The influence of shape and skeletal axis structure on texture perception

Sarah J. Harrison   State College of Optometry, Vision Sciences, State University of New York, New York, NY, USA

Jacob Feldman   Department of Psychology, Center for Cognitive Science, Rutgers University, New Brunswick, NJ, USA

We studied the relationship between texture orientation and shape skeletal axes in two tasks related to texture perception. The first series of experiments investigated discrimination of texture-defined shapes. We found that alignment between texture orientation and the skeletal axis of a figural region improved the segmentation strength, as did a perpendicular arrangement to a lesser extent. The alignment effect is attributable to the orientation of the skeletal axis itself, not the orientation of the figure edges; these two factors were deconfounded by the use of shapes whose contours undulated relative to the main axis orientation. Discrimination of multi-part shapes additionally showed that local alignment of texture with the axis of the enclosing part gave superior segmentation performance when compared to the classically optimal case of uniform texture orientation. A second series of experiments investigated sensitivity to changes in texture orientation within texture patches. Texture orientation discrimination was heightened when texture was aligned with the axis of the patch shape, demonstrating that the "axis effect" also affects the encoding of texture orientation. Taken together, these findings point to a broad role of skeletal axes in influencing the processes by which texture elements are aggregated to form the object itself.

Introduction

The study of texture segmentation, the perceptual separation of regions of the visual field that differ in the properties of their constituent elements, has an extensive literature. It is well documented that textures segment on the basis of local orientation and frequency content, properties for which V1 neurons are selective (De Valois, Albrecht, & Thorell, 1982; De Valois, Yund, & Hepler, 1982; Hubel & Wiesel, 1962). This parallel between perception and cortical properties has led to the default approach that texture segmentation is unavoidably constrained by early neural circuitry performing local analysis, with minimal influence from global or configural aspects of the stimulus. Reflecting this premise, the basic architecture of many models of texture segmentation follows a succession of serial processing stages whereby first-order oriented elements are detected and signals rectified, then sections of contour, corresponding to alignment of neighboring local orientation discontinuities, are extracted by second-order filters (Bergen & Adelson, 1988; Graham, 1994; Landy & Bergen, 1991; Malik & Perona, 1990). Although beyond the scope of such models, boundaries delineating entire coherent shapes would presumably be created by linking of these smaller segments. This type of scheme is described as "hierarchical",

meaning that local elements are analyzed first ("low-level" processing) and their properties consequently determine global aspects of the visual input ("higher level" processing). In this view of the sequence of events, individual texture elements are assessed within a restricted local context and earlier than global scene properties. While some models have proposed influences on texture elements from parts of the image that are further afield, these interactions are strictly at the level of the local elements themselves (e.g., Wolfson & Landy, 1999), not from global, configural, properties of the image.

Despite the prevalence of local-to-global approaches in texture segmentation (and in visual perception in general), there are many examples elsewhere in the literature where supposedly later processes influence supposedly earlier ones. For instance, seminal work by Navon (1977), Palmer (1977), and Pomerantz, Sager, and Stoever (1977) firmly established the existence of "global precedence" effects, in which the global property of configuration not only takes precedence in perception but also affects the salience of local elements from which the configuration is formed. Subjects' ability to discriminate local elements is altered, suggesting that global configuration has impacted on the process of local feature representation. In support of this interpretation, Enns (1986) showed that pop-out of local features depends on the "context", i.e., global configuration, in which they appear. More recently,

global-to-local aspects of visual processing have been increasingly discussed and investigated (for a review, see Hegde, 2008).

A number of findings in the texture segmentation literature itself likewise intimate that perception of separable texture regions may not be a strictly local-to-global process. For instance, it appears that local orientation signals are only a starting point—not necessarily the deciding factor—in the segmentation of textures: Disparity cues to surface completion override local, first-order, orientation signals in determining segmentation borders (He & Nakayama, 1994). Other contextual manipulations, such as presenting only the portions of texture adjacent to the orientation contrast border, without the figure interior, also critically affect the nature of the texture segmentation percept (Motoyoshi & Nishida, 2001). These studies showed that subjects' holistic interpretation of the texture stimulus, in terms of global figure and ground regions, was decisive in constructing the final segmentation percept. In broader terms, global form may play a more substantial role in texture perception than has been generally recognized, potentially combining with early inputs in determining the ultimate percept. In any case, it is clear that the issue of configural influences on texture perception deserves more vigorous study.

In this paper, we investigate the influence of the global *shape* of the enclosing region on the interpretation of texture within it. In the case of a shape defined by contrasting texture orientation, of course the shape does not exist in the absence of its constituent texture elements. Here, the question we pose is whether a reverse influence also acts: Does the shape of the enclosing region influence the perceptual organization of the elements within it? This issue goes beyond the overriding of local orientation signals by unrelated global manipulations: Any such influence of global shape on the percept of local elements would suggest a cyclical interaction between local and global percepts, rather than a strictly hierarchical construct of the texture percept as is often assumed. Hence, evidence for such influences would have important theoretical implications, challenging standard assumptions about the architecture of the system and the flow of information within it.

