Orientation- and frequency-modulated textures at low depths of modulation are processed by off-orientation and off-frequency texture mechanisms

Nicolaas Prins *, Frederick A.A. Kingdom

McGill Vision Research, McGill University, Room H4-14, 687 Pine Avenue West, Montral, Que., Canada H3A 1A1

Received 30 August 2001; received in revised form 3 December 2001

Abstract

Intuitively it may seem likely that orientation-modulated (OM) and frequency-modulated (FM) textures are processed utilizing the first-order channels that are most responsive to the first-order (luminance) information contained in the textures. This assumption would imply that the detection or segmentation of OM or FM textures is accomplished by second-order mechanisms that receive their first-order input from neurons tuned to either the center, or to the peaks in the orientation and spatial-frequency distribution of the texture. Here we show that at low depths of modulation this is not the case. Using an adaptation paradigm, we show that the first-order filters involved in the perception of OM and FM textures are those which maximize the *differential* response between the different texture regions. Our explanation of this result is similar to that made by Regan and Beverley [J. Opt. Soc. Am. 73 (1983) 1684; J. Opt. Soc. Am. A 2 (1985) 147] for simple grating stimuli. However, we show that whereas Regan and Beverley's results could be accounted for on the basis of the tuning functions of the putative mechanisms involved, our results can be explained in terms of the characteristics of the textures themselves. Some implications of our finding are discussed. © 2002 Elsevier Science Ltd. All rights reserved.

Keywords: Texture perception; Second-order; Adaptation; Spatial-frequency discrimination

1. Introduction

The detection and segmentation of texture regions differing from surrounding regions only in contrast, orientation, or spatial frequency cannot be accomplished using linear filters alone. Instead, the perception of texture discontinuities is believed to be mediated by filter–rectify–filter (FRF) mechanisms (e.g., Malik & Perona, 1990). The idea behind the FRF mechanism is simple (Fig. 1): the texture is first filtered by a bank of filters selective for the orientation and spatial-frequency of luminance modulations in the texture (e.g., cortical simple cells). The outputs of these filters are then subjected to a non-linearity (typically half- or full-wave rectification or a squaring operation) and consequently integrated by a larger second-order filter. Any discontinuity in orientation, spatial frequency, or contrast in the texture will lead to differential activation of the excitatory and inhibitory regions of the second-order filter and hence will lead to activation of the FRF unit. The inherent flexibility of the FRF mechanism—it will respond to a variety of different types of texture discontinuities—has made it possible to relate it to the perception of texture modulations in contrast (Sutter, Sperling, & Chubb, 1995), orientation (Bergen & Landy, 1991; Kingdom & Keeble, 1996; Landy & Bergen, 1991; Malik & Perona, 1990), and spatial frequency (Arsenault, Wilkinson, & Kingdom, 1999).

The second-stage filters of the mechanisms for detecting orientation change are believed to be tuned for the orientation (Kwan & Regan, 1998) and spatial frequency (Kingdom & Keeble, 1996) of the texture modulation, and are positionally labeled (Prins & Mussap, 2000, 2001). Psychophysical studies have suggested that the inhibitory surround is exceptionally broad and shallow (Kingdom & Keeble, 1996; Prins & Mussap, 2000).

* Corresponding author. Tel.: +1-514-843-1690; fax: +1-514-843-1691.

*E-mail addresses:* nicolaas.prins@mcgill.ca, nprins@po-box.mcgill.ca (N. Prins).

Fig. 1. The FRF mechanism. Two schematic FRF mechanisms are shown atop an OM texture. The amplitude of modulation is 7°. The excitatory centers of the FRFs shown cover a texture region in which the center orientation is 7° ccw from horizontal, the inhibitory surrounds cover texture regions in which the center orientation is 7° cw from horizontal. The texture is first filtered with simple first-order luminance filters which are selective for orientation and spatial frequency. Two example Gabor-shaped filters are shown on the right of the figure, one tuned to 7° (the peak orientation of the texture region presented to the FRFs center), the other tuned to 30° from horizontal. The rectified output from the first-stage filters is consequently integrated by the larger second-stage filter. In the example shown, the center region leads to activation of the FRF mechanism (+), whereas the surround leads to inhibition of the FRF mechanism (−). Shown within the RFs of the second-stage filter is the full-wave rectified output from the filters shown on the right of the figure. It should be noted that for purposes of illustration the contrast in the filtered regions in the figure has been maximized individually for each of the FRFs. In reality, the overall first-stage activation in the lower FRF would be of much lesser magnitude than that in the upper FRF. However, as is clear from the figure the activation of the lower FRF would exceed that of the upper FRF since the *difference* in first-stage activation between excitatory and inhibitory regions is larger in the former compared to the latter.

