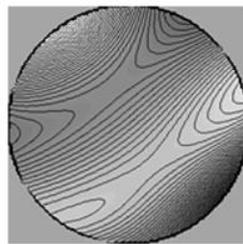


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



VISION RESEARCH

An International Journal for Functional Aspects of Vision



Biochemistry & Cell Biology • Molecular Biology & Genetics
Anatomy, Physiology, Pathology & Pharmacology • Optics, Accommodation & Refractive Error
Circuitry & Pathways • Psychophysics • Perception • Attention & Cognition
Computational Vision • Eye Movements & Visuomotor Control



ISSN 0042-6989 | Volume 47 | Number 18 | August 2007

This article was published in an Elsevier journal. The attached copy is furnished to the author for non-commercial research and education use, including for instruction at the author's institution, sharing with colleagues and providing to institution administration.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Detecting low shape-frequencies in smooth and jagged contours

Nicolaas Prins^{a,*}, Frederick A.A. Kingdom^b, Anthony Hayes^c

^a *Department of Psychology, University of Mississippi, Oxford, MS 38677, USA*

^b *McGill Vision Research Unit, McGill University, Montréal, QC, Canada*

^c *School of Psychology, University College Dublin, Dublin 4, Ireland*

Received 21 November 2006; received in revised form 6 June 2007

Abstract

It is often assumed that a two-stage filter-process is involved in the coding of visual contours. In the first stage the contour is coded by localized luminance filters selective for, among other dimensions, orientation and spatial frequency. A second stage then integrates the outputs of these local luminance filters over space. In the experiments described here we address the issue of which spatial scales of early luminance filters are involved in the detection of a contour's deviation from linearity ('co-linearity failure'), especially at low contour frequencies, for both smooth contours and jagged edges. We also address the question whether it is the orientation or position of the first-stage luminance filters that is used by the second stage. We report two main conclusions. Firstly, across a wide range of shape frequencies, we find that detection thresholds are relatively independent of the spatial scale of the luminance information present in the contour, indicating that detection of co-linearity failure can be effectively mediated by luminance filters tuned to a range of spatial scales. Secondly, we find that detection of co-linearity failure in low shape-frequency contours is primarily based on the local positions, not orientations, of the first-stage luminance filters. As our results suggest that the contour's local orientation may nevertheless play a role, we hypothesize that the local orientation of the contour is not signaled by luminance filters directly but rather by second-order filters acting on the local positions of the contour.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Curvature; Contour; Fractal edge; Shape frequency; Edge; Shape; Psychophysics

1. Introduction

Marr (1982) was not the first to recognize the importance of the defining boundaries of objects and to address seriously how human vision codes those contours: Attneave (1954) famously raised the issue some thirty years before Marr. Marr was, however, perhaps the first to pose the question, as part of a grand theory of visual computation, of what mechanisms underlie the detection of contour shape, and he was the first to recognize fully that a complete model of contour coding needed to address the role of spatial scale. In recent years these issues have prompted considerable research effort. It is often assumed that a two-

stage process is involved in the coding of contours. In the first stage, the contour is coded by localized luminance filters selective for, among other dimensions, orientation and spatial frequency (e.g., DeValois & DeValois, 1988). A second stage then integrates the outputs of these local luminance filters over space. Important questions are, *which* luminance filters are employed and *how* they are integrated.

These questions have particular relevance for edges that are jagged, such as those in the Mandelbrot image shown in Fig. 1. As with many natural contours, the contour in this image contains different information at different spatial scales. For example, the pattern of orientations of the coarsest scale image in Fig. 1 represent the global shape of the Mandelbrot edge, but this pattern is likely to have imprecise positional information. The opposite is true for the finest scale image: the local position of the edge at each point is coded precisely, but the local orientations do not

* Corresponding author. Fax: +1 662 915 5398.
E-mail address: nprins@olemiss.edu (N. Prins).

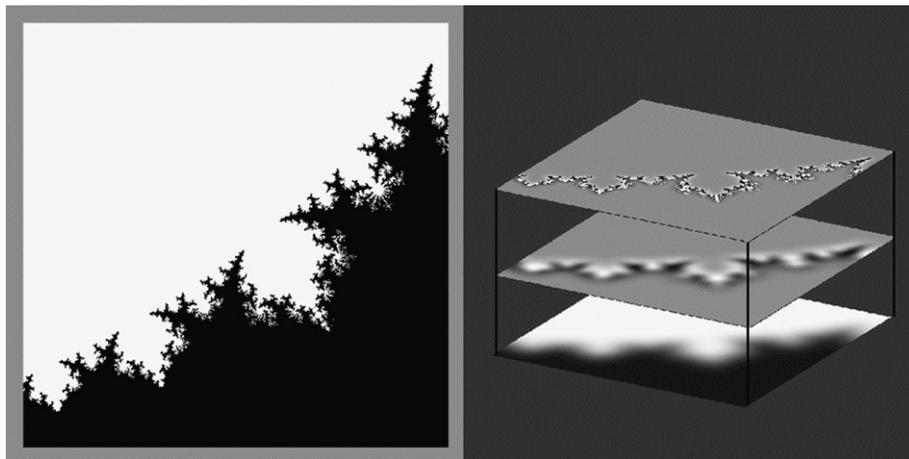


Fig. 1. Mandelbrot fractal edge. The image shown on the left is filtered at three different scales; the results are shown on the right. See text and Field et al. (1993), p. 175, for details. Reprinted from *Vision Research*, 33, Field, D.J., Hayes, A. & Hess, R.F., Contour integration by the human visual system: Evidence for a local “association field”, pp. 173–193, Pergamon Press Ltd., 1993, with permission from Elsevier.

correspond to those of the global edge shape. The question therefore arises; in coding the global position and orientation of a contour, does the visual system use the outputs of coarse-scale luminance filters, which capture the orientations of the global edge shape but whose positional information is imprecise, or does it use the outputs of fine-scale luminance filters whose orientations poorly represent the global edge shape yet whose positional information is precise? A closely related question is which information, position or orientation, is employed for global shape analysis? The aim of this study is to answer these two questions, for both smooth and jagged contours.

With regard to the first question, namely which spatial scale(s) of luminance filters are involved in global contour-shape processing, the studies by Wilson (1985) and Wilson and Richards (1989), and more recently Gheorghiu and Kingdom (2006) are, to our knowledge, unique in addressing the issue. Wilson measured curvature discrimination thresholds for line contours, and found that low-frequency grating masks had little or no effect on discrimination thresholds, whereas high-frequency masks resulted in elevated discrimination thresholds. Wilson and Richards confirmed these findings using filtered line contours. They, like Wilson, measured curvature discrimination thresholds for line contours, and found that high-pass filtering the lines had little effect on discrimination thresholds, especially at low degrees of curvature, whereas low-pass filtering degraded performance at all curvatures. Both these studies suggest that the processing of curvature relies primarily on high-frequency luminance-filters. Gheorghiu and Kingdom considered the question of the luminance spatial-frequency inputs to contour-shape processing in the context of the ‘shape-frequency after-effect’, or SFAE (Kingdom & Prins, 2005). The SFAE is the apparent shift in shape frequency of a sinusoidally shaped contour following adaptation to a contour of slightly different shape frequency (analogous to the well-known luminance frequency after-effect demonstrated by

Blakemore and Sutton (1969)) and appears to be mediated by mechanisms sensitive to local curvature (Gheorghiu & Kingdom, 2007). Gheorghiu and Kingdom (2006) found that the SFAE was both tuned to luminance spatial frequency as well as supported by a range of luminance spatial frequencies. Importantly, however, they found that in contours that were broadband in luminance spatial frequency, fine luminance scales contributed more to the SFAE than coarse luminance scales, in keeping with the results of Wilson and colleagues.

