PERGAMON

Vision Research 42 (2002) 2013–2020

Vision Research

www.elsevier.com/locate/visres

Summation of concentric orientation structure: seeing the Glass or the window?

S.C. Dakin *, P.J. Bex

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

Received 19 November 1999; received in revised form 31 October 2001

Abstract

Rotational Glass patterns are discriminable from noise at substantially lower signal-to-noise levels than translational patterns, a finding that has been attributed to the operation of concentrically tuned units in cortical area V4 (Wilson, Wilkinson, & Asaad, Vis. Res. 37 (17) (1997) 2325; Wilson & Wilkinson, Vis. Res. 38 (19) (1998) 2933). Under experimental conditions similar to Wilson et al. we found this advantage to be largely contingent on the pattern being viewed through a circular aperture. Because rotation of a random dot set cannot lead to the presence of unmatched dots at the boundary of a circular aperture, the integrity of low spatial frequency information at the boundary reliably indicates the presence of rotational, but not translational, structure. When we removed this cue, either using a square aperture or surrounding a round aperture with noise dots, none of the nine subjects tested showed any statistically significant advantage for rotational Glass patterns (although at least two did take longer to master the task with translational compared to rotational patterns). We go on to show generally similar patterns of global integration for both rotational and translational patterns. We conclude that this paradigm presently offers no concrete psychophysical evidence for specialised concentric orientation detectors.

Keywords: Form vision; Pattern discrimination; Glass patterns; Global form

1. Introduction

Glass patterns are composed of a field of dot pairs (or dipoles) whose orientations are determined by some geometrical transformation (Glass, 1969). The impression gained from inspecting these patterns is of orientation structure corresponding to the transformation (e.g. rotation in Fig. 1b) indicating that the visual system is grouping members of the same dipole. For high-density patterns this grouping problem is compounded by the fact that dots will typically have a large number of dots closer to them than their dipole correspondent (Stevens, 1978). Various manipulations of the spacing, density, and contrast of Glass patterns have allowed the visual grouping processes underlying this phenomenon to be probed. Results are largely consistent with structure being derived not by specialised symbolic token matchers, but from the output of spatial filters (Dakin, 1997a,b, 1999; Prazdny, 1986; Zucker, 1985). In particular one of us has shown that observer's precision at judging the orientation of translational Glass patterns requires that they can access the output of oriented spatial filters at a narrow band of spatial frequencies (Dakin, 1997a). This shift in theoretical perspective is unsurprising given the success with which a variety of similar correspondence problems have been recast in terms of spatial filtering, (e.g., stereo, Ohzawa, DeAngelis, & Freeman (1990), and motion, Adelson & Bergen (1985)).

Filtering models focus on the idea that it is the local statistics of Glass patterns that limit subjects' performance on this task. However, Wilson and co-workers (Wilson & Wilkinson, 1998; Wilson, Wilkinson, & Asaad, 1997) have recently reported a finding that challenges the sufficiency of such an explanation. These authors showed that subjects' ability to report the presence of circularly windowed Glass patterns (composed of a large number of widely separated dot-pairs) depended on the type of orientation structure present in the pattern. Specifically, subjects' threshold signal to

* Corresponding author. Tel.: +20-7608-6988; fax: +20-7608-6850. *E-mail address:* s.dakin@ucl.ac.uk (S.C. Dakin).

0042-6989/02/$ - see front matter © 2002 Elsevier Science Ltd. All rights reserved.
PII: S0042-6989(02)00057-3

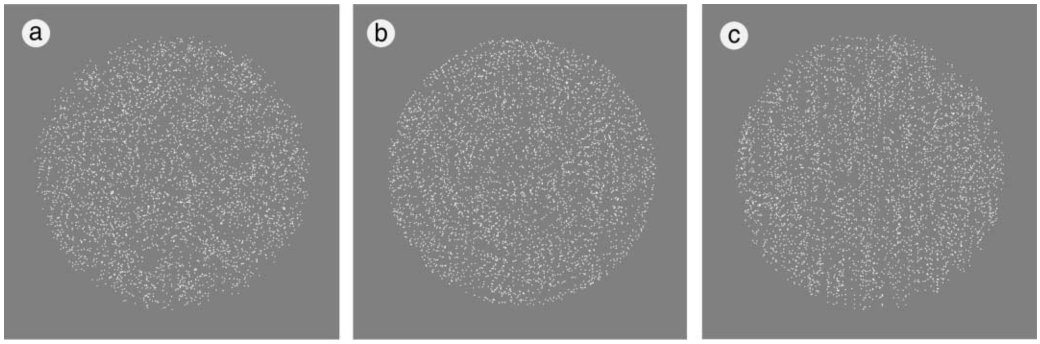

Fig. 1. (a–c) Glass patterns composed of 1966 dot pairs with a separation of 17.6 pixels. Dipoles are oriented according to (a) a random distribution (b) a rotation and (c) a 90° translation. It was observers task to discriminate between unstructured patterns (e.g. (a)) and structured patterns (e.g. (b)).