Presumably, textures are processed by a range of FRF mechanisms which differ with regard to the tuning properties of their first-order input. Intuitively it may seem likely that the first-order orientation and spatial-frequency preferences of the most active FRF mechanisms would match those of the stimulus. An example would be the upper FRF mechanism in Fig. 1 which receives its first-stage input from filters tuned to the orientation presented to the excitatory region of its receptive field (RF). However, it is quite possible, and perhaps even likely, that the detection of texture modulations is mediated by FRF mechanisms whose inputs are first-order stage filters which optimize the *differential* response between the different texture regions. At least at low modulation depths, these first-stage filters will have their peak response at orientations or spatial frequencies which are quite different from the peaks contained within the orientation-modulated (OM) or frequency-modulated (FM) modulated texture.

The argument is similar to that made by Regan and Beverley (1983, 1985). Regan and Beverley (1983) determined that the mechanisms involved in the spatial-frequency discrimination of two sequentially presented simple luminance gratings were optimally tuned to frequencies roughly one octave above and below the spatial frequency of the test stimuli. Similarly, Regan and Beverley (1985) determined that the mechanisms involved in an orientation discrimination task were those which were optimally tuned to orientations which differed by about 10–20° from that of the test stimuli. They argued that discrimination thresholds are determined by those mechanisms which maximize the differential response between the two test stimuli. Fig. 1 illustrates that the first-order channels which optimize the differential response between two texture regions that differ in orientation are tuned to orientations different from those of the peak or the trough orientation of the OM texture. Shown schematically are two FRF mechanisms, both positioned optimally atop an OM texture. Within each RF are the full-wave rectified outputs of simple Gabor-shaped luminance filters. The optimal frequency response for both sets of luminance filters lies at the peak luminance spatial frequency of the texture. However, the upper FRF mechanism compares the rectified activation of first-order filters tuned to the peak orientation (7° from horizontal) of the texture bar on which the FRF mechanism is centered, whereas the lower FRF mechanism compares the rectified outputs of first-order filters tuned to an orientation 30° from horizontal. What is clear from this simple demonstration is that the difference in activation between center and surround regions of the FRF mechanism is larger for the FRF mechanism tuned to an orientation which differs from the peak of the orientation contained in the texture by 23°. A similar argument holds for FM textures: the first-order channels which optimize the differential response between the two texture regions will be most sensitive to spatial frequencies that differ quite substantially from either the peak or the trough of the FM modulation. A more comprehensive description of the above argument is given in Section 5 and in Appendix A.

In the current series of experiments we sought to determine empirically the spatial-frequency and orientation tuning preferences of the first-stage filters of mechanisms underlying the processing of OM/FM texture modulations at low modulation depths. We presented observers with square-wave modulated textures. The textures consisted of oriented Gabor micropatterns which were modulated either in spatial frequency or orientation. The task of the observers was to perform a discrimination with regard to the spatial frequency of the texture modulation. Two modulation frequencies were used which differed by an octave such that the discrimination was trivial once the modulation was visible. At the start of a session, observers adapted to a

simple sinusoidal luminance grating whose spatial frequency and orientation were varied between different conditions. Following adaptation, we measured the orientation or spatial-frequency modulation threshold for making the second-order spatial-frequency discrimination. The adaptation would be expected to suppress the sensitivity of the neurons selective to the properties (i.e., spatial frequency and orientation) of the adapting grating. Diminished performance in the test phase (conveyed by elevated discrimination thresholds) signifies that these neurons are involved in the discrimination task.

If discrimination is achieved by FRF mechanisms which receive first-order input from the most active channels, threshold elevations (TEs) should be most pronounced when the spatial frequency and orientation of the adapting grating match those of the peak or trough of the test stimulus. If, on the other hand, the FRF mechanisms receive their first-order input from channels which maximize the differential response to the two texture regions, TEs should be most pronounced at adapting gratings whose spatial properties lie outside the peak or trough of the test texture.

To pre-empt, our results support the latter prediction: TEs for the OM textures peaked at adapting orientations which differed by about 30° from the center orientation of the text textures, and TEs for the FM textures peaked at adapting frequencies which differed by about one octave from the center frequency of the test textures. We show that these results are not so much a result of the tuning characteristics of the first-order filters as they are of the spectral distributions of the stimuli themselves.

