Information limit on the spatial integration of local orientation signals

Steven C. Dakin

Institute of Ophthalmology, University College London, 11-43 Bath Street, London EC1V 9EL, UK

Received August 1, 2000; revised manuscript received November 14, 2000; accepted November 15, 2000

Channel-based models of human spatial vision require that the output of spatial filters be pooled across space. This pooling yields global estimates of local feature attributes such as orientation that are useful in situations in which that attribute may be locally variable, as is the case for visual texture. The spatial characteristics of orientation summation are considered in the study. By assessing the effect of orientation variability on observers' ability to estimate the mean orientation of spatially unstructured textures, one can determine both the internal noise on each orientation sample and the number of samples being pooled. By a combination of fixing and covarying the size of textured regions and the number of elements constituting them, one can then assess the effects of the texture's size, density, and numerosity (the number of elements present) on the internal noise and the sampling density. Results indicate that internal noise shows a primary dependence on texture density but that, counterintuitively, subjects rely on a sample size approximately equal to a fixed power of the number of samples present, regardless of their spatial arrangement. Orientation pooling is entirely flexible with respect to the position of input features.

OCIS codes: 330.5000, 330.5510, 330.6110.

1. INTRODUCTION

A. Background

Cells in primary visual cortex respond selectively to the presence of image attributes such as orientation, spatial frequency, and direction.[1] The cornerstone of computational theories of spatial vision is that one can model this selectivity using spatial filtering. Filter kernels are constructed to reflect the known structure of cortical-cell receptive fields[2] and are then convolved with a stimulus to produce a response image. The response image is the output of a matrix of overlapping receptive fields where every pixel reflects the response of an independent mechanism centered at that location. To make use of such information, computational models of early spatiotemporal vision pool values from the response image within a locale. For example, so-called "back-pocket" models of visual texture segregation employ banks of Gabor filters whose (squared) outputs are integrated over space by some second-stage mechanism.[3–6] Although a great deal of psychophysical evidence pertains to the filters, less is known of the nature of the spatial integration process that follows it. Probably the most intensively studied issue is how filters near their threshold are combined to predict contrast sensitivity. Observers' contrast sensitivity improves with the area of the target. Initial improvements at small stimulus areas are attributed to improved summation by a single filter, but continued improvement can be modeled by assuming a probabilistic combination of the outputs of multiple filters.[7–10] Such models are generally able to successfully predict subjects' detection of stimuli composed of multiple targets[11] (although there may be exceptions[12]). Recently it has also been suggested that the spatial extent of this pooling is anisotropic and dependent on the shape of the filters being combined.[13]

The functional significance of pooling is to increase sensitivity to near-threshold, spatially distributed signals and to confer resistance to local noise. With respect to other image attributes, such as orientation and direction, it would seem likely on similar grounds that some form of integration across space takes place. In the motion domain, subjects can detect extended motion trajectories, which implicates the pooling of local motion detectors across space.[14] Orientation integration has been observed by using several paradigms; observers are capable of detecting extended visual contours that are invisible to the outputs of single linear spatial filters.[15,16] It is also known that estimates of orientation can be combined in the absence of contour structure. Specifically, observers confronted with a spatially unstructured texture composed of oriented elements are able to makes estimates of the mean orientation and orientation variance with great accuracy.[17] For example, orientation thresholds for textures composed of elements with considerable orientation variability are comparable to thresholds measured with sinusoidal gratings[17] [on the order of 1–2°; Fig. 1(a)], even though the latter is entirely free of extrinsic orientation variation. This suggests that internal noise limits performance on all orientation judgments. Dakin and Watt[17] showed that observers overcome this noise by relying on the average orientation pooled over a series of local estimates and that observers also had access to second-order (i.e., variance) but not third-order (i.e., skew) orientation statistics. This paper addresses the issue of how the visual system determines the sample it uses to compute such orientation statistics.

Does the extent of spatial pooling depend on the size of an object or on the density of features covering it? With respect to texture processing there is some indication that subjects' performance exhibits scale invariance; that is, it

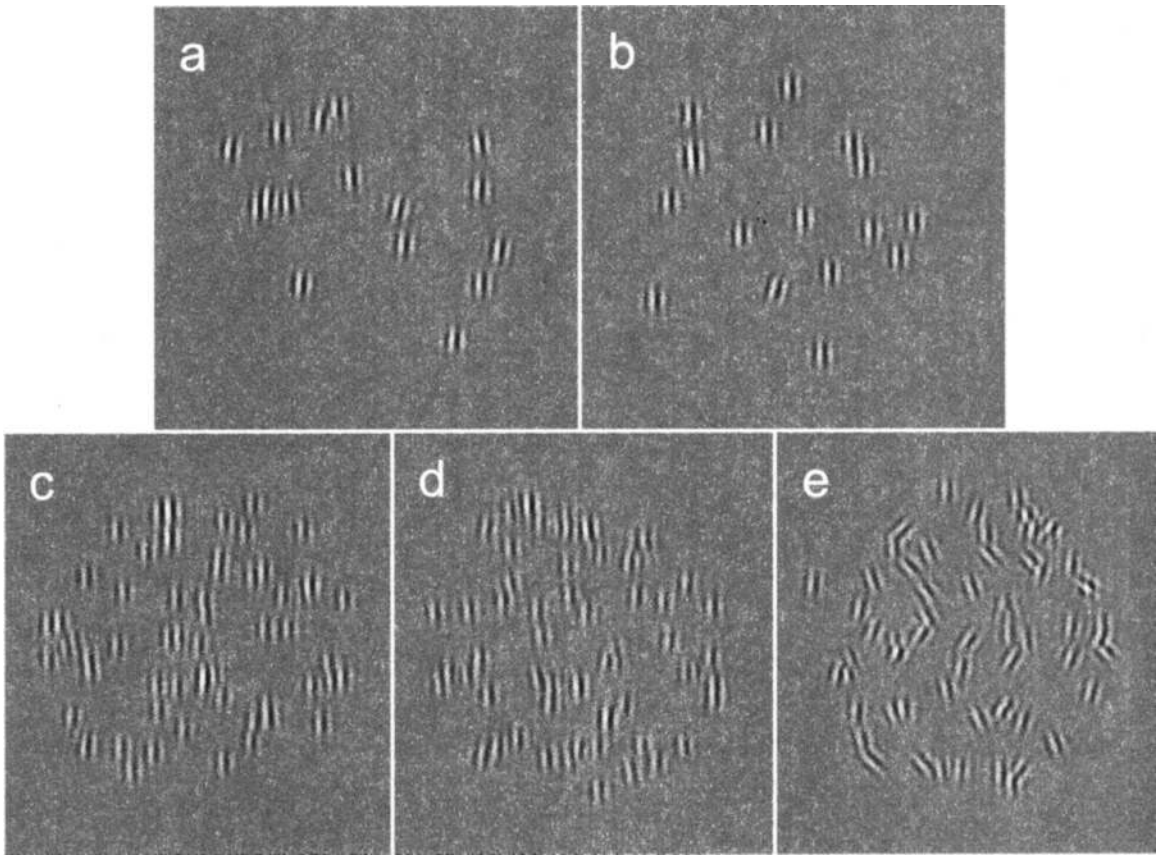

Fig. 1. (a) Texture composed of 16 Gabor micropatterns with orientations drawn from a Gaussian random distribution with a mean of 92° (2° clockwise from vertical) and a standard deviation of 4°. Observers' ability to accurately discriminate mean orientation in the presence of considerable orientation variance necessitates pooling of multiple orientation estimates. (b) A "pop-out" stimulus. A single element, tilted 15° from vertical, is embedded in a field of vertical distractors. The presence of the discrepant element could be signaled by a simple local orientation statistic (such as the mean). (c–e) Examples of the stimuli used in the experiments (contrast enhanced for the purpose of reproduction). Each is composed of 64 Gabor elements randomly distributed within a circular region with orientations drawn from a Gaussian distribution with a mean of 90° and a standard deviation equal to (c) 0.5°, (d) 4°, and (e) 23°. Subjects' ability to estimate mean orientation deteriorates with increasing orientation variance, allowing one to estimate the effective internal noise and the number of orientation samples being employed. The textures shown in (c–e) fall in the midrange of the patch sizes, densities, and numerosities tested.