noise ratio for discrimination of Glass patterns from a field of randomly oriented dipoles was lowest for rotations, and highest for translations. Wilson et al. went on to model these data using concentric orientation summation units inspired by the response properties of cells in areas V4 of the macaque (Gallant, Braun, & Van Essen, 1993).

Previously (Maloney, Mitchison, & Barlow, 1987) used a similar experimental paradigm (estimation of threshold S/N ratios) but did not report substantial differences in performance that were dependent on transformation type. Dakin (1999) modelled these data using simple orientation statistics derived from the output of spatial filters, but also reported data from a similar experiment showing a small advantage for rotational patterns over translations. It is informative to note that the main difference between studies that have shown any effect (Dakin, 1999; Wilson & Wilkinson, 1998; Wilson et al., 1997) and those that have not (Maloney et al., 1987) is that the former used round, and the latter square, stimulus windows. Below, we show that windowing Glass patterns introduces low spatial-frequency artefacts near the pattern boundary that could confer a substantial advantage for rotational Glass patterns under the experimental conditions of Wilson et al.

Fig. 1 shows examples of the stimuli used in the experiments reported. Following Wilson and co-workers we used dense patterns (6% coverage $=$ 3932 dots for a 256 pixel radius pattern with 2×2 pixel elements), with wide separations between members of each dipole (17.6 pixels $= 10.0'$ under experimental viewing conditions), and a circular stimulus window. There is an issue with the generation of these patterns similar to that encountered with random-dot motion stimuli: what does one do with elements that fall-off the edge of the display? One can either plot them regardless (as Wilson et al. did; Wilson, personal communication) or leave unmatched/singleton dipole elements at the pattern edge. Fig. 2a–c illustrates why this is never an issue for rotational patterns, placing dipoles in a circular region simply cannot lead to individual dots falling outside the delineated region. As a consequence this type of pattern will have more clearly defined edges than either a translational pattern or a Glass pattern composed of randomly oriented dipoles. Clearly such boundary cues will be stronger in the unusually dense patterns used in the Wilson et al. studies. This difference in edge-integrity is highlighted in Fig. 2d–f. Here we have convolved the Glass patterns shown in Fig. 1, with an isotropic spatially band-pass filter to highlight information at low spatial frequencies. Notice that the blobs around the edge of Fig. 2e are longer and of higher contrast than corresponding features around the edge of either the translational or random orientation textures (Fig. 2d and f). Given that observers' task is to discriminate between structured and unstructured patterns (e.g. Fig. 1a versus b) we sought to test if this edge integrity cue could confer an advantage for rotational compared to translational patterns.

In the following sections we describe the results from various experiments bearing on the edge integrity hypothesis. The first uses stimuli closely matched to Wilson et al. and shows that the advantage for rotational patterns is contingent on the stimulus window being circular. When gross edge effects are corrected, we go on to report that 9/9 subjects show no significant advantage for rotational Glass patterns. Finally we show that reported differences in spatial summation between translations and rotations, cannot explain these findings. Subjects' summation shows some variation but is basically similar for both classes of patterns.

