PERGAMON

Vision Research 39 (1999) 1477–1489

Vision Research

On the mechanism for scale invariance in orientation-defined textures

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

[a] *McGill Vision Research Unit, Department of Ophthalmology, Pine Avenue West, Rm. H4–14, Montreal, QC H3A 1A1, Canada*
[b] *Department of Optometry, University of Bradford, Richmond Road, Bradford BD7 1DP, UK*

Received 12 January 1998; received in revised form 4 May 1998

Abstract

Texture perception is generally found to be scale invariant, that is, the perceived properties of textures do not change with viewing distance. Previously, Kingdom, F. A. A., Keeble, D. R. T., & Moulden, B. (Vision Research, 1995, 35, 79–91) showed that the orientation modulation function (OMF), which describes sensitivity to sinusoidal modulations of micropattern orientation as a function of modulation spatial frequency, was scale invariant—peak sensitivity occurred at a modulation spatial frequency which was invariant with viewing distance when modulation frequency was plotted in object units, e.g. cycles cm^{-1}. We have attempted to determine the mechanism underlying the scale invariant properties of the OMF. We first confirmed that the OMF was scale invariant using Gabor-micropattern textures. We then measured OMFs at a number of viewing distances, while holding constant various stimulus features in the retinal image. The question was which stimulus feature(s) disrupted scale invariance when manipulated in this way. We found that the scale (size) of the micropatterns was a critical factor and that the most important scale parameter was the micropatterns' carrier spatial frequency. Micropattern length and density were shown to have a small influence on scale invariance, while micropattern width had no influence at all. These results are consistent with the idea that scale invariance in orientation-defined textures is a consequence of 'second-stage' texture-sensitive mechanisms being tied in spatial scale selectivity to their 'first-stage' luminance-contrast-sensitive inputs.

Keywords: Scale invariance; Texture gradients; Orientation

1. Introduction

Scale invariance describes the situation in which the perceptual properties of an object are unaffected by viewing distance. Broadly speaking, one finds that the more complex the stimulus attribute, the more likely it is to show scale invariance. Thus while detection thresholds for luminance sine-wave gratings do not show scale invariance (Howell & Hess, 1978), thresholds for detecting 'second order' gratings, such as in dot density (Van Meeteran & Barlow, 1981), contrast amplitude and frequency (Jamar & Koenderink, 1983) and orientation (Kingdom, Keeble & Moulden, 1995), do show scale invariance. Scale invariance is also found for texture segregation thresholds (Nothdurft, 1985; Landy & Bergen, 1991; Bergen, 1991; Bach & Meigen, 1992). The orientation grating stimulus employed by Kingdom, Keeble & Moulden (1995) is shown in Fig. 1. The stimulus consists of a dense array of Gabor micropatterns, whose orientations vary sinusoidally across the display. We refer to the spatial frequency of the Gabor carrier as luminance spatial frequency, while that of the orientation modulation as texture spatial frequency. For both Gabor- and line-micropattern stimuli, Kingdom, Keeble & Moulden (1995) found that the texture spatial frequency to which subjects were most sensitive was invariant with viewing distance when texture spatial frequency was measured in object units, e.g. cycles cm^{-1}. This is a signature of scale invariance. Kingdom, Keeble & Moulden (1995) suggested that the scale invariant properties of the orientation modulation function (OMF) were a consequence of the changes in the retinal-image scale (size) of the micropatterns that occurred with viewing distance. This followed from the finding that a physical change in the

* Corresponding author. Fax: +1 514 8431693; e-mail: fred@jiffy.vision.mcgill.ca.

0042-6989/99/$ - see front matter
PII: S0042-6989(98)00217-X

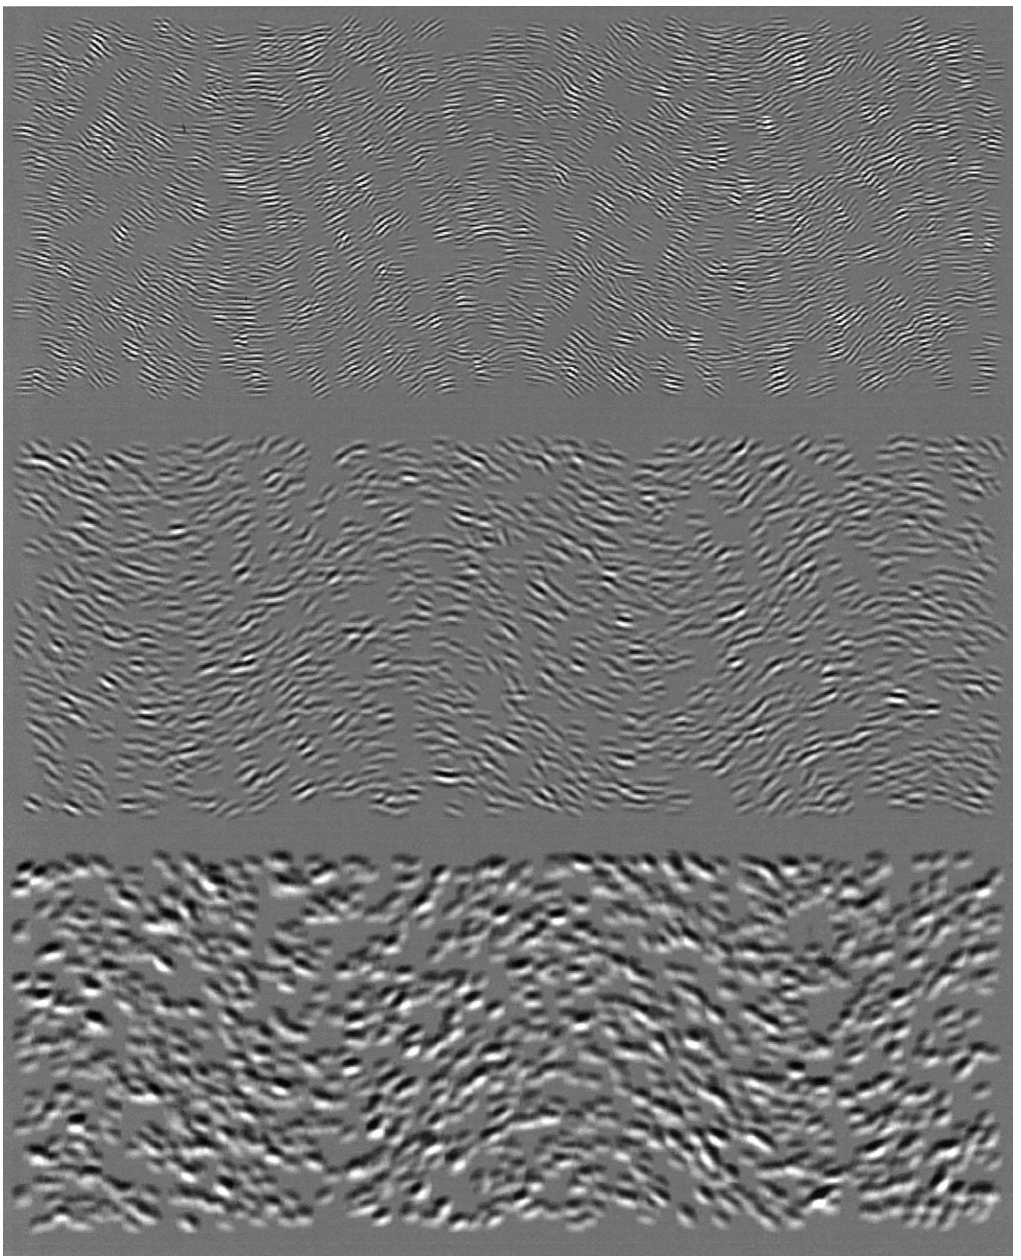

Fig. 1. Stimuli used for the constant retinal micropattern spatial frequency experiment. The stimuli are shown as if viewed from the same distance, i.e. they were physically the same size on the monitor screen. When the top, middle and bottom stimuli were viewed at octave interval viewing distances (VD1, VD2 and VD4 in text), the micropatterns all had the same retinal luminance spatial frequency. The middle stimulus is also the standard stimulus.

