PERGAMON

Vision Research 40 (2000) 1077–1087

Vision Research

www.elsevier.com/locate/visres

Luminance spatial frequency differences facilitate the segmentation of superimposed textures

Frederick A.A. Kingdom [a,*], David R.T. Keeble [b]

[a] *Department of Ophthalmology, McGill Vision Research Unit, 687 Pine Avenue West, Rm. H4-14, Montréal, Que., Canada H3A 1A1*
[b] *Department of Optometry, University of Bradford, Richmond Rd., Bradford BD7 1DP, UK*

Received 26 February 1999; received in revised form 20 October 1999

Abstract

Do superimposed textures segregate on the basis of a difference in their luminance spatial frequency? We addressed this question using orientation-gratings, which consist of dense arrays of Gabor micropatterns whose orientations vary sinusoidally across space. Two orientation gratings of the same texture spatial frequency were combined in anti-phase, to produce a 'dual-modulation' orientation grating. Thresholds for detecting the dual-modulation gratings were measured as a function of the difference in Gabor spatial frequency between the two grating components. When the two components were made from the same Gabors, thresholds were relatively high. However a one octave difference in Gabor spatial frequency between the components caused thresholds to fall close to those of single-modulation orientation gratings. The fall in threshold was accompanied by a change in appearance of the stimulus; to that of two transparent, interwoven, flow patterns. We show that these results are incompatible with current Filter–Rectify–Filter models of 'second-order' pattern detection. Rather, they favour the idea that feature analysis precedes texture analysis, with the visual system encoding local orientation content prior to the texture stage.

Keywords: Segmentation; Texture gradients; Orientation gratings; Spatial frequency

1. Introduction

Two abutting textures often appear to segment without need for scrutiny. Besides abutting textures, however, natural scenes often contain textures that overlap, or are superimposed; for example a field of corn viewed through the leaves of a forest. Watanabe and Cavanagh (1996) recently showed that overlapping textures often appeared as distinct surfaces, with their overlapping parts seen through each other, or in transparency, an effect the authors termed texture 'laciness'. The presence of non-overlapping parts of the stimulus presumably helped to define each texture surface, facilitating the appearance of transparency where those surfaces overlapped. What of two textures that are spatially superimposed however, such that no non-overlapping regions are present to uniquely define each texture component? Do they appear to segregate, and if so, what factors facilitate this segregation?

In this study we have examined whether luminance spatial frequency differences facilitate the segmentation of two, superimposed, orientation-defined textures. Our motivation is 2-fold. First, as we have said, superimposed textures occur in natural scenes, and it is important to know what factors help to segregate them perceptually. Second, a number of studies testify that texture mechanisms are sensitive to luminance spatial frequency, even under conditions when luminance spatial frequency is not itself the basis for texture segmentation (e.g. Sutter, Sperling & Chubb, 1995; Kingdom & Keeble, 1996, 1999). For our study we have employed the 'orientation grating' stimulus used in many previous studies (Zucker, 1983; Kingdom, Keeble & Moulden, 1995; Kingdom & Keeble, 1996, 1999; Gray & Regan, 1998; Kwan & Regan, 1998). An example is shown in Fig. 1a, together with its unmodulated comparison in Fig. 1b. Orientation gratings fall within the rubric of what for many are 'second-order' stimuli, stimuli defined by modulations in something other than

* Corresponding author. Fax: +1-514-8431691.
E-mail address: fred@jiffy.vision.mcgill.ca (F.A.A. Kingdom)

0042-6989/00/$ - see front matter
PII: S0042-6989(99)00233-3

luminance, here in orientation. Our orientation gratings consist of dense arrays of Gabor micropatterns whose orientations vary sinusoidally across space. We refer to the spatial frequency of the Gabor carrier as Gabor, or luminance, spatial frequency, and the spatial frequency of the orientation modulation as texture spatial frequency. The detectability of an orientation grating is typically defined as the amplitude of orientation modulation required to discriminate a modulated (Fig. 1a) from an unmodulated (Fig. 1b) grating at threshold.

In order to use the orientation grating stimulus to study the segregation of overlapping textures, we designed a novel stimulus called a 'dual-modulation' orientation grating. Examples are shown in Figs. 2–4. Each dual-modulation grating consists of *two*, superimposed, orientation gratings, with identical texture spatial frequencies, orientations, and amplitudes, but 180° out-of-phase. The two 'component' gratings can be constructed from the same Gabors, as in Fig. 2, or from different Gabors, as in Fig. 3. Note that superimposing two, opposite-in-phase, orientation gratings constructed from the same Gabors does not result in the *physical* cancellation of the two grating components, as would happen for example if two opposite-in-phase *luminance* gratings were added. This is because the two components in the dual-modulation stimulus are constructed from Gabors whose positions in each component are randomised. Thus each component has its 'own' micropattern structure. Furthermore, even if by chance two Gabors fall on top of one another, their luminance modulations (but not dc levels) are added, so their individual contribution to the structure of each grating component is never lost.

Consider now Figs. 1–3. In the single-modulation grating in Fig. 1a, one sees a vivid 'flow' pattern. This flow pattern is absent in Fig. 2a. Although one can see that Fig. 2a is non-uniform, it appears to vary in orientation *variance* rather than orientation per se. The flow pattern is however restored in the dual-modulation grating in Fig. 3a, where the two grating components are made from Gabors that differ by a factor of two in spatial frequency. In Fig. 3a one sees two textures weaving in and out of one another, giving rise to a sensation of texture transparency, similar to that described by Watanabe and Cavanagh (1996) for their overlapping textures. The question we now ask is whether the differences in appearance between Fig. 1a, Fig. 2a and Fig. 3a are paralleled by a difference in their detectability.