2. Methods

2.1. 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.) using code from

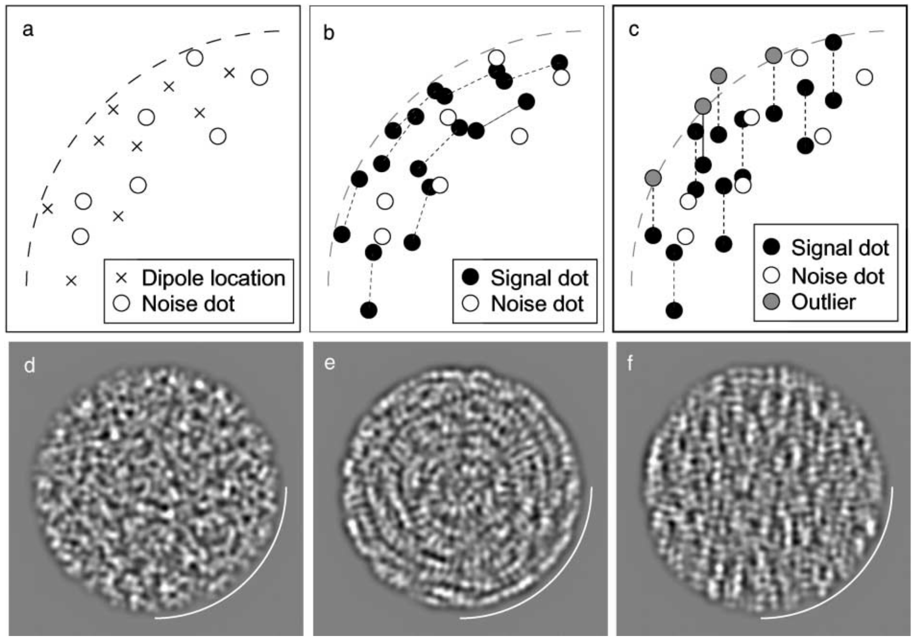

Fig. 2. (a–c) Schematic illustrating how dipole positioning can fragment the pattern boundary. (a) Crosses and circles show the locations of dipoles and noise dots respectively (the dashed line indicates the limit of the circular region within which dipole centres are constrained to fall). (b,c) Dipole orientations are arranged according to (b) a rotation and (c) a vertical translation. Only in the case of the translation can some elements (outliers) fall outside of the pattern boundary. These elements can either be retained (as was the case in Wilson et al.) or not plotted (as we did in Experiments 2 and 3) to leave unpaired dots (or singletons). The effects of this on the integrity of the pattern boundary is visible when typical Glass patterns, taken from Fig. 1, are filtered to highlight low spatial frequency information. (d–f) Glass patterns with dipole orientations determined by (d) a random, (e) a rotational or (f) a translational pattern, filtered with an isotropic Laplacian-of Gaussian filter (with $\sigma = 8.0$ pixels). A white arc is superimposed in the lower right-hand corner to assist the reader in judging edge smoothness.

the Psychophysics Toolbox (Brainard, 1997) and the Videotoolbox (Pelli, 1997) packages. Stimuli were displayed either on a 19″ Sony Multiscan 400PS colour monitor, or a 22 La Cie Electron Blue monitor. Both were driven by a Mac Picasso 850 graphics card (Villagetronic Ltd.) and the screen had a resolution of 1280×1024 pixels and operated at a frame refresh rate of 85 Hz. 12-bit contrast accuracy was achieved by electronically combining the RGB outputs from the graphics card using a video attenuator (Pelli & Zhang, 1991). A monochrome signal was generated by amplifying and sending the same attenuated signal to all three guns. The output luminance was linearised using a look-up table. The screen was viewed binocularly at a viewing distance such that 1 pixel on the screen subtended 0.57′ of visual angle. The display had a background luminance of 48 cd/m^2.

2.2. Stimuli

Unless stated otherwise, stimulus parameters were set to agree as closely as possible with those described by Wilson and Wilkinson, 1998. Stimuli were 768 pixel (7.3°) square images containing a texture composed of a mixture of dipoles and randomly positioned dots. Dots measured 2×2 pixels ($1.14' \times 1.14'$) and we used a dipole separation of 10′ throughout. 3932 dots were used in the circular window condition, and 5006 in the square window condition, maintaining a dot density of 6% across conditions. [1] Stimulus patterns were generated at various signal-to-noise ratios by replacing a proportion of the dipoles by randomly positioned dots. Noise patterns were generated in an identical manner, and were composed of a mixture of randomly oriented dipoles and randomly positioned dots (in the same ratio as the signal-to-noise ratio of the cued interval). This procedure matches intervals for the mean dot separation and forces subjects to use orientation cues. Dipoles were constrained to fall in a circular region with radius 2.43° (256 pixels). In Experiment 1, dots falling outside of this region were plotted regardless (i.e. a radial pattern would actually have a greater overall extent than a rotation). In Experiments 2 onwards dots falling outside the circular region were not plotted. In all experiments, stimuli contained a small circular hole at the centre (with diameter equal to the dipole separation—10′) that contained randomly positioned dots matched to the density of the rest of the pattern. We did this to avoid the problem that dipole correspondence is defined for translational, but for neither rotational nor radial

[1] True density will be slightly lower due to overlaps, but we, like Wilson et al., did not attempt to correct this.

patterns, within this region. A centrally presented fixation marker (a small white cross) was visible throughout the course of the experiment.