scale of the micropatterns caused peak OMF sensitivity to shift along the texture frequency axis. If the micropatterns were re-scaled by a change in viewing distance, the shift in peak sensitivity was perfectly counterbalanced by the shift that also occurred in retinal texture spatial frequency. However, the evidence that micropattern scale was the critical factor underlying scale invariance was only indirect in the study by Kingdom, Keeble & Moulden (1995), as the shift in peak OMF sensitivity was produced by changing the physical size of the micropatterns at a given viewing distance, rather than by changing viewing distance itself. It is quite possible that scale invariance in textures is due to reasons other than the changes in micropattern scale that occur with viewing distance. One obvious possibility is that observers simply take into account the estimated distance to the texture when computing its properties. The first experiment of this study tests for this possibility. The remaining experiments consider the role of micropattern scale, as well as other stimulus features, in producing texture scale invariance. We measured OMFs at various viewing dis-

tances while holding constant a particular feature of the stimulus in the retinal image. The rationale is that if this manipulation disrupts the scale invariant properties of the OMF, then the feature is implicated as critical for scale invariance.

2. Method

2.1. Subjects

The two authors FK and DK acted as observers for all experiments. For two conditions we employed a third observer, AW, who was a paid undergraduate volunteer. All had normal or corrected-to-normal vision and were highly experienced psychophysical observers.

2.2. Stimuli

2.2.1. Generation

Example stimuli are shown in Fig. 1. They were usually generated on a Macintosh Quadra 950 computer and displayed on a SuperMac Trinitron monitor. Several of DK's data sets were collected using stimuli produced by a Macintosh Quadra 840AV computer and a Nanao FlexScan T660i-J monitor. The displays were all monochrome (black–white) and were gamma corrected by suitable selection of intensity levels from an eight-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:

$$L(x, y) = M + A\cos[2\pi f_l(x\cos\theta - y\sin\theta)] \times \exp[-((x\cos\theta)^2/2\sigma_w^2 + (y\cos\theta)^2/2\sigma_l^2)] \quad (1)$$

with M mean luminance, A amplitude, f_l luminance spatial frequency, θ orientation, σ_w and σ_l standard deviations of the width and length of the Gaussian envelope, respectively. In the standard condition, σ_w and σ_l were set equal, producing a circularly symmetric envelope. In other conditions, σ_w and σ_l were unequal, producing fattened or elongated micropatterns. In the standard condition, f_l was set to three cycles per degree (cpd) at the standard viewing distance of 64.7 cm, and both σ_w and σ_l were set to 0.2°, giving the micropatterns a spatial frequency bandwidth at half height of 0.94 octaves. The micropatterns had a contrast, defined as peak amplitude A divided by the mean M of 23.6%. M and the background were set to 35 cpd m^{-2}. Details of micropattern parameters other than the standard condition are given with each experiment.

2.2.3. Orientation modulated textures

The standard stimulus shown in Fig. 1b was a display $11.3 \times 28.2°$ containing 1537 micropatterns. The position of each micropattern within the display window was completely randomised. The orientation of the micropatterns was constrained, in that the nominal mean orientation of the micropatterns varied along the horizontal axis of the display sinusoidally. The amplitude of orientation modulation determined how much the orientation of the micropatterns changed throughout one complete cycle of orientation modulation. For example, an amplitude of orientation modulation of 10° implied that the micropattern orientation changed by 20° throughout one complete cycle of orientation modulation. All micropatterns at a given horizontal position were given an orientation drawn from a Gaussian distribution with a specified mean (determined by the point on the waveform) and S.D. The S.D. 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 combined additively. Six texture spatial frequencies were employed in all conditions: 0.5, 1, 2, 4, 8, and 16 cycles per screen (cps), corresponding respectively to 0.017, 0.033, 0.067, 0.133, 0.267, and 0.533 cpd at the standard viewing distance.

2.3. Procedure

A two-interval-forced-choice procedure was used throughout to measure the threshold amplitude of orientation modulation. On each trial two displays were presented, each for 147 ms and separated by a 500 ms inter-stimulus-interval (ISI). The task for the subject was to decide which interval contained the stimulus with the orientation modulation. The only difference between the two stimuli presented on a given trial was in their amplitude of orientation modulation, which in the comparison stimulus was zero. The method of constant stimuli was used with five amplitudes of orientation modulation for each texture spatial frequency, the magnitudes being determined from pilot studies. In a given session all six texture spatial frequencies were employed and were displayed in random order. The phase of orientation modulation was also randomised on each trial. Feedback in the form of a tone was given for an incorrect decision.

2.4. Analysis

We fitted Weibull functions to the psychometric functions using a maximum likelihood method. This produced a threshold at the 82% correct level and a 67% confidence interval which is the error bar on each data point.