Lines are not strong stimuli for low luminance spatial-frequency filters, because even though they are broadband in luminance spatial-frequency, lines only cover a fraction of the filter’s receptive field, even when oriented appropriately. An edge, on the other hand, is a stronger stimulus for an appropriately oriented low-frequency filter, and thus low-frequency luminance filters may play a role in the processing of low shape-frequency edges. Moreover, with jagged edges such as the Mandelbrot edge, the onus for the visual system to use low-frequency luminance filters might be even greater, for reasons outlined above. For these reasons we considered the effect of luminance scale on contour-shape processing using two quite different classes of contour: smooth lines, which from previous studies appear to implicate predominantly high luminance frequencies, and jagged edges, which as we have argued above are the most likely type of contour to implicate low-frequency luminance filters.

With regard to the second issue, namely whether orientation or position is the feature involved in contour-shape processing, a number of studies are pertinent. The first parametric study on contour-shape processing was arguably that of Tyler (1973), who measured sensitivities to sinusoidally modulated contours over a range of shape frequencies. Tyler found that the sensitivity function for shape frequency is bandpass, with a peak around 3 c/deg and a sharp decline in sensitivity at both lower and higher shape frequencies. For the low-frequency portion of his data,

Tyler concluded that sensitivity was limited by the biggest difference in orientation along the contour, around $20'$, suggesting that orientation information serves as the basic substrate for the processing of contour shapes. Similar conclusions were drawn by Kramer and Fahle (1996), who measured curvature detection thresholds for sinusoidal lines, chevrons and trapezoids, and Wilson and Richards (1989), who measured curvature discrimination thresholds. Using a different paradigm, a similar conclusion was drawn by Blakemore and Over (1974) who found evidence for a curvature after-effect: after prolonged inspection of curved contours, observers perceived a straight line to be curved in the opposite direction to the adapting contours. Importantly, this effect was eliminated when observers were instructed to scan the adapting stimulus so that the local orientation in any given retinal region varied during the adaptation phase, suggesting again that contour shape is processed by comparing local orientations along the contour.

On the other hand, Andrews, Butcher, and Buckley (1973) found that for long lines curvature detection was far superior (in terms of efficiency) than predicted by orientation discrimination between two straight lines. The authors argued that detecting 'co-linearity failure' provides a better explanation for curvature detection than orientation discrimination. This idea was based on their finding that efficiency for curvature detection equaled that for detection of other deviations from linearity (e.g., the bend in chevrons and vernier misalignments). However, no specific mechanism for detection of co-linearity failure was proposed.

Watt and Andrews (1982) proposed that the processing of curvature involves two, or possibly three, separate mechanisms. The first, 'orthoaxial position system' acts on positional information directly, and registers orthoaxial position (i.e., position in a direction orthogonal to the primary axis of the stimulus) with high precision. Whereas such a mechanism is highly efficient, it can only process contours that are nearly straight (i.e., very shallow curvatures) in order to detect, for example, deviations from linearity. This mechanism cannot code the shape of curved lines in order to, for example, mediate performance in a curvature *discrimination* task, nor can it code exact shapes of contours, since both these tasks require position information in two orthogonal directions. The second mechanism acts on a representation of local orientations and their relative positions: processing of curvature proceeds by comparison of local orientation across position. Since a mechanism of this type codes the shape of contours, it can mediate performance in curvature discrimination tasks. Watt and Andrews also discuss the possibility of a third mechanism which operates on short, highly curved lines.

More recently, the context for the issue of orientation versus position has shifted from the detection and discrimination of curves to the detection of non-circularity in radial-frequency patterns (Hess, Wang, & Dakin, 1999; Loffler, Wilson, & Wilkinson, 2003; Wang & Hess, 2005;

Wilkinson, Wilson, & Habak, 1998). For example, Wang and Hess used micropattern-sampled radial-frequency patterns in which either the peaks in curvature, where position but not orientation signals the shape, or the curvature minima, where orientation but not position signals the shape, were sampled. They found that although orientation and position were both employed for detecting radial-frequency patterns, orientation was the more important cue. In summary, whether orientation or position subserves contour-shape processing, there appears to be conflicting evidence, especially for shallow curvatures.

2. Rationale and general method

In the following three experiments we address the issue of which spatial scales of early luminance filters are involved in contour-shape processing, especially the detection of low shape frequencies, and we investigate whether it is the luminance filters' orientation or position information that is used by the visual system. The task of the observers in all experiments was to detect the presence of a modulation of position and/or orientation in a contour using a 2-IFC paradigm. As the individual experiments were performed in different laboratories, the methods vary somewhat between the different experiments, and they will be discussed individually for each of the experiments.

3. Experiment 1

In Experiment 1 we addressed the issue of luminance scale for smooth contours. Observers were required to detect a sinusoidal modulation of position in a line with a Gaussian cross-sectional luminance profile of various standard deviations (σ). Increasing σ progressively eliminates the higher luminance spatial frequencies present in the contour so, following Wilson and Richards (1989), we expect detection thresholds to rise with increasing σ .

3.1. Method

3.1.1. Stimuli

The stimuli consisted of lines whose vertical position was modulated sinusoidally:

$$y(x) = A \cdot \sin(2\pi xf + \varphi) \quad (1)$$

where $y(x)$ is the vertical displacement of the line as a function of horizontal position (x), A is the modulation amplitude, f is the spatial frequency of the modulation, and φ is the phase of the modulation which was randomized for each stimulus. Six shape frequencies (f) were used: 1, 2, 4, 8, 16, and 32 cycles per screen (1 c/screen = 0.12 c/deg). The cross-sectional luminance profile of the lines was a negative-polarity Gaussian with one of six values of standard deviation σ : 0.5', 1.0', 2.0', 4.0', 8.0', and 11.0'. Edges were rotated by a random angle of ± 3 deg and jittered in the vertical direction by a random value between

± 3.5 min, such that absolute local orientation and position could not be used as a cue. Michelson contrast was 0.8, with a mean luminance of 40 cd m^{-2} . Viewing distance was 2 m.

3.1.2. Procedure

On each trial two lines were presented, one of which contained a modulation. Each line was presented at full contrast for 680 ms, and was bracketed by a half-Gaussian contrast ramp, each of which lasted 160 ms. Between stimuli a blank screen was presented for 100 ms. Observers (authors T.H. and F.K.) were required to indicate through a button press whether the first or second line contained the modulation. Feedback was provided in the form of a tone following an incorrect response. The modulation amplitude (A) was varied by an adaptive (one-up and two-down) staircase.

3.2. Results and discussion

Threshold modulation amplitudes (A) for each staircase were calculated as the average amplitude of the last eight reversals of a total of 10 reversals. Average thresholds for

each condition were calculated as the geometric mean of individual staircase thresholds, and are plotted in Fig. 2a as a function of σ , and in Fig. 2b as a function of shape frequency. In Table 1, thresholds for the smallest blur condition are also presented as the maximum difference in orientation in the contour as well as the maximum curvature along the contour. Consider Fig. 2a. Blurring the line, surprisingly, had little effect on detection thresholds. That is, thresholds are relatively constant as a function of line blur, at least for σ 's up to about $4'$. This finding suggests that line-shape processing is not mediated only by fine scale luminance filters, as Wilson and Richards (1989) had concluded. Furthermore, there appears to be little or no interaction between blur and shape frequency. This lack of interaction is most evident in Fig. 2a, where it can be seen that the effect of blur on thresholds is small but, more importantly, remarkably constant across the different shape frequencies (and thus different thresholds as expressed as maximum curvatures along the contour, see Table 1). This finding argues against a model proposed in the introduction in which low curvatures are processed by low-frequency luminance filters, and high curvatures by high-frequency luminance filters, since this model would