does not vary greatly with viewing distance. Kingdom and Keeble[18] measured sensitivity to sinusoidal modulations of orientation across space and found that element spatial frequency could interfere significantly with scale invariance but that local element density did not. Density has generally been viewed as important only insofar as it yields information about local orientation gradients.[19–21] Notably, the detectability of an oriented target amidst a field of oriented distractors [e.g., Fig. 1(b)] exhibits a nonmonotonic dependency on stimulus density.[24] Performance is poor at very low distractor densities, improves with small increases in distractor number, but then deteriorates again at high distractor densities. Although such results have been interpreted as evidence for local-orientation contrast detectors, this result is also consistent with targets being signaled by the change that they produce in simple statistics, such as the mean orientation. Small numbers of distractors lead to a noisy estimate of the mean, but as the number increases this improves, making the target orientation more conspicuous, until eventually the distractors swamp the target's contribution to the statistic.[17] In the same vein, Morgan *et al.*[22] have proposed that crowding in visual search displays can be accounted for by a texture analyzer that computes mean orientation.

While the work of Sagi and Julesz[21] suggests that density is the critical parameter for pop-out tasks, these authors considered its effect only for a single region size; i.e., they systematically confounded density with numerosity. The response of neurons in cat striate cortex that are responsive to texture boundaries is known to be strongest for densely spaced elements.[23] Observers' psychophysical texture segregation shows a similar dependence. Furthermore, the maximum tolerable separation of elements depends on the orientation difference at the texture boundary; elements with large orientation differences can be more widely separated than elements with smaller orientation differences.[19] These data suggested that it is change in orientation in space (the structure gradient) that determines strength of segregation. However, because the number of elements has generally been kept constant in these experiments, increasing spacing necessarily decreases sampling density so that one cannot be sure which parameter determines performance.[20] In

summary, according to Gorea and Papathomas, "the nature and significance of the density parameter in global orientation processing remain unclear." [24]

B. Logic of the Experiments

The goal of the experiments described below is to illuminate observers' sampling strategy for local orientation and specifically to determine whether texture size, texture density, or texture numerosity determine the density and precision of orientation sampling. This presents two problems: how to estimate sampling precision and density and how to separate out the effects of size, density, and numerosity. I deal with these points in turn.

The idea of understanding a complex system by exam-

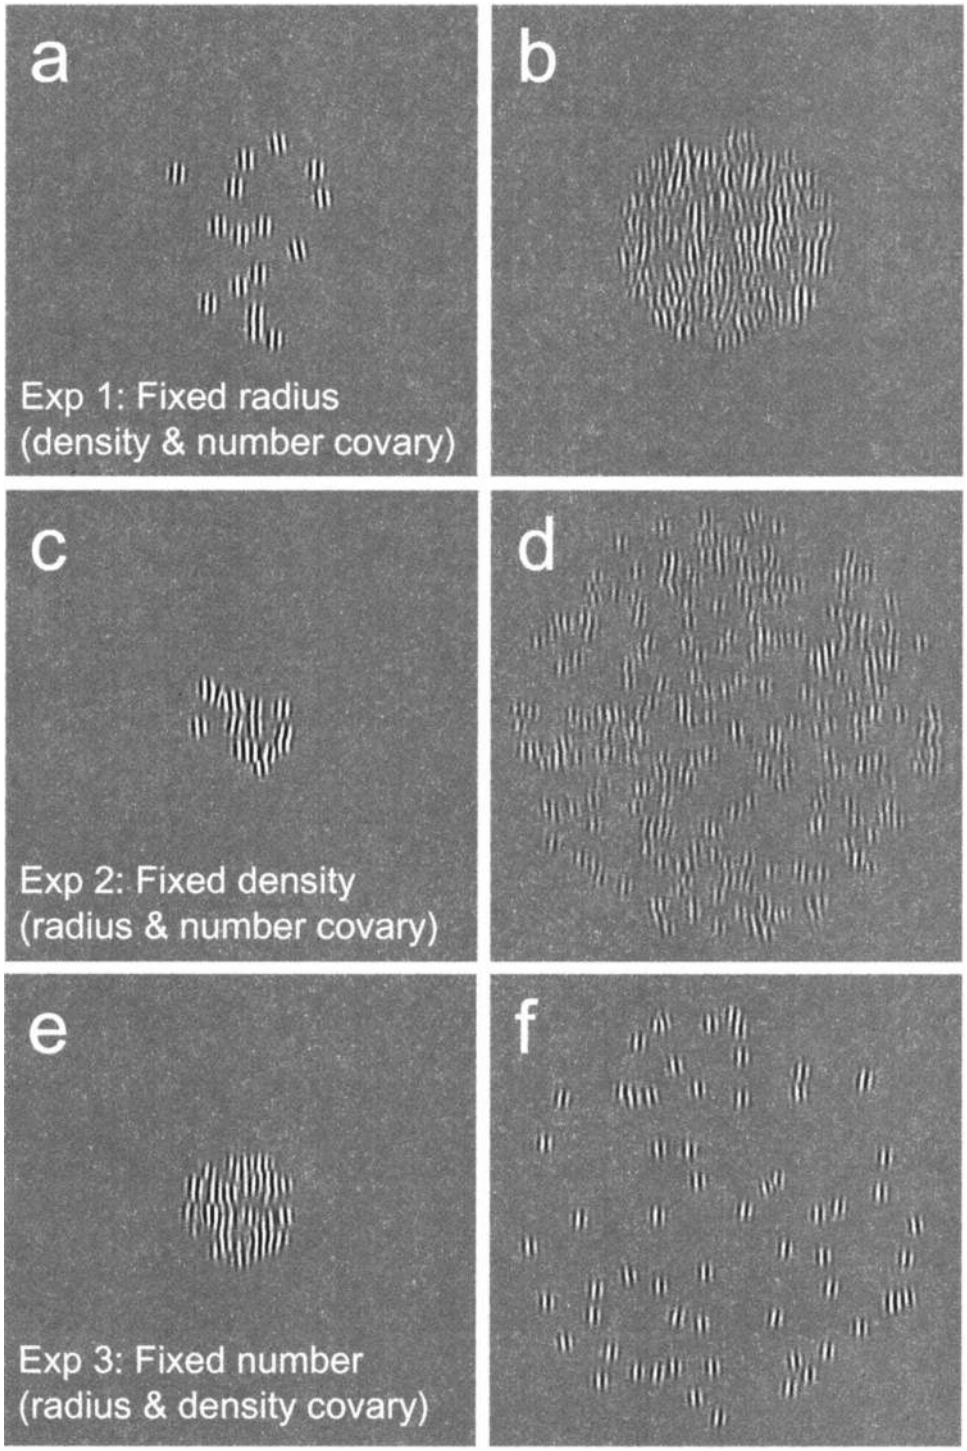

Fig. 2. Examples of the stimuli; the orientation of elements is Gaussian distributed with a mean of 90° and a standard deviation of 4°. (a, b) Examples from the fixed radius condition. Textures are composed of (a) 16 elements and (b) 256 elements falling within a circular region with a radius of 3.5°. (c, d) Examples from the fixed density condition. Textures are composed of (c) 16 elements with a 1.7° radius and (d) 256 elements with a 7.0° radius. Thus density is fixed at 5.1 elements per degree squared. (e, f) Examples from the fixed number condition. Textures are composed of 256 elements falling in a region of radius (e) 1.7° and (f) 7.0°.