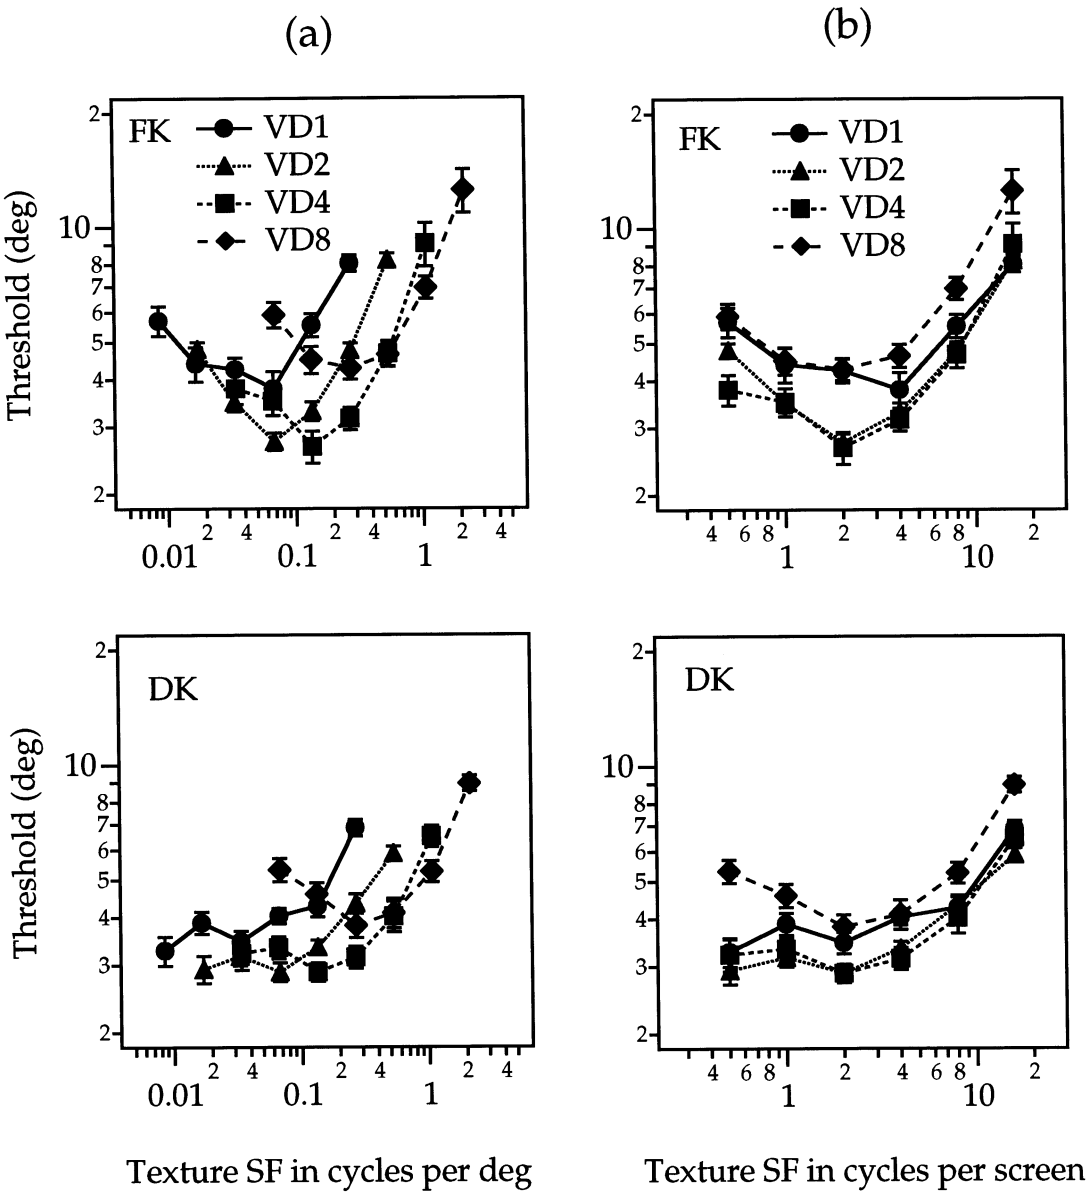

Fig. 2. OMFs for the standard condition shown in Fig. 1, obtained at four viewing distances. Each plot is the threshold amplitude of orientation modulation as a function of texture spatial frequency. The viewing distances VD1–8 are multiples of the smallest viewing distance (32.5 cm). On the left-hand side under (a), texture spatial frequency is given in retinal units of cpd, whereas on the right-hand side under (b), it is given in object units of cps.

3. Experiments and results

3.1. Experiment 1: demonstration of scale invariance

It was first necessary to confirm for the stimuli used in this study, that the shape of the OMF was scale invariant. We therefore measured OMFs for the standard stimulus [Fig. 1(b)] at four viewing distances: 31.8; 64.7; 130; and 262 cm, referred to in the figures as VD1, VD2, VD4, and VD8, respectively. These viewing distances were chosen to produce octave intervals in the visual angle subtended by fixation to a point halfway to the edges of the monitor screen. The results of this experiment are shown in Fig. 2. Each graph plots modulation amplitude thresholds as a function of texture spatial frequency. Texture spatial frequency is given in the left hand panels in retinal units of cpd (cycles per degree), and in the right hand panels in object units of cps (cycles per screen). When expressed in retinal units the curves translate leftwards with increasing viewing distance, whereas in object units the leftward translation is absent. To establish scale invariance quantitatively we performed the following analysis. After logarithmically transforming the thresholds and their associated errors, we fitted a Gaussian function to each OMF using IGOR (Wavemetrics), with the

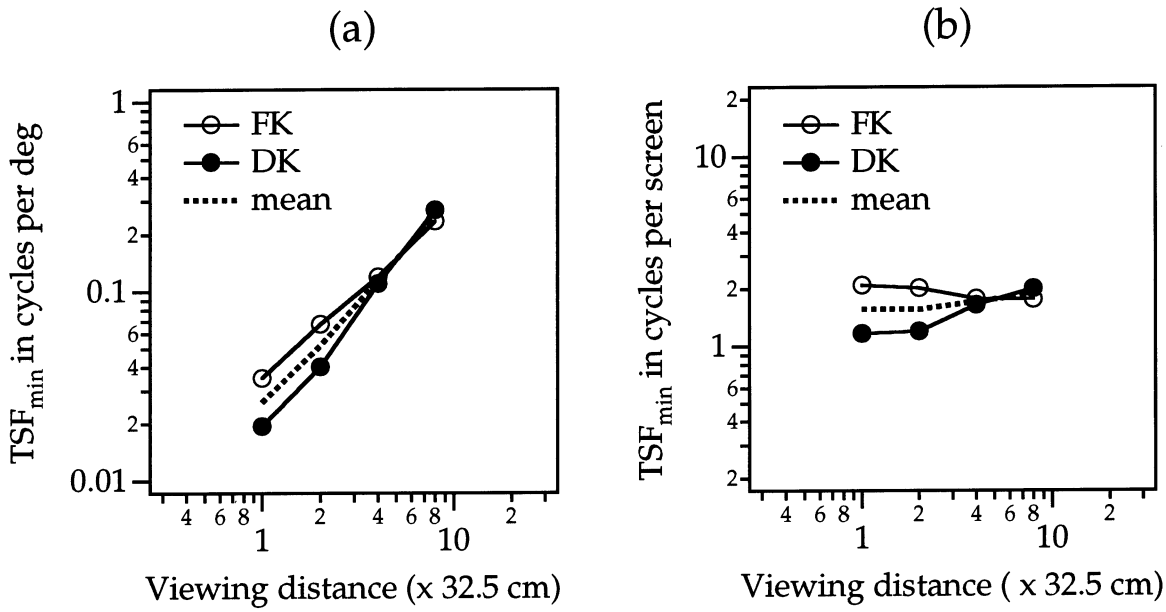

Fig. 3. Estimates of TSF_{min}, the texture spatial frequency with the lowest threshold, plotted as a function of viewing distance for the data shown in Fig. 2. The values of TSF_{min} were obtained after fitting a Gaussian to each log–log plot. In (a) texture spatial frequency is expressed in cpd. The mean slope for the two subjects is 1.1. In (b) texture spatial frequency is given in cps.

reciprocal of the errors on each data point serving as a weighting function in the fitting procedure. The Gaussian function used was:

$$a + b \cdot \exp(-(\log(f_t/c))^2/d^2) \quad (2)$$

where f_t represents texture spatial frequency, and a, b, c, d are free parameters. Note that the term $\log(f_t/c)$ in the exponent transforms texture spatial frequency into logarithmic units. The parameter which represents the overall position of the OMF along the texture spatial frequency axis is c. Although c is strictly speaking an estimate of the texture spatial frequency with the lowest threshold, a minimum threshold is obvious only in FK's data, with three out of four OMFs in DK's data being more-nearly lowpass in shape. Nevertheless the visible shifts in the rising portion of DK's data on the right of each plot in Fig. 2(a) are captured by the parameter c because the Gaussian function is constrained to be symmetrical. We will refer to c as the texture spatial frequency with minimum threshold, or TSF_{min}. TSF_{min} is plotted in cpd in Fig. 3a and cps in Fig. 3(b). Viewing distance on the abscissa is given in multiples of the nearest viewing distance (32.5 cm). Both the individual subjects' data, as well as the geometric mean across subjects is shown. In Fig. 3(a) TSF_{min} is almost exactly proportional to viewing distance—a slope of 1.1 on a double-logarithmic plot—whereas in Fig. 3(b) it is nearly flat. These results verify that the overall shape of the OMF is scale invariant.

3.2. Experiment 2: is scale invariance due to distance estimation?

It is possible that the scale invariant properties of orientation-defined textures are due to the visual system taking into account an estimate of viewing distance, rather than using features in the retinal image of the stimulus itself. To test this we measured OMFs at two viewing distances for stimuli whose retinal images were made identical. The results are shown in Fig. 4. Texture spatial frequency and TSF_{min} are both given in cpd. TSF_{min}, averaged across subjects, does not appear to change with viewing distance with this manipulation. Hence these data suggest that scale invariance is not a result of the visual system taking into account viewing distance, but is instead due to changes in the retinal image of the stimulus that occur as viewing distance is changed.