Medial axis and skeletal representations of shape

If we are to consider the influence of region shape on texture representation, we first need a system by which to quantify and manipulate shape. One useful way of capturing shape properties, first established by Blum (1967, 1973), is that of a shape's medial axes (Figure 1). Blum's implementation of medial axes used a "grassfire" algorithm, which can be visualized as the quenching points of a fire started simultaneously at all points along a shape's contour. This results in a simplified "skeletal" representation of the original shape, which summarizes
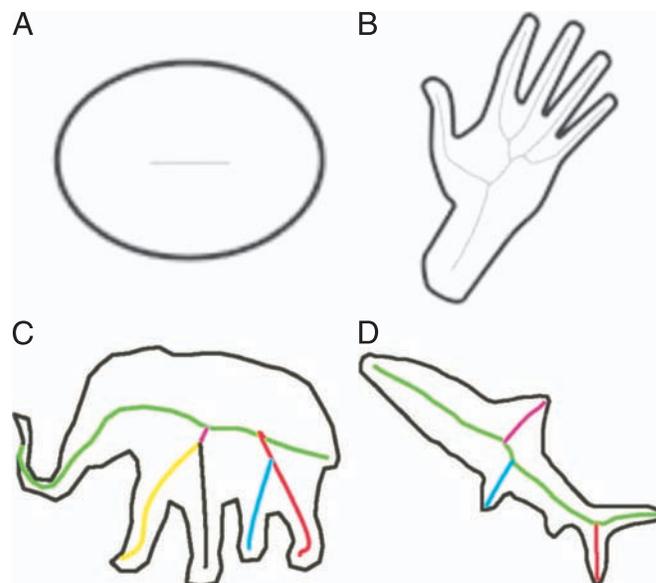

Figure 1. Examples of shapes and their medial axes. (A, B) Axes calculated using Blum's (1973) grassfire algorithm. (C, D) Axes calculated using Feldman and Singh's (2006) maximum posterior probability (MAP) algorithm.

the local axial symmetry inherent in the shape's bounding contour. Many more computationally sophisticated algorithms for extracting medial axes have since been developed (Katz & Pizer, 2003; Kovacs, Feher, & Julesz, 1998; Siddiqi, Shokoufandeh, Dickinson, & Zucker, 1999; Zhu, 1999). One recent approach (Feldman & Singh, 2006), which introduces some ideas to which we will return below, uses a Bayesian inverse probability framework to estimate the underlying skeleton (a hierarchy of the shape-part medial axes) from which the shape can be considered to have "grown." Our experimental evidence does not generally address the exact nature of the computational processes by which shape axes are determined, but rather primarily concerns the more basic question of whether such axes in any way modulate texture perception.

Shape axes are frequently presented as a computational solution to the gap between local featural representations and globally integrated object representations and have been proposed as a means of performing shape recognition and categorization (for a review, see Kimia, 2003). Additionally, a handful of studies have found psychophysical and neurophysiological evidence in support of the extraction of shape axes by the visual system: Threshold detection of Gabor patches is heightened inside a texture-defined figure as compared to outside, and this effect appears to be strongest along a closed region's axis (Kovacs et al., 1998; Kovacs & Julesz, 1993, 1994). This perceptual effect is mirrored in the responses of V1 neurons (Lamme, 1995; Lee, Mumford, Romero, & Lamme, 1998), with a heightened firing rate evolving over

time for a proportion of neurons whose receptive fields lie on the region's axis. The perceptual significance of such findings is not without debate. Most pertinent is the model of Zhaoping (2003), where the above effects are accounted for using a network of excitatory and inhibitory intracortical interactions entirely within V1, and without recourse to cortical processing of global structure or configuration. That propagation of signals between locally neighboring neurons could suffice to explain that the emergence of axial salience for certain global shape geometries implies that the psychological and physiological effects described above are not in themselves indicative of higher perceptual shape representation or processing.

Irrespective of the modulating mechanism, our study seeks psychophysical evidence that the organization of texture is influenced by shape and skeletal axes. In the following experiments, we seek to establish whether the strength of representation of a texture region is affected by the underlying skeletal structure of the region. We investigate strength of representation using two different aspects of texture perception; segmentation and orientation discrimination.