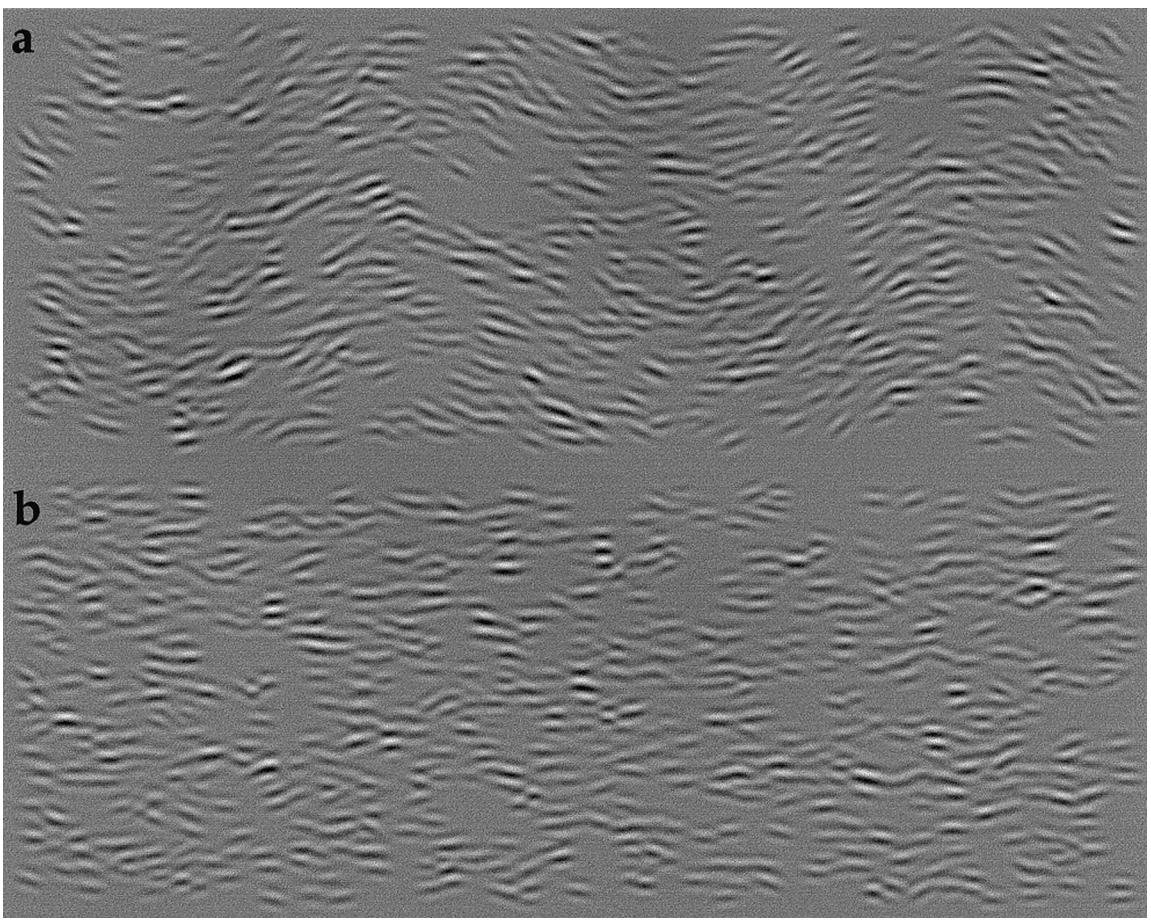

Fig. 1. Orientation grating stimulus. (a) Test stimulus with a texture amplitude of 20°, (b) unmodulated comparison stimulus.

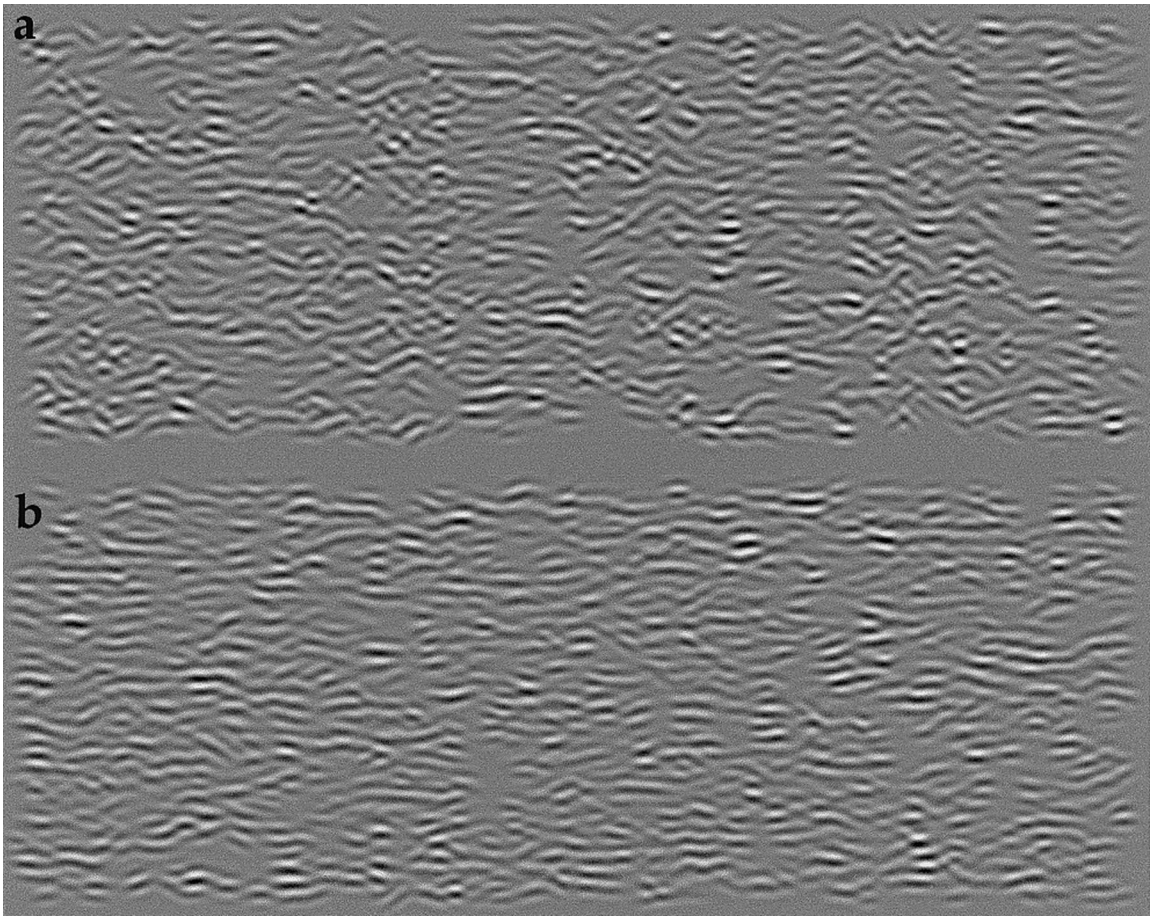

Fig. 2. Dual-modulation orientation grating made from a single spatial frequency of Gabor. (a) Test stimulus constructed from two opposite-in-phase orientation gratings with a texture amplitude of 20°, as in Fig. 1a. (b) Zero amplitude comparison stimulus. Note the absence of the perceived flow pattern in (a).

2. Method

2.1. Subjects

One of the authors FK, and a naive undergraduate volunteer AC, acted as observers. Both had normal vision and were experienced psychophysical observers.

2.2. Stimuli

2.2.1. Generation

Example stimuli are shown in Figs. 1–4. They were generated on a Macintosh Quadra 950 computer and displayed on a SuperMac Trinitron monitor. The displays were all monochrome (black–white) and were gamma corrected by selecting the appropriate intensity levels from an 8-bit (256 grey levels) look-up-table, following calibration using a UDT photometer.