2.3. Procedure

The subjects task was a two-interval, two-alternative forced choice. Two textures were presented sequentially for 147 ms, separated by a 500 ms ISI. The cued interval contained a Glass pattern, with a proportion of the dipoles replaced by randomly positioned dots. The noise interval contained a stimulus composed of randomly oriented dipoles (interspersed with the same proportion of randomly position elements). Subjects were required to indicate which interval contained the Glass pattern, which they did using the computer keyboard. QUEST (Watson & Pelli, 1983), was used to sample a range of signal to noise ratios and attempted to converge on the level giving 83% correct performance on this task. Runs consisted of blocks of 45 trials and at least three runs were undertaken for each data point plotted. Runs were not interleaved; subjects always knew the organisation for which they were looking. As discussed below, some subjects showed significant learning during the course of the experiment. For this reason data presented are always pooled across a maximum of the last four runs performed with a particular stimulus configuration. Error bars show the estimated standard error.

2.4. Subjects

The authors (SCD and PJB) and naïve observers served as subjects in the experiments. All wore optical correction as required.

3. Results

Fig. 3 shows results from the first experiment. Consider first the leftmost two bars of each graph. These data show that all four subjects show an advantage for the rotational patterns (SCD: 13.7%, PJB: 11.8%, AJ: 9.5%, RW: 15.2%) when patterns were presented using a circular stimulus window. However, the rightmost two bars show that this rotational advantage is reduced, or even reversed, when stimuli are presented within a square window (SCD: −9.5%, PJB: −1.5%, AJ: +2.0%, RW: +2.3%). Clearly the overall shape of the Glass pattern can affect performance on this task.

In the second experiment we ran a larger pool of subjects on the same task, but this time we attempted to reduce the magnitude of the edge integrity cue by not plotting dots that fell outside of the circular stimulus region (the outliers shown in Fig. 2c). All subjects were required to detect structure in vertical or rotational Glass patterns that were windowed by either a square or

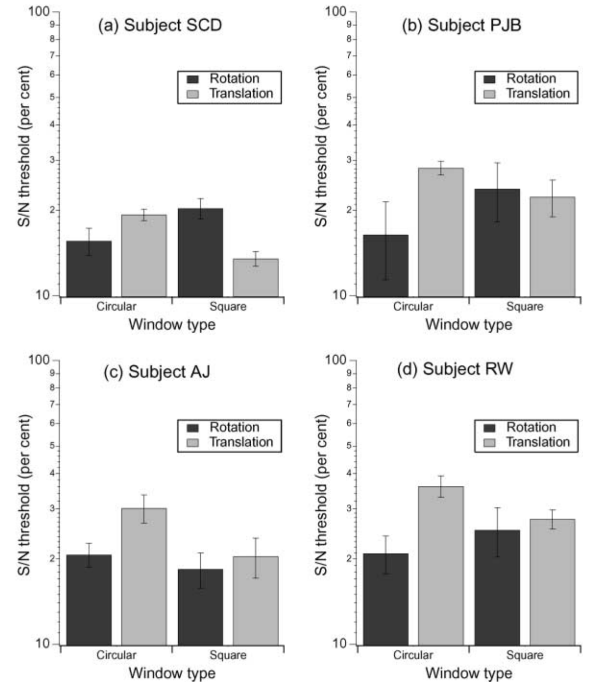

Fig. 3. Results from Experiment 1. All subjects show an advantage for rotational patterns that is contingent on the stimulus window being circular.

a round aperture. Six of the subjects also participated in a third condition where stimuli were windowed by a round aperture that was surrounded by an annulus (extending from the edge of the Glass pattern to a radius of 3.0°) filled with randomly positioned dots. Fig. 4 graphs the results. Although it seems unlikely that any one of the three manipulations should greatly affect the degree of activation of a detector tuned for circular structure, we report little or no advantage for circular structure in the nine subjects tested. Subjects whose data are plotted in Fig. 4a–c and g–i show no difference in performance with circular and translational patterns, and SCD (data plotted in Fig. 4d) shows a consistent advantage for translational patterns. Fig. 4e and f, shows that subject PJB and DBL both showed a small advantage for rotational patterns but that this was contingent on the pattern window being circular.

We noted that two of the subjects (IE and JS) showed substantial learning in the course of the experiment. A trial-by-trial plot of thresholds collected from subject IE in the second experiment is depicted in Fig. 5, and shows that this learning was greatest for translational patterns. It is unclear why some subjects should take longer to learn how to detect translational Glass patterns but it is certainly the case that their performance with both types of patterns converges after sufficient training.

