants had normal or corrected-to-normal vision. All experiments conformed to the ethical standards of the Federal Code of Regulations Title 45 (Public Welfare) and Department of Health and Human Services, Part 46 (Protection of Human Subjects). Institutional Review Board-approved (McGill University) informed written consent was obtained.

Experiment 1: Stimulus generation

The individual Gabor micropatches in the stimuli were filtered white noise to which a spatial circular Gaussian envelope was applied. Filtering took place in

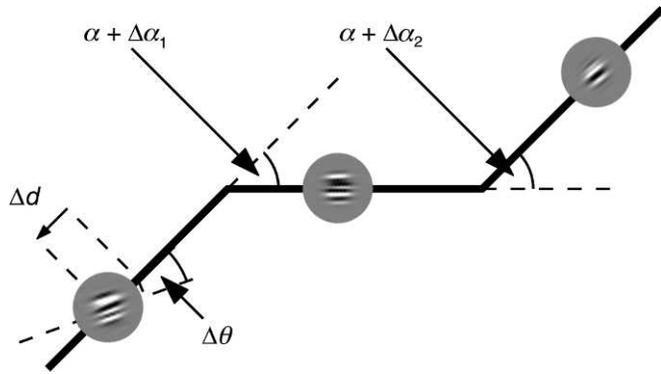


Figure 2. Schematic for building contours for all experiments in the current study. See text for further details.

the Fourier domain. The spatial frequency filter took the form of an annulus in Fourier space, with a Gaussian radial cross-section, which peaked at a spatial frequency of $f_0 = 4.5$ cycles per degree (cpd), and had a standard deviation of $\sigma_f = 2.25$ cpd. Thus, if points in Fourier space are defined in Polar coordinates (f, θ) , where f is spatial frequency and θ is angle (in degrees) from the horizontal axis in Fourier space, then the spatial frequency filter had the form

$$w_f(f, \theta) = e^{\left[\frac{-(f-f_0)^2}{2\sigma_f^2} \right]} \quad (1)$$

The orientation filter took the form of a pair of wedges in Fourier space, with amplitude given by a Gaussian function of orientation:

$$w_\theta(f, \theta) = e^{\left[\frac{-(\theta-\theta_0)^2}{2\sigma_\theta^2} \right]} + e^{\left[\frac{-(\theta-\theta_0+180)^2}{2\sigma_\theta^2} \right]} \quad (2)$$

where θ_0 is the orientation of the element from vertical (in stimulus space), and σ_θ , which we call the *orientation bandwidth SD*, controls the element's orientation bandwidth. The filter described in Equation 2 had to contain two wedges to maintain a complex-conjugate relationship between points on opposite sides of Fourier space, so that the image obtained after inverse-Fourier transforming did not contain any imaginary components.

White noise was generated and filtered with the orientation and frequency filters described in Equations 1 and 2 in Fourier space. If $n(f, \theta)$ is the noise value at position (f, θ) in Fourier space, then the filtered noise was given by $n(f, \theta) \times w_f(f, \theta) \times w_\theta(f, \theta)$. The combined orientation and spatial frequency filter has the property that the spatial frequency bandwidth is constant across orientation, and the orientation bandwidth is constant across spatial frequency. The direct current (DC) component was then set to zero, and the filtered noise was inverse-Fourier transformed, and windowed with a circular Gaussian spatial envelope

with standard deviation 0.17° visual angle. Finally, the element contrast was scaled to give a root-mean-square (RMS) contrast of 0.12, and was set to possess a mean pixel luminance that matched the background.

Stimuli were generated using a similar procedure to that of Hansen and Hess (2006), which was based on Field et al.'s (1993) original derivation. Briefly, the stimulus area was divided into an invisible 16×16 square grid. For stimuli containing no contour, each grid square was filled with one randomly oriented filtered noise element, placed at a random location within the square, subject to the constraint that the elements did not overlap. For stimuli containing a contour, the contour was first positioned randomly within the grid and then each remaining empty grid square was filled with one element with random orientation and random position within the grid square.² Part of an example contour is represented schematically in Figure 2. The contour was constructed along an invisible backbone of nine line segments, joined end to end. A contour element was placed at the center of each segment. The absolute difference in orientation between adjacent segments is referred to as the *path angle*, α (which controls contour curvature), and was varied systematically in our experiments. In order to reduce any chance alignments of three or more elements along the contour paths, we added $\pm 10^\circ$ of random (uniform distribution) path angle jitter, $\Delta\alpha$, to all contour elements in all experiments of the current study (e.g., Hansen & Hess, 2006). Contour elements were further randomly jittered by ± 5 pixels along the contour path (i.e., Δd). Finally, the *element angle* (i.e., the difference in orientation between a given path segment and the central orientation of a given element) was varied by adding a value, $\Delta\theta$, sampled from a zero-mean Gaussian distribution whose standard deviation (the angle *SD*) was varied systematically in the experiment. Experiment 1 was designed to test the effects of varying the parameters of (a) orientation bandwidth *SD* of the stimulus field elements, (b) element angle *SD*, and (c) the path angle on the ability of humans to detect contours made up of orientation-filtered random noise elements. All three parameters were varied in a factorial design. In each trial, the stimulus element orientation bandwidth, σ_θ , was 2° , 10° , 20° , 30° , or 40° . For stimuli containing contours, the angle *SD* was 0° , 10° , 20° , 30° , or 40° , and the path angle was 0° , 10° , 20° , or 30° . Refer to Figure 3 for stimulus examples.

Experiment 2: Stimulus generation

Individual stimulus elements were constructed with the same protocol as described above for Experiment 1, except that a notch in the orientation filter was

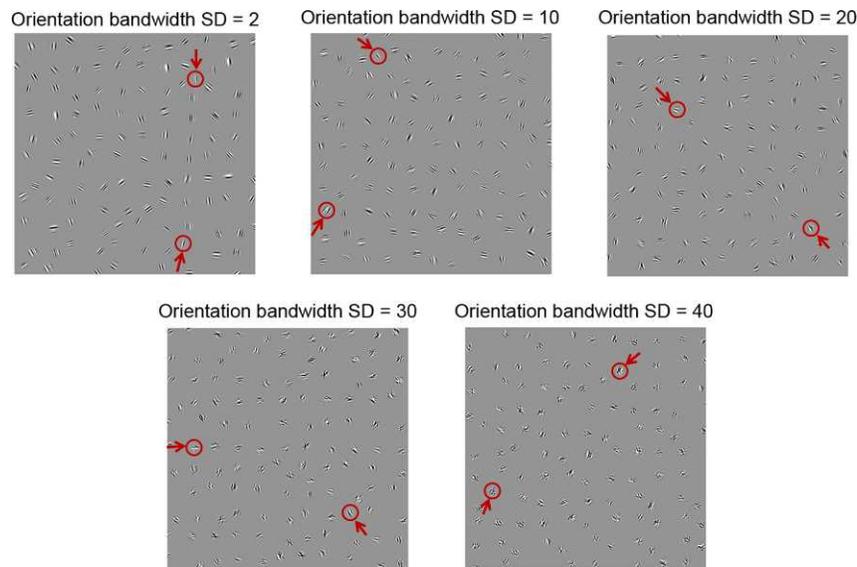


Figure 3. Example stimuli from Experiment 1. Stimulus contours have been highlighted to facilitate viewing (i.e., contours were not highlighted during the actual experiment).

introduced by subtracting from w_θ in Equation 2 a narrower wedge shape, w_{notch} , given by

$$w_{notch}(f, \theta) = e^{\left[\frac{-(\theta - \theta_0)^2}{(2\sigma_{notch}^2)} \right]} + e^{\left[\frac{-(\theta - \theta_0 + 180)^2}{(2\sigma_{notch}^2)} \right]} \quad (3)$$

where σ_{notch} , which we call the notch SD, controls the width of the notch. The orientation bandwidth SD, σ_θ , was set to 25° for Experiment 2, and σ_{notch} took values of 2°, 6°, or 12°. As for Experiment 1 stimuli, the angle SD could take values of 0°, 10°, 20°, 30°, or 40°, and the path angle took values of 0°, 10°, 20°, or 30°. See Figure 4 for stimulus examples. All three variables were varied in a factorial design.

Experiment 3: Stimulus generation

Experiment 3 was designed to test the effects of alternating either the orientation bandwidth or notch

width for successive elements along the contour. For stimuli containing contours, the elements that made up the contours were made to possess path angles of 0°, 10°, 20°, or 30°. Individual stimulus elements were constructed using either the same protocol described for Experiment 1 with orientation bandwidths of either 2°, 30°, or 40° (full-width at half-height) or Experiment 2 with orientation notch bandwidths of either 6° or 12° (the notches were taken out of Gabor filters with a bandwidth of 25°). Stimulus element fields and contours were constructed in an identical fashion as described for Experiments 1 and 2, but with the following modifications. For the alternating orientation bandwidth condition, the stimulus element fields were constructed with half of the elements consisting of orientation filtered bandwidths of 2° and the other half 30°, or half of the elements consisting of 2° and the other half 40°. The contours themselves were made up of elements that alternated along the contour between either 2° and 30° or 2° and 40°,

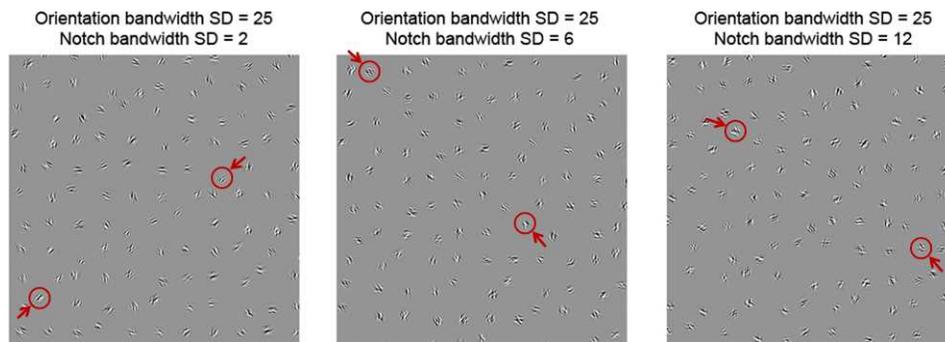


Figure 4. Example stimuli from Experiment 2. Stimulus contours have been highlighted to facilitate viewing (i.e., contours were not highlighted during the actual experiment).

respectively. For the alternating orientation notch filtered element condition, the stimulus element fields were constructed with half of the elements consisting of orientation notch filtered bandwidths of 0° (i.e., no notch) and the other half 6° , or half of the elements consisting of 0° and the other half 12° . The contours themselves were made up of elements that alternated between either 0° and 6° or 0° and 12° , respectively. For Experiment 3, all contour elements were aligned with the contour path (i.e., element $SD = 0^\circ$) since it is with that particular parameter setting that performance is at its highest (therefore allowing a full range of performance modulation).