ining how its behavior deteriorates with the addition of noise is an engineering technique introduced to vision sci ence by Barlow.[25,26] The equivalent-noise approach has been applied to various visual acuities that require integration, including luminance offset detection,[26] coding of spatial position,[27,28] discrimination of edge blur[29] spatial frequency acuity,[30] contrast detection,[31,32] contour integration,[33] orientation discrimination,[34,35] and the characterization of attentional strategies.[36] In words, the approach is as follows: Given that thresholds are estimates of response variance, observers' nonideal behavior when presented with noiseless stimuli can be expressed as additive, internal noise. The simplest way to measure the amount of internal noise is to add increasing levels of external noise to a stimulus and determine the point at which subjects' performance begins to deteriorate. If the task requires integration, subjects' resistance to increasing levels of external noise will depend decreasingly on internal noise and increasingly on how many samples they are averaging over. This system view is usually expressed in the form of a variance-summation model,

$$\sigma_{\mathrm{obs}} = \left(\sigma_{\mathrm{int}}^2 + \sigma_{\mathrm{ext}}^2/n\right)^{1/2}, \tag{1}$$

where σ_{obs} is the observed threshold, σ_{ext} is the external noise, σ_{int} is the equivalent intrinsic or internal noise, and n is the number of samples being employed. In terms of an orientation discrimination task, σ_{obs} corresponds to the threshold orientation offset, σ_{ext} corresponds to the orientation bandwidth of the stimulus, σ_{int} corresponds to the noise associated with the measurement of each orientation sample and their combination, and n corresponds to the number of orientation samples being combined. Note that σ_{int}^2 is the total internal noise for a given condition. It is likely that this value is actually the sum of noise resulting from the decision process as well as sampling uncertainty. The experiments described are unable to separate the effects of these multiple sources of uncertainty, so later I treat them as a single noise source whose value is fitted for each condition. It is usual to assess these parameters by measuring the threshold orientation offset for a mean orientation judgment (σ_{obs}) as a function of the orientation bandwidth of the texture (σ_{ext}) and then to fit these data with the variance-summation model giving estimates of σ_{int} and n. Heeley *et al.*[35] used exactly this approach to determine whether the oblique effect (meridional anisotropies for various psychophysical tasks) is due to undersampling or to increased intrinsic uncertainty. They showed that subjects' average intrinsic uncertainty was higher and that their average sample size was smaller for oblique stimuli ($\sigma_{\mathrm{int}} = 3.26°$, $n = 18.9$) compared with vertical stimuli ($\sigma_{\mathrm{int}} = 2.08°$, $n = 28.8$), so that the oblique effect could be wholly attributed to neither.

Given that the equivalent-noise paradigm may be used to determine the sampling characteristics of an orientation judgment, consider the second problem: how to determine which stimulus factors determine these values. The first departure from the technique used by Heeley *et al.*[35] is in the type of stimulus used. Dense, filtered noise patterns do not give one control over the local density of the patterns. Watt[34] used line patterns that allow one to control density but are spatially broadband. I chose to use stimuli composed of a field of discrete narrow-band Gabor micropatterns [for examples, see Fig. 1(c) and 1(d)]. Both the spatial bandwidth and the density of these patterns may be controlled, and, because overlapping elements are summed, these textures are indistinguishable from filtered noise textures at sufficiently high texture element densities.

I consider three spatial attributes of a visual texture that are likely to influence orientation integration: a textured region size, density, and numerosity (the number of elements constituting it). These parameters cannot be varied independently; a change in one is always accompanied by a change in at least one other. For example, if one fixes the size of the textured region and alters the number of elements falling in it, both numerosity and density change. Consequently, it is not possible to determine from any one manipulation which is the parameter of importance. For example, if the number of samples the observer uses to compute mean orientation is observed to depend strongly on density in one condition, it is necessary to ascertain that it does *not* depend on numerosity in another condition, since one must have covaried the two parameters in the first condition. To allow for this, the design that I used investigates all three possible manipulations: fixed radius (numerosity and density covary), fixed density (radius and numerosity covary) and fixed numerosity (radius and density covary). Examples of these manipulations are shown in Fig. 2. If one fits the variance-summation model to the data derived from each condition, one may estimate the internal noise and the number of samples used to perform the task. If one of the three spatial parameters determines the observers' integration strategy, then one should observe a strong effect (on the best-fitting variance model parameters) of the independent variable in the two conditions where the critical parameter varied but none in the condition where it was fixed.

2. METHODS

A. Equipment

An Apple Macintosh G3 computer controlled stimulus presentation and recorded subjects' responses. The programs for running the experiment were written in the Matlab environment (Mathworks, Ltd.) with code from the Psychophysics Toolbox[37] and the VideoToolbox[38] packages. Stimuli were displayed on a 19-in. Sony Multiscan 400PS color monitor driven by a Mac Picasso 850 graphics card (Villagetronic Ltd). The screen had a resolution of 1280×1024 pixels and operated at a frame refresh rate of 85 Hz. Pseudo-12-bit contrast accuracy was achieved by electronically combining the red–green–blue outputs from the graphics card by using a video attenuator.[39] A monochrome signal was generated by amplifying and sending the same attenuated signal to all three guns of the CRT. The display output was calibrated using a Minolta CS100 photometer and output luminances that were linearized with a lookup table. The screen was viewed binocularly at a distance of 36 cm, so that 1 pixel on the screen subtended 2.4 arc min of visual

angle. The display had a background luminance of 48 cd/m^2.

B. Stimuli

Stimuli were 256-pixel (10.2°) square images containing a spatially unstructured patch of texture. Textures were generated by distributing a number of Gabor micropatterns, each with 10% Michelson contrast, throughout the image. For the experiments reported, the peak spatial frequency of the carrier was 4 cycles per degree (λ = 15 arc min), and the standard deviation of the Gaussian envelope was 10.56 arc min. This spatial frequency was chosen because it falls well within the range of spatial frequencies that produce optimal orientation discrimination.[40,41]

Overlapping patches were added. Gray levels falling outside the permissible range were clipped at the maximum or minimum gray level accordingly. Depending on the condition, 4–1024 elements were used. Two positioning strategies were used in two separate conditions. In the first, Gabors were positioned using a Gaussian random distribution with a mean at the center of the image and a standard deviation of 0.5–8.0°. The second positioning strategy was similar except that a uniform random distribution was used with ranges of ±0.9 to ±13.8° (chosen so that spatial standard deviations were matched to the Gaussian condition). Elements falling outside of a circular window (with a radius equal to the positional range) were repositioned within the circular region. There are several reasons for having two positioning strategies. First, the problem with uniform random distributions is that, assuming the integration processes operate outwardly from fixation, some patterns may be sparse at their centers, producing poor performance. Second, although Gaussian distributions alleviate this problem at low densities, at high densities the nonuniform distribution leads to outliers which subjects could rely on. Finally, by comparing the estimates for sampling precision and density for one subject across the two positioning conditions, one may be able to infer the underlying spatial distribution of the local mechanisms; the sampling strategy closer to this spatial distribution should lead to greater efficiency.

The orientation of elements was selected from a Gaussian distribution with a mean equal to the cued orientation [i.e., 90°± the cue generated by the adaptive probit estimation (APE) procedure] and a standard deviation equal to 0.5°, 1.0°, 2.0°, 4.0°, 8.0°, 16.0° or 23.0°. As described above, the logic of the experiments is to examine the effects of density, size, and numerosity by fixing each parameter in three separate conditions and by allowing the other two parameters to vary. Figure 3 gives the specific spatial parameters of the textures used. If one considers the effects of fixing the texture radius, for example, notice that this design yields not only the five subconditions associated with the fixed-radius experiment (the horizontal row of five boxes), but also the four pairwise estimates of the effect of fixing the radius at four different values. For example, data measured with a patch radius of 0.9° from the fixed-density experiment can be compared with data measured with a similar patch radius in the fixed-numerosity condition. The results of these pairwise comparisons are given in Fig. 7.