Wilson et al. claimed that lower S/N thresholds for rotational patterns are attributable to differences in the way orientation is integrated over space for the various types of pattern. They reported that S/N thresholds for

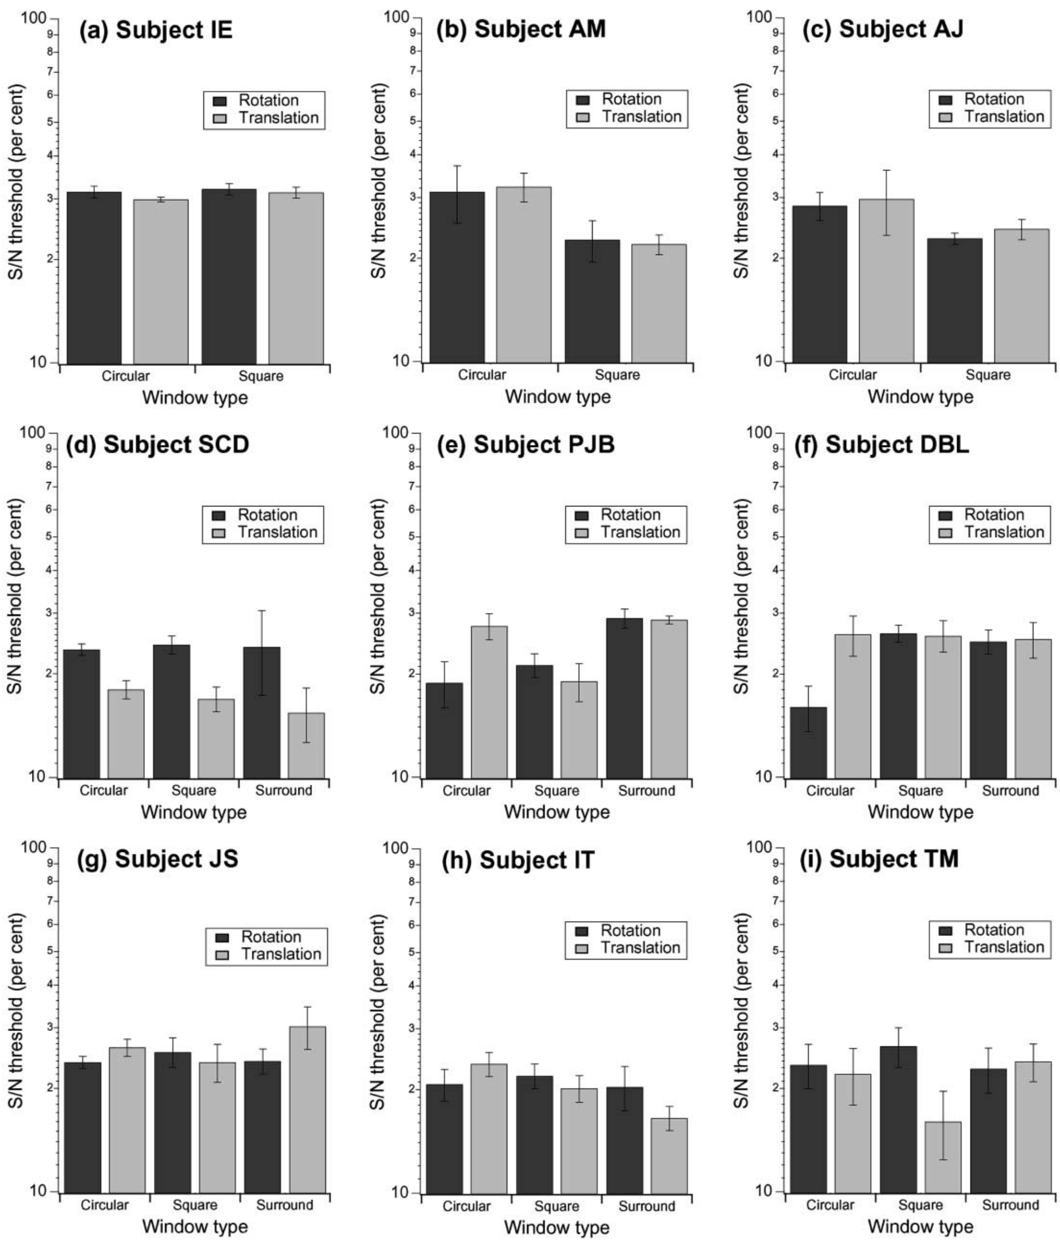

Fig. 4. Summary graphs for nine subjects. Patterns in these experiments were generated with singletons. (a–d) The first four subjects show no advantage for rotational patterns (in fact SCD shows a systematic advantage for translational patterns). (e) PJB shows an advantage that is contingent on the stimulus window being circular.