Psychophysical procedure

The psychophysical procedure employed for all experiments in Part I of the current study consisted of a two-interval forced choice (2IFC) method of constant stimuli paradigm. The stimulus intervals in the 2IFC paradigm contained either a field of randomly oriented filtered noise elements or a field of randomly oriented filtered noise elements containing a contour. Each trial began by presenting the participant with a fixation point (500 ms) followed by Stimulus Interval 1 (1000 ms) followed by the fixation point (200 ms), followed by Stimulus Interval 2 (1000 ms), after which the display was set to mean luminance. At this point the participant indicated which interval contained the contour (the participant had unlimited time to respond).

For Experiment 1, the three independent variables (stimulus element orientation bandwidth, angle SD , and path angle) were blocked by noise element orientation bandwidth and angle SD ; within each block the path angle of the contour varied between one of the four values mentioned above. Each block had 20 trials per path angle, resulting in 80 trials per block. Each block was repeated five times, resulting in 100 trials per level of path angle for each noise element orientation bandwidth and element angle (total number of trials for Experiment 1 was 10,000).

For Experiment 2, the three independent variables (notch SD , angle SD , and path angle) were blocked by notch SD and element angle SD ; within each block, path angle varied between one of the four curvatures mentioned above. Each block had 20 trials for each level of path angle, resulting in 80 trials per block. Each block was repeated five times, resulting in 100 trials per level of curvature for each noise element orientation notch bandwidth and angle SD (total number of trials for Experiment 2 was 6,000).

Lastly, for Experiment 3, the three independent variables (alternating noise element orientation bandwidth, alternating notch bandwidth, and path angle)

were blocked by alternating noise element orientation bandwidth or alternating notch bandwidth. Within each block the path angle of the contour varied between one of the four curvatures mentioned above. For each level of curvature, one of 20 different contours was embedded in the stimulus field, resulting in 80 trials per block. Each block was repeated five times resulting in 100 trials per level of curvature for each alternating bandwidth or notch bandwidth conditions (total number of trials for Experiment 3 was 1,600).

Thus, each participant completed 17,600 trials (~16.5 hrs of psychophysics each) across all three experiments.

Two-stage filter-overlap contour integration model

To model the psychophysical results from Experiments 1 through 3, we employed May and Hess' (2008) two-stage filter-overlap model for contour integration which operates by filtering a given stimulus image with a set of oriented filters. The resulting filtered images from the different orientation channels are then stacked on top of each other to create a three-dimensional (3-D) space, where the dimensions are filter orientation and the two spatial dimensions of the image. The filter output values in this 3-D representation are thresholded to 1 or 0, to create blobs or zero-bounded response distributions (ZBRs) that can extend across space and orientation. Each ZBR is taken to be a contour. The contours can be rated on some property that indicates strength or goodness of the contour, and the stimulus with the higher rating on a particular trial is selected as the target by the model. In our current simulations, each contour rating was given by the longest distance in the image plane between any two pixels in the ZBR. The stimulus containing the highest-rated contour was selected by the model as being the target.

May and Hess' (2008) model is entirely feed forward. It contains orientation channels, spaced at 10° intervals from 0° to 170° (each filter kernel was even-symmetric, with a complete cycle every 180° of rotation). Each orientation channel contains a filter-rectify-filter (FRF) mechanism. A first-stage filter sensitive to the elements is followed by a squaring operation, followed by a second-stage filter, oriented identically to the first-stage filter, with scale (i.e., size) large enough to bridge the gap between successive contour elements. The first-stage filter gives a strong positive or negative response to individual elements of the same orientation as the filter; the squaring operation makes the responses all positive; and the second-stage filter blurs this positive signal along the length of the contour, giving rise to a

Parameter	Value (degrees visual angle)	Rationale
$\sigma_{u,1}$	0.0629	Fit to the data
$\sigma_{u,1}/\lambda_1$	0.2906	$\sigma_u/\lambda = 0.2906$ is the median value of this ratio from Jones and Palmer's (1987) physiological data set.
$\sigma_{v,1}/\sigma_{u,1}$	1.0714	Aspect ratio $\sigma_v/\sigma_u = 1.0714$ is the smallest value from Jones and Palmer's (1987) data set, giving the largest physiologically plausible orientation bandwidth for the first-stage filter (full-width at half-height was 70.5°). The need for a wide orientation bandwidth for the first-stage filter was indicated by the relatively mild effect of the notch in Experiment 2.
$\sigma_{u,2}$	0.5032	Fit to the data
$\sigma_{u,2}/\lambda_2$	0.33964	$\sigma_u/\lambda = 0.33964$ is the value of this ratio for Jones and Palmer's (1987) cell with the highest ratio of σ_v/λ . May and Hess (2008) argued that a high value of σ_v/λ would benefit contour integration.
$\sigma_{v,2}/\sigma_{u,2}$	2.407	Aspect ratio $\sigma_v/\sigma_u = 2.407$ is the aspect ratio of Jones and Palmer's (1987) cell with the highest ratio of σ_v/λ , as described above. The orientation bandwidth (full-width at half-height) was 26.5° .

Table 1. The filter parameters used in the simulation of the experiments shown in Figures 10 through 13. *Notes:* Two were fit to the data. The others were chosen intuitively on the basis of the requirements of the contour integration task, but constrained to be physiologically plausible. While some of the filter parameters fall at the extremes of the physiologically plausible range, they are nevertheless biologically plausible; the brain is likely to use the neurons that are most useful for the task at hand, even if these fall at the extremes of the available range of parameters.

continuous ZBR that delineates the contour. The rationale for the FRF mechanism was May and Hess' (2008) finding that contour integration performance was relatively unperturbed by increases in separation between the elements, even for elements with high-frequency carriers. May and Hess showed that this finding could not be accounted for by a single filter because a physiologically plausible filter kernel that was large enough to bridge the gap between widely spaced elements would not be sensitive to high-frequency elements (a kernel sensitive to the high-frequency elements would have needed an implausibly high number of carrier cycles to bridge the gap between the elements). In the FRF mechanism, the first-stage filter scale can be tuned to the frequency of the stimulus elements, and the second-stage filter scale can be tuned to a size appropriate for the gaps between the elements, allowing integration of elements with any carrier frequency across any gap size. We do not currently know how the visual system selects the filter scales but, since all the stimuli in every experiment in the current study used the same element spatial frequency distribution and interelement spacing, we could just treat the filter scales as free parameters, which were kept constant across all conditions.

The model has seven parameters: three parameters describe the first-stage Gabor filter kernel, three parameters describe the second-stage Gabor filter kernel, and the final parameter is the threshold on the final filter output that creates the ZBRs. The three Gabor filter parameters are as follows: λ , the carrier wavelength; σ_u , the envelope standard deviation in a direction perpendicular to the bars of the

carrier (i.e., the receptive field width); σ_v , the envelope standard deviation in a direction parallel to the bars of the carrier (i.e., the receptive field length). These three parameters (λ , σ_u , σ_v) can alternatively be presented as (σ_u , σ_u/λ , σ_v/σ_u). The σ_u parameter describes the overall scale (i.e., size) of the filter kernel, while the σ_u/λ and σ_v/σ_u parameters describe the *shape* of the kernel. To constrain the fit, we set the two shape parameters, σ_u/λ and σ_v/σ_u , at each filter stage to sensible values based on physiological findings and the demands of the contour integration task, and fit only the scale parameter, σ_u , of each stage, and the threshold. The filter parameter values that we used, along with the rationale, are given in Table 1, with the extra subscript "1" or "2" indicating first- or second-stage filter kernel. The fitting process itself was a manual procedure, guided by intuition. Simulating the experiments with just one set of parameter values was immensely time-consuming, requiring many computers running in parallel for many months, so it was not possible to fit the parameters using standard fitting methods. Instead, we simulated some selected conditions with a few sets of Gabor filter parameters, and then, having found the set of filter parameters that fit best to the selected conditions, we proceeded to simulate every condition of every experiment with these filter parameters. The only parameter that was fit in the conventional sense was the threshold since, having obtained the filter outputs, we could easily set the threshold at many different levels and map out the contours for each threshold level. The threshold is expressed in standard deviations above the mean

second-stage filter response across all orientation channels.

Results

Part I

The data for all experiments take the form of proportion correct on each condition. The confidence intervals displayed on all the graphs are 95% Bayes credible intervals, which show the region in which there is a 95% chance that the true probability of a correct response lies, assuming a flat prior. It can be shown that, if we assume a flat prior on the probability of a correct response, then the probability density function (PDF) of the true probability of a correct response is the beta distribution, $\text{Beta}(X + 1, N - X + 1)$, where N is the total number of trials on that condition, and X is the observed number of correct trials (Nicholson, 1985). The lower confidence limit is the 0.025 quantile of this distribution, and the upper limit is the 0.975 quantile:

$$\text{Lower Limit} = B^{-1}(0.025; X + 1, N - X + 1),$$

$$\text{Upper Limit} = B^{-1}(0.975; X + 1, N - X + 1),$$

where B^{-1} is the inverse cumulative distribution function of the beta distribution. These can be calculated in MATLAB using the *betainv* function in the Statistics Toolbox.

In order to assess whether or not the two participants performed similarly, we calculated the absolute difference between the two subjects' proportion of correct scores for each condition of each experiment. The mean absolute difference (MAD) across all conditions of all experiments was only 0.0585, which is a small proportion of the total range for two-alternative forced choice (0.5 to 1). The standard error of the mean was 0.00345. The MADs for the individual experiments are as follows (standard errors in parentheses): Experiment 1: 0.0624 (0.00431); Experiment 2: 0.0593 (0.00661); Experiment 3, alternating bandwidth: 0.0175 (0.00675); Experiment 3, alternating notch width: 0.0425 (0.0126). So for all experiments, the two subjects performed similarly.