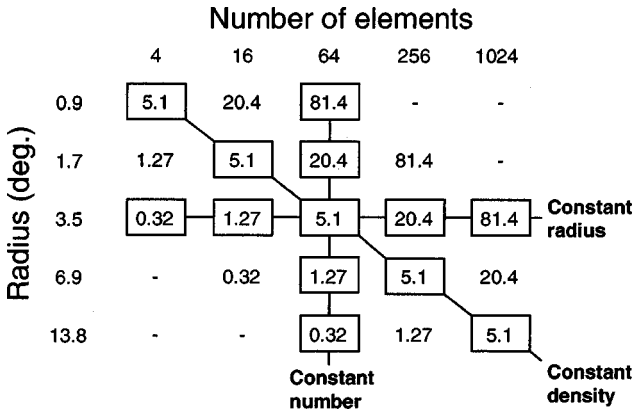

Fig. 3. Texture density (measured in elements per square degree of visual angle) as a function of the patch radius and the number of elements. The boxed cells give the spatial parameters of the textures used in the 13 subconditions tested.

C. Procedure

The subjects' task was a single-interval, binary forced choice. An oriented texture was presented in the center of the display for 100 ms, and the observer judged whether the overall orientation of the texture was tilted clockwise or counterclockwise, compared with their internal standard for vertical. Subjects signaled their responses on the computer keyboard. APE, an adaptive method of constant stimuli[42] was used to sample a range of orientations around vertical. A session consisted of seven interleaved runs of 64 trials, corresponding to the seven levels of orientation variability tested. At least three runs were undertaken for each data point plotted. Data were pooled across all runs performed with a particular stimulus configuration, and a bootstrapping procedure was used to fit a cumulative Gaussian function to the results.[43] This procedure yields estimates of the standard deviation (the reciprocal of slope) and bias parameters of the fitting function as well as estimates of their associated confidence intervals. The term "orientation threshold" will be used throughout to refer to the standard deviation of the best-fitting psychometric function. Generally, subjects showed little systematic bias on this task, and data reported are the standard deviation of the best-fitting cumulative Gaussian function. Error bars show the associated 95% confidence intervals.

D. Subjects

The author (SCD) and two naïve observers (ACM and JS) served as subjects in the experiments. All wore optical correction as required.

3. EXPERIMENTS

A. Experiment 1: Effect of Numerosity and Density on Orientation Estimation of a Texture of Fixed Size

The first experiment considered the effect of covarying the numerosity and density of a fixed-size texture patch on observers' ability to combine local orientation estimates across space. Approximately circular patches of texture were used. For subjects AM and SCD, elements were uniformly randomly distributed throughout a circular region with a radius of 5.1°. For subject JS, a two-

dimensional Gaussian spatial distribution of elements was used, with a standard deviation of 2.0°. The number of texture elements varied between 4 and 1024 so that, for the uniformly distributed textures, density varied between 0.32 and 81.4 elements per square degree. (The Gaussian distribution always leads to nonuniform density). Examples of the uniform spatially distributed stimuli from this experiment are shown in Figs. 2(a) and 2(b).

Results. Data from this experiment are graphed in Fig. 4. Consider data from the 256-element condition (▽), which are typical. Low levels of orientation variability ($\sigma < 8°$) produce little discernable effect on orientation thresholds that remain low (1–2°) and comparable with thresholds that are estimated with gratings. Beyond this level of variability, performance deteriorates steadily. The variance-summation model (Eq. 1) has been fitted to these data, and its predictions are shown as thin curves in Fig. 4. The model captures this trend well, and subjects' performance is generally consistent across runs. The parameters derived from the variance summation fit are shown in the legend for each plot. Results

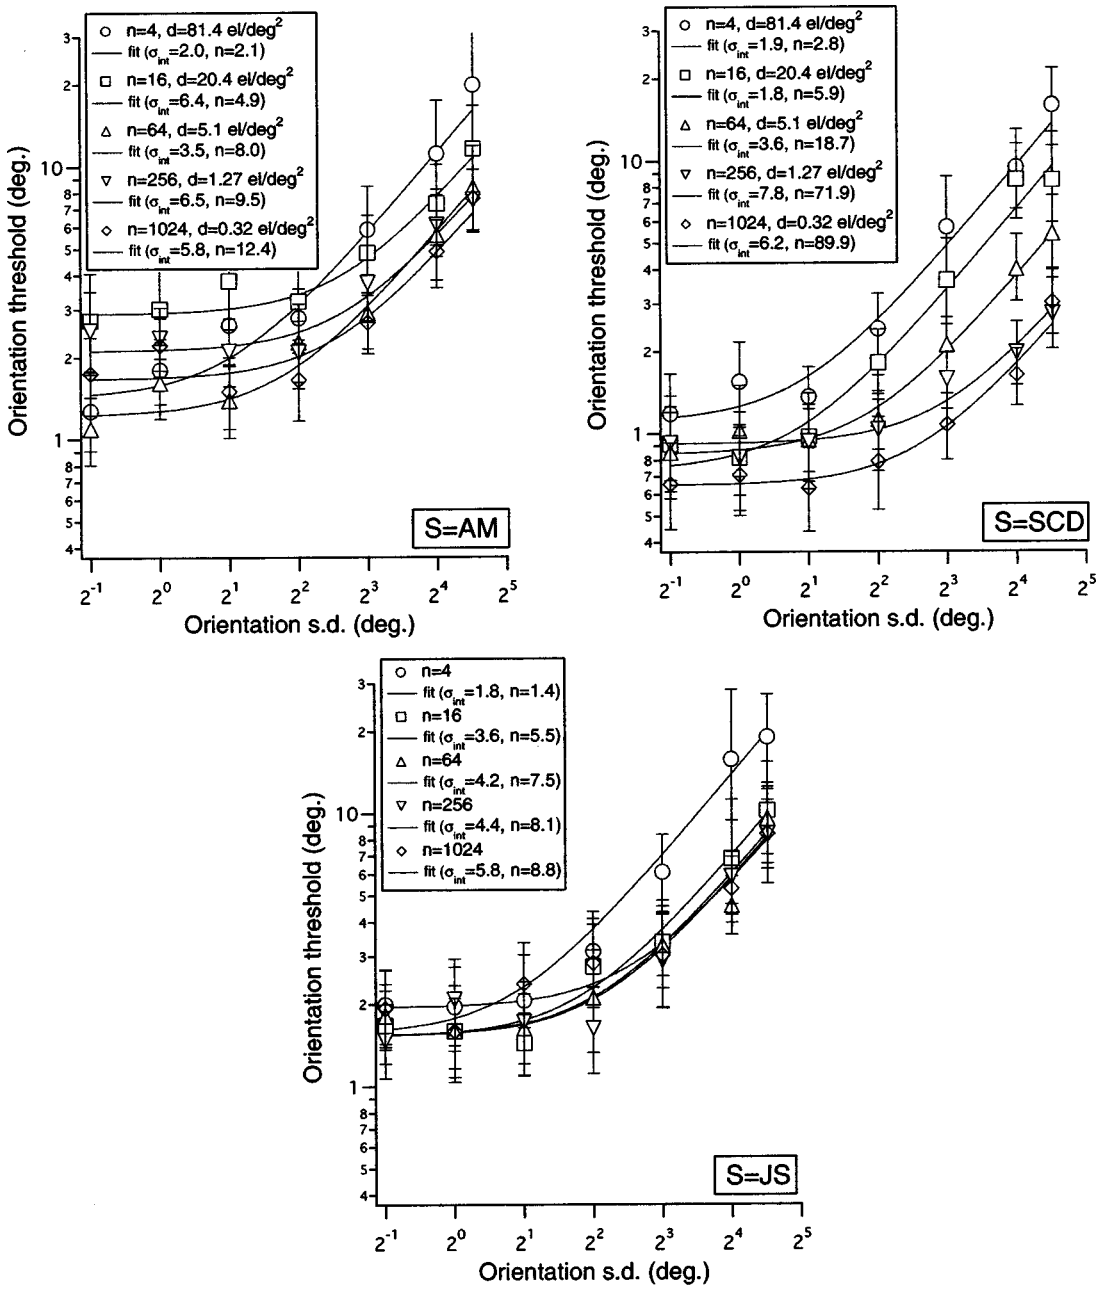

Fig. 4. Data from the fixed-size experiment (experiment 1). Each graph shows the effect of orientation variability (bandwidth) on one subject's judgment of the mean orientation of a texture patch as a function of the number of Gabor microelements constituting the texture. For subjects AM and SCD, Gabor elements were uniformly randomly distributed within a circular window with a radius of 3.5°. For subject JS, Gabors were distributed with a two-dimensional Gaussian spatial distribution. Density covaried with numerosity in this condition. The fits shown are for the variance model described by Eq. 1. Notice that the parameters derived from the fit (shown in the legend) indicate that both the number of orientation samples used and the internal noise increase with the numerosity and density, although efficiency drops.