rotational patterns show a stronger dependency on signal area than performance with translational patterns; in particular that for parallel Glass patterns, restricting the signal dots to 31% of the total area did not increase the threshold for any subject, (Wilson et al., 1997, p. 809). These authors conclude that "Surprisingly, no such global pooling was found for Glass patterns with parallel structure". There are, however, considerable differences between Wilson et al. (1997) and Wilson and Wilkinson (1998) as to the reported extent of this effect since the latter paper indicates that 3/4 subjects show substantially better performance with larger rather than smaller translational Glass patterns, apparently indicating considerable pooling. Indeed Fig. 6 of Wilson and Wilkinson (1998) suggests that translational pooling is only marginally less efficient than pooling with radial patterns. We sought to measure this difference in pooling at a range of stimulus areas (Wilson & Wilkinson (1998) used only 25% and 100%) using our procedure.

In Experiment 3, we first ran three subjects on our task as a function of the area of a circularly windowed Glass patterns (outliers not plotted), sampled at areas of 25%, 35.3%, 50%, 70.7% and 100% of the original patterns area (18.55° of arc^2). Results are presented in

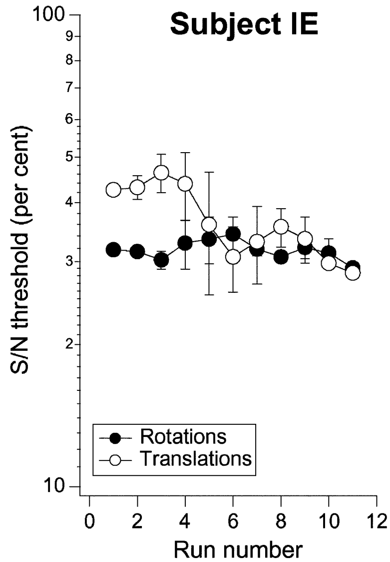

Fig. 5. Some subjects showed substantial learning in the course of the experiment. The graph shows thresholds from subject IE for the detection of translational and rotational Glass patterns presented within a circular aperture. Outliers were not plotted.

Fig. 6b–d. Note that all subjects show robust and improvements in performance with increasing stimulus area for all pattern types. There are some individual differences; subject TM has problems with radial patterns, and subject SCD shows generally superior performance with translational patterns. PJB shows near identical summation throughout. These data are fundamentally at odds with those presented by Wilson et al. who found shallower summation functions for translational compared to radial patterns (which they took to indicate differences in pooling). We propose that these relatively shallower summation functions may have a much simpler explanation involving two factors. First, most of their subjects showed relatively poor performance with translational patterns at 100% signal area (the edge effects described earlier are the likely cause of this discrepancy). Second, 3/4 of their subjects actually showed superior performance with translational patterns at small signal areas (the authors do not comment on this although it runs contrary to their explanation in terms of inferior pooling for translational patterns). This is a simple consequence of the structure of Glass patterns: as signal area decreases an increasingly larger proportion of the rotational/radial patterns falls within the central region (with radius equal to the dipole length) where correspondence is undefined. Although small, this undefined region always falls at subjects' fixation. Clearly, this limitation does not apply to translational patterns. Thus thresholds for translational patterns are elevated at large signal areas due to edge artefacts, and are (relatively) depressed at small stimulus areas because rotational/radial patterns become difficult to see. To summarise, we suggest that statistics of dipole pairing at both the edge and centre of rotational patterns conspire to steepen the summation function for rotations but not translations. This has nothing do with neural machinery. Fig. 6b–d demonstrate that when these factors are controlled, although summation functions tend to show variation in overall levels of performance they show similar slopes across transformations.