2.2.2. Gabor micropatterns

These were generated by multiplying a sine function by a two-dimensional Gaussian envelope:

$$L(x,y) = M + A\ \cos[2\pi\ f_l(x\ \cos\ \theta - y\ \sin\ \theta)] \times\ \exp[-(x^2+y^2)/2\sigma^2]$$

with M, mean luminance, A, amplitude, f_l luminance spatial frequency, θ, orientation and σ, the standard deviation of the Gaussian envelope, which was circularly symmetric. Their contrast, defined as peak amplitude A divided by the mean M, was 23.6%. M, or in other words the background, was set to 35 cd/m. Details of actual micropattern parameters are given with each experiment.

2.2.3. Orientation gratings

Fig. 1 shows an example of a single-modulation orientation grating; Figs. 2–4 examples of dual-modulation orientation gratings. Each stimulus subtended $11.3 \times 28.2°$ at the viewing distance of 64.7 cm, and contained 1537 micropatterns. The position of each micropattern within the stimulus window was randomised. The orientations of the micropatterns were constrained in that their nominal mean orientations

varied sinusoidally along the horizontal axis of the display for each grating component. The amplitude of orientation modulation, or texture amplitude, determined by how much the orientation of the micropatterns changed throughout one complete cycle of orientation modulation. For example, a texture amplitude of 10° implied that the micropattern orientations changed by 20° throughout one complete cycle of orientation modulation. All micropatterns at a given horizontal position had an orientation drawn from a Gaussian distribution with a specified mean (determined by the point on the waveform) and standard deviation. The standard deviation of the Gaussian distribution in all conditions was 10°, and this represented the amount of orientation 'noise' in the stimulus. Where Gabor micropatterns overlapped, their amplitudes, though not dc levels, were additively combined. Texture spatial frequency was fixed throughout at 0.067 c/deg, producing two complete cycles of orientation modulation across the stimulus. This texture spatial frequency was close to that with peak sensitivity for gratings made from the Gabor micropatterns used here (Kingdom et al., 1995). The phase of orientation modulation was randomised on each stimulus presentation.

2.3. Procedure

A two-interval-forced-choice (2IFC) procedure was used throughout to measure the threshold amplitude of orientation modulation, or 'texture' threshold. On each trial two stimuli were presented, each for 147 ms and separated by a 500 ms inter-stimulus-interval. The task for the subject was to decide which interval contained the stimulus with the orientation modulation. The only difference between the two intervals was in the amplitude of orientation modulation, which in the comparison stimulus was zero (see Fig. 1b, Fig. 2b, Fig. 3b, Fig. 4b). The method of constant stimuli was used with five texture amplitudes for each texture spatial frequency, the magnitudes being determined from pilot studies. Feedback in the form of a tone was given for an incorrect decision.

2.4. Analysis

Weibull functions were fitted to the psychometric functions using a maximum likelihood method. This produced a threshold at the 82% correct level, and a

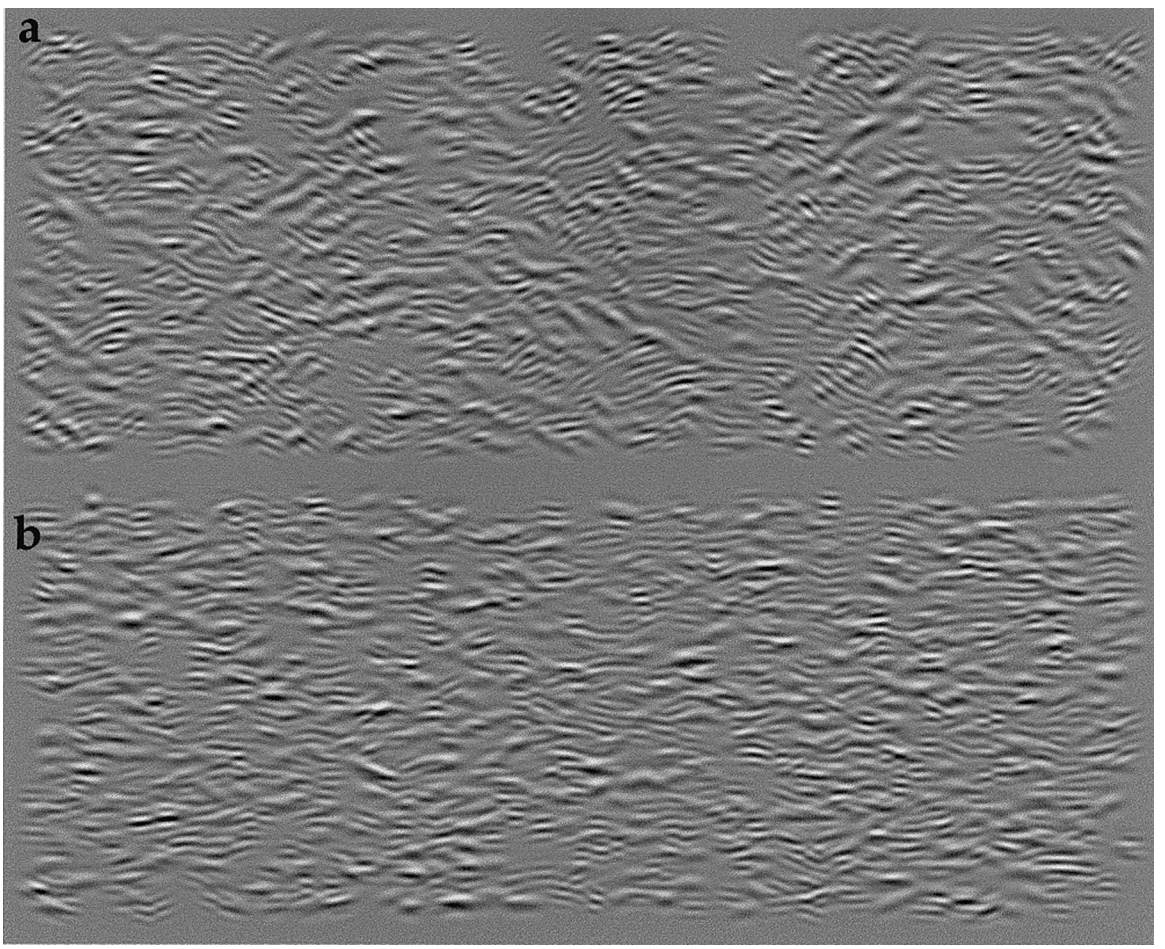

Fig. 3. Dual-modulation grating stimulus in which the two components are made from *different* Gabors a factor two apart in spatial frequency. (a) 20° amplitude test stimulus, (b) zero amplitude comparison stimulus. Note that one can now see two flow patterns in transparency in (a).