2. Stimuli

Textures consisted of randomly positioned Gabor-micropatterns:

$$L(x,y) = L_0 + L_m \cos(2\pi f(x\sin(\theta) + y\cos(\theta)) + \varphi)$$
$$\times \exp(-(x^2 + y^2)/(2\sigma_e^2)) \quad (1)$$

$$\sigma_e = (f\pi)^{-1}\sqrt{0.5\ln[2]}(2^{1.5}+1)(2^{1.5}-1)^{-1} \quad (2)$$

where L_0 is mean luminance (124.0 cd m^{-2}), L_m is the luminance modulation amplitude (61.9 cd m^{-2}), f is spatial frequency, θ is orientation, φ is the phase of the cosine component ($\varphi = \pi/2$ or $3\pi/2$), and σ_e is the standard deviation of the Gaussian envelope. The value of σ_e covaried with the value of f so as to keep the spatial-frequency bandwidth (full-width at half-height) of the micropatterns constant at 1.5 octaves across different frequencies.

The stimulus area was circular with a radius of 4.2°. The positioning of the Gabor-micropatterns within the stimulus area was random under the constraint that the minimum center-to-center separation between any two micropatterns was equal to the standard deviation of the micropattern envelope (σ_e in Eq. (1)). This constraint ensured approximately equal 'coverage' across the stimulus area and avoided excessive luminance summation (which would otherwise occur where several micropatterns happened to overlap). The luminance modulations, but not dc components, of the Gabor-micropatterns were summed where they overlapped. The number of micropatterns per unit area varied as a function of the spatial frequency of the micropatterns such that the number of micropatterns per square degree stimulus area was $0.6/\sigma_e^2$. For example, at the center spatial frequency of 5 cpd, the number of micropatterns per square degree was equal to 97. Implementing this density constraint ensured that both 'coverage' and rms contrast between regions of differing spatial frequency were, at least statistically, equal.

The spatial frequency or orientation of the Gabor-micropatterns was square-wave modulated around a center spatial frequency of 5 cpd and a center orientation of 0° (horizontal). In the case of spatial-frequency modulation the two spatial frequencies present in the texture differed from the center spatial frequency by an equal number of log-frequency units. The phase of the texture modulation was randomized across trials. The stimulus area contained either two or four full cycles of the square-wave modulation, corresponding to bar widths of 1.05° and 0.53° respectively. The bars were oriented vertically. Fig. 2a displays an example OM texture, with a modulation amplitude of 8° (peak-to-trough difference of 16°). Fig. 2b displays an example FM texture, with a modulation amplitude of 0.2 octave.

The adapting gratings were simple sine-wave luminance modulations. The orientations and spatial frequencies of the adapting gratings varied between blocks and are as specified in Section 3. Adapting gratings were presented in a square area (9.1° × 9.1°) which was centered on the display.

Textures were generated on-line (trial-by-trial) in computer memory and were presented on a Clinton Monoray monitor controlled by a Cambridge Research Systems VSG 2/5 graphics board. At the employed viewing distance of 100 cm, the resolution of the monitor was equal to 42.3 pixels per degree.

3. Procedure

Within each block of trials, only one type of texture modulation (orientation or spatial frequency) was presented. At the start of each block and between any two trials, an adapting sine grating was presented. In blocks in which the texture was orientation modulated the adapting grating always had a spatial frequency of 5 cpd

Fig. 2. Example test stimuli. (a) Example of an OM texture. The amplitude of modulation equals 8° (peak-to-trough difference = 16°). The modulation frequency is 2 cycles per texture. (b) Example of a FM texture. The amplitude of modulation equals 0.2 octaves. The modulation frequency is 4 cycles per texture. Details are given in Section 2. Textures as shown here are not gamma-corrected.

and between blocks the orientation was varied across the full (180°) range of orientations, sampled at intervals of 10°. In blocks in which the texture was modulated in spatial frequency the adapting grating had an orientation of 0° (horizontal) and between blocks the spatial frequency was varied from −3 to 2 octaves relative to the texture's center spatial frequency (5 cpd) in steps of one half octave.[1] The different adaptation conditions were presented in quasi-random order.

At the start of each block, the adapting grating was presented for 120 s, counter phasing at 0.5 Hz (i.e., 1 contrast reversal/s). Between any two presentations of the test stimulus, the adapting grating was presented for 4 s, again counter phasing at 0.5 Hz. After each adaptation series a blank (mean luminance) screen was presented for 300 ms followed by the test texture. The test textures were presented for 200 ms. A new adaptation series started immediately after the response of the observer. Observers were instructed to respond to each test texture without delay. The task of the observers was to indicate through a button press whether the orientation or frequency modulation contained in the texture was of a low or high spatial frequency (2 or 4 cycles per stimulus). Feedback was provided after each trial by means of a tone following an incorrect response. A block consisted of a total of 120 trials, consisting of three randomly interleaved adaptive staircases (the BEST Pest; Pentland, 1980) of 40 trials each. The BEST Pest uses a maximum likelihood procedure to estimate the OM or FM threshold amplitude based on all preceding trials in a staircase and presents the stimulus on the next trial at the current threshold estimate. The observers were the authors.