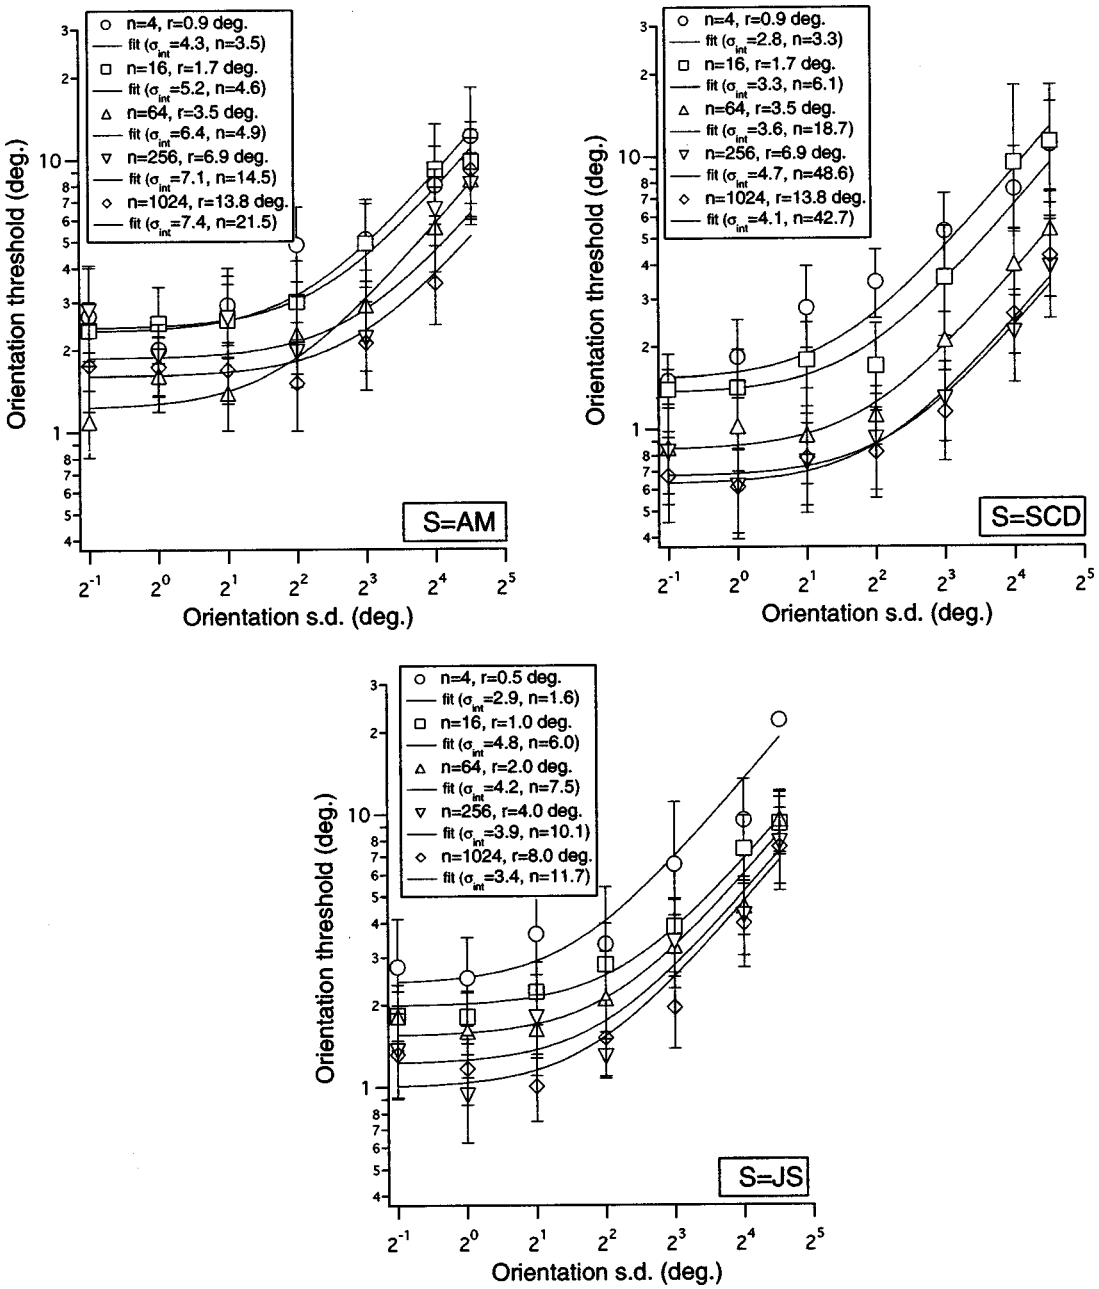

Fig. 5. Data from the fixed-density experiment (experiment 2). Each graph shows the effect of orientation variability on an observer's orientation thresholds measured with textures composed of various numbers of elements. In this condition the size of patches was covaried with numerosity so that the densities of all patches tested were similar. Note that texture density was uniform for subjects AM and SCD but that subject JS was tested with Gaussian-distributed elements. Parameters derived from the variance-summation model are given in the legend to each graph and indicate that, as for the fixed patch-radius condition, progressively more samples are employed as the numerosity and patch size increase. However, there appears to be much less variation in the estimated internal noise across subconditions, compared with experiment 1.

indicate that both the estimated internal noise (σ_{int}) and the number of samples (n) show a strong dependence on the changes in patch density and numerosity.

B. Experiment 2: Effect of Numerosity and Size on Orientation stimation of Textures of Fixed Density

The goal of this experiment was to determine the effect of fixing texture density, while covarying texture patch size and numerosity, on observers' integration of local orientation. Texture density was fixed at 5.1 elements per square degree. The number of elements varied between 4 and 1024, and the patch size varied: The radius fell between 0.32° and 81.4° (in the case of uniformly distributed elements) or the standard deviation fell between 0.5° and 8.0° (in the case of Gaussian element distribution). Examples of the stimuli are shown in Figs. 2(c) and 2(d). Data from this experiment, along with fits from the variance-summation model, are shown in Fig. 5. Curves now appear to be roughly parallel to one another, being shifted with respect to one another on the ordinate axis. This pattern of results indicates a change in the number of samples being used, accompanied by little change in

the internal noise parameter. The number of samples used by subjects depends largely on the sample and region size presented. Internal noise seems to be determined largely by the texture density.