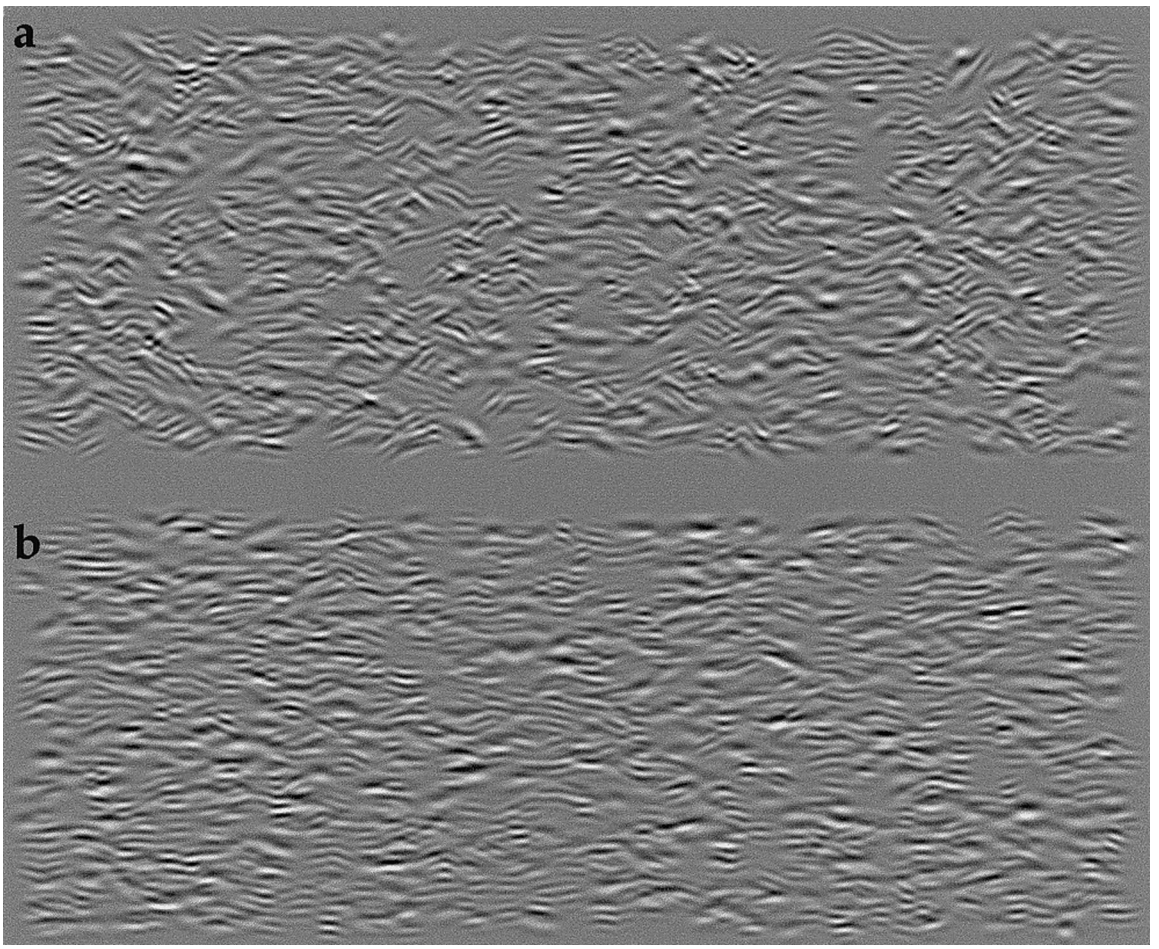

Fig. 4. Dual-modulation grating stimulus in which each component is constructed from both Gabors used in Fig. 3. (a) 20° amplitude test stimulus, (b) zero amplitude comparison stimulus. Note that the perception of a flow pattern is once again missing.

67% confidence interval, the error bar on each data point.

3. Experiments and results

In order to measure the detectability of a dual-modulation grating whose components were made from different Gabors, we set the relative amplitudes of the two grating components to take into account any difference in their individual detectabilities. To do this we first measured thresholds for the single-modulation grating shown in Fig. 1, at a number of Gabor spatial frequencies. In this, as well as the main experiment below, the Gabor standard deviation σ was set to 0.2°. The results are shown in Fig. 5. Thresholds appear to vary only slightly across a 6-fold range of Gabor spatial frequency, roughly between 4° and 7°. There is a hint, particularly in FK's data, that thresholds are lowest around 2–4 c/deg. In order to obtain threshold estimates for Gabor spatial frequencies between those measured, we took the geometric average of the two thresholds on either side. To equate the detectability of the two grating components we set their amplitudes to a constant ratio — the ratio of their respective single-modulation grating thresholds.

For the main experiment, we measured thresholds for dual-modulation gratings as a function of the difference in Gabor spatial frequency between the two grating components. We fixed the Gabor spatial frequency for one grating component, and varied the Gabor spatial

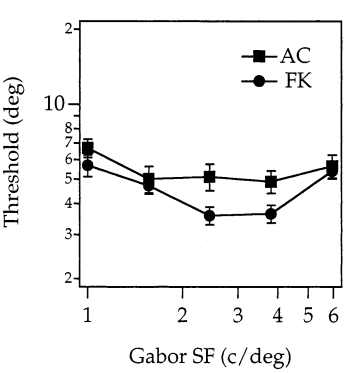

Fig. 5. Effect of Gabor spatial frequency on thresholds for a single-surface grating. Data are shown for two subjects.

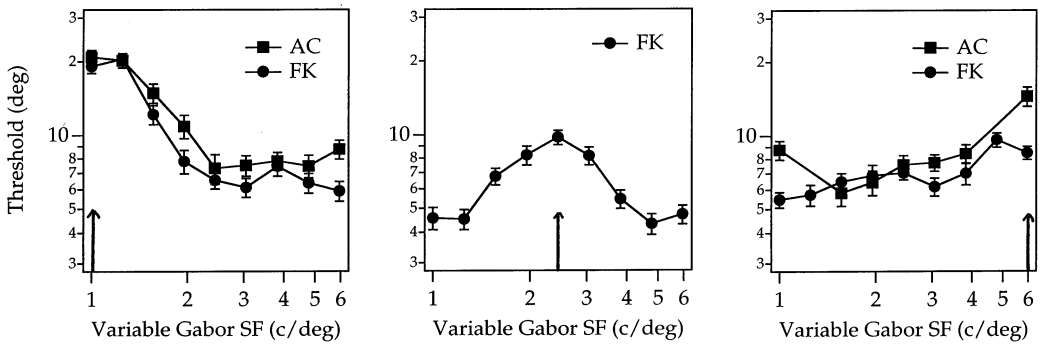

Fig. 6. Thresholds for dual-modulation orientation gratings as a function of the Gabor spatial frequency of one of the component gratings, for three fixed Gabor spatial frequencies: 1.0, 2.44 and 5.96 c/deg. The fixed Gabor spatial frequency is indicated by the arrow on each graph.