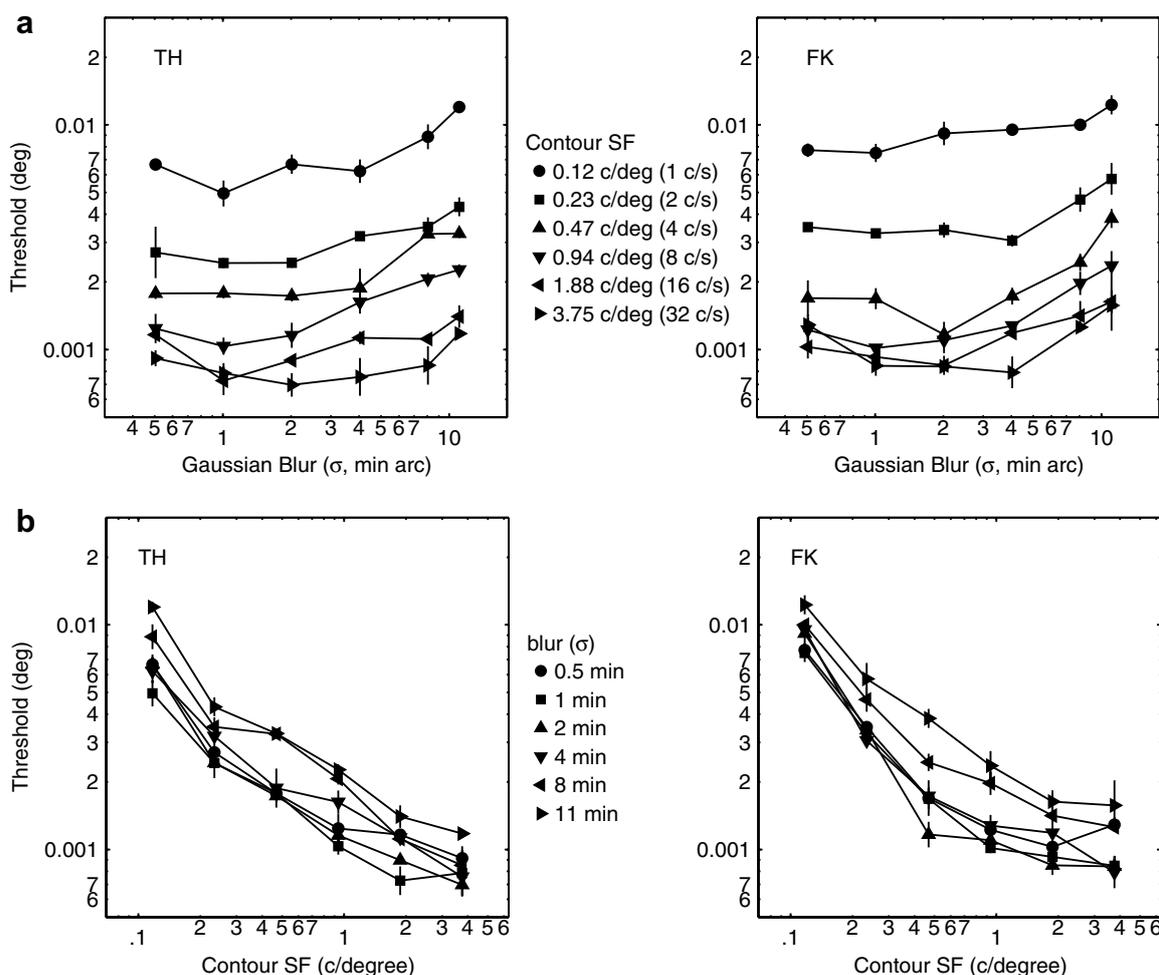


Fig. 2. Experiment 1 results. Threshold modulation amplitudes are plotted as a function of blur (a) and spatial frequency (b) of the contour.

Table 1
Detection thresholds in Experiment 1 expressed as the range of orientations in the stimulus ($\Delta\theta$; min), and as maximum curvature along the edge (κ_{\max} ; deg^{-1}) in the lowest blur condition, for both observers

Shape SF (c/deg)	$\Delta\theta$		κ_{\max}	
	TH	FK	TH	FK
0.12	33.5	38.8	0.004	0.004
0.23	27.2	35.2	0.006	0.008
0.47	35.7	34.0	0.015	0.015
0.93	49.9	49.1	0.043	0.043
1.86	93.4	82.5	0.162	0.143
3.72	146.8	207.4	0.508	0.717

predict an interaction between blur and shape frequency, with sharper increases in thresholds as a function of blur for higher shape frequencies. Instead, the model that is supported by the data is one in which all scales of luminance filter are capable of signaling curvature information with more-or-less equal precision. We will return in a later section to a discussion of this issue.

Our observed relationship between shape frequency and detection threshold (Fig. 2b) closely mirrors that obtained by Jeffrey, Wang, and Birch (2002) in the context of radial-frequency detection in a circular contour. That is, thresholds initially decrease with shape frequency but flatten out for shape frequencies above approximately 2 c/deg.

What do the results of this experiment reveal about the second of the two issues addressed by this study, namely whether position or orientation mediates contour-shape processing? As discussed earlier, Tyler (1973) argued from his data that thresholds for detecting sinusoidal lines were determined by the maximum orientation difference in the waveform. For a sinusoid, the maximum orientation difference is between the orientations at the d.c. positions, and the difference is given by $2 \tan^{-1}(2\pi Af)$, where A and f are amplitude and spatial frequency. Table 1 shows maximum orientation differences for the lowest σ edges at detection threshold. Between shape frequencies of 0.12 and 0.45 c/deg this measure is more or less constant at around 30'–40', after which it rises sharply. The values are close to the 20'–30' reported by Tyler for sinusoidal lines over the same range of shape frequencies. The small difference between the results may be due to the fact that, unlike with our data, Tyler's data were collected using the method of adjustment with unlimited viewing time. As such, however, the data from this experiment are consistent with orientation being the feature mediating contour-shape detection. We will return to this issue in the general discussion.

Also listed in Table 1 are detection thresholds expressed as the maximum curvature (κ_{\max} ; for a sinusoid κ_{\max} is given by $\kappa_{\max} = 4\pi^2 f^2 A$, where f and A are as defined above) along the contour, from which it is apparent that higher contour frequencies require greater curvature before the modulation is detected. Thus, it appears unlikely that task performance is mediated by a mechanism which signals curvature *per se*.

4. Experiment 2

In Experiment 2, we investigated the issue of the response scale of the “front-end luminance filters” for contour-shape detection using jagged edges, in which observers were required to detect the fundamental shape harmonic. The edge was either presented unfiltered, high-pass filtered (eliminating low spatial-frequency content) or low-pass filtered (eliminating high spatial-frequency content). Furthermore, we investigated whether local position or orientation serves as the cue to detect curvature by varying the relative amplitudes of the lower- and higher-frequency shape harmonics. The lower-frequency harmonics chiefly affect the relative local orthoaxial position of the edge while having a relatively small effect on the local orientation of the edge. The higher-frequency harmonics mainly affect local orientation while having a relatively small effect on local orthoaxial position of the edge.

4.1. Method

4.1.1. Stimuli

The edge profiles consisted of a fundamental shape frequency plus 64 shape harmonics,

$$y(x) = A_F \times \sin(2\pi x f_F + \varphi_F) + \sum_{N=2}^{65} [A_H \times N^{-\omega} \times \sin(2\pi x f_F N + \varphi_N)] \quad (2)$$

where $y(x)$ is the vertical displacement of the edge as a function of horizontal position (x), A_F is the amplitude of the fundamental, f_F is the spatial frequency of the fundamental (1 cycle per screen, 0.12 c/deg), φ_F is the phase of the fundamental which was randomized for each stimulus, A_H dictates the amplitudes of the harmonics, ω is the fractal slope which dictates the distribution of energy among the harmonics, and φ_N is the phase of the N th harmonic which was random for all harmonics. The value of A_H covaried with the fractal slope (ω) in order to keep the total amount of energy, proportional to, $\sum_{N=2}^{65} (A_H \times N^{-\omega})^2$, constant. The values of the fractal slope (ω) and A_H that were used are given in Table 2. A condition in which only the fundamental frequency was presented (i.e., $A_H = 0$) was included as a control. Edges were rotated by a random angle between ± 4 deg and jittered in the vertical direction by a random value between

Table 2
Values of the fractal slope and base amplitude of harmonics (A_H) used in Experiment 2

Fractal slope (ω)	A_H (deg)
No harmonics	0
-1	2.79×10^{-4}
-0.5	1.85×10^{-3}
0	1.07×10^{-2}
0.5	0.44
1	1.08
2	2.98

± 0.32 deg, such that the local orientation and position of the edge could not be used as a cue. A simple algorithm was used to achieve subpixel resolution.