Experiment 1

Experiment 1 examined the effect of varying the orientation bandwidth of the stimulus elements on contour integration performance. Dakin and Hess (1998) found that contour integration performance was fairly narrowly tuned to the spatial frequencies of the

contour elements. Here, we investigated whether a similar trend can be observed with respect to contour element orientation bandwidth. Specifically, we were interested in measuring proportion correct as a function of element orientation bandwidth, and from that function, calculate an estimate of orientation tuning for contour integration. We also varied two other stimulus properties: (a) the amount of curvature (i.e., *path angle*) exhibited by the contour in a field of distracter elements, and (b) the contour element-to-path orientation alignment, i.e., the amount of noise in the relative orientation alignment of the elements making up the contour to that contour's path trajectory (referred to as *angle SD*). There were two reasons for these other manipulations. Firstly, they would enable us to see whether the influence of orientation bandwidth interacted with the other stimulus properties; secondly, the *angle SD* provided a manipulation that could be compared in size with the orientation bandwidth manipulation, in the same units.

Proportion correct on each condition of Experiment 1 for each participant is plotted in Figures 5 and 6. These two figures plot identical data but in Figure 5, each curve plots the performance as a function of orientation bandwidth *SD*, whereas in Figure 6, each curve plots the performance as a function of angle *SD*. It is clear that performance drops much faster with increasing angle *SD* (Figure 6) than with increasing orientation bandwidth *SD* (Figure 5). This is a meaningful comparison because both angle *SD* and orientation bandwidth *SD* are expressed in the same units. To quantify this comparison, we calculated the drop in proportion correct as orientation bandwidth *SD* increased from 10° to 40°. Table 2 shows the mean drop across subjects for each path angle and each angle *SD* ≥ 10°. We also calculated the drop in proportion correct as angle *SD* increased from 10° to 40°. Table 3 shows the mean drop across subjects for each path angle and each orientation bandwidth *SD* ≥ 10°. By comparing corresponding cells in the two tables, we can see how much more performance was impaired by the increase in angle *SD* than by the increase in orientation bandwidth *SD* (we omitted the smallest orientation bandwidth *SD* and angle *SD* from this analysis, as they were not identical, being equal to 0° and 2°, respectively, so these conditions were not directly comparable). Table 4 shows the result of subtracting each cell of Table 2 from the corresponding cell of Table 3. In every case, the difference was positive, indicating that angle *SD* had a greater effect. We subjected the data in Table 4 to a one-sample *t* test against a hypothesized mean of zero, and found a highly significant difference, $t(15) = 8.34, p = 5.15 \times 10^{-7}$. This confirms that when the orientation bandwidth *SD* is held at σ while the angle *SD* increases from 10° to 40°, the drop in performance is much greater than when the angle *SD* is held at σ

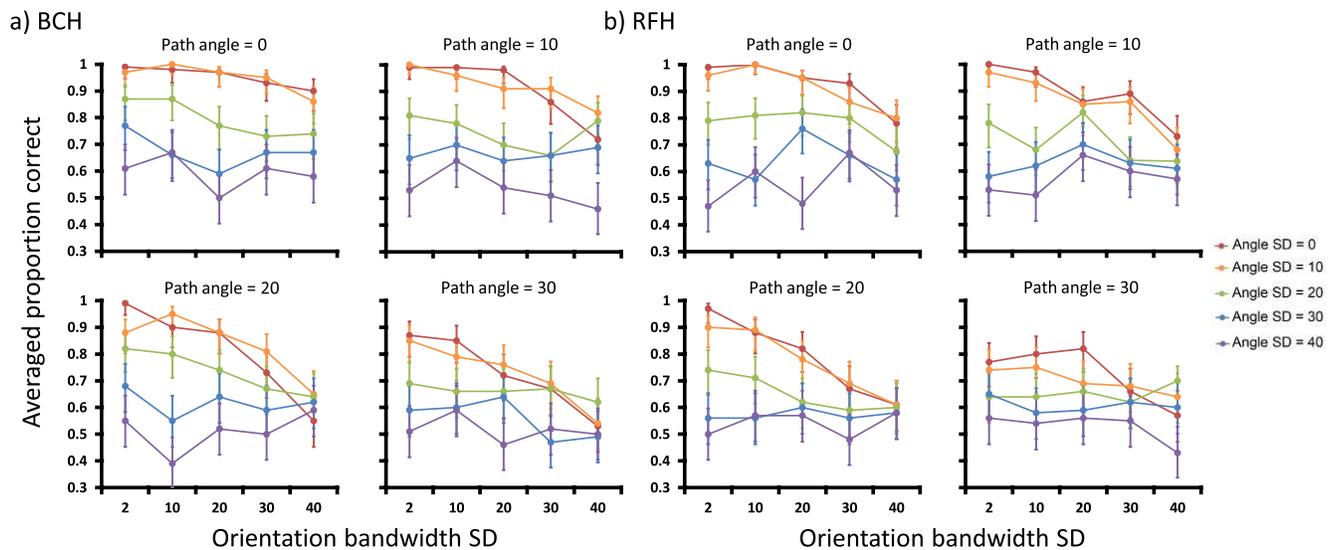


Figure 5. Proportion correct (ordinate) for each subject in each condition of Experiment 1. Each curve within a panel plots proportion correct as a function of orientation bandwidth *SD* (abscissa) for a particular angle *SD* (see legend to the far right). Error bars represent 95% Bayes credible confidence intervals (refer to the text for further detail).

while the orientation bandwidth *SD* increases from 10° to 40°.

While the orientation bandwidth *SD* controls the spread of contrast energy across orientation for an individual element, the angle *SD* controls how much the peak of this distribution deviates from the element’s contour path segment. It is clear that for fairly straight paths, performance depends critically on alignment of the peak of the energy distribution with the path segment, and is relatively insensitive to the spread of this distribution across orientation. This is still true to some extent for the larger path angles (indicated by the positive values in every cell of Table 4), but the effects

of angle *SD* and element orientation bandwidth *SD* are more similar to each other, both converging at chance performance. It therefore seems as though contour integration performance is more narrowly tuned for orientation bandwidth *SD* for larger path angles, as long as the angle *SD* is small enough to avoid floor effects (i.e., angle *SD* ≥ 20). To verify this quantitatively, the participant-averaged performance curves from Figure 5 for the 0° and 10° path angle conditions (angle *SD* = 20 to avoid ceiling effects) were best fit by a truncated Gaussian function [−90 90] with a bandwidth (half-width at half-height) of ~40°. Using the fitting procedure described above, the participant-averaged

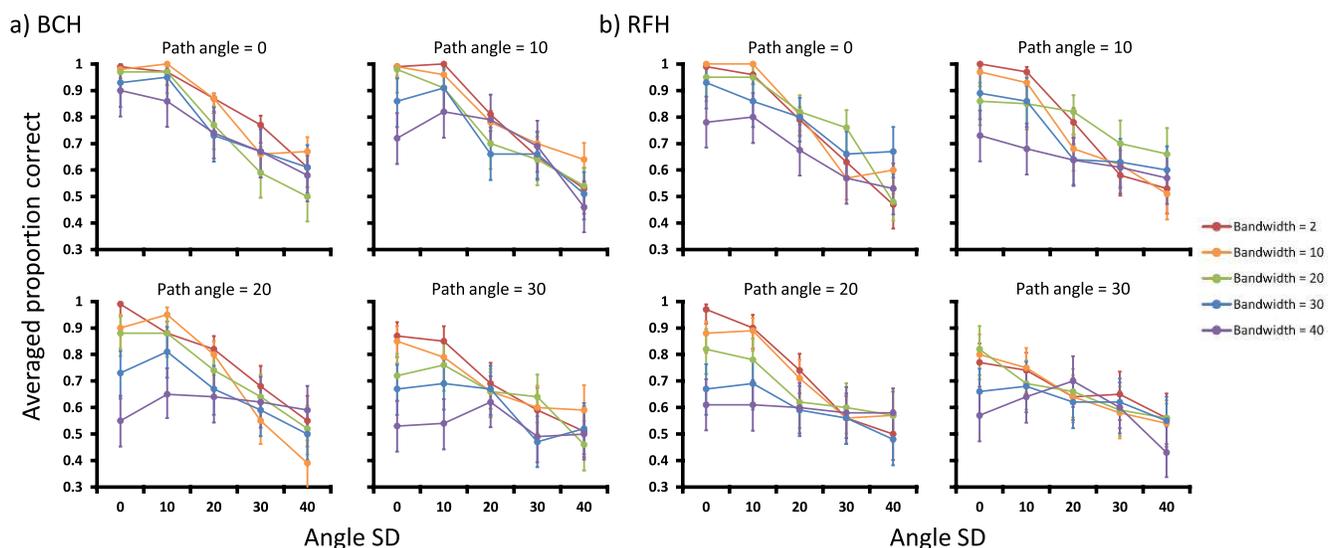


Figure 6. The same data as in Figure 5, but each curve within a panel plots proportion correct as a function of angle *SD* for a particular orientation bandwidth *SD*.

	Path angle 0°	Path angle 10°	Path angle 20°	Path angle 30°
Angle SD 10°	0.17	0.195	0.29	0.18
Angle SD 20°	0.1325	0.08625	0.135	−0.01
Angle SD 30°	−0.005	0	−0.045	0.045
Angle SD 40°	0.08	0.015	−0.105	0.1

Table 2. Drop in proportion correct as orientation bandwidth SD increases from 10° to 40°. Notes: The scores are the mean across the two subjects. Negative values indicate that performance improved over this step in orientation bandwidth SD.

performance curves for the 20° path angle conditions yielded a bandwidth (half-width at half-height) of ~25°. This difference in estimated bandwidth might seem to imply a difference between the underlying mechanisms for curved and straight contours, but in fact we show later that a single contour integration mechanism applied to all stimuli can account for this effect.

Experiment 2

It is difficult to draw any direct conclusions about the bandwidth of the contour integration mechanisms from Experiment 1. It is not immediately obvious whether the observed effects of orientation bandwidth are consistent with a narrowband contour integration mechanism, a broadband contour integration mechanism, or both. In Experiment 2 we sought to address the issue by introducing a notch into the orientation filter, centered on the center orientation. This manipulation would clearly defeat a narrowband contour integration mechanism, because a narrowband filter would show a very weak response at the orientation of each contour element; the response to the contour element would be a bimodal distribution with peaks that deviated from the element orientation, and the result would be similar to manipulating the element-to-path angle, which Experiment 1 showed had a very disruptive influence on performance (the conditions with high angle SD). On the other hand, a filter broadly tuned for orientation would bridge the notch, resulting in a unimodal distribution of energy centered on the element orienta-

	Path angle 0°	Path angle 10°	Path angle 20°	Path angle 30°
Bandwidth SD 10°	0.365	0.375	0.44	0.205
Bandwidth SD 20°	0.47	0.33	0.285	0.215
Bandwidth SD 30°	0.265	0.3	0.26	0.15
Bandwidth SD 40°	0.275	0.195	0.045	0.125

Table 3. Drop in proportion correct as angle SD increases from 10° to 40°. Note: The scores are the mean across the two subjects.