frequency of the other. The results are shown in Fig. 6, for three fixed Gabor spatial frequencies: 1.0, 2.44 and 5.96 c/deg. Consider the left-hand graph which shows the data for the fixed 1.0 c/deg Gabor spatial frequency. When the two grating components had the same Gabor spatial frequency (indicated by the arrow), thresholds were around 20°. With the introduction of just over a one octave difference in Gabor spatial frequency however, thresholds fell to near-asymptotic levels of around 7°, a change of almost a factor of three. For the fixed Gabor spatial frequency of 2.44 c/deg (middle graph) thresholds fell by a factor of two for the subject tested, and for the 5.96 c/deg condition, thresholds fell on average by a factor of 1.6. In this last case the data appear somewhat noisy at the ends of the functions, which might render this figure an underestimate of the true value.

3.1. Spatial frequency differences, or feature enrichment?

It is possible that the fall in thresholds following an introduction of a Gabor spatial frequency difference between the two grating components was due to the stimulus becoming richer in features, rather than because segmentation was facilitated. To test this possibility we measured thresholds for dual-modulation orientation gratings composed of two types of Gabors, as before. The two Gabor types were either segregated between the two grating components — the 'segregated' condition, or distributed equally between them — the 'non-segregated' condition. Examples of these two conditions are shown in Figs. 3 and 4. Since both the segregated and non-segregated conditions had the same overall Gabor composition, any difference in their thresholds could not be due to any differences in overall feature richness. For this experiment we used Gabors with spatial frequencies 1.56 and 4.77 c/deg (about a factor of three apart) with a space constant σ of 0.2°. The results are shown in Fig. 7a, along with single-modulation grating thresholds constructed from each type of Gabor alone. In Fig. 7a thresholds for the segregated condition are around a factor of two lower than for the non-segregated condition, and close on average to those of the single-modulation thresholds. The finding that thresholds for segregated dual-modulation gratings were significantly lower than for non-segregated dual-modulation gratings made from the same two types of Gabor, allows us to reject the 'feature-enrichment' hypothesis for the results of the earlier experiment.

3.2. Spatial frequency or bandwidth?

Fig. 7b shows results obtained using Gabors of the same spatial frequency as in Fig. 7a, but this time with

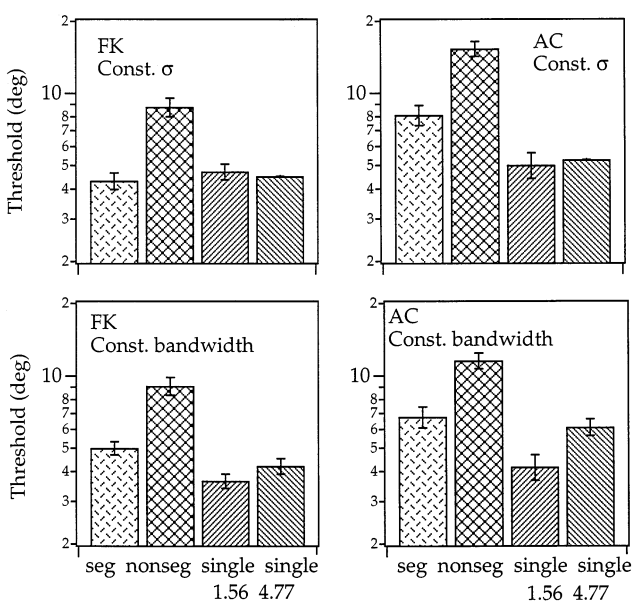

Fig. 7. Thresholds for dual-modulation gratings constructed from Gabors a factor of three apart in spatial frequency (1.56 and 4.77 cpd). (a) Constant Gabor σ condition, (b) constant Gabor bandwidth condition. Seg: segregated condition in which the two Gabors are separated in the different grating components. Nonseg, non-segregated condition in which the two Gabors are present in equal amounts in both grating components. Single, single-modulation orientation grating. The numbers below the single conditions refer to their Gabor spatial frequencies in cpd.

the same bandwidth (same number of carrier cycles) rather than the same envelope size. Equal bandwidths were achieved by setting the envelope standard deviations of the 1.56 and 4.77 c/deg micropatterns to 0.305 and 0.1°. The results in Fig. 7b are similar to those of Fig. 7a. This confirms that the critical parameter underlying the beneficial effects on dual-modulation grating thresholds is luminance spatial frequency rather than bandwidth.

4. Discussion

We have demonstrated that luminance spatial frequency differences facilitate the segmentation of superimposed orientation-defined textures. When two orientation gratings made from Gabors were combined in opposite spatial phase, the normal flow pattern disappeared, and thresholds for detecting the gratings were significantly elevated. However, the introduction of just over a one octave difference in Gabor spatial frequency between the two grating components created the impression of two flow patterns in transparency, and this was paralleled by a fall in thresholds to their original level.

Why were the effects most dramatic in the 1.0 c/deg fixed Gabor spatial frequency condition, and least dramatic in the 5.96 c/deg condition (see Fig. 6)? At present we can only speculate that this is due to the differential strength of the variance-modulation cue in the dual-modulation grating conditions made from a single type of Gabor (Fig. 2). Inspection of Fig. 6 shows that for subject FK, who completed all of the three fixed micropattern spatial frequency conditions, 1.0, 2.44 and 5.96 c/deg, thresholds were respectively 20.0, 10.0 and $\sim 10.0°$. The Gabors in these conditions had a fixed envelope size, and therefore their orientation bandwidth increased with spatial frequency. Orientation-variance discrimination may be better with narrow orientation bandwidths, having the effect of reducing thresholds.

4.1. Relation to other form tasks

The results of this study parallel similar findings with other 'form' tasks. Julesz and Chang (1979) demonstrated that when two mirror-symmetric noise patterns, one symmetric about its vertical axis, the other about its horizontal axis, were superimposed, mirror symmetry was not perceived in either component. However, with a two octave difference in luminance spatial frequency between the components, symmetry perception was achieved. Recently, Kingdom, Ziegler and Hess, (2000) found similar results to those here for detecting the orientation of dual-modulation *disparity* gratings. These findings from symmetry and stereopsis domains, together with those from the present study, reinforce the idea that luminance spatial scale differences are potent in facilitating the perceptual segmentation of overlapping textured surfaces.