3.3. Experiment 3: is scale invariance due to micropattern scale?

In our previous study we found that changing the scale of the micropatterns shifted the OMF along the texture spatial frequency axis (Kingdom, Keeble & Moulden, 1995). We therefore suggested that the rescaling of the micropatterns in the retinal image that accompanied a change in viewing distance underlay the scale invariant properties of the OMF. In this experiment we directly test this hypothesis by measuring OMFs at different viewing distances while physically re-scaling the micropatterns to keep them identical in

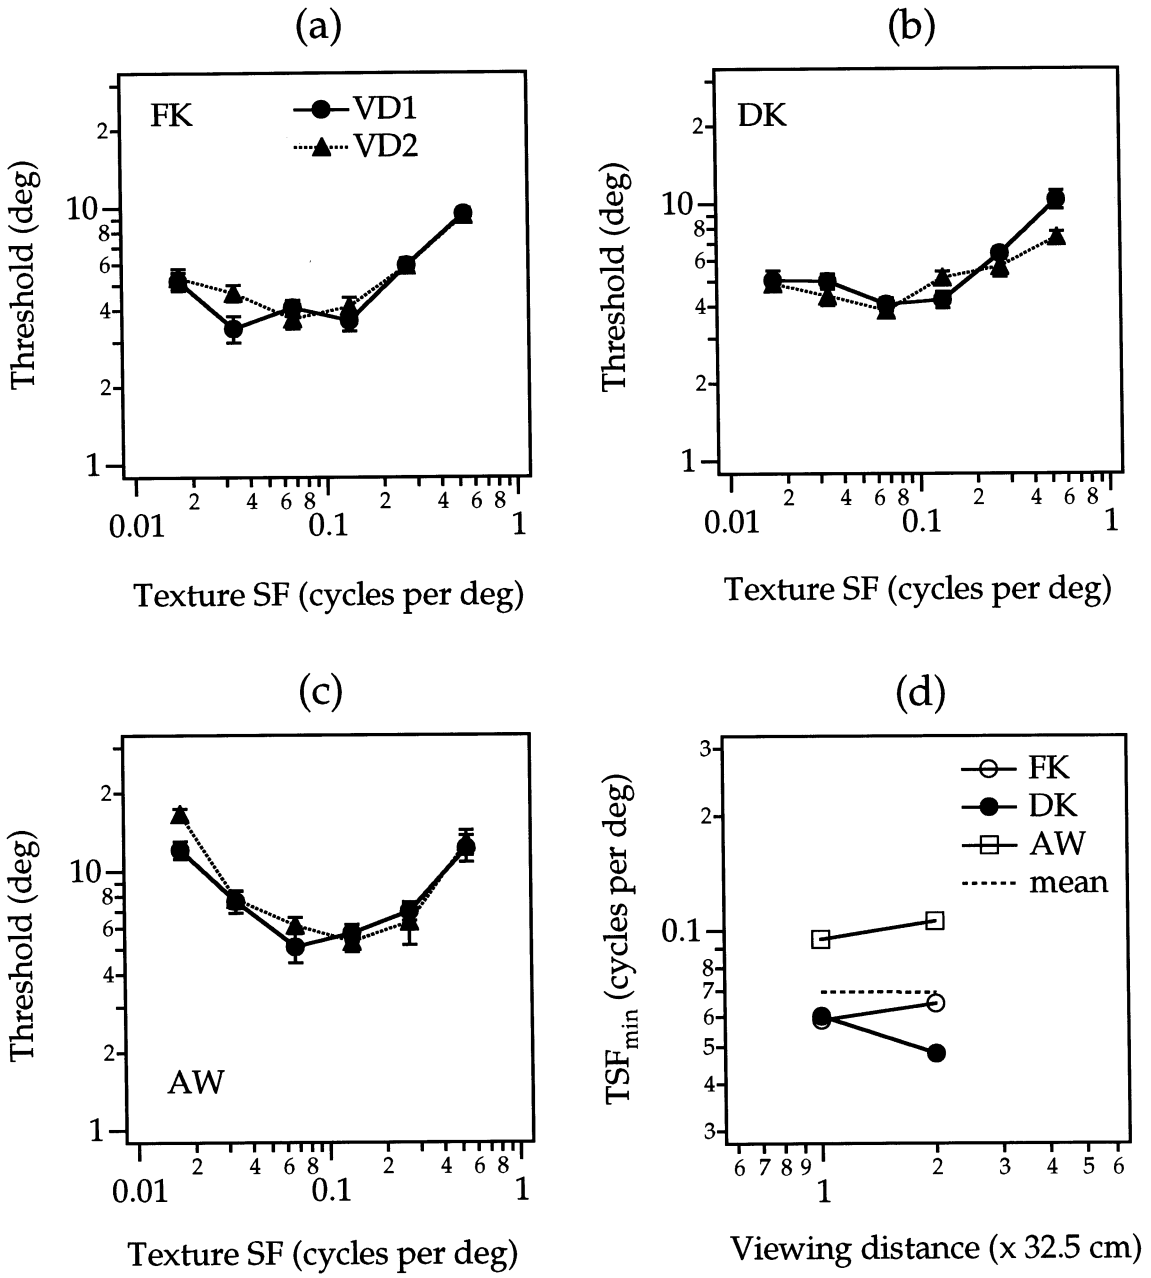

Fig. 4. Effect of holding the retinal image of the stimulus constant across two viewing distances. (a–c) Shows data for the three subjects. Note that the abscissa in all three figures is given in cpd. (d) Shows TSF_{min} in cpd as a function of viewing distance for each subject. The dotted line gives the geometric mean values across subjects.

the retinal image. Thus the VD1 condition differed from the VD2 condition in that the former's micropatterns were physically halved in scale, while the micropatterns in the VD4 condition were doubled in scale. If micropattern scale is the critical feature for producing scale invariance, then this manipulation should disrupt it. We included FK's standard (VD2) condition data from the previous experiment, while for DK this condition was measured anew as the experiment was carried out on different equipment (see Section 2). The OMFs for the VD1 and VD4 conditions are shown in Fig. 5(a) (left hand panels), with texture spatial frequency plotted in cps. Fig. 6(a) shows the estimates of TSF_{min} for all three viewing distances. As can be seen, the scale invariance demonstrated by the flat function in Fig. 3(b) is now absent in both subjects in Fig. 6(a). The mean slope of the function describing log TSF_{min} versus log viewing distance is -0.53. A complete disruption of scale invariance would produce a slope of -1.0. This suggests that holding micropattern scale constant in the retinal image significantly, though not completely, disrupts scale invariance.