C. Experiment 3: Effect of Density and Size on Orientation Estimation of Textures of Fixed Numerosity

The final experiment determined the effect of covarying patch size and density (such that numerosity was kept constant) on observers' ability to integrate texture orientation information. The number of elements was fixed at 64, and the patch radius and standard deviation varied in a range similar to that of experiment 2. Figures 2(e) and 2(f) show examples of the stimuli from this condition. Results, shown in Fig. 6, show that for low levels of orientation variability, curves are shifted on the ordinate axis with respect to one another (indicating a change in the internal noise parameter). However, subjects' performance tends to converge across the various sizes and densities tested at higher levels of orientation variability. The similar slopes of the fits indicate that similar sample

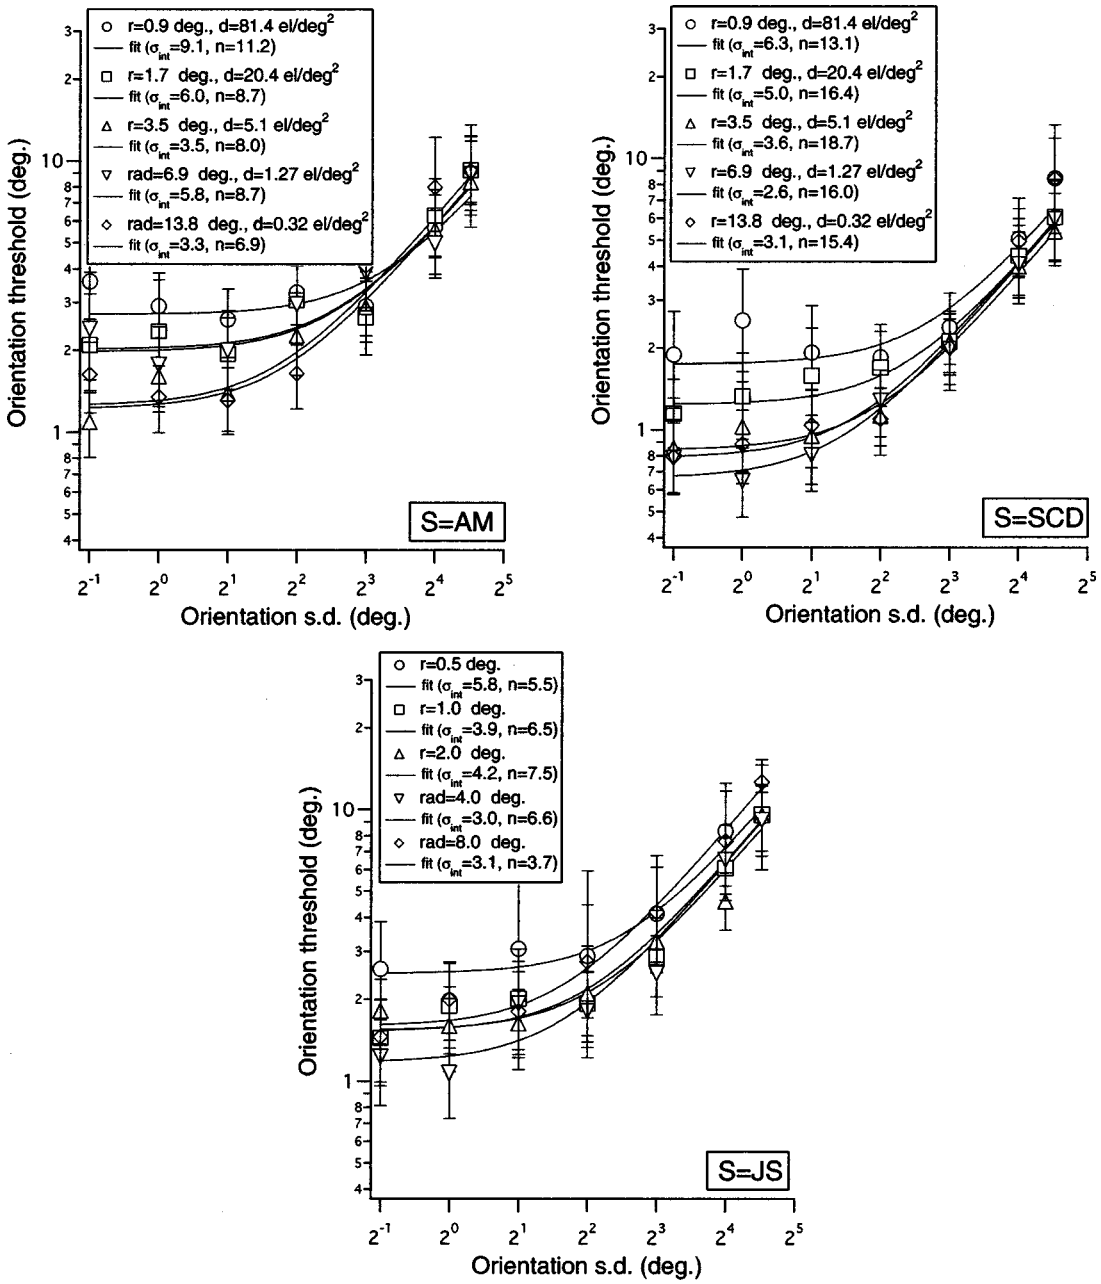

Fig. 6. Data from the fixed-numerosity experiment (experiment 3). Graphs show the effect of orientation variability on a mean orientation judgment as a function of the size of the texture patch. Because the number of texture elements was fixed, patch density covaried with patch size. Unlike previous conditions, parameters derived from fitting the variance model indicate that subjects are now employing a broadly similar number of orientation samples over approximately eight octaves of densities and patch sizes. This suggests that it is the number of samples presented to subjects that determines their sampling strategy regardless of size or density.