We wondered if differences in pooling might be manifest by a more subtle manipulation of dipole positioning than alteration of the overall patch size. Wilson measured performance with translational patterns that were restricted to a single centrally located elongated strip and compared this to performance with radial patterns that were restricted to wedge-shaped regions. Although this procedure produces different amount of summation this in unsurprising since dipoles were not evenly distributed throughout the translational pattern. In the second part of Experiment 3 we conducted what we consider to be a fairer comparison between rotational and translational patterns. Dipoles were positioned either in parallel strips or in concentric rings (white regions in Fig. 6a) and interleaved with density-balanced noise (black regions in Fig. 6a). Regions where dipoles were permitted to fall were generated using linear or radial ("bulls-eye") sine-wave grating with identical wavelength ($0.25 \times$ patch radius $= 0.61'$), which were then differentially thresholded to give patterns covering 37%, 50%, 63%, 75%, 87% or 100% of the pattern area. We then used a 2×2 design and measured summation for translational and rotational patterns, where dipoles were positioned according to either the translational or rotational positioning scheme. Results are shown in Fig. 6e–g. They indicate first that, contrary to Wilson et al., all subjects showed robust improvement in threshold with increasing stimulus area for both transformation types. Second that it is primarily the transformation type that determines signal-to-noise thresholds, and that slopes for these pooling functions are broadly similar across all conditions. This is consistent with the notion that orientation pooling across space is flexible with respect to spatial variables such as density, and is more limited by the informational content of texture (Dakin, 2001).

4. Summary

We have presented evidence challenging the generality of the finding that detection of globally organised rotational structure in Glass patterns is generally more resistant to the intrusion of noise than translational patterns (Dakin, 1999; Wilson & Wilkinson, 1998; Wilson et al., 1997). Rather we have used stimulus generation techniques similar to Wilson et al., to show that these effects seem to be contingent on the stimuli being presented within a circular aperture. Moreover,