	Path angle 0°	Path angle 10°	Path angle 20°	Path angle 30°
Angle & bandwidth SD 10°	0.195	0.18	0.15	0.025
Angle & bandwidth SD 20°	0.3375	0.24375	0.15	0.225
Angle & bandwidth SD 30°	0.27	0.3	0.305	0.105
Angle & bandwidth SD 40°	0.195	0.18	0.15	0.025

Table 4. Result of subtracting each value in Table 2 from the value in the corresponding cell in Table 3. Note: The positive values in this table indicate that increasing angle SD impairs performance more than increasing orientation bandwidth SD by the same amount.

tion, so a contour integration mechanism that was broadly tuned for element orientation should still perform reasonably well in the presence of a notch, as long as the angle SD and path angles were low.

Averaged data from each participant are plotted in Figure 7a (BCH) and 7b (RFH). The notch manipulation led to an averaged performance reduction of only 20% and 25% (from ceiling) at angle SDs of 0° and 10°, respectively. As argued above, this suggests that the elements are processed with a mechanism that is not simply narrowly tuned for orientation; otherwise one would expect to find chance (or near chance) performance at the 6° and 12° notch conditions. We return to this notion in the modeling section (Part II).

Experiment 3

Experiment 3 further assessed the orientation properties of contour integration for contours of different curvature by alternating the contour elements with respect to either two different element orientation bandwidth SDs for two different orientation notch filter bandwidth SDs. Dakin and Hess (1998) employed a similar paradigm to assess the spatial frequency tuning of contour integration performance. Thus, the idea here was to provide an alternative paradigm to that explored in Experiment 2. However, it is worth noting that the spatial frequency alternation of contour elements in Dakin and Hess (1998) does not cleanly map onto a paradigm that alternates the orientation bandwidth of contour elements. Specifically, mechanisms tuned for spatial frequency will respond to some spatial frequencies but not others. Thus, if the spatial frequencies of the alternating elements differ by more than the tuning width of the contour integration mechanism, then contour integration will break down (leaving the critical spatial frequency difference as a measure of the tuning width of the contour integration mechanism). However, it's not clear whether such a breakdown process could be expected for alternating orientation bandwidths. Therefore, the primary motivation for conducting Experiment 3 was to provide

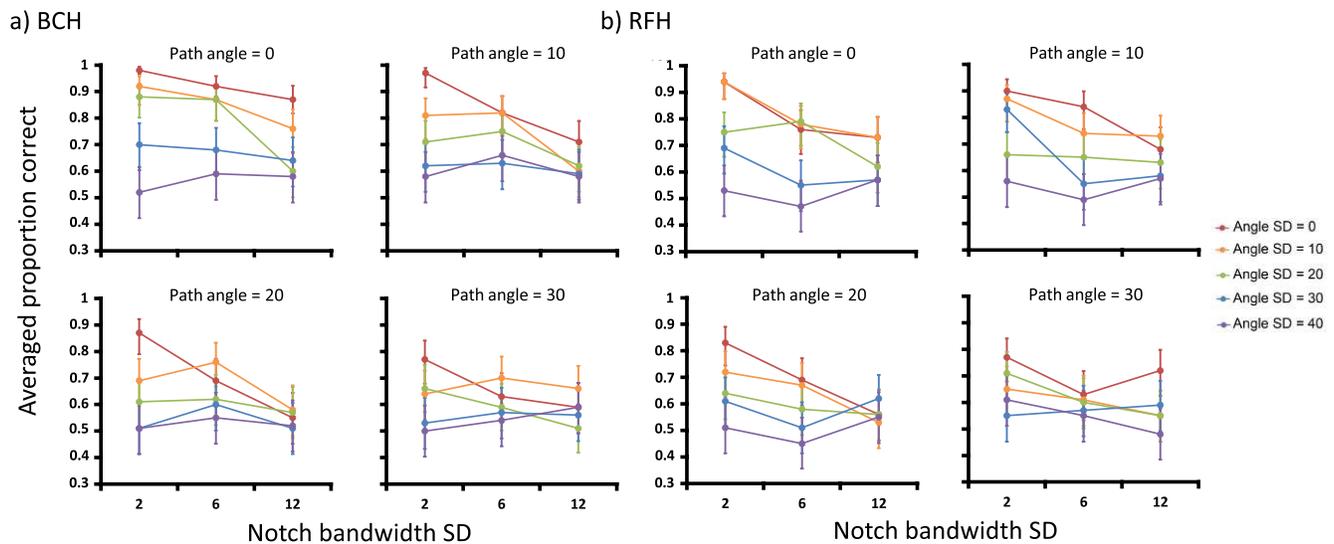


Figure 7. Individual data from Experiment 2, averaged across repetition. Each subpanel shows averaged proportion correct (ordinates) for the four different path angles as a function of element orientation bandwidth (abscissa). Error bars are calculated in the same way as in Figures 5 and 6.

additional data to test the two-stage filter-overlap contour integration model in order to develop a more robust estimate of orientation bandwidth of the front-end processes underlying the contour integration mechanism.

Data from each observer (along with the data averaged across observers) are shown in Figure 8. Figure 8a plots performance for detecting contours consisting of elements with orientation bandwidth SD alternating between 2° and 30° , or alternating between 2° and 40° , with data replotted from Experiment 1 for comparison.

For relatively straight contours, performance is virtually identical to that obtained in Experiment 1 for element bandwidths of 2° (probably resulting from a ceiling effect). As path angle increases, performance for the alternating contour falls to a level between Experiment 1 performances for the two bandwidths used in the Experiment 3 contour.

Figure 8b plots performance for detecting contours consisting of elements with orientation notch bandwidth SD alternating between 0° (i.e., no notch) and 6° , or alternating between 0° and 12° , with data replotted from Experiment 2 for comparison. Interestingly, performance in both notch conditions (i.e., 0° – 6° or 0° – 12°) is much higher in the alternating no-notch and notch bandwidths than in the notch conditions measured in Experiment 2.

Part II

Computational modeling of Experiments 1 through 3

Contour integration is a complicated process; it is difficult to know for certain what our data imply about

contour integration mechanisms without carrying out a full implementation of a contour integration mechanism. To this end, we simulated all our experiments using the two-stage filter-overlap contour integration model described by May and Hess (2008). This model takes the stimulus images, processes them, and delineates contours in the images. For each condition in our experiment, we simulated 200 trials by generating 200 pairs of stimuli (target and nontarget) using the same software that we had used to generate stimuli in the psychophysical experiments, and running the model on these stimuli to extract contours. On each trial, we generated a psychophysical response from the model's output to the two stimuli, and we recorded the proportion correct for each condition. The results are described below.

Figure 9 shows the MAD between mean human performance and model performance as a function of the threshold. To find the MAD, we averaged BCH and RFH's scores together, to produce a single proportion correct for each condition; then, for each condition, we found the absolute difference between the averaged human score and the model's score, and then we took the mean absolute difference across conditions. Experiments 1 and 2 were fairly insensitive to the threshold level: Only Experiment 3, in which the contour elements alternated in orientation bandwidth SD or notch width, showed a substantial effect of the threshold. This shows that, for more ordinary contours, which do not fluctuate wildly in their properties from element to element, May and Hess's (2008) contour integration algorithm is quite robust to changes in the threshold level, so the threshold level does not have to be set to a precise value for the mechanism to work well. This suggests that the model is