---

[1] Actually, the highest spatial frequency of the adapting grating was +2.08 octaves relative to the texture's center spatial frequency (21.1 cpd), which corresponded to the Nyquist frequency of the monitor at the employed viewing distance. It should be noted that at the Nyquist frequency the actual contrast of the grating is likely not at the intended level due to non-linearities of the monitor.

4. Results

Within each condition, responses were combined across the different staircases and fitted with a logistic function using a maximum likelihood criterion. Standard errors of the thresholds (at 75% correct) were determined by bootstrap analysis using 400 repetitions for each standard error (e.g., Efron & Tibshirani, 1986; Foster & Bischof, 1991). Modulation amplitudes are expressed as half the peak-to-trough difference in orientation or spatial frequency between the two different texture regions.

4.1. Orientation-modulated textures

In the conditions where the texture modulation was defined by a difference in orientation between the two texture regions, no systematic differences were obtained between negative and positive adapting orientations (clockwise (cw) versus counterclockwise (ccw) from horizontal) for either observer. Hence it was decided to collapse the data across negative and positive angles of the adapting grating. In Fig. 3A are plotted the thresholds for both observers as a function of the orientation (deviation from horizontal) of the adapting grating. The ordinates on the right-hand side of the graphs show threshold elevations (TE), defined as

$$TE = Th_a / Th_u - 1,$$

where Th_a is the threshold obtained in an adaptation condition, and Th_u is the threshold obtained in the no-adaptation condition.

As can be seen from the figure, thresholds in the condition where the observers were adapted to a sine grating of the texture center orientation are only slightly elevated compared to those obtained in the no-adaptation condition. Rather, TEs peak at adapting orientations which are about 30° from the center orientation of the orientation, displaying a broad tuning function.
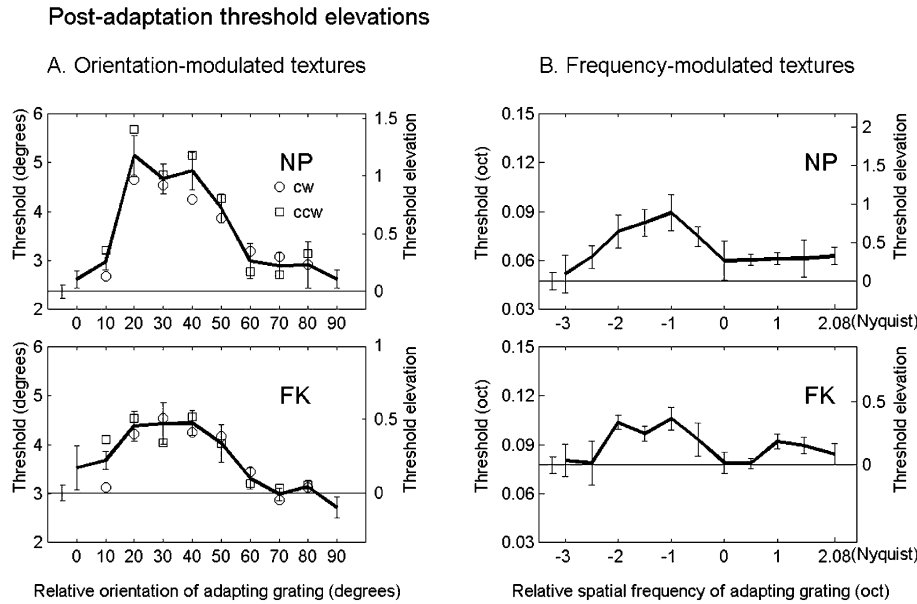
Post-adaptation threshold elevations

Fig. 3. Post-adaptation TEs. Discrimination thresholds (left ordinate) and TEs (right ordinate) for both observers for the (A) OM textures as a function of the orientation of the adapting grating and (B) FM textures as a function of the relative spatial frequency of the adapting grating, where 0 corresponds to the center spatial frequency of the test textures (5 cpd). Also shown in (A) are the TEs separately for cw and ccw angles of the adapting grating (relative to horizontal). The thresholds in the no-adaptation conditions are indicated by the horizontal lines. Modulation amplitudes correspond to half the difference between the peak and trough orientation/spatial-frequency value. Error bars indicate ±1 standard error. Each data point is determined by a minimum of 240 observations.