Experimental approach

The first series of experiments presented here investigates the interaction, should it exist, between texture segmentation and shape. If global shape percepts interact with the processing of local texture elements, we can expect this to manifest in the phenomenon of texture segmentation, where regions of local elements are grouped and separated on the basis of their encoded properties. Orientation-based texture segmentation (OBTS) has traditionally been studied using simple geometric shapes such as circles or squares. Within this constraint, the strength of texture segmentation does not appear to be affected by the global shape of the figural region. However, OBTS of more complex shapes has only been considered in passing (Nothdurft, 1985b), rather than as a parameter of interest in its own right, and so the influence of shape has not been extensively studied. Certainly, while it seems intuitive that the extraction of shape borders would increase in difficulty as shapes increase in complexity (compare stimuli appearance and results in Nothdurft, 1985a, 1985b), the interaction between shape and OBTS has not, to our knowledge, been further explored.

Manipulation of shape itself simultaneously co-varies a broad collection of parameters (e.g., number of vertices, degree of curvature, total boundary length) and so could be predicted to affect segmentation strength for many reasons that are not "shape" itself as summarized by a skeletal structure. Rather than compare segmentation strengths of dissimilar shapes, our first series of experiments sought a link between OBTS and axial shape

representations more subtly, by testing whether the strength of OBTS was affected by the structural organization of texture within the shape. The basis of the structural organization that we use in all of the following experiments is the shape's underlying skeletal axes: The orientation of texture within shapes was defined relative to the orientation of the shape axis, where a "texture–axis offset" of 0 degrees means that texture and axis orientations are congruent and an offset of 90 degrees means that texture and axis are perpendicular. Segmentation strength for a range of texture–axis offset conditions was quantified by finding the threshold texture coherence in two shape discrimination tasks.

Our second series of experiments addressed the representation of texture orientation within the context of a shaped texture patch. An effect of shape on segmentation, if indeed it were found in our first series of experiments, would not necessarily imply a parallel effect on resolution of texture orientation; segmentation is known to occur even when subjects are unable to report the composite texture orientations (see, for instance, Ben-Shahar, 2006, Experiment 2; see also, Rogers-Ramachandran & Ramachandran, 1998). However, given that OBTS requires that elements in different regions be encoded as having different orientations, it follows that if OBTS is affected by region shape then the strength of texture orientation representation might also be affected. In the second series of experiments, this aspect of shape influence on texture perception was probed directly.

The perceived orientation of a texture region can be thought of as a "statistical summary" of the composite signals: While the average orientation of a group of elements can be judged almost as precisely as that of a single element presented in isolation, this accuracy entails a loss of precision in judging individual orientations (Ariely, 2001; Dakin & Watt, 1997; Parkes, Lund, Angelucci, Solomon, & Morgan, 2001). Importantly, this integration of individual signals occurs not only over a fixed spatial extent but also within perceived groups of signals (Livne & Sagi, 2007). Hence, the question posed in Experiment 2 was given as follows: is texture orientation discrimination influenced by the shape of the texture patch or "group"? As in our segmentation experiments, this question was investigated by altering the structural organization of texture within an unchanging shape, and structural organization of texture was defined by its orientation relative to that of the region shape axis. Orientation discrimination thresholds were found for different angular offsets between the texture orientation and shape axis orientation.

General methods

Stimuli were generated on a Macintosh computer via Matlab (The MathWorks, Natick, MA) utilizing functions

provided in the Psychtoolbox (Brainard, 1997; Pelli, 1997) and presented on a ViewSonic monitor at 60 Hz. Texture stimuli in all experiments consisted of black lines presented on a mid-gray background, and experiments were performed in low lighting conditions. All stimuli were viewed at a distance of 50 cm, which resulted in a monitor pixel resolution of 25 pixels deg^{-1} and total screen dimensions of 30×40 degrees. Anti-aliased texture lines were 12 pixels in length (30 min of visual angle).

Data from each observer were analyzed separately, and thresholds were estimated by curve fitting with a Weibull function, using the Psychtoolbox. Subsequent analysis of threshold values was carried out using SPSS.

All subjects were naive to the purposes of the experiments and completed preliminary practice trials to ensure their comprehension of the tasks. Some subjects completed more than one experiment section. All subjects were paid for their participation.

Experiment 1

Experiment 1a

We investigated the effect on segmentation strength, of varying texture orientation relative to the shape axis. "Peanut" shapes were used in a shape discrimination task, where subjects judged which of two shapes was presented, a "thin" peanut or a "fat" peanut (Figure 2A). These shapes were chosen so as to deconfound shape axis orientation and shape edge orientation, as texture edges are known to have increased saliency when texture elements are locally aligned with the edge (Nothdurft, 1992; Wolfson & Landy, 1995). Texture shapes were presented against a background of randomly oriented texture, so that all local sections of the shape boundary were equally salient in terms of their orientation contrast with the adjacent background region.