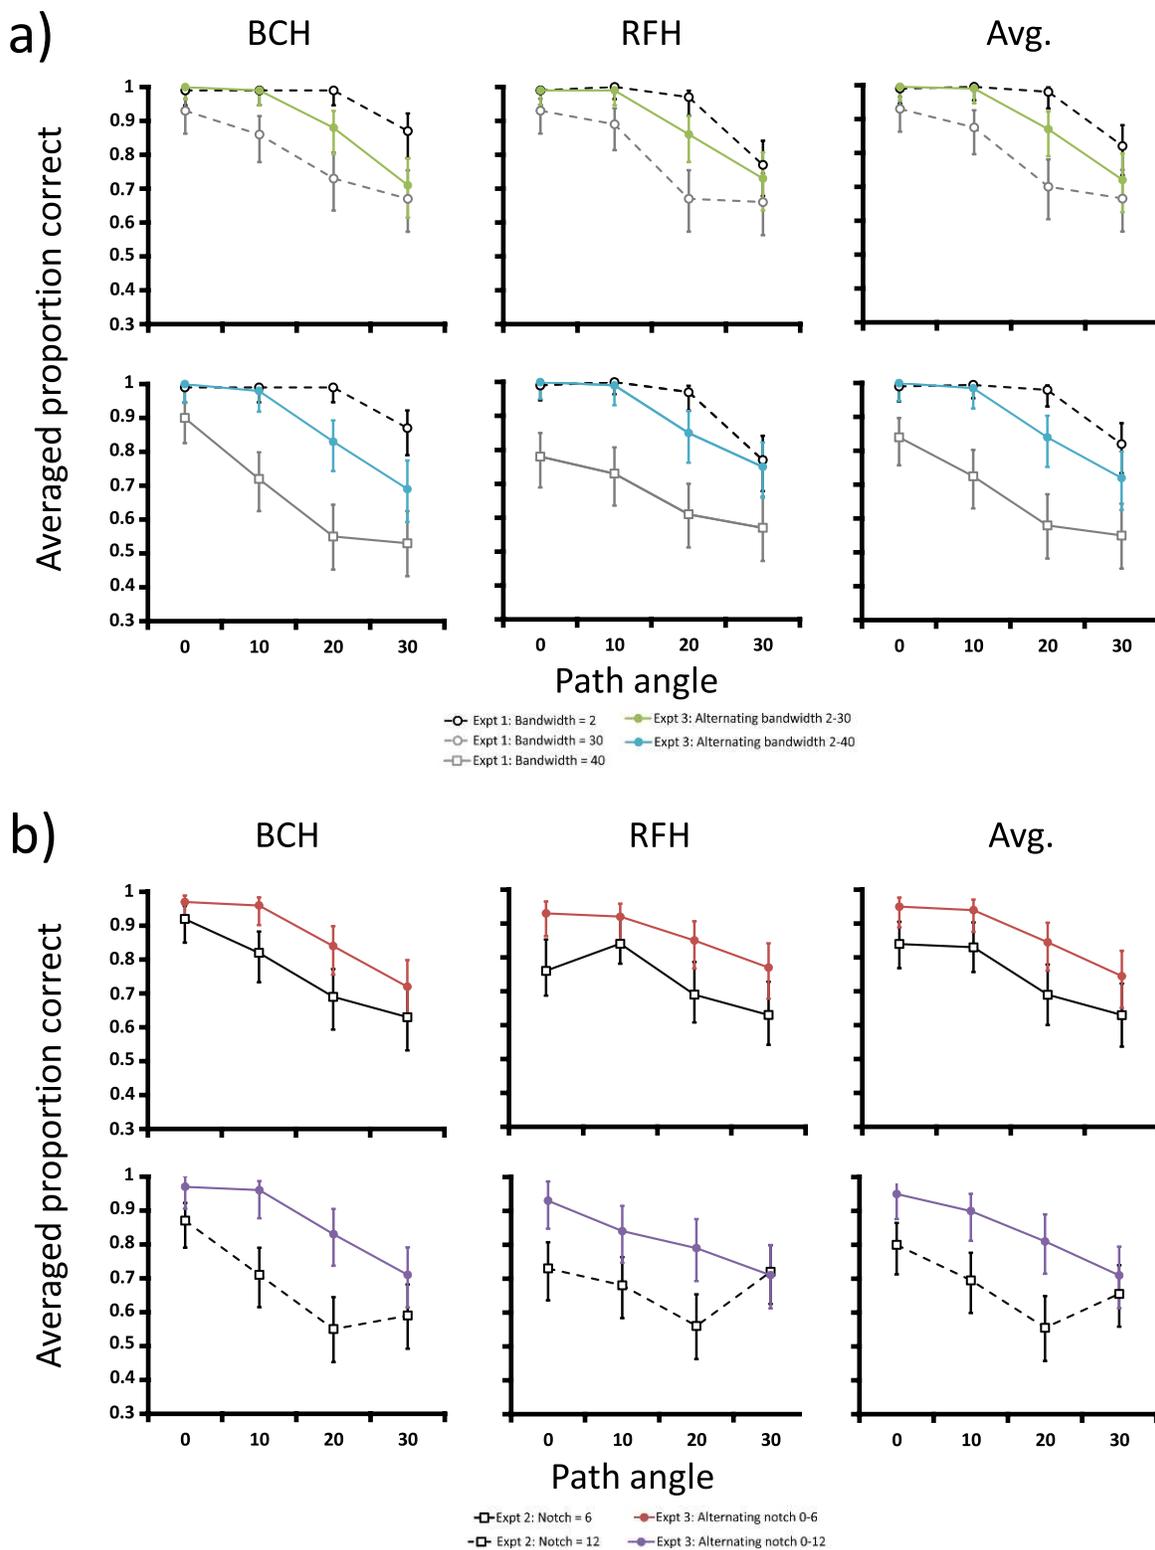


Figure 8. (a) Far left and middle panels: Individual data from Experiment 3, averaged across repetition, far right panel, data averaged across participants. On the ordinates is averaged proportion correct, and on the abscissas is path angle. Gray and black traces are data replotted from Figure 5 (i.e., Experiment 1). (b) Far left and middle panels: Individual data from Experiment 3, averaged across repetition, far right panel, data averaged across participants. On the ordinates is averaged proportion correct, and on the abscissas is path angle. Black traces are data replotted from Figure 7 (i.e., Experiment 2). Error bars are calculated in the same way as in Figures 5 through 7.

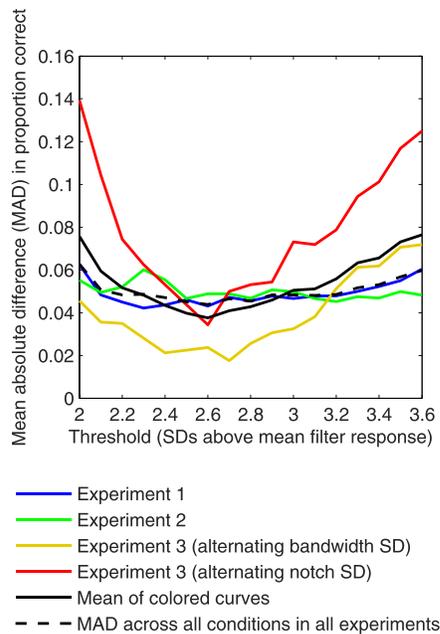


Figure 9. MAD between mean human performance (across the two subjects) and model performance as a function of the threshold. All other model parameters were set to the values given in Table 1. The colored lines plot the MADs for the individual experiments. The solid black line plots the mean of these values. The dashed black line plots the MAD across all conditions of all the experiments. In the measurement shown by the dashed line, each condition has equal weight. In the measurement shown by the solid black line, each experiment has equal weight. Both the dashed and solid lines have their minimum value at a threshold value of 2.6. We also carried out a similar analysis using root mean square error instead of MAD, and obtained similar results, again with 2.6 as the best-fitting threshold.

a robust algorithm that could be implemented in imperfect biological hardware.

The best-fitting threshold across all experiments was 2.6 SDs above the mean second-stage filter response. Figures 10 through 13 plot the performance of the model with this threshold level, and all the filter parameters fixed at the values given in Table 1. This model, with a single set of parameters across all conditions of all experiments, gives a surprisingly good fit to the whole data set: Mean absolute difference in proportion correct between the model and the averaged human data was only 0.0438 (this is the value plotted with the dashed line for threshold 2.6 in Figure 9). Thus, the data for all the different contours in our experiment can be accounted for by a single contour integration mechanism. Given the limited search of parameter space that we could feasibly conduct, it is almost certain that we missed a set of parameters that would have fit the data even better, but the fits that we obtained with the parameter values in Table 1 are very satisfactory.

Figure 14 takes the data from Figures 10 through 13 and plots the model's performance against the mean human performance for each condition of each experiment. The Pearson correlation coefficient, r , between these two datasets is 0.934, so the model explains 87.4% of the variance in the psychophysical data (this is $100 \times r^2$). The t statistic corresponding to $r = 0.934$ is given by $t = 48.2$ ($df = 174$), which is so high that p evaluates to zero in MATLAB, suggesting that p is too small to represent using a 64-bit floating point number.

We reported above that the model gives a good fit to the psychophysical data: The mean absolute difference in proportion correct between the model and human data was only 0.0438. But is the fit of the model to the psychophysical data good enough? We now demonstrate that it is.

Since the model is a model of human performance, the most we could reasonably expect of the best-fitting model is that it fits to the data of BCH and RFH as closely as BCH's data and RFH's data fit to each other. We found the MAD between BCH and RFH, and compared that with the MAD between the human subjects and the model. In finding the MAD between the model and each subject, we needed to split the 200-trial modeling dataset into two sets of 100 trials; otherwise the MAD between model and human would be unfairly reduced (compared with the MAD between the two humans) due to the lower noise in the model's data. We therefore set up two simulated subjects, S1 and S2, one with the first 100 modeling trials on each condition, and the other with the second 100 modeling trials. This gave us two human subjects, BCH and RFH, and two simulated subjects, S1 and S2, each of whom performed 100 trials on the experiment. Figure 15 plots each of these four subjects' scores against each other subjects' scores. Table 5 shows the MAD between each set of scores, and Table 6 shows the Pearson correlation between each set of scores. The MAD between BCH and RFH (0.0585) falls within the range of MADs between the simulated and human subjects (0.0539–0.0625). Similarly, the correlation between BCH and RFH (0.889) falls within the range of correlations between the simulated and human subjects (0.874–0.905). Thus, the fit between the two human subjects is generally no better than the fit between the model and the human subjects. This, as stated above, is the best we can reasonably expect of the model fit.

Examining the effect of first-stage orientation bandwidth

Earlier, we gave an intuitive argument that the weak effect of the notch found in Experiment 2 implied that the elements were processed with filters that were not narrowly tuned for orientation. Here, we provide concrete support for this assertion. In May and Hess' (2008) model, the stimulus elements are initially

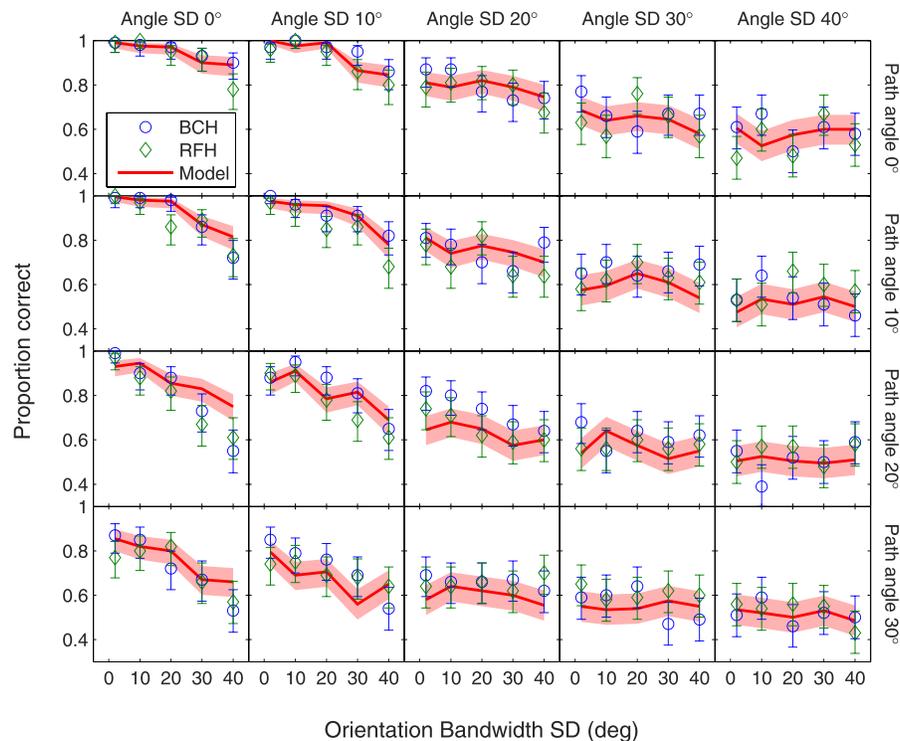


Figure 10. Model's performance on Experiment 1. Symbols show human psychophysical data from subjects BCH (blue circles) and RFH (green diamonds). The red lines show the performance of the model with filter parameters given in Table 1, and a threshold of 2.6. Each column of panels gives data from a different angle SD , which controls the element-to-path alignment. Each row of panels gives data for a different path angle, which controls the contour curvature. Within each panel, the proportion correct is plotted as a function of orientation bandwidth SD , which controls the orientation bandwidth of the stimulus elements. The pink regions show 95% Bayes credible intervals for the model's scores. The error bars show 95% Bayes credible intervals for the human subjects' scores.