4.2. Frequency-modulated textures

Thresholds for the FM textures are displayed in Fig. 3B as a function of the spatial frequency of the adapting frequency. The thresholds in the condition where observers were adapted to a sine grating of a spatial frequency equal to the center spatial frequency of the texture are slightly elevated compared to those obtained in the no-adaptation condition. For NP a clearly pronounced elevation of thresholds occurs only at adapting spatial frequencies below the center spatial frequency of the test texture, peaking at about 1 octave below the center spatial frequency of the test texture. FK's thresholds show two elevation peaks, one at about 1–1.5 octaves below the center spatial frequency of the test texture and another of lesser magnitude peaking at about 1 octave above the center spatial frequency of the test texture.

5. Discussion

We have determined that, at low depths of modulation, the most active second-order texture mechanisms receive their first-order input from channels which have peak sensitivities at orientations and spatial frequencies which do not match the center orientation or spatial frequency of either the center, peak or trough of the texture modulation. Rather, based on the pattern of TEs as a function of the orientation of the adapting grating,

we conclude that in the case of an OM texture the second-order spatial-frequency discriminations are mediated by mechanisms which receive their first-order input from filters tuned to a broad range of orientations centered at about 30° from the center orientation of the OM texture. Similarly, in the case of FM textures the discriminations are mediated by mechanisms which receive their first-order input from filters tuned to a broad range of spatial frequencies centered at about 1 octave above and below the center spatial frequency of the FM texture.

Our results extend those obtained by Regan and Beverley (1983, 1985) into the texture domain. Regan and Beverley (1985) found that TEs in an orientation discrimination task using simple luminance gratings peaked when observers had adapted to a grating with an orientation differing by about 10–20° from that of the test stimuli. They argued that performance in the orientation discrimination task was mediated by channels which optimize the differential response between the two test stimuli, i.e., those channels that have the steepest slope in their orientation tuning curve at the orientation of the test stimuli. Earlier, using a similar procedure, Regan and Beverley (1983) had argued that performance in a spatial-frequency discrimination task is mediated by channels which have the steepest slope in their spatial-frequency tuning curves at the spatial frequency of the test stimuli.

There is, however, an important difference between the origin of Regan and Beverley's (1983, 1985) effect and the origin underlying our effect. The difference lies

in the stimulus characteristics. The test stimuli Regan and Beverley used were simple sine gratings which contained several cycles of the waveform. Hence, the spatial-frequency and orientation bandwidths of their stimuli were very narrow. Because of this their results can almost entirely be ascribed to the spatial-frequency and orientation response characteristics of the filters/channels involved in the task. Our stimuli, on the other hand, are relatively broadband both in spatial frequency and in orientation. The full-width at half-height spatial-frequency bandwidth of our textures is approximately 1.5 octaves and the full-width at half-height orientation bandwidth is approximately 60° (details as to how these values were determined can be found in Appendix A). In Fig. 4A we plot the amplitude spectra of the two different texture regions of an OM texture as a function of orientation. The amplitude of the modulation corresponding to Fig. 4A is NPs discrimination threshold in the no-adaptation condition. Also plotted in the figure is the difference between the two spectra. Analogously, plotted in Fig. 4B are the amplitude spectra for the two texture regions of an FM texture as a function of spatial frequency, again at NPs discrimination threshold in the no-adaptation condition. As is clear from Fig. 4, the differences in spectral amplitude between the two texture regions are largest at orientations/spatial frequencies which are quite different from the center orientation/spatial frequency of the stimulus. Also, mainly as a result of the broad orientation and spatial-frequency bandwidths of our stimuli, the function describing the difference in spectral amplitude between the two texture regions is very shallow and broad.