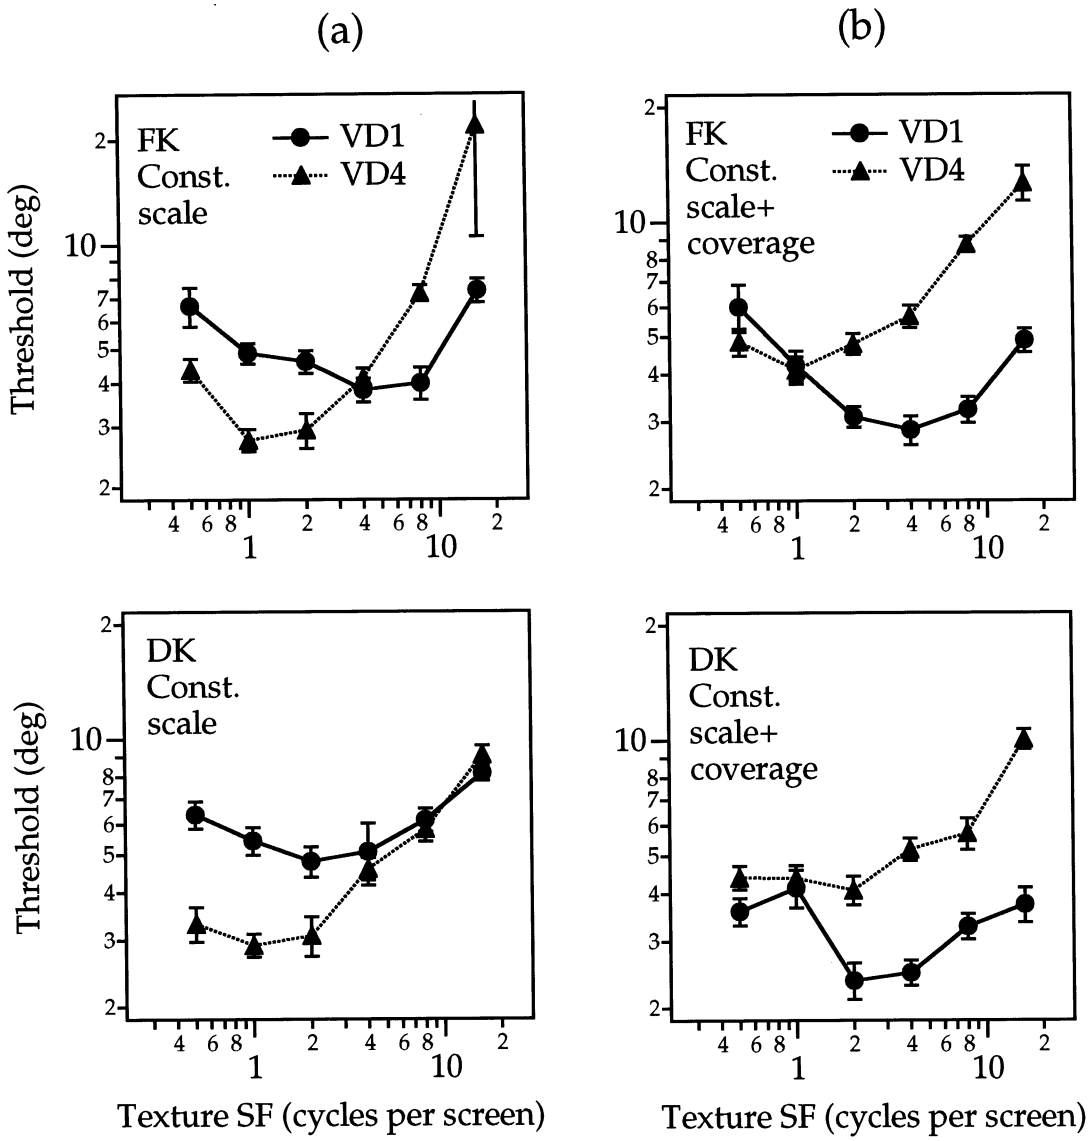

Fig. 5. OMFs at two viewing distances a factor of four apart with the Gabor micropatterns physically re-scaled such that each micropattern's retinal image was the same as in the standard condition. In (a) the number of Gabor micropatterns was identical across viewing distance. In (b) the number of micropatterns was set to keep the proportion of the display covered by micropatterns constant.

It is possible that the disruptive effect on scale invariance resulting from holding micropattern scale constant in the retinal image is not due to the effects of micropattern scale per se, but instead due to the effects of micropattern coverage. In the above experiment, more and more of the available stimulus area was occupied by micropatterns as viewing distance was increased, because the micropatterns were physically enlarged to preserve their size in the retinal image. By the same token, the average distance between the edges of the micropatterns diminished (though note that the total number of micropatterns remained constant). An increase in micropattern coverage, and/or reduction in between-micropattern-edge distance could selectively impair the processing of high texture frequencies. This would cause TSF_{min} to shift leftwards on the texture spatial frequency (cps) axis with viewing distance. To test this possibility we repeated the above experiment while holding constant micropattern coverage as micropattern scale was physically changed. This required decreasing the number of micropatterns as viewing distance increased. The raw OMFs are shown in Fig. 5(b), and TSF_{min} in Fig. 6(b). As the figure shows scale invariance was still disrupted even though micropattern coverage did not vary with viewing distance. The average slope of the function in Fig. 6(b) is -0.76. We conclude, therefore, that the change in micropattern coverage that accompanied the physical change in micropattern scale in the first part of experiment 3 was not responsible for disrupting the scale invariance in the OMF. The factor that was responsible was micropattern scale.

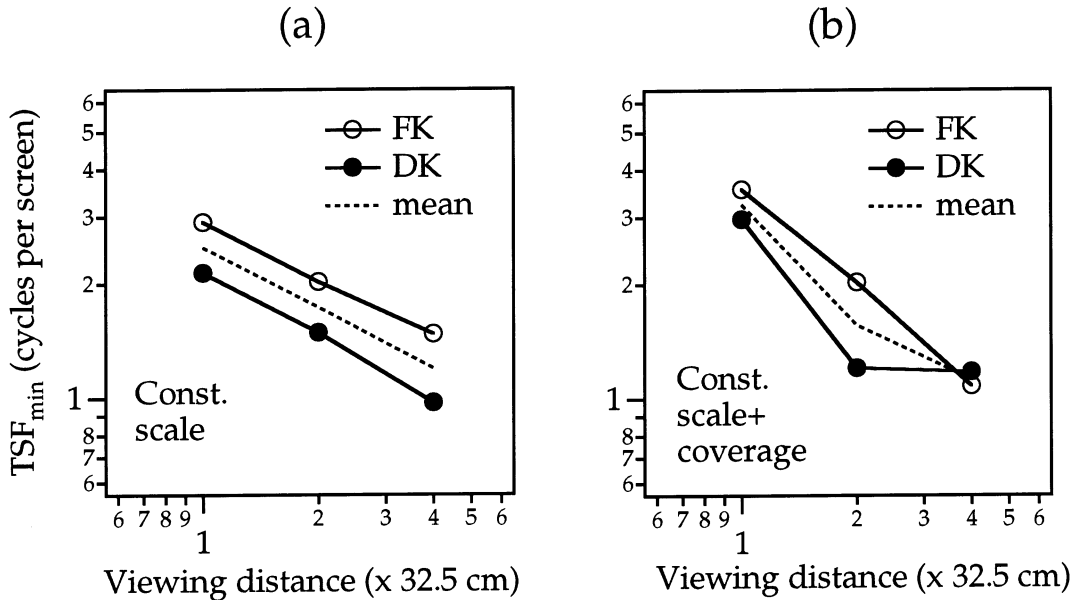

Fig. 6. TSF_{min} in cps, plotted as a function of viewing distance for the data shown in Fig. 5, plus the standard condition (VD2). (a) constant number of micropatterns; (b) constant coverage. Note that for DK, the standard condition is different for (a) and (b), as explained in the text.

3.4. Experiment 4: which scale feature causes scale invariance?

Although we have demonstrated that micropattern scale is a critical factor underlying scale invariance in the OMF, the question arises as to which micropattern scale feature(s) is critical. Is it the carrier spatial frequency, the envelope size, or perhaps even the length of the micropatterns? To answer this question, we repeated the main condition of experiment 2, but instead of holding micropattern scale constant in the retinal image, we held one of three factors constant: the carrier spatial frequency f_l, the size (S.D.) of the envelope σ, the length of the micropatterns, defined as the S.D. of the envelope along the axis of the carrier σ_l. The stimuli employed for the first manipulation, carrier spatial frequency, are shown in Fig. 1. Note that in the figure the stimuli are shown as if viewed from the same viewing distance. The results of this experiment are shown in Figs. 7 and 8 for the raw OMFs and TSF_{min}.

It is clear from Fig. 8 that the most critical feature is carrier spatial frequency. The average (geometric) mean slope of the carrier frequency condition is -0.53. Envelope length appears to be a factor for one subject FK (slopes: FK = -0.46; DK = -0.09; mean = -0.27). Envelope size does not appear to be a factor at all (mean slope = -0.06).