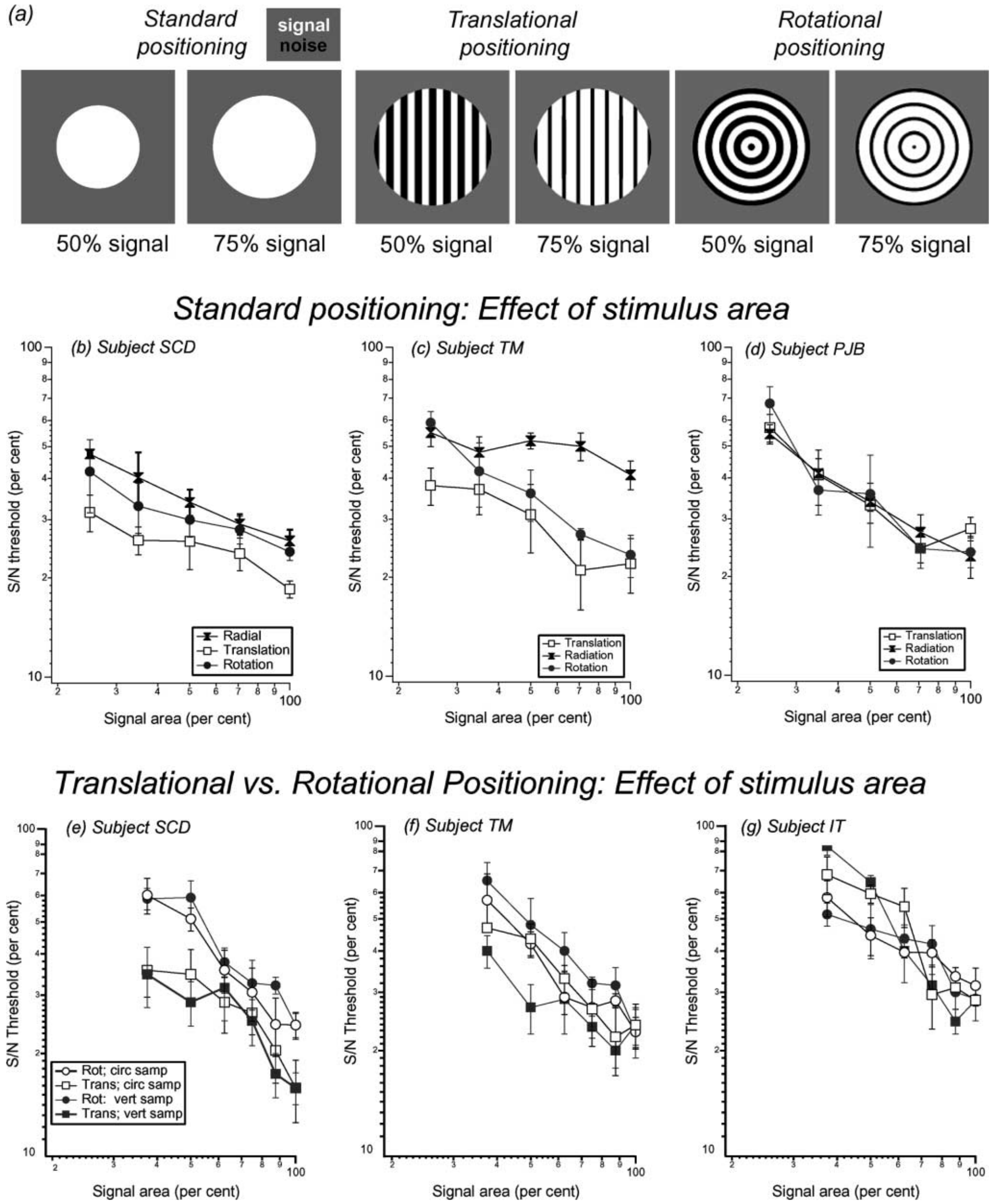

Fig. 6. (a) Schematic representation of the dipole-positioning schemes used in Experiment 3. In the first part of the experiment (standard positioning) we simply varied the radius of the texture. In the second part (translational and rotational positioning) we positioned signal dots in strips or rings (white regions) and interleaved them with noise dots (positioned in the black regions). (b–d) S/N thresholds for detection of translational, rotational and radial patterns, as a function of the overall size of the stimulus. All three subjects show similar patterns of summation for translation and rotations, although TM performs poorly with radial patterns. (e–g) Results as a function of the signal area of rotational and translational patterns where dots have been positioned according to either the translational or rotational scheme. Subjects' performance is largely determined by the transformation type (symbol shape) rather than the positioning strategy (filled versus open symbols).

we have presented a simple explanation for the rotational advantage under these conditions; subjects are able to use the integrity of the pattern edge as an additional cue to the presence of organised structure. This cue is clearly visible in the low-frequency structure of Glass patterns. When such cues are absent we report no substantial differences in the processing of translational and rotational Glass patterns, either in terms of basic detectability or spatial summation. Although we note that some observers take longer to learn how to detect translational compared to rotational structure it is our conclusion that that this paradigm offers, as yet, no substantial psychophysical evidence for the presence of dedicated higher-order structure detectors in human vision.

Acknowledgements

SCD was supported by a fellowship from the Wellcome Trust.

References

Adelson, E. H., & Bergen, J. R. (1985). Spatio-temporal energy models for the perception of motion. *Journal of the Optical Society of America A, 2*, 284–299.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*, 433–436.

Dakin, S. C. (1997a). The detection of structure in Glass patterns: psychophysics and computational models. *Vision Research, 37*, 2227–2259.

Dakin, S. C. (1997b). Glass patterns: some contrast effects re-evaluated. *Perception, 26*, 253–268.

Dakin, S. C. (1999). Orientation variance as a quantifier of structure in texture. *Spatial Vision, 12*, 1–30.

Dakin, S. C. (2001). Information limit on the spatial integration of local orientation signals. *Journal of the Optical Society of America A, 18*(5), 1016–1026.

Gallant, J. L., Braun, J., & Van Essen, D. C. (1993). Selectivity for polar, hyperbolic, and cartesian gratings in macaque visual cortex. *Science, 259*(5091), 100–103.

Glass, L. (1969). Moire effects from random dots. *Nature, 243*, 578–580.

Maloney, R., Mitchison, G., & Barlow, H. (1987). Limit to the detection of Glass patterns in the presence of noise. *Journal of the Optical Society of America A, 4*, 2336–2341.

Ohzawa, I., DeAngelis, G. C., & Freeman, R. D. (1990). Stereoscopic depth discrimination in the visual cortex: neurons ideally suited as disparity detectors. *Science, 249*(4972), 1037–1041.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming number into movies. *Spatial Vision, 10*, 437–442.

Pelli, D. G., & Zhang, L. (1991). Accurate control of contrast on microcomputer displays. *Vision Research, 31*, 1337–1350.

Prazdny, K. (1986). Some new phenomena in the perception of Glass patterns. *Biological Cybernetics, 53*, 153–158.

Stevens, K. (1978). Computation of locally parallel structure. *Biological Cybernetics, 6*, 19–28.

Watson, A. B., & Pelli, D. G. (1983). QUEST: a Bayesian adaptive psychometric method. *Perception and Psychophysics, 33*(2), 113–120.

Wilson, H. R., & Wilkinson, F. (1998). Detection of global structure in Glass patterns: implications for form vision. *Vision Research, 38*(19), 2933–2947.

Wilson, H. R., Wilkinson, F., & Asaad, W. (1997). Concentric orientation summation in human form vision. *Vision Research, 37*(17), 2325–2330.

Zucker, S. W. (1985). Early orientation selection: tangent fields and the dimensionality of their support. *Computer Vision, Graphics and Image Processing, 8*, 71–77.