It should be stressed that our results only apply to stimuli in which the depth of modulation is low. At high modulation depths (relative to the bandwidth of the stimulus) the first-order filters which maximize the differential response between two texture regions will have preferences for orientations/spatial frequencies more closely matching the peak and trough orientation/spatial frequency of the stimulus. To illustrate, in Fig. 4C the difference between amplitude spectra is plotted for two texture regions of an OM texture at different depths of modulation. Similarly, in Fig. 4D the difference between amplitude spectra is plotted for two texture regions of an FM texture, again at different depths of modulation. It is interesting to note that the shape of the difference distribution is remarkably constant across a wide range of modulation depths. This would suggest that the distribution of activation across different FRF mechanisms would be largely constant across a wide range of modulation depths. It is important to realize that the exact shape of the difference functions will depend on the orientation and spatial-frequency bandwidths of the particular stimuli used.
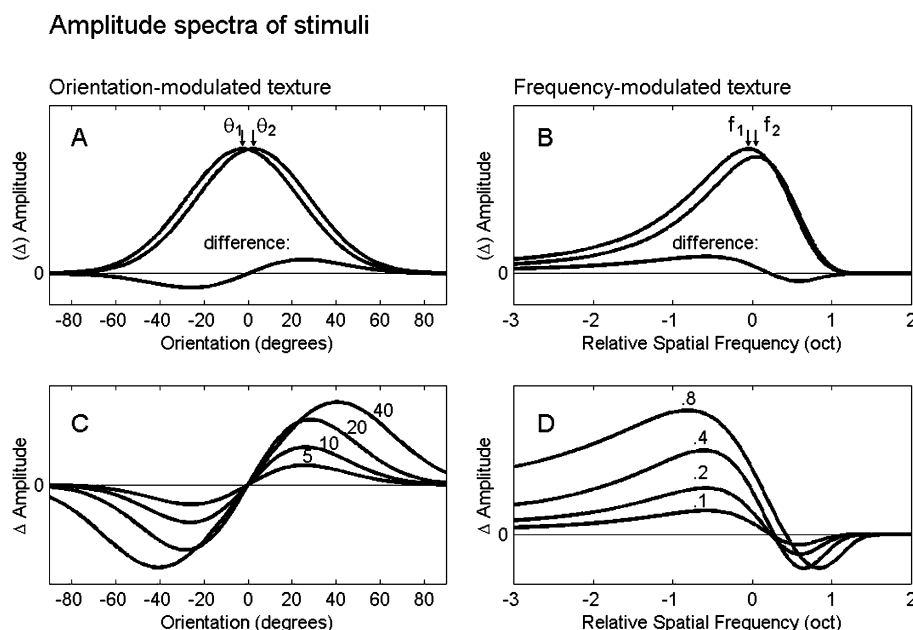

Fig. 4. Amplitude spectra of stimuli. Shown are 1-D 'slices' through the (2-D) idealized amplitude spectra (see Appendix A) of the two different texture regions together with the difference in amplitude between the two regions. (A) The amplitude spectra of the two regions of an OM texture at an amplitude of modulation corresponding to NPs threshold in the no-adaptation condition ($-\theta_1 = \theta_2 = 2.4°$). (B) The amplitude spectra of the two regions of an FM texture at an amplitude of modulation corresponding to NPs threshold in the no-adaptation condition ($-f_1 = f_2 = 0.047$ octaves (relative to the texture's center spatial frequency of 5 cpd)). (C) The spectral difference functions for an OM texture at depths of modulation of 5°, 10°, 20° and 40°. (D) The spectral difference functions for an FM texture at depths of modulation of 0.1, 0.2, 0.4 and 0.8 octaves. Note that spatial-frequency spectral content comes out asymmetric when plotted in terms of octaves in a way which mirrors Fig. 3. Graphic representations in 2-D Fourier space of the difference between amplitude spectra may be found in Fig. 5.

As discussed above, our results are in large part determined by the characteristics of the stimulus rather than those of the filter mechanisms. Varying the orientation and spatial-frequency bandwidths of the stimulus textures would result in different patterns of TE across the orientations and spatial frequencies of the adapting grating. For example, if the orientation and spatial-frequency bandwidths of the textures were reduced the peak TEs would presumably be shifted closer to the center orientation and spatial frequency of the texture. With decreasing orientation and spatial-frequency bandwidths of the texture the tuning properties of the filters will play an ever-increasing role in determining the pattern of TEs. In the extreme (and purely theoretical) case where the textures have infinitesimally small bandwidths, the pattern of TEs would be determined entirely by the tuning properties of the filters involved. Based on the work by Regan and Beverley (1983, 1985) one might expect in this case that TEs peak at around 10–20° or 1 octave from the center orientation/spatial frequency of the texture.

6. Conclusion

In conclusion, our results are consistent with the idea that the first-order filters involved in the processing of texture modulations at low depths of modulation are those which maximize the differential response between the two texture regions. This is indicated by the general similarity between the obtained TE functions and the functions describing the difference in spectral amplitude between the two texture regions.

An interesting implication of these results is that the distribution of activity between FRF mechanisms with differing first-order tuning functions might be used to determine the nature of the texture modulation. As discussed in Section 1, any individual FRF mechanism will be indiscriminate as to the nature of the texture discontinuity which triggered its response. Activation of any given FRF mechanism may arise from a difference between texture regions in orientation, spatial frequency, or contrast. However, as our results indicate, the pattern of activation across FRF mechanisms differing with regard to the orientation and spatial frequency preferences of their first-order filters, will be unique for each of the three mentioned discontinuities. A discontinuity in orientation between two texture regions will lead to two groups of active FRFs, both of which will be centered at the stimulus spatial frequency, but at two distinct orientations. Conversely, a spatial frequency discontinuity will again lead to two groups of active FRFs, however in this case both will be centered at the peak orientation of the stimulus, but each group will have different preferences for spatial frequency. A discontinuity in contrast will arguably lead only to the activation of one group of FRFs, centered at the peak orientation and spatial frequency of the texture.