4. Discussion

There are two principle findings of this study. First, scale invariance in textured stimuli results from changes in the retinal-image properties of the stimulus as viewing distance is changed, rather than by the visual system incorporating estimates of viewing distance in the computation of texture properties. This is consistent with Gibson's early notion of the primacy of intrinsic image information for vision (Gibson, 1979). Second, scale invariance appears to be due primarily to changes in micropattern scale in the retinal image, and specifically micropattern spatial frequency. When micropattern scale or spatial frequency was held constant in the retinal image as viewing distance was changed, scale invariance was severely disrupted. Micropattern length was found to have a small disruptive effect on scale invariance when manipulated in this way.

The first question that needs to be addressed is why none of the stimulus manipulations fully disrupted scale invariance. Full disruption would have been evidenced by a slope of -1 on the plots of TSF_{min} versus viewing distance, when TSF_{min} was expressed in cps. For the feature manipulation producing most disruption of scale invariance, namely micropattern carrier spatial frequency, the slope was -0.53. Two additional factors besides luminance spatial frequency could contribute to the scale invariance of the OMF and would operate to reduce the potentially disruptive effect of holding micropattern spatial frequency constant in the retinal image. The first is sampling. The finite number of micropatterns in the stimuli inevitably places an upper limit on the physical representation of the high texture frequencies, and this upper limit would be invariant with viewing distance. In other words, we would expect a scale invariant rise in thresholds beyond a certain physical texture spatial frequency due to sam-

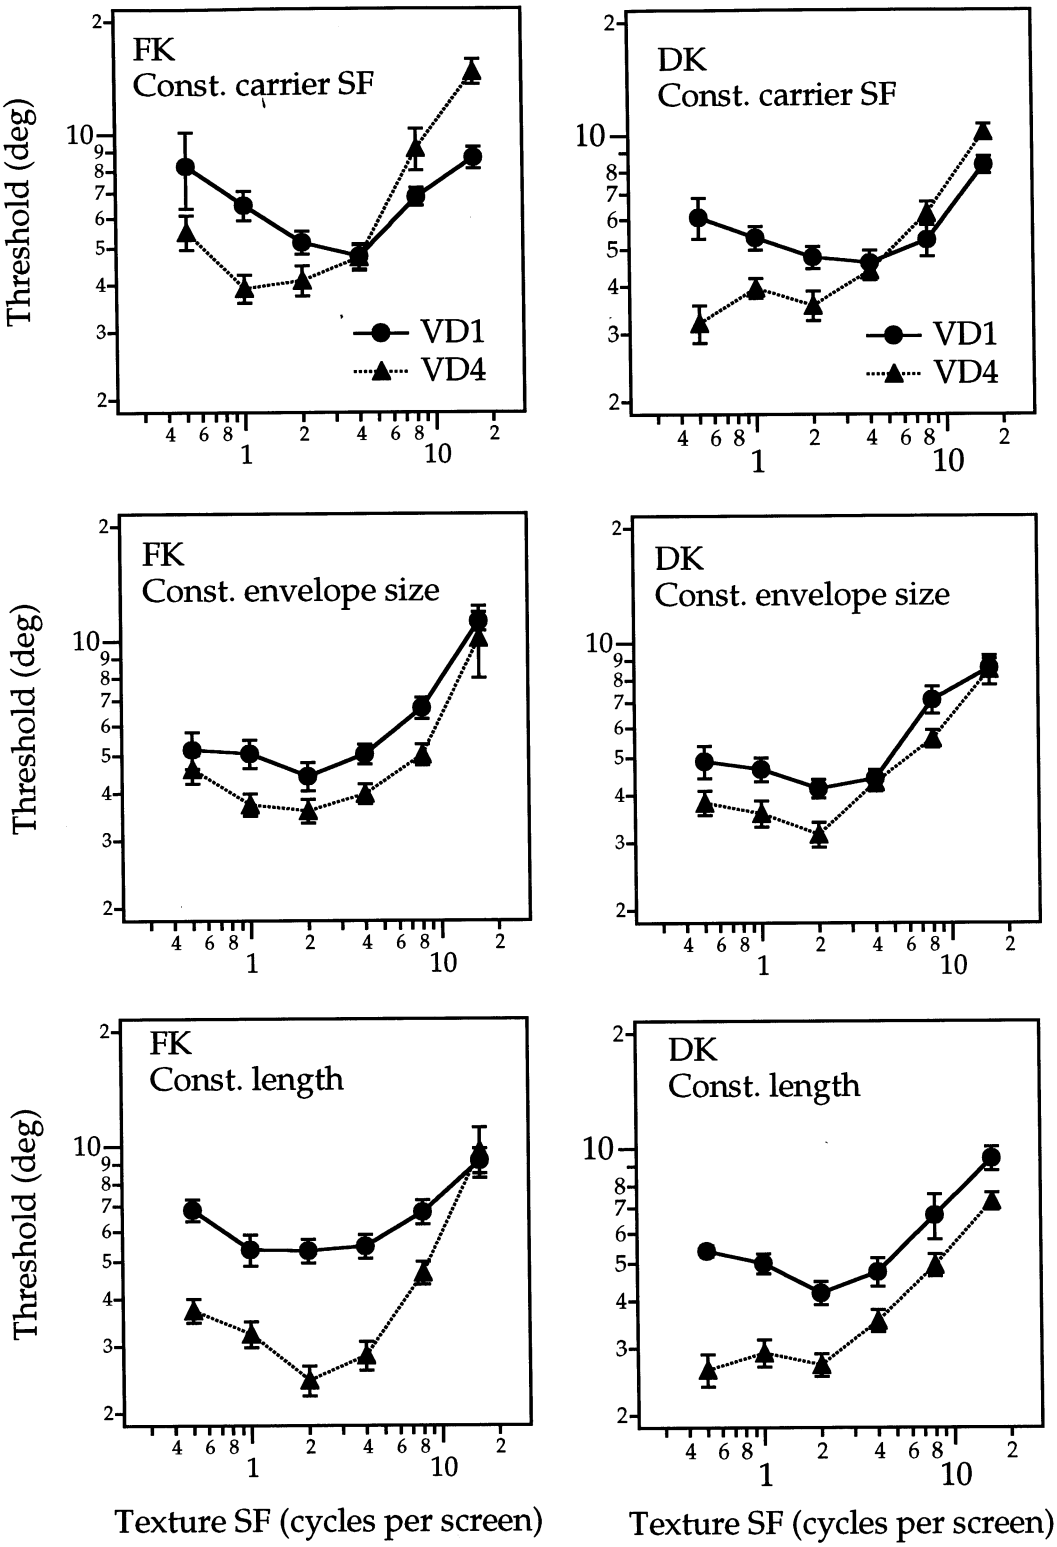

Fig. 7. OMFs at two viewing distances with the micropatterns physically altered in one of three ways to maintain a particular feature constant in the retinal image. Top, constant carrier spatial frequency f_1; middle, constant envelope size σ; bottom, constant envelope length σ_1.

pling limitations. To test this idea we performed an additional experiment in which we measured OMFs at three viewing distances, while holding the density of micropatterns, defined as the number of micropatterns per unit visual angle, constant in the retinal image at 0.41 micropatterns degree^{-2}. The results are shown in Fig. 9 for three subjects. The mean slope across subjects is -0.13. This slope is very small, and we must there-