processed by the first-stage filters, so we would expect that these filters would have to be fairly broadband for the model to be able to perform as well as humans when the notch SD was large. We took the conditions from the top-right panel of Figure 11 (notch SD 12° , path angle 0°) and simulated these conditions with first-stage filters set to have a narrower orientation bandwidth. We refer to this new parameterization as the “narrowband model”; the parameterization shown in Table 1 (and plotted in all the other figures) is referred to as the “broadband model.” All the filter parameters in the narrowband model were as in Table 1, except for the first-stage aspect ratio, $\sigma_{v,1}/\sigma_{u,1}$ which was set to 2. The higher aspect ratio elongated the Gabor envelope, resulting in an orientation bandwidth (full-width at half-height) of 37° . We varied the threshold from 1.7 to 3.6 in steps of 0.1, and selected the threshold that gave the lowest MAD between the narrowband model and the human data (averaged across subjects) from the five conditions in the top-right panel of Figure 11. The best-fitting threshold for the narrowband model was 1.8. Figure 16a plots the mean of the human data from the top-right panel of Figure 11, along with the performance of the narrowband

model with a threshold of 1.8. For the smallest two angle SD s (i.e., those with substantially above-chance human performance), the confidence intervals for the narrowband model and human data do not overlap. As a comparison, Figure 16b replots the broadband model data from the top-right panel of Figure 11, along with the mean human data from that panel: The confidence intervals of the broadband model and human data overlap considerably. In summary, we have shown that the model with the broadband first-stage filter is able to perform as well as humans when the notch SD was large, whereas the model with the more narrowband first-stage filter cannot reach human performance. It should be noted that this comparison was greatly biased in favor of the narrowband model, because the threshold in this model was fitted just to the data in Figure 16, while the broadband model had its threshold set to give the best fit across all 176 conditions in the study. Yet, despite having this massive advantage, the narrowband model was unable to reach human performance levels on these conditions. We conclude that a broadband first-stage filter is necessary for the model to be able to integrate contours that contain a 12° orientation notch.

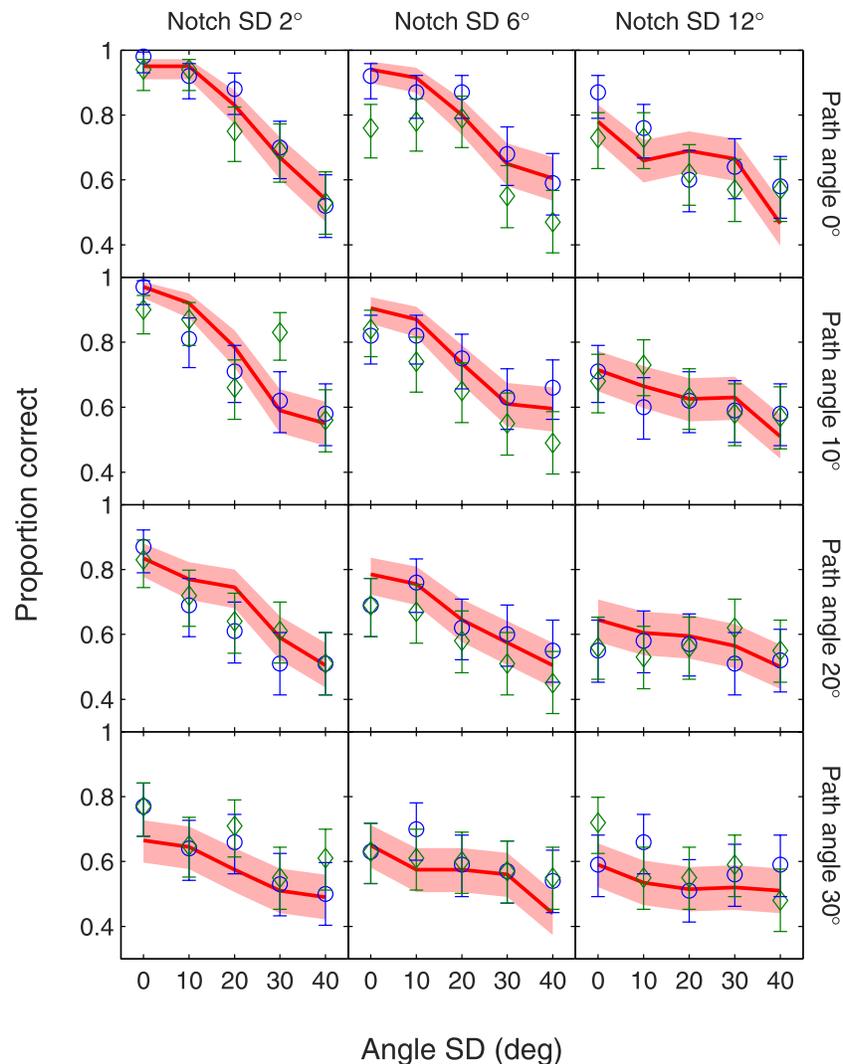


Figure 11. Model's performance on Experiment 2. Each column of panels gives data for a different notch SD , which controls the width of the notch in the distribution of energy across orientation for the individual elements. Each row of panels gives data for a different path angle. Within a panel, proportion correct is plotted as a function of angle SD , which controls the element-to-path alignment. Symbols, lines, and confidence intervals follow the same plotting conventions as in Figure 10.

Discussion

Experiment 1 suggests that contour integration performance in human vision is relatively broadly tuned for orientation bandwidth SD for small path angles, and more narrowly tuned for larger path angles. Experiment 2 showed that notch-filtering the stimulus elements with a notch centered on the element orientation impaired performance to some extent, but did not completely disrupt the contour integration process. The latter finding suggests that, in contour integration, the stimulus is initially processed with filters that have sufficiently wide orientation bandwidth to be able to bridge the notch. This notion was supported by the results of Experiment 3. As expected,

element-to-path orientation misalignment (i.e., increased angle SD) had a detrimental impact on integration performance, regardless of element orientation bandwidth or path angle.

In order to provide a quantitative account of the qualitative description given above regarding contour integration performance, we simulated all the experiments with the two-stage model of contour integration presented by May and Hess (2008). The model takes the stimulus images as input, and outputs ZBRs that each correspond to a contour. The model has seven parameters: three parameters for the first-stage Gabor filters, three parameters for the second-stage Gabor filters, and a threshold. The four filter parameters that determined the *shapes* of the Gabor kernels (i.e., shapes of the receptive field profiles) were set to values

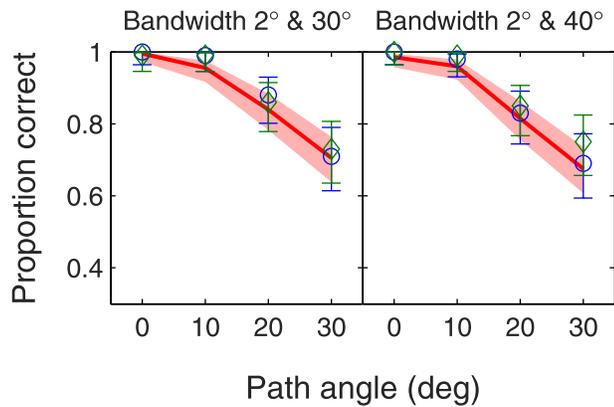


Figure 12. Model's performance on the conditions of Experiment 3 where the stimuli consisted of alternating orientation bandwidth SD s. Symbols, lines, and confidence intervals as in Figure 10.

constrained by taking into account both physiological data and the task requirements. The two parameters that determined the overall scales (i.e., sizes) of the first- and second-stage filter kernels ($\sigma_{u,1}$ and $\sigma_{u,2}$) were fit to the data, as was the threshold.

As shown by May and Hess (2008), the appropriate filter scales depend on the stimulus properties: The first-stage scale should be appropriate for the element spatial frequency, and the second-stage scale should be appropriate for the element spacing. We envisage that the human visual system would select these scales automatically. The mechanism by which this scale selection occurs is currently unknown, but might be similar to that which occurs in edge processing, whereby feature detection at each point in the image appears to be mediated by the filter with the strongest response (Georgeson, May, Freeman, & Hesse, 2007; May & Georgeson, 2007a, 2007b; McIlhagga & May, 2012). Given that the scale selection process in contour integration is currently unknown, we simply fit the filter scales, $\sigma_{u,1}$ and $\sigma_{u,2}$, as free parameters in the current

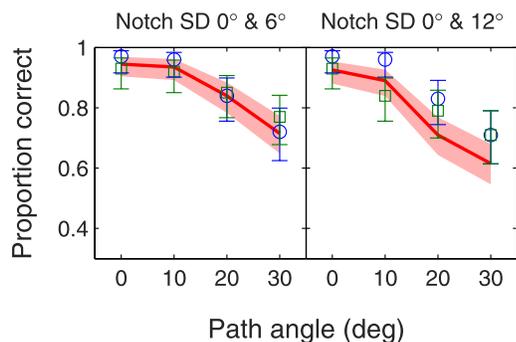


Figure 13. Model's performance on the conditions of Experiment 3 where the stimuli consisted of alternating notch SD s. Symbols, lines, and confidence intervals as in Figure 10.

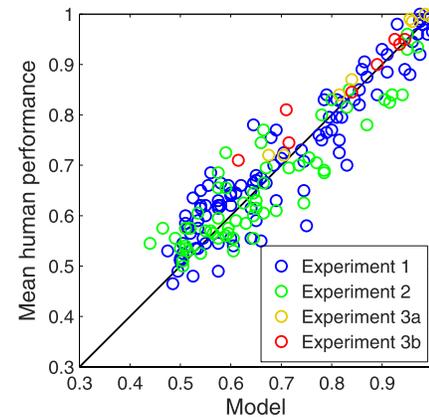


Figure 14. Plot of model performance against the mean human performance. Each symbol represents one condition of one experiment. The conditions are color-coded according to which experiment they are from (see legend in the bottom right). Experiment 3a is the part of Experiment 3 with alternating bandwidth SD ; Experiment 3b is the part of Experiment 3 with alternating notch SD . The Model data are the modeling data plotted in Figures 10 through 13, i.e., proportion correct over 200 trials. The mean human performance is the mean of BCH and RFH's data, giving a total of 200 trials per condition across the two subjects.

study. Since all conditions of all the experiments in this study had the same element spatial frequency and element spacing, all the conditions required the same filter scale parameters, so these parameters could be kept constant across all conditions. Keeping the element spatial frequency and element spacing constant, we varied several key stimulus parameters that should reveal characteristics of the contour integration mechanisms: These stimulus parameters were path angle, angle SD , element orientation bandwidth, and element orientation notch width. Each stimulus parameter was varied over a very wide range. Across this large range of conditions, 100 (Experiment 1) + 60 (Experiment 2) + 8 (Experiment 3a) + 8 (Experiment 3b) = 176 conditions, a single set of model parameters fit well to all the data.