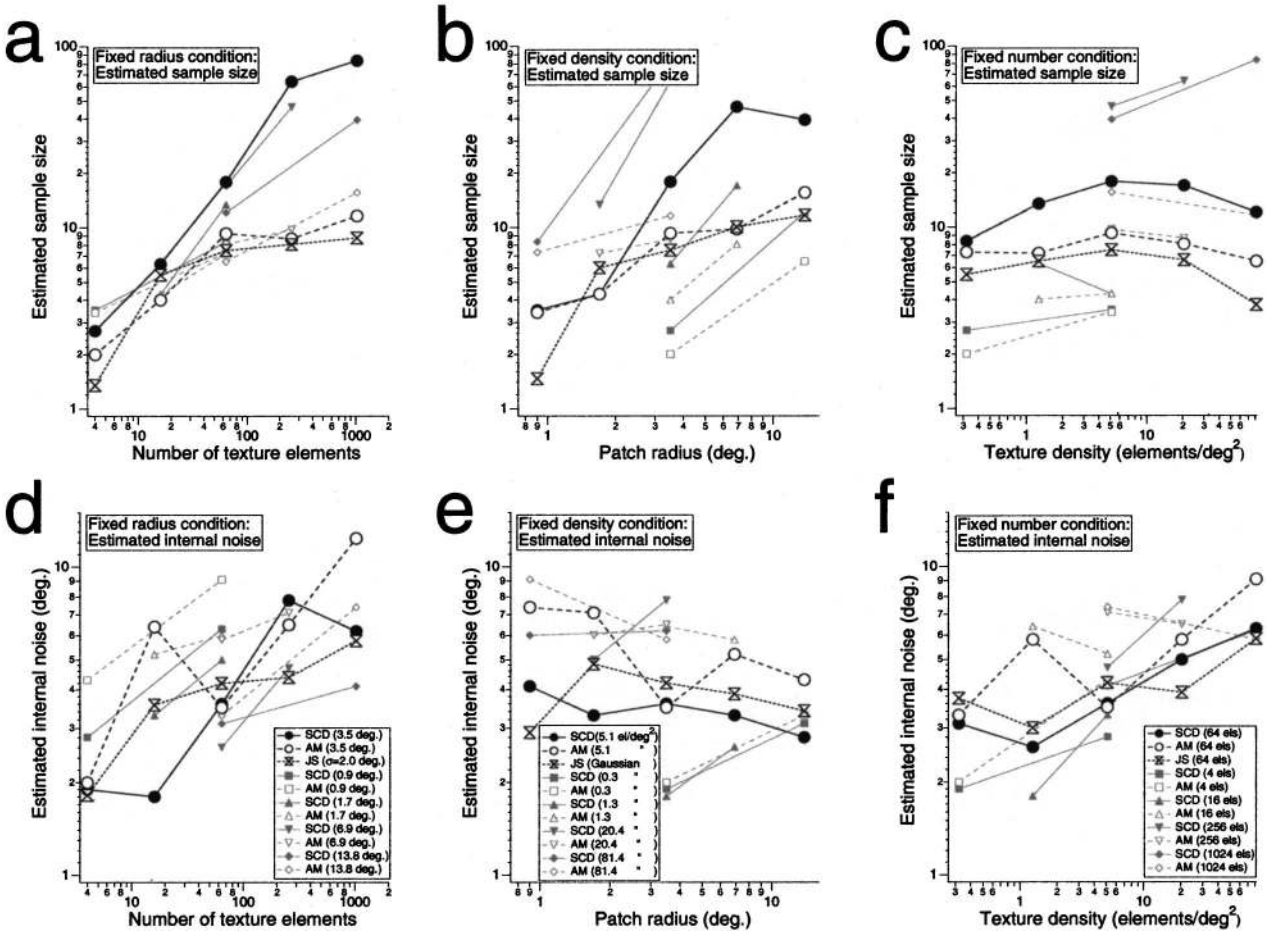

Fig. 7. Summary of parameters derived from the model fits. The top row shows estimates of the number of samples employed by observer, and the bottom row shows the associated internal noise parameter. (a) and (d) fixed size condition, (b) and (e) fixed density-condition, and (c) and (f) the fixed-numerosity condition. Notice in (c) that there is little change in the number of samples employed for a fixed number of elements, and in (e) that there is little change in estimated internal noise for a fixed texture density.

sizes are being employed across all subconditions. The number of texture elements presented determines the number of samples used by the observer.

4. DISCUSSION

Figure 7 summarizes the parameters derived from the variance-summation model. All estimates were made with stimuli generated through the use of uniform random distributions of element positions, apart from the "bow-tie" symbols that show results from subject JS obtained with Gaussian spatial distributions. Results are similar for the two types of distributions, and for this reason I have not presented the additional four subconditions per graph for subject JS.

Consider first Figs. 7(d)–7(f), which show the effect of fixing each of the three spatial parameters on the estimated internal noise on the orientation combination process. Internal noise shows some dependency on all of the texture parameters tested, although there appears to be the least amount of change in estimated noise under conditions of fixed density. There are a number of features of the task-stimulus combination that could be interacting to determine this pattern of results. First, a primary dependence on density is to be expected if the main contribution to internal noise comes from constraints on local sampling. Specifically, as density increases, the increasing number of overlaps between elements leads to an increase in the local orientational bandwidth of the textures that could in turn inflate any estimate of effective local noise. To test this directly I measured the orientational bandwidth of textures (generated in the same manner as in the experiments) as a function of the number of elements constituting the textures and those elements' orientation standard deviation. Results are graphed in Fig. 8 and show that bandwidth is constant across a wide range of numerosities. Any reduction in bandwidth is likely to be significant only at combinations of the lowest numerosities (<8 elements) and highest orientation standard deviations tested ($s \geq 8°$). Changes in bandwidth therefore cannot wholly account for a primary dependence on texture density. A second, more likely explanation is visual crowding. As density increases, elements will tend to have increasing numbers of elements located near them that could corrupt a local orientation estimate.[22] The mechanisms underlying such disruptions of feature

processing are unclear, but it has been suggested that crowding is due to limitations of the spatial resolution of visual attention,[44] which is known to be required for orientation judgments.[45]

Figures 7(a)–7(c) show the effects of the estimated number of orientation samples used by subjects for the various spatial configurations tested. It is evident from these data that sampling density is almost entirely independent of the density and size of a texture patch and is determined by the texture's numerosity. This is a counterintuitive result: The spatial configuration of a texture cannot disrupt performance, provided that the texture's density and radius conspire to maintain the number of samples present. The finding that local density has little effect on the task is contrary to the accepted interpretation of previous findings[21] measured using "pop-out" stimuli [e.g., Fig. 1(b)], but is consistent with recent findings on the effect of density on other texture tasks (such as symmetry detection)[46] and with the general notion of a visual attentive processor limited by a low data capacity.[47] As for this capacity, the curves shown in Fig. 6(a) suggest that the number of samples used is a linear function of the number of samples present, on log–log axes, indicating a simple power-law relationship. Assumption of a zero offset on the fit (i.e., that subjects use one sample when only one sample is presented) produces a one-parameter fit, giving powers of 0.61 and 0.52 for subjects SCD and ACM, respectively. Clearly this relationship breaks down at higher signal densities and numerosities. This could have two explanations: first, that there is a genuine upper limit on the number of samples that can be processed; second, that only so many oriented elements can be presented within a limited area before overlaps begin to reduce the number of samples that are effectively available. The present study is unable to distinguish between these two explanations.

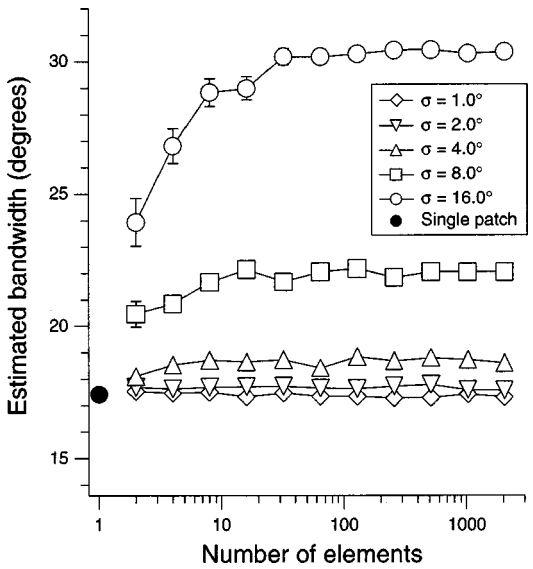

Fig. 8. Orientation bandwidth of the textures used (standard deviation of best-fitting Gaussian fit to Fourier power spectra) as a function of the orientation standard deviation and the number of elements (2–2048) constituting each texture. Note that bandwidth is largely constant as a function of the number of elements, except for the lowest density/highest orientation standard deviation combinations.

That observers' sampling is contingent only on the sample size presented means that the spatial pooling of orientation detectors across the visual field is flexible. This flexibility seems to incur a low attentional load; subjects are quite capable of performing these tasks with a 100-ms presentation time (although stimuli in this study were unmasked, which may have increased the effective exposure duration somewhat). It is interesting to consider these results in the context of filter-rectify-filter models of second-order and texture processing.[48,49] Such systems propose that the second-stage pooling mechanisms are themselves filters operating on the rectified output of banks of first-stage filters. Observers are known to be sensitive to the orientation structure of typical second-order stimuli (e.g., contrast-modulated noise[50]), suggesting that the notional first-stage filters are oriented. There is physiological evidence for responsiveness of cells in primary visual cortex to contrast-defined structure,[51] and it has been further suggested that responsiveness to second-order structure could be achieved by using center–surround selectivity to orientation, which is found in area *V*1 of the macaque monkey.[52] The results reported here suggest that our ability to integrate the output of local first-order detectors is independent of the spatial distribution of their activation. It is difficult to envision how later filtering mechanisms could be organized in such a manner as to confer this degree of flexibility. That is not to say that filter-rectify-filter mechanisms might not be implicated in other tasks (involving boundary localization, for example) but rather that performance on the task described here suggests the existence of a pooling process operating on a description of image structure that almost completely disregards spatial position.

ACKNOWLEDGMENTS

This work was funded by a fellowship from the Wellcome Trust. Thanks to Anna Morgan and Jenny Sloper for acting as subjects and to Mike Morgan and Josh Solomon for their helpful comments on this project.

Steven Dakin can be reached at s.dakin@ucl.ac.uk.