The (unfiltered) luminance edge is defined by a step function and thus contains luminance energy at all frequencies. The resulting edges were then either high-pass filtered, low-pass filtered, or not filtered. The filter profile used to create low-pass filtered edges is given by,

$$F = e^{-\left(\frac{f}{5.48}\right)^4} \quad (3)$$

where f is spatial frequency in c/deg. This filter has a 50% cut-off at 5 c/deg. The cut-off value of 5 c/deg was based on the consideration that the human contrast sensitivity function peaks around this value and is sensitive to frequencies both above and below this value. The filter profile used to create high-pass filtered edges is given by the complement of the above filter. Filtering was conducted in the Fourier domain. Michelson contrast for the unfiltered and low-pass filtered edges was set to 50%. In order to increase the visibility of the high-pass filtered edges, luminance values for these edges were rescaled to result in $\sim 100\%$ Michelson contrast. Mean luminance was 145 cd m^{-2} . Example stimuli are shown in Fig. 3.

4.1.2. Procedure

On each trial two edges were presented sequentially; one contained the fundamental shape frequency plus the

harmonics, the other contained only harmonics (i.e., $A_F = 0$). Each edge was presented at full contrast for 167 ms, bracketed by linear contrast ramps, each of which lasted 167 ms. Between stimuli a blank screen was presented for 500 ms. Observers (authors N.P. and F.K., and one naïve observer, L.H.) indicated through a button press whether the first or the second edge contained the fundamental shape frequency. Feedback was provided in the form of a tone following an incorrect response. Each block of trials consisted of three adaptive staircases, interleaved in random order (best PEST, Pentland, 1980) of 40 trials each, in which the modulation amplitude [A_F] of the fundamental was varied. Each block of trials contained edges from one fractal slope and filtering condition only. The different types of stimuli were run in quasi-random order. Stimuli were generated on-line, and stimulus presentation was controlled by a Cambridge Research VSG2/5 graphics board. The resolution of the monitor was set to 800×600 pixels (96 pixels/deg). Viewing distance was 230 cm.

4.2. Results and conclusion

Trials from all blocks in any given condition were combined and fit with a logistic function using a maximum likelihood criterion. Standard errors were obtained using a bootstrap procedure ($N = 400$, e.g., Efron & Tibshirani,



Fig. 3. Experiment 2 example stimuli. (A) Fundamental frequency only, unfiltered. (B) Fundamental plus harmonics, fractal slope = 2, unfiltered. (C) Fundamental plus harmonics, fractal slope = 1, unfiltered. (D) Fundamental plus harmonics, fractal slope = -1, unfiltered. (E) Fundamental plus harmonics, fractal slope = -1, low-pass filtered. (F) Fundamental plus harmonics, fractal slope = -1, high-pass filtered. For the purpose of comparison, the amplitude and phase of the fundamental is identical in each of the examples.

1986). Detection thresholds and standard errors are given in Fig. 4. Each threshold was based on at least 240 trials.

Two of the three observers (F.K. and L.H.) displayed higher thresholds for the high-pass filtered stimuli compared to low-pass and unfiltered stimuli when the fractal slope was shallow. However, due to its jaggedness, the unfiltered edge contains relatively little energy at intermediate spatial frequencies in the shallow fractal slope conditions. As a result, the RMS contrast of these contours, when high-pass filtered, was very low and these two observers reported that on a significant proportion of trials the contour was not perceived in its entirety. No other consistent differences were found between the different filtering conditions.

Since neither the removal of high frequencies, nor the removal of the complementary low frequencies appreciably affects detection performance, these results suggest that the processing of the global shape of jagged edges is not limited by any one scale of luminance filter, and in this respect the results are consistent with the results of Experiment 1. We argue that the results, contrary to what we predicted, do not provide compelling evidence that low luminance spatial frequencies are necessary for the processing of jagged edges. It appears that the mechanism underlying the detec-

tion of low-frequency edge shapes can accept input from either fine- or coarse-scale luminance filters.

Inspection of Fig. 4 reveals that detection thresholds for the fundamental frequency are elevated by the presence of irrelevant harmonics. Low-frequency harmonics are more detrimental than high-frequency harmonics, as evidenced by the increase in thresholds with increasing fractal slope. Since the low-frequency harmonics primarily affect the local position of the edge, but leave the local orientations relatively unaffected, the results of Experiment 2 suggest that the detection of low shape frequencies is based on the local orthoaxial position of the edge. Since no consistent threshold differences are observed between the different filtering conditions, it can be concluded that the position of the contour may be processed by fine- or coarse-scale luminance filters with approximately equal precision.

In Experiment 1, our evidence was consistent with Tyler's (1973) finding that for shallow-curvature detection performance is limited by the maximum orientation difference in the contour, consistent with orientation being the feature underlying contour-shape processing. The results of Experiment 2, however, do not support this conclusion, at least in so far as it implicates luminance filters (see Sec-

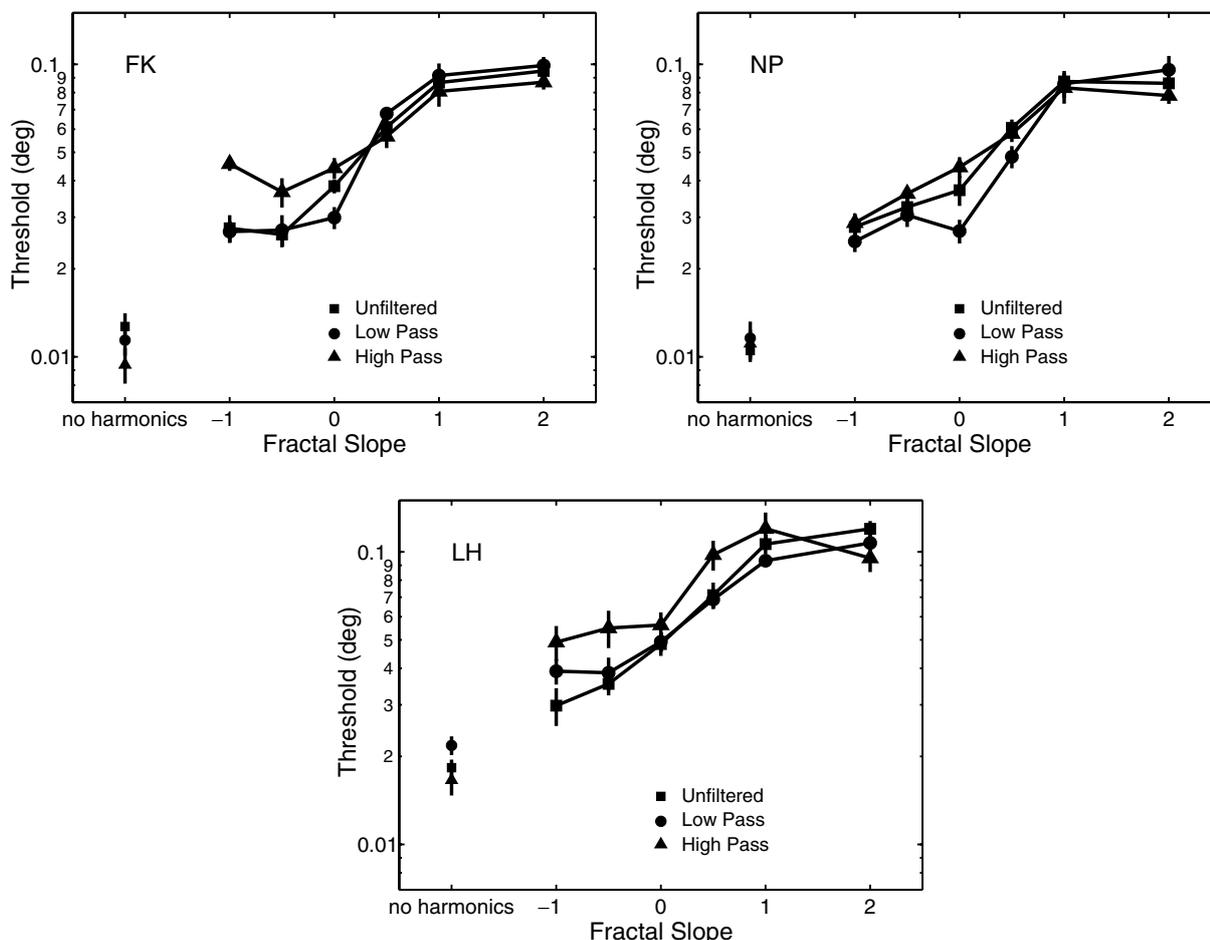


Fig. 4. Experiment 2 results. Thresholds are plotted as a function of fractal slope for each filter condition.

tion 6). It seems highly implausible that when fractal slopes are shallow, the low-frequency fundamental is detected on the basis of local luminance orientation, since local luminance orientation will rarely be even close to the overall orientation of the contour (see Fig 3C–F). Whereas one might argue that local luminance orientation could be encoded by a relatively coarse-scale luminance filter in the unfiltered and low-pass filtered conditions, these filters would not be able to determine the orientation of the edge in the high-pass filtered condition, which contains no energy at low spatial frequencies.