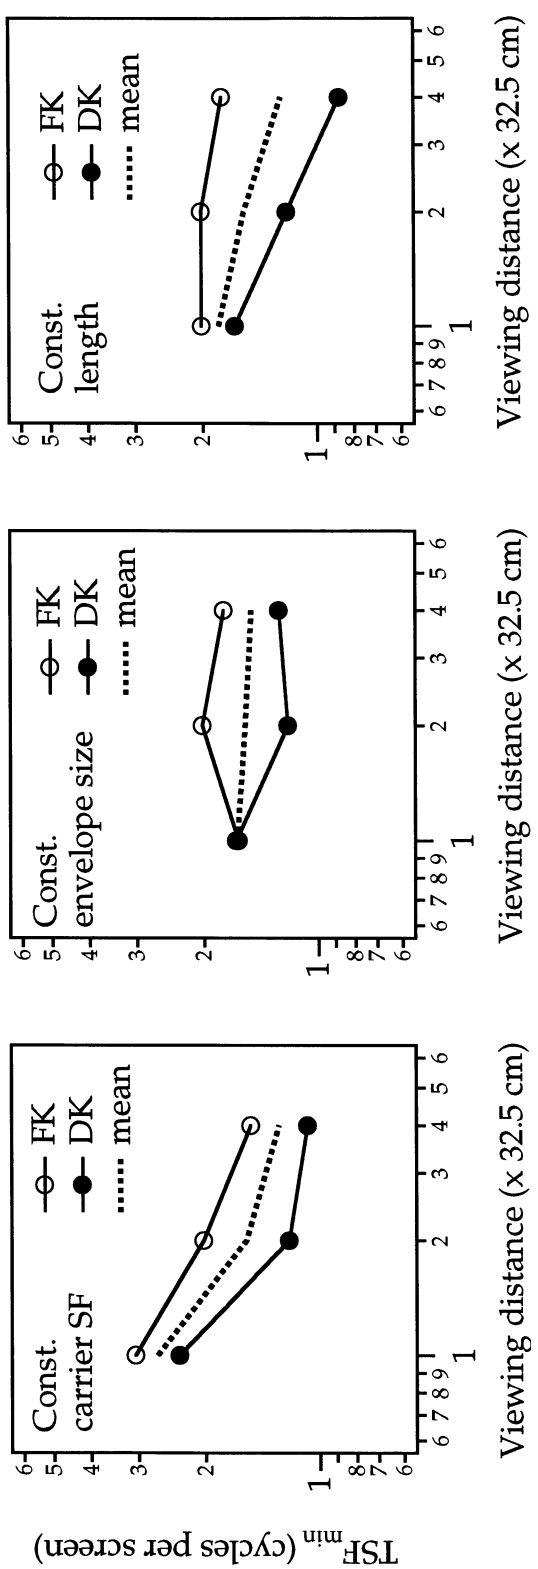

Fig. 8. TSF_{min} as a function of viewing distance for the data in Fig. 7 plus the standard condition (VD2).

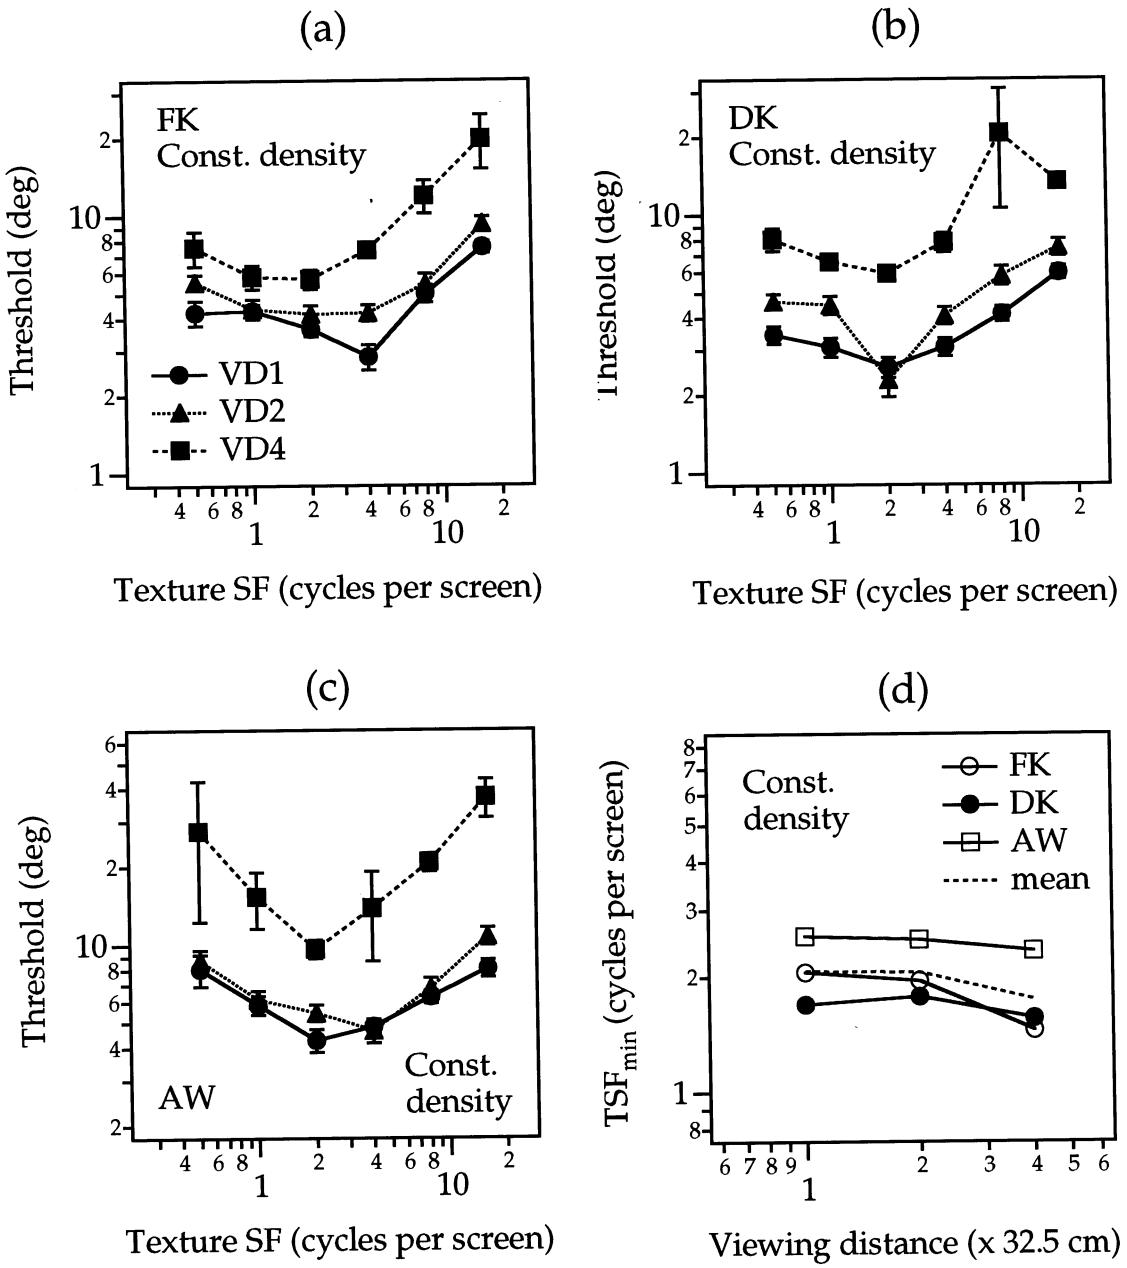

Fig. 9. OMFs at three viewing distances with a constant micropattern density in the retinal image. (a–c) Data for three subjects. (d) Gives TSF_{min} as a function of viewing distance.

fore conclude that any sampling limitation resulting from the finite number of micropatterns appears to make only a very small contribution towards scale invariance. A second factor besides micropattern scale that could contribute to scale invariance is the finite size of the stimulus. If the mechanisms processing orientation-defined texture gradients were to summate over a fixed number of texture cycles, texture frequencies with fewer cycles than the summation limit would be relatively impaired and this impairment would result in a scale invariant, low texture frequency decline in performance. Although we have evidence from our first study (Kingdom, Keeble & Moulden, 1995) that the physical size of the stimulus sets the summation limit rather than the number of texture cycles, we cannot rule out a contribution of this factor to scale invariance. We can legitimately speculate however that were we to repeat our measurements using both very dense and very large displays, in which the two additional factors mentioned here would be expected to have little effect, the disruptive effects of micropattern spatial frequency on scale invariance would be even more dramatic than observed here.