Measuring thresholds for dual-modulation gratings is a simple and objective way of revealing the consequences of perceptual segmentation, yet there are undoubtedly many other perceptual consequences. For example, adding random orientation noise to orientation gratings would presumably increase their thresholds and obscure their appearance as a flow pattern, but less so when target and noise differed in luminance spatial frequency. A recent study by Arsenault, Wilkinson and Kingdom (1999) is relevant to this conjecture. They found that thresholds for detecting *frequency*-modulated patterns in masking noise were lowered when the test and mask differed in carrier orientation. Thus, just as an orientation-defined texture can be segmented from another by a difference in luminance spatial frequency (the present result), so a spatial-frequency-defined texture can be segmented from masking noise via a difference in its carrier orientation. Interestingly, Arsenault et al. also found that when their noise mask was *modulated* in spatial frequency, thresholds were elevated even when the mask carrier orientation differed from the test. Further experiments measuring orientation-grating thresholds in the presence of *modulated*, as well as unmodulated, masks of various luminance spatial frequencies, will establish if the analogous result occurs for orientation gratings.

On the other hand our results are at odds with some other paradigms of complex pattern discrimination. Olzak and colleagues (e.g. Olzak & Wickens, 1997; Olzak & Thomas, 1999) have provided evidence that when discriminating subtle differences in orientation in certain types of complex pattern, we tend to pool information across spatial scale, even if doing so is deleterious to performance. In their experiments, subjects were required to discriminate the orientation of luminance grating patches superimposed upon grating patches of various spatial frequencies and orientations. The stimulus and task used by Olzak and colleagues are very different from here and their paradigm may be more relevant to the perceptual properties of features such as edges and bars, rather than textures.

4.2. Relevance to current models of texture discrimination

It is widely believed that 'second-order' patterns are detected by mechanisms which can be approximated as Filter–Rectify–Filter (FRF) models (see Wilson, 1999, for a recent review). There are many types of FRF model, and while they differ somewhat in detail, they all embody the same basic principle. The stimulus is first filtered by a bank of '1st stage' linear filters tuned

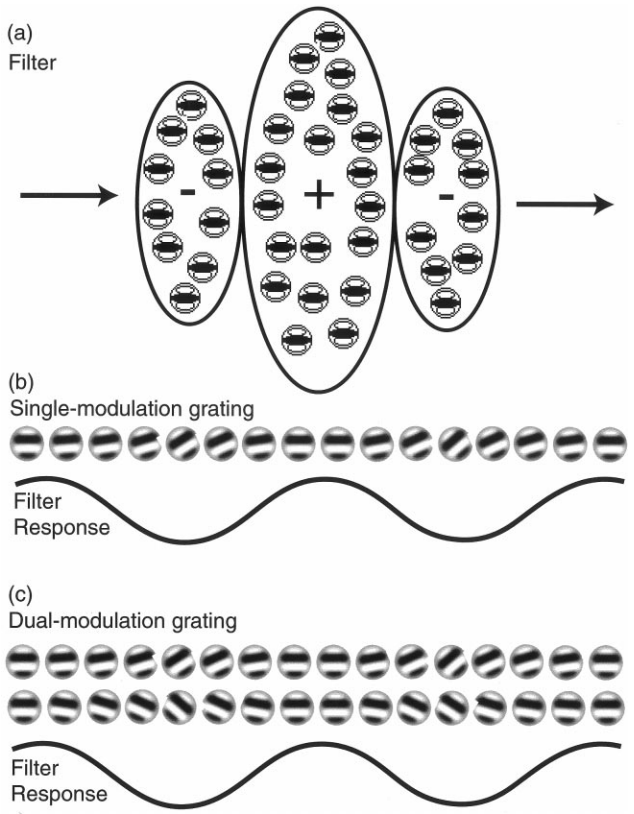

Fig. 8. Complex channel model responses to both single-modulation and dual-modulation orientation gratings made from one spatial frequency of Gabor. The two continuous lines in the figure show schematic responses of the complex channel in (a) to both the single-modulated (b), and dual-modulated (c), orientation gratings shown in Fig. 1aFig. 2a. The stimuli in (b) and (c) are represented as rows of Gabors whose orientations are similar to the mean orientation of the Gabors along the horizontal axis of the actual stimuli. In (c) the two rows of Gabors represent the two grating components.

to various orientations and spatial frequencies; this part of the model approximates a wavelet transform. The outputs of these filters are then subject to a nonlinearity (squaring, full-wave rectification, etc.), that results in filter-response energy maps at each scale/orientation. Finally, '2nd stage' filters, with larger receptive fields than their 1st stage counterparts, filter each scale/orientation energy map. In some FRF models, the 2nd stage filters detect differences in energy *between* two, abutting, texture regions; this type of model has been used to account for the segmentation of 'first-order' or 'Fourier' textures (Fogel & Sagi, 1989; Malik & Perona, 1990; Bergen & Landy, 1991; Landy & Bergen, 1991; Rubenstein & Sagi, 1993). In other FRF models, the 2nd stage filter captures the pattern of energy *within* a texture region; this type of model has been used to account for the segmentation of 'second-order' or 'non-Fourier' textures (Sutter & Graham, 1995; Graham & Sutter, 1998). This second class of FRF model also lends itself naturally to the detection of textures which vary sinusoidally (Sutter et al. 1995, for contrast modulation; Kingdom & Keeble, 1996; Gray & Regan, 1998, for orientation modulation; Arsenault et al. 1999, for spatial frequency modulation). We now consider whether such a model is consistent with our present data. Following Graham et al., we will refer to this type of FRF model as a 'complex channel' model.

Fig. 8 shows the hypothetical convolution responses of a complex channel to both a 'single-modulation' and a 'dual-modulation' orientation grating of appropriate texture spatial frequency, both made from the same spatial frequency of Gabor. The schematic 2nd stage filter in Fig. 8a has both excitatory and inhibitory sub-regions, each receiving inputs from an array of 1st stage, simple-cell-like filters tuned to relatively high spatial frequencies and horizontal orientations. The nonlinearity, e.g. rectification, imposed on the 1st stage filter responses prior to their integration by the 2nd stage filter is not shown for simplicity. The orientation-grating stimuli in (b) and (c) are shown schematically as rows of Gabors whose orientations show the mean orientation at each horizontal location in the stimulus. In Fig. 8b, the complex channel gives a strong modulated response to the single-modulation grating, as illustrated in the response profile below the stimulus. When the centre of the complex channel receptive field falls onto the dc level of the grating, where the Gabors are horizontal, the channel gives a strong positive response, since the non-horizontal Gabors falling in the inhibitory flanks do not elicit such a strong response. On the other hand, when the complex channel falls on the peak of the waveform, it gives a strong negative response, as the preferred horizontal Gabors now fall within its flanks.