On the other hand, it is possible that the introduction of noise in the form of higher-frequency shape harmonics incapacitated a mechanism acting on the local orientation of the contour, and forced the visual system to use an alternative mechanism based on orthoaxial position. To test this interpretation, we conducted an experiment to decouple local orientation and local orthoaxial position information in a contour presented without the addition of noise harmonics.

5. Experiment 3

The stimuli for Experiment 3 were contours constructed of strings of Gabor micropatterns. This type of construction allowed us to manipulate local orthoaxial position and local orientation independently. The task of the observers was again to detect a sinusoidal modulation of the contour. In separate conditions, the modulation was defined either by, (a) both the orientation and position of the micropatterns, (b) orientation only, or (c) position only. If the low-frequency shape modulation is detected by comparing local positions along the contour, performance should be identical for the condition in which the modulation is defined by both position and orientation and the condition in which the modulation is defined solely by position. Performance for the orientation only condition, however, should be worse.

5.1. Method

5.1.1. Stimuli

Contours were defined by the orientation and/or position of 36 Gabor micropatterns. The luminance profile of the Gabor micropatterns was defined as follows:

$$L(x, y) = M + A \times \cos(2\pi f_M (x \sin(\theta) + y \cos(\theta)) + \varphi_M) \times e^{-0.5 \left(\frac{x^2 + y^2}{\sigma^2} \right)} \quad (4)$$

where M is mean luminance (38.6 cd m^{-2}), A is luminance modulation amplitude of the sine component (18.1 cd m^{-2}), f_M is spatial frequency (8 c/deg), θ is orientation, φ is the phase of the cosine component ($\varphi_M = \pi/2$ or $3\pi/2$) and σ is the spatial constant of the Gaussian envelope ($\sigma = 0.049 \text{ deg}$, such that the spatial-frequency bandwidth of the Gabor micropatterns was 1.5 octaves).

The 36 Gabor micropatterns were spaced equally along the entire width of the screen (10.9 deg) to form a contour. Contours were rotated by a random angle between $\pm 5 \text{ deg}$ and jittered in the vertical direction by a random value between $\pm 0.14 \text{ deg}$, such that local orientation and position of the contour could not be used as a cue.

In the position and orientation condition, the vertical position of the micropatterns was given by

$$y = A \times \sin(2\pi f_C x + \varphi_C) \quad (5)$$

and the orientation of each micropattern was tangential to the contour defined by Eq. (5),

$$\theta = \tan^{-1} (A \times 2\pi f_C \times \cos [2\pi f_C x + \varphi_C]) \quad (6)$$

In the orientation-only condition, the vertical position of the micropatterns was constant, but the orientation of the micropatterns was as in Eq. (6). In the position-only condition, the vertical position of each micropattern was as in Eq. (5), but the orientation of each micropattern was either randomized individually across 360 deg or fixed at horizontal.

Two values of contour spatial-frequency (f_C) were used: 1 cycle per screen (0.092 c/deg) or 8 cycles per screen (0.74 c/deg). The phase of the orientation and/or position modulation (φ_C) was randomized individually for each stimulus. The phase of each micropattern (φ_M) in a contour was either identical for all micropatterns in the contour (either $\pi/3$ or $2\pi/3$, randomly determined for each contour) or was randomly set to $\pi/3$ or $2\pi/3$ for each micropattern individually. Example stimuli are shown in Fig. 5.

5.1.2. Procedure

On each trial two contours were presented sequentially for a duration of 500 ms each separated by a blank screen of 250 ms. One of the two contours contained a modulation of the position and/or orientation of the micropatterns. The other contour was unmodulated (i.e., A in Eqs. (5) and (6) was set to zero). The observers (authors N.P. and F.K. and a naive participant, K.D.) indicated by key press the interval containing the modulation. Feedback was provided in the form of a tone following an incorrect response. Modulation amplitude was varied using an adaptive staircase (best PEST, Pentland, 1980) in blocks of 50 trials. Each block of trials contained only one type of contour modulation. The different types of contour modulation were run in quasi-random order. Stimuli were generated on-line and stimulus presentation was controlled by a Cambridge Research VSG2/5 graphics board. The resolution of the monitor was set to 800×600 pixels (73 pixels/deg). Viewing distance was 200 cm.

5.2. Results and discussion

Trials from all blocks in any given condition were combined and collectively fit with a logistic function using a

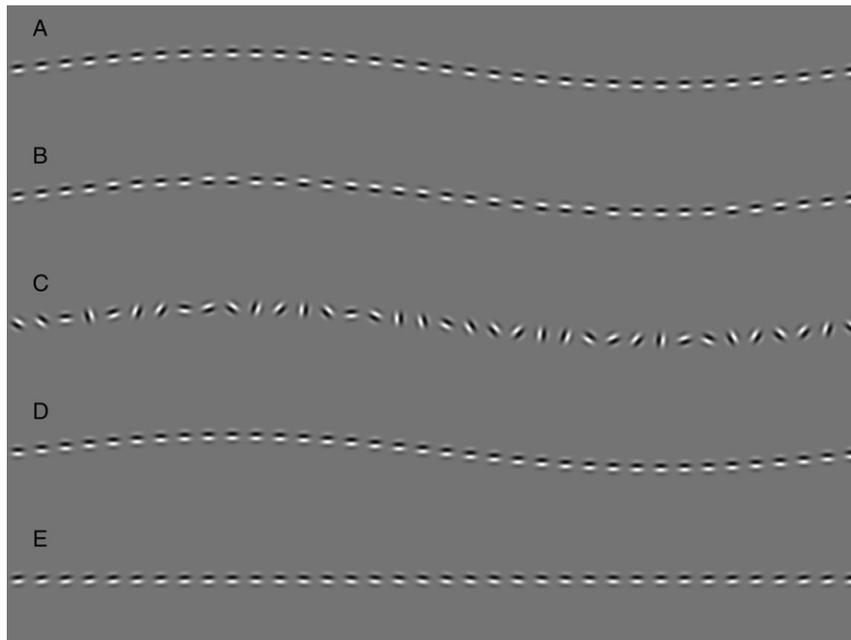


Fig. 5. Experiment 3 example stimuli. (A) Modulation is defined by the position and orientation of micropatterns. (B) As in (A), but phases of the carrier of the Gabor micropatterns are randomized. (C) Modulation is defined by positions of micropattern only, orientations are random. (D) Modulation is defined by positions of micropattern only, orientations are fixed at horizontal. (E) Modulation is defined by orientations of micropatterns only, positions are fixed.

maximum likelihood criterion. Standard errors were obtained using a bootstrap procedure ($N = 400$, e.g., Efron & Tibshirani, 1986). Detection thresholds and standard errors are given in Fig. 6. Each threshold is based on at least 150 trials.

Inspection of Fig. 6 reveals that, for the low-frequency contour (1 cycle per screen, 0.09 c/deg), detection thresholds are approximately equal in those conditions where micropattern position was available as a cue ('position' and 'both') but were much higher in conditions in which only micropattern orientation was the cue ('orientation'). These findings provide strong evidence that shape detection for low-frequency shape modulations is based on a comparison of local positions along the shape contour, and that local luminance orientations by themselves are not only poor cues to the presence of the modulation, but that local orientation also does not seem to aid detection when added to an existing position cue. No consistent differences between conditions are evident for the high-frequency contour. Detection of the contour modulation in these conditions can thus be based on either local position or local orientation. It should come as no surprise that at higher shape frequencies (i.e., greater modulation of local orientation) orientation becomes the more reliable cue, since for a given modulation amplitude, the range of orientations in the contour increased with increasing shape frequency.