REFERENCES

1. D. H. Hubel and T. N. Wiesel, "Receptive fields, binocular interaction and function architecture in the cat's visual cortex," J. Physiol. **160**, 106–154 (1962).
2. A. B. Watson, "Summation of grating patches indicates many types of detectors at one retinal location," Vision Res. **22**, 17–25 (1982).
3. I. Fogel and D. Sagi, "Gabor filters as texture discriminator," Biol. Cybern. **61**, 103–113 (1989).
4. A. Bovik, M. Clark, and W. Geisler, "Multi-channel texture analysis using localised spatial filters," IEEE Trans. Pattern Anal. Mach. Intell.**12**, 55–73 (1990).
5. J. Malik and P. Perona, "Preattentive texture discrimination with early visual mechanisms," J. Opt. Soc. Am. A **7**, 923–932 (1990).
6. J. R. Bergen and M. S. Landy, "Computational modeling of visual texture segmentation," in *Computational Models of Visual Processing*, M. S. Landy and J. A. Movshon, eds. (MIT Press, Cambridge, Mass., 1991).

7. E. R. Howell and R. F. Hess, "The functional area for summation to threshold for sinusoidal gratings," Vision Res. **18**, 369–374 (1978).
8. J. G. Robson and N. Graham, "Probability summation and regional variation in contrast sensitivity across the visual field," Vision Res. **21**, 409–418 (1981).
9. N. V. S. Graham, *Visual Pattern Analyzers* (Oxford U. Press, New York, 1989).
10. M. J. Mayer and C. W. Tyler, "Invariance of the slope of the psychometric function with spatial summation," J. Opt. Soc. Am. A **3**, 1166–1172 (1986).
11. T. S. Meese and C. B. Williams, "Probability summation for multiple patches of luminance modulation," Vision Res. **40**, 2101–2113 (2000).
12. P. Verghese and L. S. Stone, "Perceived visual speed constrained by image segmentation," Nature **381**, 161–163 (1996).
13. U. Polat and C. W. Tyler, "What pattern the eye sees best," Vision Res. **39**, 887–895 (1999).
14. P. Verghese, S. N. J. Watnmaniuk, S. P. McKee, and N. M. Grzywacz, "Local motion detectors cannot account for the detectability of an extended trajectory in noise," Vision Res. **39**, 19–30 (1999).
15. D. J. Field, A. Hayes, and R. F. Hess, "Contour integration by the human visual system: evidence for a local 'association field'," Vision Res. **33**, 173–193 (1993).
16. R. F. Hess and S. C. Dakin, "Absence of contour linking in peripheral vision," Nature (London) **390**, 602–604 (1997).
17. S. C. Dakin and R. J. Watt, "The computation of orientation statistics from visual texture," Vision Res. **37**, 3181–3192 (1997).
18. F. A. Kingdom and D. R. Keeble, "On the mechanism for scale invariance in orientation-defined textures," Vision Res. **39**, 1477–1489 (1999).
19. H. C. Nothdurft, "Sensitivity for structure gradient in texture discrimination tasks," Vision Res. **25**, 1957–1968 (1985).
20. M. S. Landy and J. R. Bergen, "Texture segregation and orientation gradient," Vision Res. **31**, 679–691 (1991).
21. D. Sagi and B. Julesz, "Short-range limitation on detection of feature differences," Spatial Vision **2**, 39–49 (1987).
22. M. J. Morgan, R. M. Ward, and E. Castet, "Visual search for a tilted target: tests of spatial uncertainty models," Q. J. Exp. Psychology A **51**, 347–370 (1998).
23. H. C. Nothdurft and C. Y. Li, "Texture discrimination: representation of orientation and luminance differences in cells of the cat striate cortex," Vision Res. **25**, 99–113 (1985).
24. A. Gorea and T. V. Papathomas, "Local versus global contrasts in texture segregation," J. Opt. Soc. Am. A **16**, 728–741 (1999).
25. H. B. Barlow, "Retinal noise and absolute threshold," J. Opt. Soc. Am. A **46**, 634–639 (1956).
26. H. B. Barlow, "Increment thresholds at low intensities considered as signal/noise discrimination," J. Physiol. (London) **136**, 469–488 (1957).
27. Y. Y. Zeevi and S. S. Mangoubi, "Vernier acuity with noisy lines: estimation of relative position uncertainty," Biol. Cybern. **50**, 371–376 (1984).
28. R. J. Watt and R. F. Hess, "Spatial information and uncertainty in anisometropic amblyopia," Vision Res. **27**, 661–674 (1987).
29. R. J. Watt and M. J. Morgan, "The recognition and representation of edge blur: evidence for spatial primitives in human vision," Vision Res. **23**, 1465–1477 (1983).
30. D. W. Heeley, "Spatial frequency discrimination for sinewave gratings with random, bandpass frequency modulation: evidence for averaging in spatial acuity," Spatial Vision **2**, 317–335 (1987).
31. D. G. Pelli, "The quantum efficiency of vision," in *Vision Coding and Efficiency*, C. Blakemore, ed. (Cambridge U. Press, Cambridge, UK, 1990).
32. A. J. J. Ahumada and A. B. Watson, "Equivalent-noise model for contrast detection and discrimination," J. Opt. Soc. Am. A **2**, 1133–1139 (1985).
33. R. F. Hess and S. C. Dakin, "Contour integration in the peripheral field," Vision Res. **39**, 947–959 (1999).
34. R. J. Watt, *Understanding Vision* (Academic Press, London, UK, 1991).
35. D. W. Heeley, H. M. Buchanan-Smith, J. A. Cromwell, and J. S. Wright, "The oblique effect in orientation acuity," Vision Res. **37**, p. 235–242 (1997)
36. Z. L. Lu and B. A. Dosher, "External noise distinguishes attention mechanisms," Vision Res. **38**, 1183–1198 (1998).
37. D. H. Brainard, "The Psychophysics Toolbox," Spatial Vision **10**, 433–436 (1997).
38. D. G. Pelli, "The VideoToolbox software for visual psychophysics: transforming number into movies," Spatial Vision **10**, 437–442 (1997).
39. D. G. Pelli and L. Zhang, "Accurate control of contrast on microcomputer displays," Vision Res. **31**, 1337–1350 (1991).
40. D. W. Heeley and H. M. Buchanan-Smith, "Recognition of stimulus orientation," Vision Res. **32**, 719–743 (1990).
41. D. C. Burr and S. A. Wijesundra, "Orientation discrimination depends on spatial frequency," Vision Res. **31**, 1449–1452 (1991).
42. R. J. Watt and D. Andrews, "APE: Adaptive probit estimation of psychometric functions," Current Psychol. Rev. **1**, 205–214 (1981).
43. D. H. Foster and W. F. Bischof, "Bootstrap estimates of the statistical accuracy of thresholds obtained from psychometric functions," Spatial Vision **11**, 135–139 (1997).
44. S. He, P. Cavanagh, and J. Intriligator, "Attentional resolution and the locus of visual awareness," Nature **383**, 334–347 (1996).
45. J. S. Joseph, M. M. Chun, and K. Nakayama, "Attentional requirements in a 'preattentive' feature search task," Nature **387**, 805–807 (1997).
46. S. J. M. Rainville and F. A. A. Kingdom, "The functional role of oriented spatial filters in the perception of mirror symmetry," Vision Res. **40**, 2621–2644 (2000).
47. P. Verghese and D. G. Pelli, "The information capacity of visual attention," Vision Res. **32**, 983–995 (1992).
48. C. Chubb and G. Sperling, "Drift-balanced random stimuli: a general basis for studying non-Fourier motion perception," J. Opt. Soc. Am. A **5**, 1986–2007 (1988).
49. H. R. Wilson, P. Ferrera, and C. Yo, "A psychophysically motivated model for two-dimensional motion perception," Vision Res. **9**, 79–97 (1992).
50. S. C. Dakin and I. Mareschal, "Sensitivity to contrast modulation depends on carrier spatial frequency and orientation," Vision Res. **40**, 311–329 (2000).
51. I. Mareschal and C. L. Baker, "Cortical processing of second-order motion," Visual Neurosci., **3**, 527–540 (1999).
52. J. J. Knierim and D. C. van Essen, "Neuronal responses to static texture patterns in area V1 of the alert macaque monkey," J. Neurophysiol. **67**, 961–980 (1992).