Consider now the response of the complex channel to the dual-modulation grating, also made from just one spatial frequency of Gabor (see Fig. 2a), shown schematically in Fig. 8c. Remember that for this stimulus, thresholds were significantly elevated compared to those of the single-modulation grating. Contrary to this result however, the complex channel would give as strong a response to the dual-modulation as to the single-modulation grating. When the complex channel is centred on the dc of the stimulus, its receptive-field centre is again strongly stimulated by the horizontal Gabors, while its flanks are again only weakly stimulated by the non-horizontal Gabors. The net positive response will be of the same magnitude as with the single-modulation grating, as the differential change in orientation is the same in both cases. Complex channels receiving inputs from 1st stage filters with other-than-horizontal preferred orientations would also respond equally well to both single-modulation and dual-modulation orientation gratings. It is therefore difficult to see how the elevated thresholds resulting from the superimposition of two opposite-in-phase orientation gratings made from the same spatial frequency of Gabor can be explained using a conventional complex channel model.

Other varieties of complex channel model aimed specifically at modelling the detection of orientation gratings would fare no better. Kingdom and Keeble (1996) proposed that the 2nd stage filters might receive differently-oriented 1st stage inputs to the centre and surround of their receptive fields, with both sets of inputs being excitatory. The idea was inspired somewhat loosely by neurophysiological data from V1 (Knierim & Van Essen, 1992; Gallant, Van Essen & Nothdurft, 1995) and MT (Olavarria, DeYoe, Knierim, Fox & Van Essen, 1992). These studies had shown that responses to a line falling within a cell's 'classical' receptive field were inhibited when the line was surrounded by lines outside the receptive field. However, when the lines outside the receptive field were orthogonal to that in the centre, the inhibition was minimised (see also Nothdurft, 1997). Gray and Regan (1998) suggested a 2nd stage filter which pooled the outputs of a number of circularly-symmetric versions of the filter shown in Fig. 8a (termed 'double-opponent' neurones), each receiving inputs from differently oriented 1st stage filters. However, both the Kingdom and Keeble (1996) and Gray and Regan (1998) versions of complex channel model would have the same problem as the model illustrated in Fig. 8; they would respond with equal strength to the single-modulation and dual-modulation gratings in Fig. 1a and Fig. 2a. They are therefore inconsistent with the present data.

4.3. What mechanism for orientation-grating detection?

Consider first the value to vision of detecting textural variations in orientation. Orientation gradients occur in the retinal image of any non-planar textured surface, such as one folded in depth, even if the surface itself consists of physically non-oriented features (Gibson, 1979; Cutting & Millard, 1984; Knill, 1998). Such gradients provide a powerful cue to surface shape (Stevens, 1988; Knill, 1998). Although orientation variance is also a salient texture cue (Dakin & Watt, 1997), it provides different surface-shape information; a change in orientation variance is typically associated with a compression gradient, such as occurs when viewing a surface that curves away from the viewer. It would seem sensible to have mechanisms selectively responsive to gradients in average orientation, as well as in orientation variance, for use in surface-shape analysis. A general-purpose FRF mechanism would not be appropriate for surface-shape analysis, because FRF models respond indiscriminately to virtually any type of second-order pattern (e.g. contrast modulated, mean-orientation modulated, orientation-variance modulated, spatial-frequency modulated, etc.). One presumably would not want a contrast-modulated pattern to elicit a response in a shape-from-texture mechanism. Our finding that thresholds for dual-modulation gratings made from a single spatial frequency of Gabor (whose gradients are in orientation variance) are higher than for single-modulation gratings (whose gradients are in mean orientation) is consistent with there being specific mechanisms for detecting variations in particular types of orientation content. This line of reasoning leads us to speculate that for our textures, a stage of feature analysis, in which the orientations of individual elements are explicitly derived, precedes texture analysis. This is not a new idea; it was central to the thinking of Julesz (1981), Beck (1982), and Marr (1982) (see Bergen, 1991 for a review). We are currently conducting experiments to test directly whether this is indeed the case for orientation grating detection.

A second likely feature of orientation-grating detection mechanisms is that they incorporate orientation grouping processes. In Fig. 1a the strong sense of a flow pattern follows the contours defined by fortuitous co-alignments of neighbouring Gabors. Early work by Link and Zucker (1988) and Or and Zucker (1989) suggested that 'linking processes' between neighbouring co-aligned texture elements played an important role in the perception of orientation-defined texture structure, and our own evidence has been consistent with this (e.g. Kingdom et al., 1995). Dakin and Hess (1998) have recently shown that the detection of strings of co-aligned Gabors in noise is easiest when the Gabors have the same spatial frequency. If the fortuitous contours in our orientation gratings are critical for the grating's detection (as opposed to merely necessary for its appearance as a flow pattern), then it is possible that the main result of this study is a consequence of the spatial frequency selectivity of *contour* detection mechanisms. Experiments directly testing the role of contour detection in orientation-grating detection are presently underway.

A third feature of orientation-grating detection mechanisms is that they produce scale-invariance. Previously we measured orientation-grating detection at various *texture*, as well as luminance scales or spatial frequencies (Kingdom et al., 1995; Kingdom & Keeble, 1999). We termed the function relating orientation-grating thresholds to texture spatial frequency the orientation modulation function, or OMF. We found the OMF to be typically bandpass. The OMF shifted rightwards along the texture spatial frequency axis when the luminance spatial frequency of the stimulus was increased, whether the increase was produced by increasing the spatial frequency of the micropatterns, or by increasing viewing distance. This result is important because it shows that the shape of the OMF is 'scale invariant', i.e. invariant with viewing distance when texture spatial frequency is plotted in object units, e.g. cycles/cm. The scale invariant property of the OMF is most readily explained by supposing that orientation-gradient mechanisms are tied in their spatial scale selec-

tivity to the luminance-contrast-sensitive mechanisms from which they receive their inputs. This in turn carries the implication that orientation-gradient mechanisms are tuned to luminance spatial frequency. The present study provides direct evidence of this tuning. A similar conclusion was reached by Sutter et al. (1995) for the mechanisms involved in detecting contrast modulation. Therefore whatever type of operators underlie the detection of orientation gradients, their receptive field sizes most likely scale with those of their luminance spatial frequency inputs. Kingdom et al. (1995) found similar shaped OMFs for orientation gratings made both of broadband-in-luminance and narrowband-in-luminance micropatterns. Dakin (1996) has argued that with broadband-in-frequency orientation-defined textures, there is only a narrow range of spatial frequencies at which the key textural information is most efficiently represented, and thus it is sensible for the visual system to be scale-selective. The results of this study forcefully demonstrate the value of such scale selection.