Thresholds were higher when the carrier phases of the Gabor micropatterns were randomized in both the low and high contour frequency conditions, especially when the position of the micropatterns was consistent with the contour modulation (i.e., in the position and position-and-orientation conditions). This finding is interesting

and somewhat surprising since alternating micropattern phases has been shown to have no effect on contour integration *per se* (Kuai & Yu, 2006). This finding suggests that local positions are more easily integrated between luminance filters of same polarity or phase. It is not entirely clear why this might be, but, to speculate, it may suggest that position information is more easily integrated between luminance filters of like phase. This arrangement would make sense from an ecological perspective: in natural scenes, whereas luminance polarity would rarely reverse sign along a contour (see Field, Hayes, & Hess, 2000). An alternative explanation is that the phase of the carrier of the micropattern affects its perceived location. Randomizing the phase of the micropattern carriers would then, in effect, introduce position noise and would thus increase thresholds for a mechanism acting on (perceived) position (see also, Hayes, 2000). Whitaker, McGraw, Keeble, and Skilken (2004) indeed found evidence that a Gabor's carrier phase affects its perceived location, although it should be noted that Whitaker et al. observed no bias when the micropatterns were odd-symmetric, as were all of our micropatterns.

Finally, in the condition where the contour was defined solely by the position of micropatterns, thresholds were markedly higher when the orientations of the micropatterns were randomized (as in Fig. 5C) relative to the condition in which the orientations of the micropatterns were fixed at horizontal (as in Fig. 5D). When one considers that a Gabor micropattern can be more precisely localized along the axis perpendicular to its orientation compared to the axis parallel to its orientation (since the sine compo-

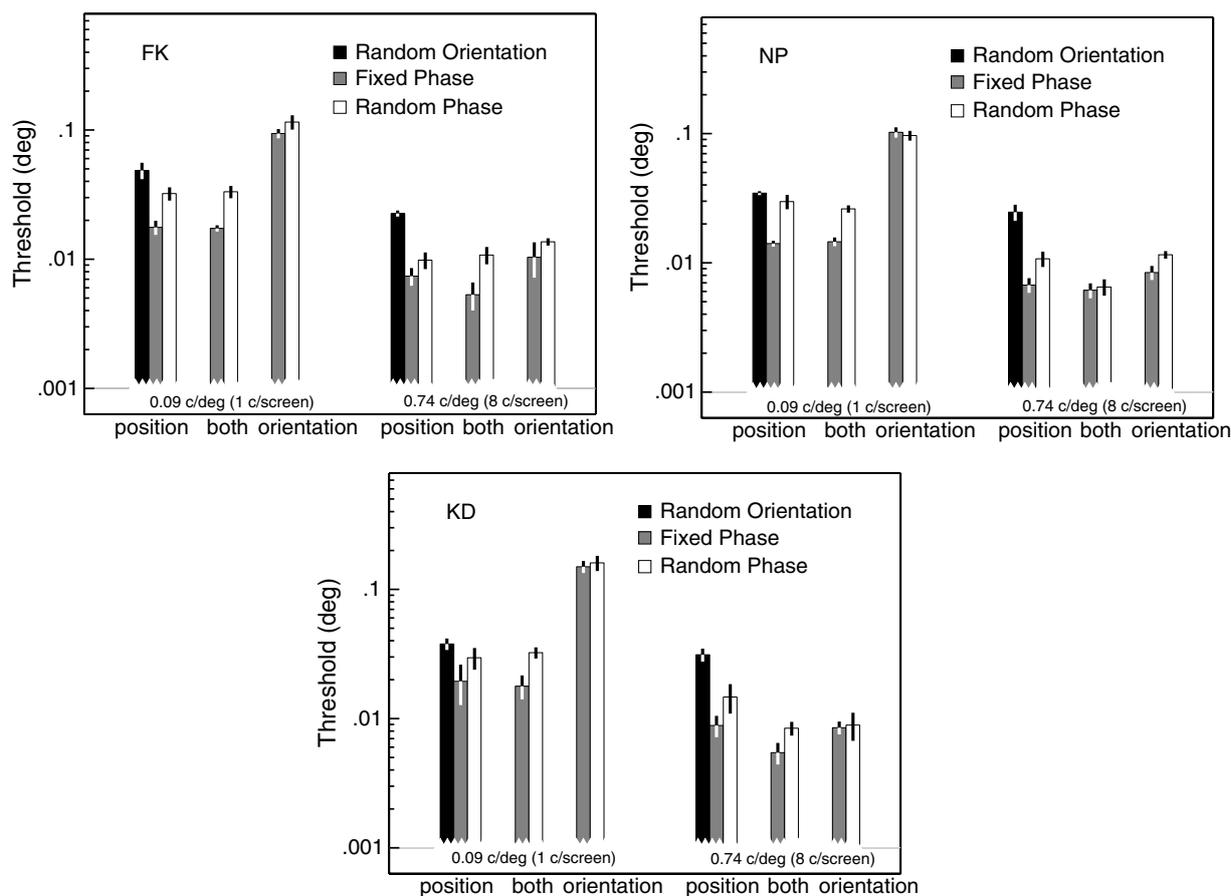


Fig. 6. Experiment 3 results. Thresholds are plotted for each condition and observer. In the ‘position only’ condition, the orientations of the Gabor micropatterns were either randomized (black bars) or fixed at horizontal (gray and white bars).

ment of the micropattern contains positional information only along the axis perpendicular to its orientation), this finding is consistent with the conclusion that shape modulation detection is indeed based on local position information.

It is possible that the lack of effectiveness of local orientation in our low shape-frequency condition is due to the relatively broad orientation bandwidth of our micropatterns. That is, our micropatterns contain relatively little orientation information. Indeed, Wang and Hess (2005) found that the bandwidth of micropatterns affected whether position or orientation information was utilized to detect radial-frequency modulation in a circular contour. Thus, in a control experiment, suggested by an anonymous reviewer, we replicated the main findings using micropatterns with narrower orientation bandwidths. In the control experiment, we varied the orientation bandwidth of our micropatterns by varying the spatial constant of the Gaussian contrast envelope (σ in (1) above). The values of σ used were 0.049 (as in the main experiment), 0.082, and 0.098 deg (example micropatterns are shown in Fig. 7). The orientation bandwidth of our narrowest bandwidth micropatterns is well below the orientation bandwidths used by Wang and Hess. Contours were defined by 36, 27, or 18 micropatterns, respectively. Phases of the micro-

patterns were identical within any given contour. Author N.P. and two naïve observers (A.H. and B.B.) participated in the control experiment. The results, shown in Fig. 7, replicate the finding from the main experiment that thresholds do not differ systematically between the position-only and position-and-orientation conditions, while the thresholds in the orientation-only condition are again significantly higher.