Why might luminance spatial frequency underlie scale invariance in the detection of orientation-defined texture gradients? To answer this question it is worth

considering some current ideas on textures processing. One influential view is the 'back-pocket model' of human texture segregation (Chubb & Landy, 1991). This model has two principle filtering stages. The first convolves the image with a set of orientation-selective linear filters, whose outputs are then subject to a non-linearity (such as squaring or rectification) to produce an orientation-energy map. The second-stage convolves this energy map with a local spatial derivative filter which detects sharp changes in texture energy at each orientation. This model has support primarily from studies on texture segregation, which typically use stimuli with sharp, abutting texture boundaries. Both the model and the data which support it have led to the notion that texture segregation mechanisms are 'edge-based' (Nothdurft, 1985; Landy & Bergen, 1991; Nothdurft, 1991; Wolfson & Landy, 1995, 1998). A different paradigm has dealt with the discriminability of textures, in which subjects are required to compare pairs of textures which do not necessarily abut. These studies have emphasised the encoding of texture properties over relatively large areas, with the implication that the mechanisms involved are 'region-based' (Gurnsey & Laundry, 1992; Keeble, Kingdom, Moulden, & Morgan, 1995; Keeble, Kingdom & Morgan, 1997; Wolfson & Landy, 1998). Wolfson & Landy (1998) have attempted to determine the conditions under which texture processing is edge- or region-based. They contend, for example, that the detection of sharp, abutting changes in texture mean orientation involves an edge-based mechanism, while the detection of sharp, abutting changes in texture orientation variance involves a region-based mechanism. Our previous studies using textures which vary continuously in their mean orientation across space (e.g. Fig. 1), suggest a mechanism which is both edge- and region-based. We found that the detection of a variety of orientation-defined texture waveforms could be modelled by a second-stage texture filter whose receptive field was somewhat bandpass, as in an edge-based model, but which was maximally sensitive to relatively low texture spatial frequencies (typically 50 times lower than the centre luminance spatial frequency of the micropatterns), consistent with the large receptive field of a region-based model (Kingdom & Keeble, 1996). Moreover, in our first study we provided evidence that co-operative interactions between co-aligned first-stage orientation-selective filters might occur prior to the second-stage (having the beneficial effect of reducing local orientation noise), again in keeping with a region-based mechanism (Kingdom, Keeble & Moulden, 1995). Thus, a possible framework for the findings of the present study is a two-stage model, consisting of first-stage (luminance-contrast-sensitive) orientation-selective filters and second-stage texture filters with large, centre-surround receptive fields tuned to relatively low texture spatial frequencies. Our results suggest a range of second-stage filter sizes, each selective for the luminance spatial frequency of its first-stage inputs. In other words, the detection of fine-scale texture variations is subserved predominantly by small second-stage filters with high spatial frequency first-stage inputs, while the detection of coarse-scale texture variations is subserved predominantly by large second-stage filters with low spatial frequency first-stage inputs. The idea that second-stage texture filters are tied in spatial scale to their first-stage inputs has been suggested previously by Bergen (1991), ourselves (Kingdom, Keeble & Moulden, 1995) and most recently by Sutter, Sperling & Chubb (1995). We are currently conducting experiments to establish directly the spatial scale selectivity of the second-stage mechanisms involved in detecting orientation-defined texture gradients.

We can only speculate here on the reason for the small disruptive effect on scale invariance found for subject FK when micropattern length was held constant in the retinal image. As Fig. 7 (bottom) shows, the effect resulted from a relative lack of improvement in thresholds when going from the VD1 to VD4 condition for the high but not low texture spatial frequencies. In the VD4 condition the micropatterns were most physically elongated and thus most likely to overlap along the direction of their axes. This could have had a selectively disruptive effect on the high texture spatial frequencies, because the overlap between micropatterns may have partially blurred the relatively rapid changes in their orientation.

What of orientation-defined textures made up of spatially broadband micropatterns, such as simple line elements? Kingdom, Keeble and Moulden (1995) found both line segment and Gabor micropattern OMFs to be scale invariant. If texture scale invariance is primarily due to changes in luminance spatial frequency with viewing distance, then this implies that line segment textures are as selective in terms of the first-stage filters they stimulate as are Gabor micropattern textures. It has been suggested that the visual system is scale selective when processing orientation-defined textures, because there is only a narrow range of spatial frequencies at which the key textural information is most efficiently represented (Dakin, 1996). If so, then such scale selectivity would have the additional benefit of producing scale invariance.

It has been argued that scale invariance is an important property of perceptual systems because it simplifies the computation of absolute shape (e.g. Parker, 1993). The spatial changes in local orientation exemplified by our stimuli would most likely arise in natural scenes from textured surfaces folded in depth, as for example in the retinal image of an undulating field of corn. For such ground surface textures the elements themselves need not be physically oriented for them to become

oriented in the retinal image because of foreshortening. The mechanisms involved in detecting the orientation modulation in our stimuli may therefore be closely related to those concerned with processing shape-from-texture.

Acknowledgements

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

References

Bach, M., & Meigen, T. (1992). Electrophysiological correlates of texture segregation-effect of orientation gradient. *Investigative Ophthalmology and Visual Science*, *A1349*(Suppl. 33), 962.

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

Chubb, C., & Landy, M. S. (1991). Orthogonal distribution analysis: a new approach to the study of texture perception. In M. S. Landy, & J. A. Movshon, *Computational Models of visual processing* (pp. 291–301). Cambridge, MA: MIT Press.

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

Gibson, J. J. (1979). *The ecological approach to visual perception*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Gurnsey, R., & Laundry, D. S. (1992). Texture discrimination with and without abrupt texture gradients. *Canadian Journal of Psychology*, *46*, 306–332.

Howell, C. R., & Hess, R. F. (1978). The functional area for summation to threshold for sinusoidal gratings. *Vision Research*, *18*, 369–374.

Jamar, J. H. T., & Koenderink, J. J. (1983). Sine-wave gratings: scale invariance and spatial integration at suprathreshold contrast. *Vision Research*, *23*, 805–810.

Keeble, D. R. T, Kingdom, F. A. A, Moulden, B., & Morgan, M. J (1995). The detection of orientationally multimodal textures. *Vision Research*, *35*, 1991–2005.

Keeble, D. R. T., Kingdom, F. A. A., & Morgan, M. J. (1997). The orientational resolution of texture perception. *Vision Research*, *37*, 2993–3007.

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.

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

Nothdurft, H. C. (1985). Sensitivity for structure gradient in texture discrimination tasks. *Vision Research*, *25*, 1957–1968.

Nothdurft, H. C. (1991). Texture segmentation and pop-out from orientation contrast. *Vision Research*, *31*, 1073–1078.

Parker, A. (1993). Solid shape and the natural world. *Current Biology*, *3*, 401–403.

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

Van Meeteran, A., & Barlow, H. B. (1981). The statistical efficiency for detecting sinusoidal modulation of average dot density in random figures. *Vision Research*, *21*, 765–777.

Wolfson, S. S., & Landy, M. S. (1995). Discrimination of orientation-defined texture edges. *Vision Research*, *35*, 2863–2877.

Wolfson, S. S., & Landy, M. S. (1998). Examining edge- and region-based texture analysis mechanisms. *Vision Research*, *38*, 439–446.