The above also implies that a given FRF mechanism which is optimally tuned for an OM texture modulated around, say, orientation θ and containing luminance spatial frequency f, will respond little to an FM texture modulated around luminance spatial frequency f and containing orientation θ. For example, the mechanism that is most responsive to our OM textures will have first-order filter preferences for orientations about 30° from horizontal and spatial frequencies at, or close to, the center spatial frequency of the texture. Its spatial-frequency preferences are thus not well suited to detect our FM textures (a quick glance at Figs. 3B, 4B and 5b


Fig. 5. The absolute difference between the idealized amplitude spectra of the two texture regions of (a) an OM texture at an amplitude of modulation corresponding to NPs threshold in the no-adaptation condition ($-\theta_1 = \theta_2 = 2.4°$) and (b) an FM texture at an amplitude of modulation corresponding to NPs threshold in the no-adaptation condition ($-f_1 = f_2 = 0.047$ octaves (relative to the texture's center spatial frequency of 5 cpd)).

tells us that a mechanism with spatial-frequency preferences around the center spatial frequency of the texture is little, if at all, involved in detecting our FM texture). Moreover, its orientation preferences are also sub-optimal for detection of our FM textures as they lie about 30° removed from the peak orientation content of the FM texture.

It should be noted that it remains a distinct possibility that performance in our task is not mediated by FRF mechanisms at all. Many schemes can be conceived of which will be able to perform the discrimination task described. For example, the second-order spatial-frequency discrimination could have been performed by comparing the local orientation or spatial frequency between only a few regions in the texture. However, our task is typical of the type of task used to investigate the properties of FRF mechanisms. Future research into the properties of FRF mechanisms should verify the spectral content of the stimuli used before relying on the assumption that the first-order stage consists of filters which are tuned to the center or peak orientations and spatial frequencies contained in the stimulus. To give one example, in physiological investigations of second-order mechanisms, it should be noted that a neuron responsive to second-order modulations of texture may have different first-order preferences depending on whether the texture is modulated in orientation, spatial frequency, or contrast. Whether this is indeed the case will depend on the depth of modulation and the orientation and spatial-frequency bandwidths of the texture.

Acknowledgements

Appendix A

A.1. Fourier spectra of stimuli

The amplitude spectra of our stimuli $(H_{c/s}(f, \theta))$ were determined by generating unmodulated textures in the manner explained in Section 2 and subjecting these to a Fourier transform. We did so for five different spatial frequencies of the Gabor-micropatterns spanning 1 octave centered on the center spatial frequency (5 cpd) of the Gabor-micropatterns. For each of the five different frequencies, the amplitude spectra of 400 textures were determined and averaged. The resulting distributions were subsequently fitted with:

$$H(f, \theta) = A \exp(-0.5[(f - f_0)/(f_0\sigma_f)]^2)$$
$$\times \exp(-0.5[(\theta - \theta_0)/\sigma_0]^2), \tag{A.1}$$

where f_0 and θ_0 were fixed at the intended values (the values of the individual Gabor-micropatterns making up the texture), and σ_f, σ_0, and A were free to vary.[2]

The averaged amplitude spectra were described remarkably well by Eq. (A.1), ($R^2 > 0.977$ across the range of frequencies considered). The estimated values of σ_f and σ_0 were rather constant across the range of frequencies considered with average values of 0.405 (se = 0.001) and 25.48 (se = 0.07), respectively. These values correspond to a frequency bandwidth (full-width at half-height) of 1.50 octaves and an orientation bandwidth (full-width at half-height) of 60.0°. The values of A were proportional to the reciprocal of f ($R^2 > 0.999$). Such a relationship was expected as it renders the total energy $\int \int \{(H[f, \theta])^2\} \, df \, d\theta$ of the stimulus in Fourier space (and hence rms contrast) equal across different frequencies of the Gabor-micropatterns.