The results of our experiments (especially Experiment 2) suggested the involvement of a filter that was broadly tuned for orientation. This seemed to conflict with the intuition that contour integration would best be achieved with elongated (i.e., orientation-narrowband) receptive fields, which can bridge the gaps between successive contour elements without picking up too many distracter elements. The two-stage model accommodates both requirements (see Table 1), with a broadband first-stage filter that processes the elements, and a narrowband second-stage filter that integrates the elements to form a perceptually coherent contour.

There have only been a few previous attempts to conduct full, quantitative, simulations of contour

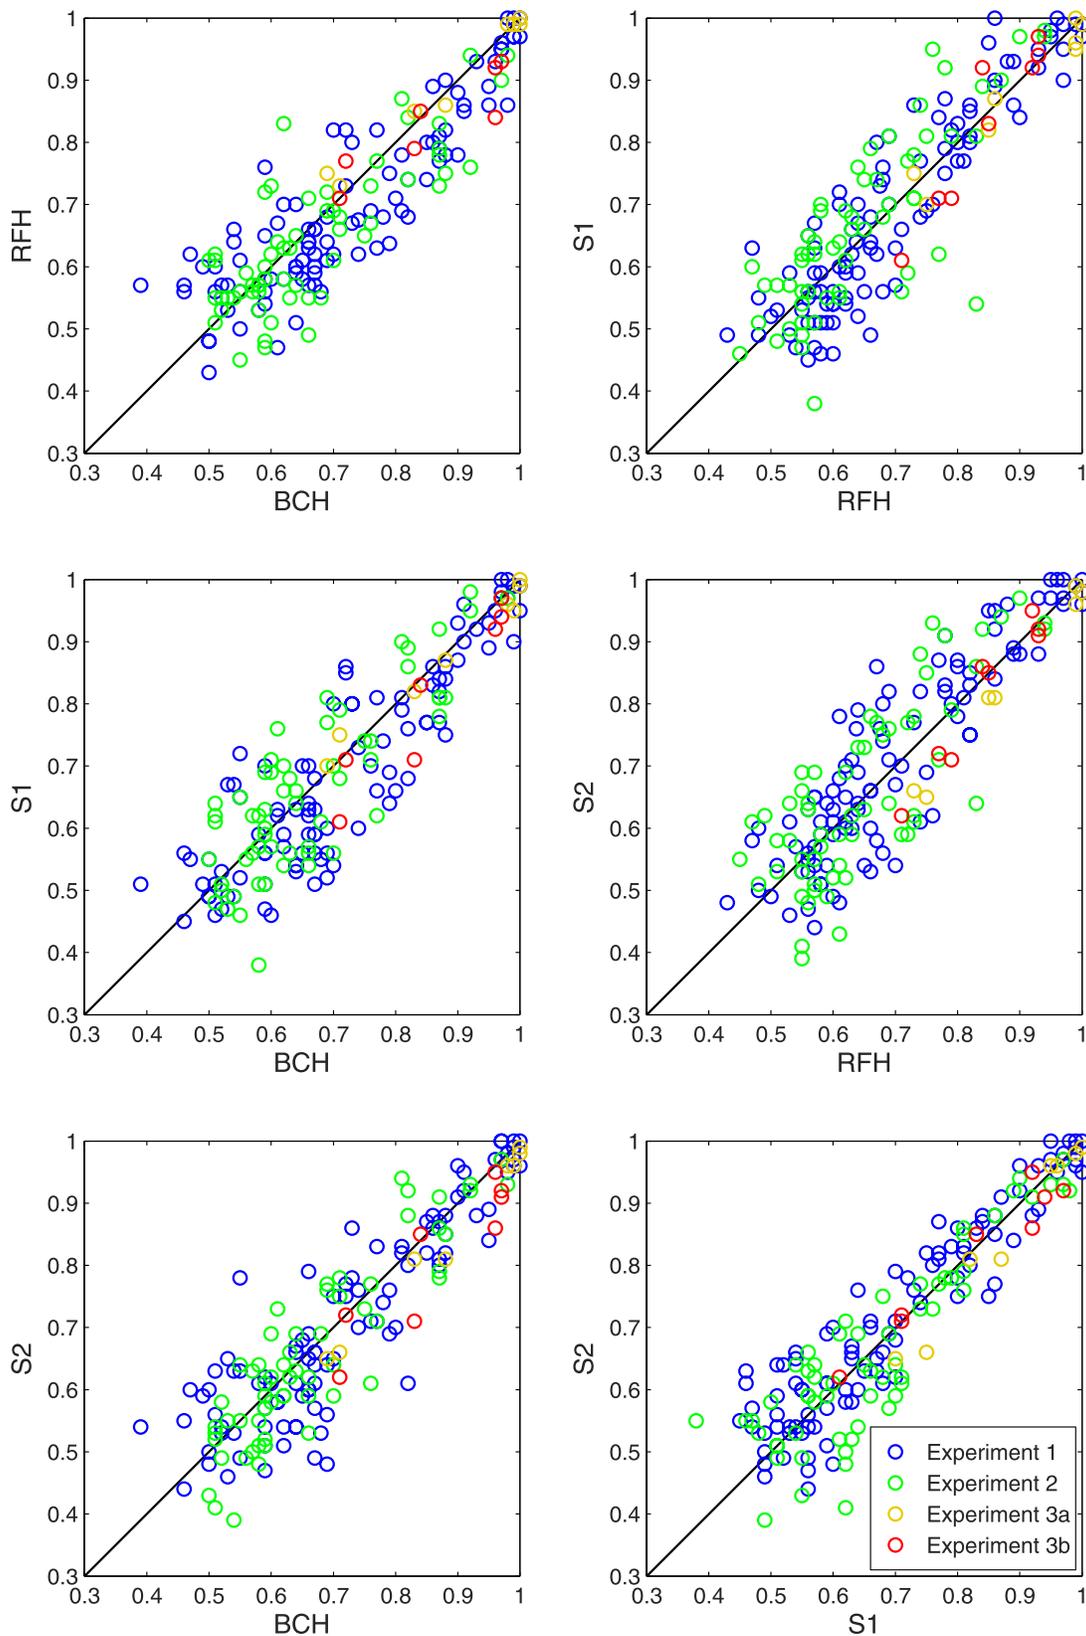


Figure 15. Scatter plots comparing the human and simulated subjects' data. Each symbol represents one condition of one experiment. The conditions are color-coded according to which experiment they are from (see legend in the bottom right panel). Experiment 3a is the part of Experiment 3 with alternating bandwidth *SD*; Experiment 3b is the part of Experiment 3 with alternating notch *SD*. Each axis label indicates which subject's data are being plotted on that axis. BCH and RFH are the human subjects. S1 and S2 are the simulated subjects, each generated by running the model on 100 stimuli on each condition.

	BCH	RFH	S1	S2
BCH	—	0.0585	0.0597	0.0539
RFH	—	—	0.0567	0.0625
S1	—	—	—	0.0482
S2	—	—	—	—

Table 5. Mean absolute difference (MAD) between each subject’s scores. Notes: BCH and RFH were the two human subjects. S1 and S2 were the two simulated subjects.

integration experiments. Without exception, these studies have focused on the linking process without much consideration of the effects of the stimulus elements themselves and what this tells us about the front-end filters used in contour integration. Some studies did not include a front-end filter at all, the input to the model being merely the positions and orientations of the stimulus elements (Ernst et al., 2012; May & Hess, 2007; Watt, Ledgeway, & Dakin, 2008); these models cannot distinguish between different element properties, such as orientation bandwidth. Other models have used front-end filters, but chosen on a fairly ad-hoc basis, simply to provide an orientation-selective response to the stimulus elements (Geisler, Perry, Super, & Gallogly, 2001; Gintautas et al., 2011; Pettet, McKee, & Grzywacz, 1998; Yen & Finkel, 1998). Our study is the first to examine in detail the properties of the orientation-selective filters used in the early stages of processing by contour integration mechanisms in human vision.

It is important to stress that our conclusions regarding the filter bandwidth are predicated on the May and Hess (2008) model. An alternative class of model is the association field model, in which local units are linked by lateral connections. It is possible that our results might be explained by an association field model with narrowly tuned front-end filters but wide orientation tuning on the lateral connections. Indeed, Schinkel, Pawelzik, and Ernst (2005) reported that disruption due to orientation jitter of the contour elements could be overcome either by increasing the input orientation uncertainty or by increasing the association field width, so that wide orientation tuning had similar effects whether it was in the front-end processing, or in the association field. However, the interpretation of Schinkel et al.’s results is not entirely

	BCH	RFH	S1	S2
BCH	—	0.889	0.892	0.905
RFH	—	—	0.892	0.874
S1	—	—	—	0.923
S2	—	—	—	—

Table 6. Pearson correlation between each subject’s scores. Notes: BCH and RFH were the two human subjects. S1 and S2 were the two simulated subjects.

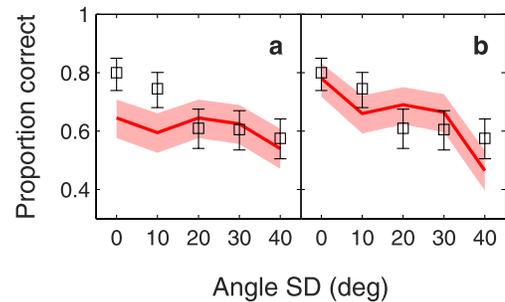


Figure 16. Failure of narrowband first-stage filter to integrate contours with a wide notch. (a) Symbols plot the mean human data (averaged across subjects) from the top right panel of Figure 11 (notch SD 12°, path angle 0°). The red line plots the performance of the narrowband model, with all filter parameters as in Table 1 except the aspect ratio, $\sigma_{v,1}/\sigma_{u,1}$, which was set to 2; the threshold for the narrowband model was 1.8, which provides the best fit to these five data points. Confidence intervals are illustrated using the conventions in Figures 10 through 13. (b) Symbols plot the same human data as in (a). The red line and corresponding confidence intervals are replotted from the top right panel of Figure 11, and show the broadband model performance. This model had the filter parameters in Table 1; the threshold was 2.6, which gave the best fit across all 176 conditions in this study.

straightforward, as explained in the next two paragraphs.