A second control experiment was carried out to rule out the possibility that orientation-selective luminance filters directly code for the orientation of the contour even when the orientations of the individual micropatterns are not consistent with the orientation of the contour. This could occur if orientation-selective luminance filters straddled two neighboring micropatterns. To examine this possibility, we considered the Fourier transform of contour fragments consisting of pairs of micropatterns. In Fig. 8A two collinear micropatterns, like those used to create our contours, are shown. Fig. 8B displays the amplitude spectrum of the Fourier transform of the pair. The orientation of the bands visible in the Fourier transform corresponds to the orientation defined by the relative position of the micropatterns (i.e., the local orientation of the contour). For example, in Fig. 8C two micropatterns are shown that form a contour fragment with an orientation of 10 deg (rel-

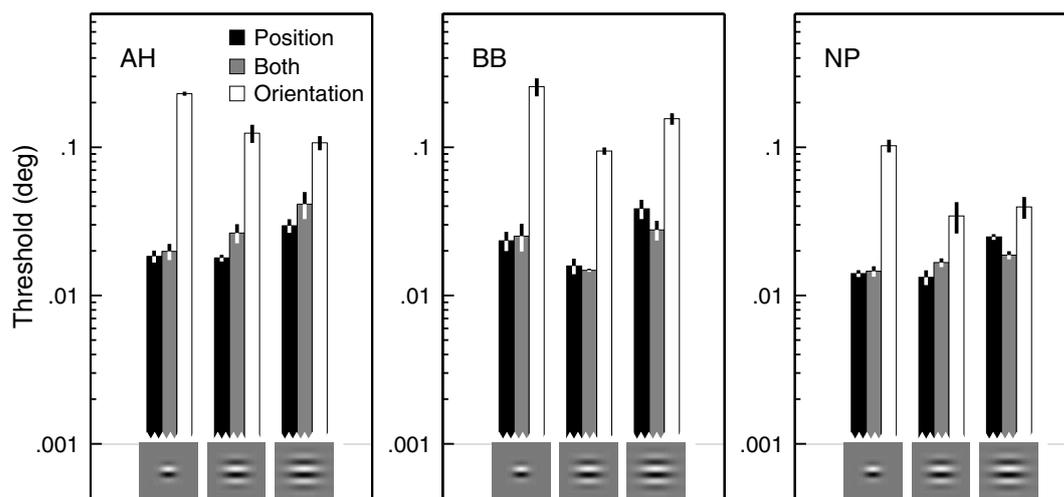


Fig. 7. Results of control experiment in which Gabor micropatterns of varying orientation bandwidth were utilized.

ative to horizontal, the orientation of the micropatterns themselves). In the Fourier transform (Fig. 8D), it can be seen that the orientation of the bands has changed as well. The bands enclosed by the gray ellipse show mutual radial symmetry and, hence, their orientation could conceivably be signaled by an orientation-selective V1 cell (only radially symmetric Fourier amplitude profiles correspond to [real-valued] Euclidian filter profiles). In order for a luminance filter to be selective for the orientation of these bands, and hence the orientation of the contour, its (Fourier) response profile may not overlap with more than a single band in the Fourier space of Fig. 8B and D. In other words, the orientation bandwidth of such a filter may not exceed the orientation bandwidths of these bands. The bands in Fig. 8B and D have an orientation bandwidth of approximately 15 deg (full-width at half-height). Based on psychophysical results, it is believed that the lower bandwidth limit of orientation-selective channels in humans is around 15 deg, a value which corresponds to the narrowest bandwidths measured physiologically in monkeys (DeValois & DeValois, 1988). Thus, in principle, it is possible that the curvature in our contours constructed from micropatterns could be detected by comparing local orientation computed from luminance filters.

To rule out the above possibility, we conducted a control experiment in which the spacing between micropatterns in the contour was doubled, as in Fig. 8E. As is shown in Fig. 8F, doubling the spacing has the effect of narrowing the orientation bandwidth of the oriented bands in the Fourier domain. The bands in Fourier space in Fig. 8F have an orientation bandwidth of approximately 6 deg only, a value below the orientation bandwidth limit of the human visual system. Author N.P. and four naïve observers, (A.H., B.B., G.P., and K.D.) were tested in two conditions. In one condition shape modulation was defined by the orientation and the position of the micropatterns, in the other condition shape modulation was defined by the position of the micropatterns only. All Gabor

micropatterns in a given contour had identical carrier phases. Thresholds and standard errors are presented in Table 3. All thresholds are based on at least 450 trials. No consistent differences between the position-only and the position-and-orientation conditions were found. The results from this control experiment thus strengthen the conclusion that shape modulation detection for low shape frequencies is based on a comparison of local luminance position, not luminance orientation, of the contour.

6. General discussion

The two main conclusions of this study are (i) the detection of contour shape modulations can be effectively mediated by luminance filters of a range of scales and (ii) whereas the detection of high-frequency shapes can be effectively mediated by either local position or local orientation, detection of low-frequency shapes is based primarily on the local positions, not orientations of contours.

The mechanism that detects deviation from linearity in low-frequency shapes apparently can accept information about the orthoaxial position of the contour from fine- or coarse-scale early luminance filters (e.g., V1 simple cells). We found that blurring a smooth edge contour had little or no effect on curvature detection performance, showing that detection performance was not limited by the finest scale luminance filter that will respond to the contour. In Experiment 2, shape modulation detection thresholds were very similar for unfiltered, high-pass filtered or low-pass filtered jagged edges, suggesting that detection performance is equally accurate whether mediated by fine- or coarse-scale luminance filters.

This result stands in apparent contrast to those of Wilson (1985) and Wilson and Richards (1989), who found that curvature discrimination performance for low curvatures was critically dependent on the availability of high-spatial luminance frequencies in the contour. In their proposed model, shallow curvatures are processed by high-fre-

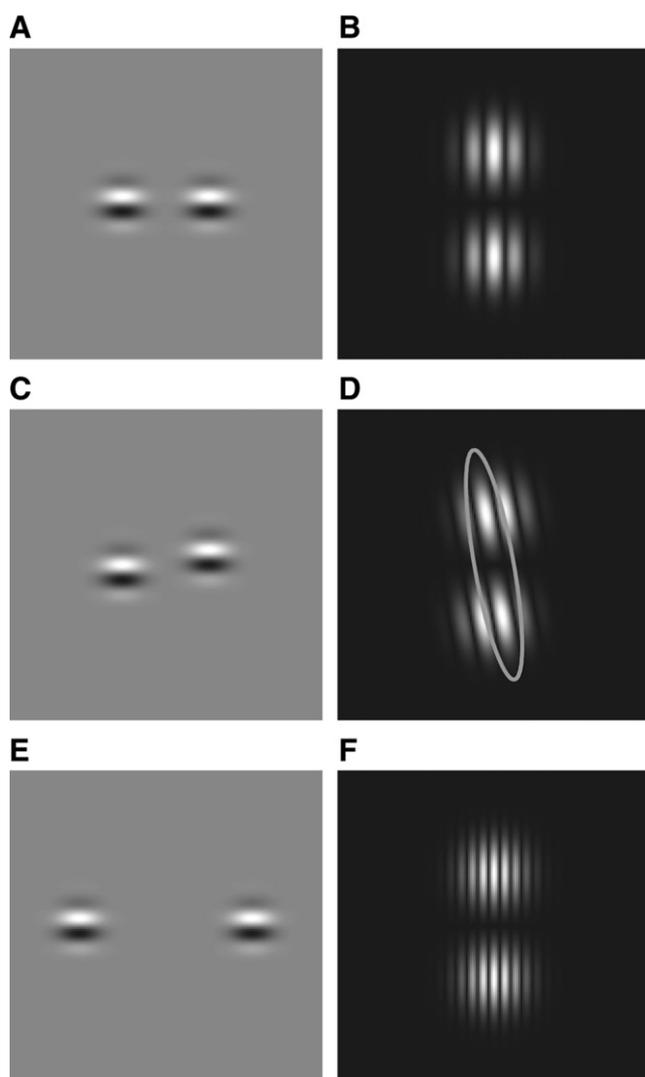


Fig. 8. Contour fragments with their Fourier amplitude spectrums. (A) Fragment of an unmodulated contour as used in Experiment 3. (B) Fragment of a modulated contour in which the modulation is defined by position of the Gabor micropatterns only (i.e., both micropatterns are horizontal). (C) Fragment of a modulated contour where spacing between micropatterns is doubled relative to the spacing in (A) and (B).

Table 3
Results of control experiment using greater separation between Gabor micropatterns

	Position only	Position and orientation
N.P.	1.58 ± 0.09	1.48 ± 0.08
A.H.	2.19 ± 0.30	2.18 ± 0.24
B.B.	1.51 ± 0.18	1.83 ± 0.26
G.P.	1.78 ± 0.19	1.82 ± 0.16
K.D.	1.50 ± 0.23	1.68 ± 0.15

Thresholds and standard errors (in min) are given for the position-only and position-and-orientation conditions for author N.P. and four naïve observers.

quency filters displaced along the contour and selective for slightly different orientations (Fig. 9). Such a mechanism essentially compares the orientation of the contour at

points along the contour. However, Wilson and Wilson and Richards utilized a curvature discrimination task, in which all stimuli had a much higher degree of curvature compared to ours. Using the measure of curvature given in Wilson, curvatures in our contours at threshold amplitude were between one and two orders of magnitude smaller than the lowest curvatures utilized by Wilson, and Wilson and Richards.