A.2. Activation of tentative FRF mechanism

The output of an FRF mechanism can be crudely modeled to be proportional to the difference between the response of the first-stage filters to the two texture regions in the stimulus. The underlying idea is that the response of an FRF mechanism is simply the difference between the rectified and summed outputs of its first-stage filters feeding into the excitatory center and the rectified and summed outputs of its first-stage filters feeding into the inhibitory surround. We assume here that the activation of an FRF mechanism's excitatory center (A_c) is proportional to the integral of the point-wise product of the Fourier transform of the first-order filter's response profile and the amplitude spectrum of the texture region covered by the excitatory center, i.e.,

$$A_c = c \int \int H_c(f, \theta) R(f, \theta) \, df \, d\theta,$$

where $H_c(f, \theta)$ is the amplitude spectrum of the texture region presented to the excitatory center and $R(f, \theta)$ is the Fourier transform of the response profile of simple cells tuned to orientation θ_t and frequency f_t. We make a similar assumption with regards to the summed activation of the inhibitory surrounds (A_s):

$$A_s = c \int \int H_s(f, \theta) R(f, \theta) \, df \, d\theta,$$

where $H_s(f, \theta)$ is the amplitude spectrum of the texture region presented to the inhibitory surrounds. Essentially, the above means that we make the assumption that the first-stage filters are linear in their response, that the rectification (the R in FRF) is either half-wave or

---

[2] We also considered fitting: $H(f, \theta) = A \exp(-0.5(\log(f/f_0)/\sigma_f)^2) \exp(-0.5((\theta - \theta_0)/\sigma_0)^2)$, which is symmetric in spatial frequency when frequency is expressed in log units. This function, however, fitted the data considerably worse.

full-wave rectification, and that the second-order mechanism sums the rectified responses of the first-stage filters in a linear fashion. Note that we have assumed that the proportionality constants are equal for the excitatory center and the inhibitory surrounds resulting in a *balanced* FRF, i.e., the response of the FRF to an unmodulated texture equals zero.

For simplicity, let us assume that the excitatory center of the FRF covers exactly the width of one of the bars in our texture and that the inhibitory surrounds fall entirely within the neighboring texture bars (as shown in Fig. 1). The net response of the FRF, then, is given by:

$$A_c - A_s = c \int \int H_c(f, \theta) R(f, \theta)\, df\, d\theta$$
$$- c \int \int H_s(f, \theta) R(f, \theta)\, df\, d\theta,$$

which can be rewritten as:

$$A_c - A_s = c \int \int [H_c(f, \theta) - H_s(f, \theta)] R(f, \theta)\, df\, d\theta.$$

In words, this means that the response of the FRF is equal to the integral of the pointwise product of the Fourier transform of the filter's response profile and the *difference* between the two stimulus amplitude spectra. Shown in polar coordinates in Fig. 5a is the absolute value of the difference in amplitude spectra between the two different texture regions $|H_c(f, \theta) - H_s(f, \theta)|$ corresponding to NPs orientation-modulation threshold in the no-adaptation condition (the amplitude spectra were modeled according to Eq. A.1 above). Clearly, the difference between the two amplitude spectra at either the orientation peak (θ_2 in figure) or the orientation trough (θ_1) of the texture is very small. Rather, the difference between the amplitude spectra peaks at around 30° inclination from the center orientation of the texture. In Fig. 5b we show the absolute difference in amplitude spectra between the two texture regions corresponding to NPs frequency-modulation threshold in the no-adaptation condition. Again, the difference between spectra is very small at either the frequency peak (f_2) or the frequency trough (f_1) of the texture. Rather, the

difference between the two amplitude spectra shows a peak at frequencies well below the center spatial frequency of the texture and another peak of lesser magnitude at spatial frequencies well above the center spatial frequency of the texture.

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.

Bergen, J. R., & Landy, M. S. (1991). Computational modeling of visual texture segregation. In M. S. Landy, & J. A. Movshon (Eds.), *Computational models of visual processing* (pp. 253–271). Cambridge, MA: MIT Press.

Efron, B., & Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science, 1,* 54–75.

Foster, D. H., & Bischof, W. F. (1991). Thresholds from psychometric functions: Superiority of bootstrap to incremental and probit variance estimators. *Psychological Bulletin, 109*(1), 152–159.

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.

Kwan, L., & Regan, E. (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.

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

Pentland, A. (1980). Maximum likelihood estimation: the best PEST. *Perception and Psychophysics, 28,* 377–379.

Prins, N., & Mussap, A. J. (2000). Alignment of orientation-modulated textures. *Vision Research, 40,* 3567–3573.

Prins, N., & Mussap, A. J. (2001). Adaptation reveals a neural code for the visual location of orientation change. *Perception, 30,* 669–680.

Regan, D., & Beverley, K. I. (1983). Spatial frequency discrimination and detection: comparison of postadaptation thresholds. *Journal of the Optical Society of America, 73,* 1684–1690.

Regan, D., & Beverley, K. I. (1985). Postadaptation orientation discrimination. *Journal of the Optical Society of America A, 2,* 147–155.

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