Firstly, their model did not actually carry out contour grouping: The model calculated the saliency of each location, and Schinkel et al. (2005) assumed the contour was detected if three of the five most salient locations corresponded to contour elements. A critical part of the contour integration process is the perceptual grouping of different sets of elements into different contours, and this part of the process seems to be missing from Schinkel et al.’s model. So the trade-off between front-end and association field orientation tuning has not been demonstrated in a model that actually carries out a contour *integration* process. It is possible that increasing the association field width would cause more excessive grouping between distracter elements and contour elements than would be caused by increasing the front-end tuning width.

Secondly, and more importantly, Schinkel et al.’s (2005) model did not actually contain any front-end filters. The input to the model (which they termed the *afferent input*) at each spatial position was a probability distribution giving the relative probability of an edge as a function of orientation. Schinkel et al. assumed that this probability distribution could be identified with the cortical orientation tuning function. However, we argue that this is a false analogy. The afferent input to their model is essentially a posterior probability distribution that would arise from decoding the local orientation from the front-end neurons. The width of

this posterior distribution does depend on neuronal tuning width, but it also depends on the number of neurons, the maximum spike rate of the neurons, the spontaneous firing rate, the Fano factor (ratio of neuronal spike variance to mean rate), and the stimulus itself. Any of these factors can trade off with each other in their contribution to decoding precision (see May & Solomon, 2014). If all we know is the posterior distribution, we cannot determine the cortical tuning width. Therefore, it cannot be correct to identify the afferent input to Schinkel et al.'s model with cortical tuning functions. We conclude that, although Schinkel et al.'s (2005) model is very interesting, it does not demonstrate conclusively that increasing the association field tuning can have similar effects to increasing the front-end filter tuning width, and further work will be needed to examine this hypothesis.

Finally, our study gives the first quantitative comparison of human data with the performance of an important new class of grouping models in which grouping occurs by virtue of the overlap between the filter responses to different elements. A key feature of the model that we used is that it represents the stimulus in a space that has a dimension representing orientation, in addition to the two dimensions of the image (May & Hess, 2008). The spread of the filter's response across image space and orientation has the effect of blurring the representation along all three dimensions so that the responses due to elements that are close in space and orientation join up to form a ZBR. Rosenholtz and colleagues (2009) extended this idea to a much broader class of grouping processes. To perform grouping on the basis of some feature dimension, f , you can create a multidimensional space (x, y, f) , where x and y are spatial coordinates, and then plot the image in this space. Then image elements with similar attributes and spatial positions will be nearby and, if you blur the representation, they join up. As noted earlier, a quantitative comparison of the performance of this class of model with that of human subjects has so far been lacking. Our quantitative evaluation of the model showed a very good fit to a large data set with relatively few free parameters, suggesting that this class of model may have an important role to play in helping us to understand the mechanisms of contour integration, and perceptual grouping more generally.

Keywords: contour integration, two-stage filter-overlap model, orientation tuning

Acknowledgments

This research was supported in part by a Colgate Research Council grant to BCH, and an NERC grant

(#46521-11) to RFH. KAM was supported by EPSRC Grant EP/H033955/1 to Joshua Solomon.

Commercial relationships: none.

Corresponding author: Bruce C. Hansen.

Email: bchansen@colgate.edu.

Address: Department of Psychology, Neuroscience Program, Colgate University, Hamilton, NY, USA.

Footnotes

¹ Given the extensively systematic nature of the current study (e.g., 176 conditions in total), it is necessary to have experienced psychophysical observers so that they did not become better psychophysical observers during the course of the study. Further, the current study is completely exploratory, thus the participants had no expectations regarding ad hoc hypotheses and can both be considered naive to the outcome of the study.

² Note that the relative contour and background elements were randomly positioned within each 16×16 grid cell (which is standard practice), and thus their distances from one another were also randomly distributed, and did not lead to any reliable detection cue when contour elements were present (note that contour detection performance in several of the conditions described later was at chance). Thus, performance cannot have been mediated by differences in density around the contours.

References

- Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., ... Rust, N. C. (2005). Do we know what the early visual system does? *Journal of Neuroscience*, 25(46), 10577–10597.
- Dakin, S. C., & Hess, R. F. (1998). Spatial-frequency tuning of visual contour integration. *Journal of the Optical Society of America, A*, 15(6), 1486–1499.
- Dakin, S. C., & Hess, R. F. (1999). Contour integration and scale combination processes in visual edge detection. *Spatial Vision*, 12(3), 309–327.
- De Valois, R. L., Albrecht, D. G., & Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research*, 22(5), 545–559.
- De Valois, R. L., Yund, E. W., & Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research*, 22(5), 531–544.

- Ernst, U. A., Mandon, S., Schinkel-Bielefeld, N., Neitzel, S. D., Kreiter, A. K., & Pawelzik, K. R. (2012). Optimality of human contour integration. *PLoS Computational Biology*, 8(5), e1002520.
- Field, D. J., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local “association field.” *Vision Research*, 33(2), 173–193.
- Field, D. J., Hayes, A., & Hess, R. F. (2000). The roles of polarity and symmetry in the perceptual grouping of contour fragments. *Spatial Vision*, 13(1), 51–66.
- Field, D. J., & Hayes, A. (2003). Contour integration and the lateral connections of V1 neurons. In L. M. Chalupa & J. S. Werner (Eds.), *The visual neurosciences* (pp. 1069–1079). Oxford, NY: MIT Press.
- Field, D. J., & Tolhurst, D. J. (1986). The structure and symmetry of simple-cell receptive-field profiles in the cat’s visual cortex. *Proceedings of the Royal Society of London B*, 228(1253), 379–400.
- Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41(6), 711–724.
- Georgeson, M. A., May, K. A., Freeman, T. C. A., & Hesse, G. S. (2007). From filters to features: Scale-space analysis of edge and blur coding in human vision. *Journal of Vision*, 7(13):7, 1–21, <http://www.journalofvision.org/content/7/13/7>, doi:10.1167/7.13.7. [PubMed] [Article]
- Gintautas, V., Ham, M. I., Kunsberg, B., Barr, S., Brumby, S. P., Rasmussen, C., ... Kenyon, G. T. (2011). Model cortical association fields account for the time course and dependence on target complexity of human contour perception. *PLoS Computational Biology*, 7(10), e1002162.
- Hansen, B. C., & Hess, R. F. (2006). The role of spatial phase in texture segmentation and contour integration. *Journal of Vision*, 6(5):5, 594–615, <http://www.journalofvision.org/content/6/5/5>, doi:10.1167/6.5.5. [PubMed] [Article]
- Hess, R. F., Hayes, A., & Field, D. J. (2003). Contour integration and cortical processing. *Journal of Physiology Paris*, 97(2–3), 105–119.
- Jones, J. P., & Palmer, L. A. (1987). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, 58(6), 1233–1258.
- Maffei, L., & Fiorentini, A. (1973). The visual cortex as a spatial frequency analyser. *Vision Research*, 13(7), 1255–1267.
- May, K. A., & Georgeson, M. A. (2007a). Blurred edges look faint, and faint edges look sharp: The effect of a gradient threshold in a multi-scale edge coding model. *Vision Research*, 47, 1705–1720.
- May, K. A., & Georgeson, M. A. (2007b). Added luminance ramp alters perceived edge blur and contrast: A critical test for derivative-based models of edge coding. *Vision Research*, 47, 1721–1731.
- May, K. A., & Hess, R. F. (2007). Ladder contours are undetectable in the periphery: A crowding effect? *Journal of Vision*, 7(13):9, 1–15. <http://www.journalofvision.org/content/7/13/9>, doi:10.1167/7.13.9. [PubMed] [Article]
- May, K. A., & Hess, R. F. (2008). Effects of element separation and carrier wavelength on detection of snakes and ladders: Implications for models of contour integration. *Journal of Vision*, 8(13):4, 1–23, <http://www.journalofvision.org/content/8/13/4>, doi:10.1167/8.13.4. [PubMed] [Article]
- May, K. A., & Solomon, J. A. (2014). Connecting psychophysical performance to neuronal response properties. Manuscript submitted for publication.
- McIlhagga, W. H., & May, K. A. (2012). Optimal edge filters explain human blur detection. *Journal of Vision*, 12(10):9, 1–13, <http://www.journalofvision.org/content/12/10/9>, doi:10.1167/12.10.9. [PubMed] [Article]
- Merigan, W. H., & Maunsell, J. H. (1993). How parallel are the primate visual pathways? *Annual Review of Neuroscience*, 16, 369–402.
- Nicholson, B. J. (1985). On the *F*-distribution for calculating Bayes credible intervals for fraction nonconforming. *IEEE Transactions on Reliability*, R-34, 227–228.
- Pettet, M. W., McKee, S. P., & Grzywacz, N. M. (1998). Constraints on long range interactions mediating contour detection. *Vision Research*, 38(6), 865–879.
- Ringach, D. L., Shapley, R. M., & Hawken, M. J. (2002). Orientation selectivity in macaque V1: Diversity and Laminar dependence. *Journal of Neuroscience*, 22(13), 5639–5651.
- Rosenholtz, R., Twarog, N. R., Schinkel-Bielefeld, N., & Wattenberg, M. (2009). An intuitive model of perceptual grouping for HCI design. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI’09*, 1331–1340.
- Schinkel, N., Pawelzik, K. R., & Ernst, U. A. (2005). Robust integration and detection of noisy contours in a probabilistic neural model. *Neurocomputing*, 65–66, 211–217.
- Shapley, R., & Lennie, P. (1985). Spatial frequency

analysis in the visual system. *Annual Review of Neuroscience*, 8, 547–583.

Watt, R., Ledgeway, T., & Dakin, S. C. (2008). Families of models for Gabor paths demonstrate the importance of spatial adjacency. *Journal of Vision*, 8(7):23, 1–19, <http://www.journalofvision.org/content/8/7/23>, doi:10.1167/8.7.23. [PubMed] [Article]

Yen, S.-C., & Finkel, L. H. (1998). Extraction of perceptually salient contours by striate cortical networks. *Vision Research*, 38(5), 719–741.