A similar mechanism to that proposed by Wilson and Richards (1989) involving interactions between V1 cells was later proposed by Field, Hayes, and Hess (1993) to underlie the detection of contours made up of oriented Gabor micropatterns embedded in a field of randomly oriented Gabor micropatterns. Field et al. showed that detection performance of such fragmented contours depended not only on the position but also on the orientation of the elements: when the elements were jittered in their orientation relative to the path, detection suffered. Consistent with this psychophysical finding is the physiological evidence for orientation-dependent lateral interactions between V1 orientation-selective cells as found by Bosking, Zhang, Schofield, and Fitzpatrick (1997) in the tree shrew. However, it must be borne in mind that detecting a contour in noise is a very different task from detecting the shape of a contour that is highly visible. It is possible that local orientation is important for one but not the other.

In Experiment 2, we found that adding noise-harmonics that mostly affected local orientation, but left local position information comparatively intact, (i.e., the small fractal slopes conditions, Fig. 3C–F) had a relatively small effect on detection of the low-frequency fundamental. However, adding noise harmonics which mostly affected local position but left local orientation relatively unaffected (i.e., the steep fractal slope conditions, Fig. 3B) had a large detrimental effect on performance. This effect held even when the contour was high-pass filtered. It seems unlikely that in the jagged high-pass filtered, low-fractal-slope conditions, the orientation information available to early (luminance) filters could be integrated in order to detect the shape modulation of the fundamental. In Experiment 3, we confirmed this conclusion using contours made of strings of Gabor micropatterns, allowing us to vary contour position and contour orientation independently. The contour modulation was defined either by both the position and orientation, only the position, or only the orientation of the

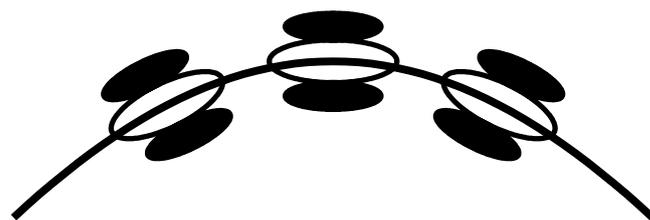


Fig. 9. Schematic diagram of curvature processing mechanism as proposed by Wilson (1985) and Wilson and Richards (1989).

micropatterns. The pattern of results differed between low shape-frequency contours and high shape-frequency contours. For high shape-frequency contours, detection performance was approximately identical whether the modulation was defined by the position of the micropatterns, the orientation of the micropatterns, or both, suggesting that either of the cues may serve as input to the mechanism which detects the modulation. However, for low shape frequencies it was found that, whereas detection thresholds were equal for the 'position and orientation' condition and the 'position only' condition, detection thresholds were much higher in the condition in which the modulation was defined by the orientation of the micropatterns only. Thus, detection of low shape frequencies appears to be mediated by mechanisms acting on the local position of the contour.

Our finding that low shape frequencies may be processed with either high- or low-frequency luminance filters, and is based on the local position of contours and not orientation, is consistent with the orthoaxial position system proposed by Watt and Andrews (1982), and suggested to be involved also in the detection of other deviations from linearity, such as the bend in chevrons and vernier acuity. We suggest that this mechanism accepts information from fine- or coarse-scale luminance filters. The manner in which this mechanism assesses the shape of the contour based on the local positions cannot be determined based on our data and remains a matter of speculation. It is, in fact, possible that the micropattern positions are integrated by 'second-order' filters, whose orientations are then used for contour-shape processing. This would explain why the modulation detection threshold for low shape frequencies in Tyler (1973) and in our Experiment 1 appears to be determined by the range of local orientations present in the contour. In both experiments the range of local orientations was perfectly correlated with the range of orientations present in the contour available to the hypothesized second-order mechanism which acts on position information. In other words, it might be that it is not the orientation of the luminance filters that is utilized, but rather the orientation of the contour *per se*.

Acknowledgments

This research was supported by a grant from the Natural Sciences and Engineering Research Council of Canada (NSERC) awarded to Frederick Kingdom (Grant No. OGP 0121713).

References

- Andrews, D. P., Butcher, A. K., & Buckley, B. R. (1973). Acuities for spatial arrangement in line figures: Human and ideal observers compared. *Vision Research*, *13*, 599–620.
- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, *61*, 183–193.
- Blakemore, C., & Over, R. (1974). Curvature detectors in human vision? *Perception*, *3*, 3–7.
- Blakemore, C., & Sutton, P. (1969). Size adaptation: A new aftereffect. *Science*, *166*, 245–247.
- Bosking, W. H., Zhang, Y., Schofield, B., & Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in Tree Shrew striate cortex. *Journal of Neuroscience*, *17*, 2112–2127.
- DeValois, R. L., & DeValois, K. K. (1988). *Spatial vision*. New York: Oxford University Press.
- Efron, B., & Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, *1*, 54–75.
- 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*, 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*, 51–66.
- Gheorghiu, E., & Kingdom, F. A. A. (2006). Luminance contrast properties of contour shape processing revealed through the shape-frequency after-effect. *Vision Research*, *46*, 3603–3615.
- Gheorghiu, E., & Kingdom, F. A. A. (2007). The spatial features underlying the shape-frequency and shape-amplitude after-effects. *Vision Research*, *47*, 834–844.
- Hayes, A. (2000). Apparent position governs contour-element binding by the visual system. *Proceedings of the Royal Society Series B*, *267*, 1341–1345.
- Hess, R. F., Wang, Y. Z., & Dakin, S. J. (1999). Are judgments of circularity local or global? *Vision Research*, *39*, 4354–4360.
- Jeffrey, B. G., Wang, Y.-Z., & Birch, E. E. (2002). Circular contour frequency in shape discrimination. *Vision Research*, *42*, 2773–2779.
- Kingdom, F. A. A., & Prins, N. (2005). Different mechanisms encode the shapes of contours and contour-textures. *Journal of Vision*, *5*, 463.
- Kramer, D., & Fahle, M. (1996). A simple mechanism for detecting low curvatures. *Vision Research*, *36*, 1411–1419.
- Kuai, S.-G., & Yu, C. (2006). Constant contour integration in peripheral vision for stimuli with good Gestalt properties. *Journal of Vision*, *6*, 1412–1420.
- Loffler, G., Wilson, H. R., & Wilkinson, F. (2003). Local and global contributions to shape discrimination. *Vision Research*, *43*, 519–530.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. New York: WH Freeman and Company.
- Pentland, A. (1980). Maximum likelihood estimation: The best PEST. *Perception & Psychophysics*, *28*, 377–379.
- Tyler, C. W. (1973). Periodic Vernier acuity. *Journal of Physiology*, *228*, 637–647.
- Wang, Y.-Z., & Hess, R. F. (2005). Contributions of local orientation and position features to shape integration. *Vision Research*, *45*, 1375–1383.
- Watt, R. J., & Andrews, D. P. (1982). Contour curvature analysis: Hyperacuities in the discrimination of detailed shape. *Vision Research*, *22*, 449–460.
- Whitaker, D., McGraw, P. V., Keeble, D. R. T., & Skillen, J. (2004). Pulling the other one: 1st- and 2nd-order visual information interact to determine perceived location. *Vision Research*, *44*, 279–286.
- Wilkinson, F., Wilson, H. R., & Habak, C. (1998). Detection and recognition of radial frequency patterns. *Vision Research*, *38*, 3555–3568.
- Wilson, H. R. (1985). Discrimination of contour curvature: Data and theory. *Journal of the Optical Society of America A*, *2*, 1191–1199.
- Wilson, H. R., & Richards, W. A. (1989). Mechanisms of curvature discrimination. *Journal of the Optical Society of America A*, *6*, 106–115.