5. Summary and conclusion

We have shown that luminance spatial frequency differences facilitate the segmentation of superimposed orientation-defined textures. Our results with dual-modulation orientation gratings are not consistent with a general-purpose Filter–Rectify–Filter model. Instead, we argue that the mechanisms for detecting orientation-gratings are (a) specific for detecting orientation gradients, which necessitatates feature analysis prior to the texture stage; (b) involve grouping-by-coalignment processes; (c) are luminance-scale selective; and (d) produce scale invariance.

Acknowledgements

This research was funded by an NSERC (Canada) grant ref: OGP 0121713 given to FK.

References

Arsenault, S. A., Wilkinson, F., & Kingdom, F. A. A. (1999). Modulation frequency and orientation tuning of second-order texture mechanisms. *Journal of the Optical Society of America A*, *16*, 427–435.

Beck, J. (1982). Textural segmentation. In J. Beck, *Organization and representation in perception*. Hillside, NJ: Erlbaum.

Bergen, J. R. (1991). Theories of visual texture perception. In D. Regan, *Vision and visual dysfunction*, vol. 10B. New York: MacMillan.

Bergen, J. R., & Landy, M. S. (1991). Computational modelling of visual texture segregation. In M. S. Landy, & J. A. Movshon, *Computational models of visual processing*. Cambridge, MA: MIT Press.

Cutting, J., & Millard, R. (1984). Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology (General)*, *113*, 198–216.

Dakin, S. C. (1996). The detection of structure in glass patterns: psychophysics & computational models. *Vision Research*, *37*, 2227–2246.

Dakin, S. C., & Hess, R. F. (1998). Spatial frequency tuning of visual contour extraction. *Journal of the Optical Society of America A*, *15*, 1486–1499.

Dakin, S. C., & Watt, R. J. (1997). The computation of orientation statistics from visual texture. *Vision Research*, *37*, 3181–3192.

Fogel, I., & Sagi, D. (1989). Gabor filters as texture discriminators. *Biological Cybernetics*, *61*, 103–113.

Gallant, J. L., Van Essen, D. C., & Nothdurft, H. C. (1995). Two dimensional and three-dimensional texture processing in visual cortex of the macaque monkey. In T. V Papathomas, C Chubb, A. Gorea, & E. Kowler, *Early vision and beyond*. Cambridge, MA: MIT Press.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Hillsdale, NJ: Erlbaum.

Graham, N., & Sutter, A. (1998). Spatial summation in simple (Fourier) and complex (no-Fourier) texture channels. *Vision Research*, *38*, 231–257.

Gray, R., & Regan, D. (1998). Spatial frequency discrimination and detection characteristics for gratings defined by orientation texture. *Vision Research*, *38*, 2601–2617.

Julesz, B. (1981). Textons, the elements of texture perception and their interactions. *Nature*, *290*, 91–97.

Julesz, B., & Chang, J. (1979). Symmetry perception and spatial-frequency channels. *Perception*, *8*, 711–718.

Kingdom, F. A. A., Keeble, D. R. T., & Moulden, B. (1995). Sensitivity to orientation modulation in micropattern-based textures. *Vision Research*, *35*, 79–91.

Kingdom, F. A. A., & Keeble, D. R. T. (1996). A linear systems approach to the detection of both abrupt and smooth spatial variations in orientation-defined textures. *Vision Research*, *36*, 409–420.

Kingdom, F. A. A., & Keeble, D. R. T. (1999). On the mechanism for scale invariance in orientation-defined textures. *Vision Research*, *39*, 1477–1489.

Kingdom, F. A. A., Ziegler, L. R., & Hess, R. F. (2000) Spatial scale facilitates stereoscopic depth segmentation. *Journal of the Optical Society of America A* (submitted).

Knierim, J. J, & Van Essen, D. C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *Journal of Neuroscience*, *67*, 961–980.

Knill, D. C. (1998). Surface orientation from texture. Ideal observers, generic observers and the information content of texture cues. *Vision Research*, *38*, 1655–1682.

Kwan, L., & Regan, D. (1998). Orientation-tuned spatial filters for texture-defined form. *Vision Research*, *38*, 3849–3855.

Landy, M. S., & Bergen, J. R. (1991). Texture segregation and orientation gradient. *Vision Research*, *31*, 679–691.

Link, N. K., & Zucker, S. W. (1988). Corner detection in curvilinear dot grouping. *Biological Cybernetics*, *59*, 247–256.

Malik, J., & Perona, P. (1990). Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America A*, *7*, 923–932.

Marr, D. (1982). *Vision*. Freeman.

Nothdurft, H-C. (1997). Different approaches to the coding of visual segmentation. In M. Jenkins, & L. Harris, *Computational and psychophysical mechanisms of visual coding*. Cambridge University Press.

Olavarria, J. F., DeYoe, E. A., Knierim, J. J., Fox, J. M., & Van Essen, D. C. (1992). Neural responses to visual texture patterns in middle temporal area of the macaque monkey. *Journal of Neuroscience*, *68*, 164–181.

Olzak, L. A., & Wickens, T. D. (1997). Discrimination of complex patterns: orientation information is integrated across spatial scale: spatial frequency and contrast information are not. *Perception*, *26*, 1101–1120.

Olzak, L. A., & Thomas, J. P. (1999). Neural recoding in human pattern vision: model and mechanisms. *Vision Research*, *39*, 231–256.

Or, Y. H., & Zucker, S. W. (1989). Texture fields and texture flows: sensitivity to differences. *Spatial Vision*, *4*, 131–139.

Rubenstein, B. S., & Sagi, D. (1993). Effects of foreground scale in texture discrimination task: performance is size, shape and content specific. *Spatial Vision*, *7*, 293–310.

Stevens, K. A. (1988). The line of curvature constraint and the interpretation of 3D shape from parallel surface contours. In W. Richards, *Natural computation*. MIT Press.

Sutter, A., & Graham, N. (1995). Investigating simple and complex mechanisms in texture segregation using speed-accuracy trade-off method. *Vision Research*, *35*, 2825–2843.

Sutter, A., Sperling, G., & Chubb, C. (1995). Measuring the spatial frequency selectivity of second-order texture mechanisms. *Vision Research*, *35*, 915–924.

Watanabe, T., & Cavanagh, P. (1996). Texture laciness: the texture equivalent of transparency? *Perception*, *25*, 293–303.

Wilson, H. (1999). Non-Fourier cortical processes in texture, form, and motion perception. *Cerebral Cortex*, *13*, 445–477.

Zucker, S. W. (1983). Computational and psychophysical experiments in grouping: early orientation selection. In J. Beck, B. Hope, & A. Rosenfield, *Human and machine vision*. New York: Academic Press.