Methods

Stimuli

Peanut shapes were constructed from a single straight axis (or "root"). The boundaries of the shape were calculated by varying the perpendicular distance (or "root width") from the axis according to one cycle of a cosine function. Root width was constant around the end of each axis, giving the shapes rounded endings. If the cosine function started at 0 degrees, then a "thin" peanut resulted, i.e., the shape was narrower at its center than at its ends. Conversely, if the function started at 180 degrees, then a "fat" peanut resulted, i.e., the shape was wider at its center than at its ends. Using these two shapes, it is not possible to simultaneously equate maximum width, length, and shape area. Hence, we approximately equated perceived
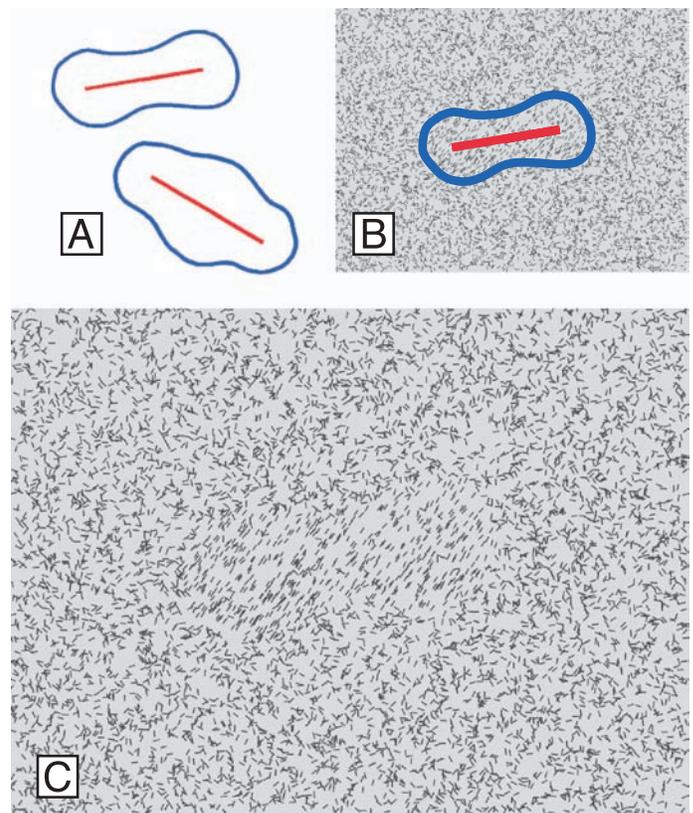

Figure 2. Examples of stimuli in Experiment 1a. (A) The two varieties of "peanut" shape; shape boundaries are depicted in blue and shape axes are in red. (B) Allocation of texture elements in a "thin peanut" stimulus; orientation of elements within the shape boundaries is constrained by the texture–axis offset condition (here 45 degrees) and coherence level (here 90%). Elements in the surround have random orientation. (C) The texture stimulus in (B) as it appeared on a trial. Total stimulus size was 30×40 degrees of visual angle. Subjects' task was to discriminate between the two possible shapes.

total spatial extent and used a small range of shape sizes so as to discourage a strategy of attending to a single geometric dimension. Shape axes were of length 270, 280, 290, or 300 pixels, corresponding root widths were 90, 93, 97, and 100 pixels, and cosine amplitude was always $0.2 \times$ root width. These parameters generated four shapes that subtended a length of between 18 and 20 degrees and a central width of between 8.64 and 9.6 degrees ("fat" peanuts) or 5.76–6.4 degrees ("thin" peanuts) of visual angle. To further ensure that attending to a single dimension or fixed monitor location could not inform correct shape discrimination, shapes were presented at random orientations, and the location of the midpoint of each shape was drawn randomly from a central 50×50 pixel (2.0×2.0 degree) window. Hence discrimination of isolated shape dimensions could only proceed consequent to recognition of the shape's location and global orientation.

Textures were created by randomly positioning 7000 line elements (Figures 2B and 2C). Texture elements inside the shape boundary were oriented according to one of five texture–axis offset conditions; the elements could be parallel to the direction of the shape axis, i.e., 0 deg offset, ±22.5, ±45, ±67.5, or perpendicular, i.e., 90 deg offset. The absolute orientation of each element was drawn from a Gaussian distribution centered on the above orientations and with a standard deviation of 6 degrees. Note that texture–axis offset condition was not confounded with absolute texture orientation, because shapes were presented at random orientations. Texture elements located outside the shape boundary were assigned random orientations. Stimuli were pregenerated, along with mask screens consisting of 7000 randomly oriented line elements.

The strength of segmentation of the shapes from the background was quantified by the threshold coherence level: The coherence of texture within the shape region in each stimulus, i.e., the proportion of texture elements conforming to the orientation stipulated by the texture–axis offset condition, was 50%, 60%, 70%, 80%, 90%, or 100%. Due to random placement of line elements, the occurrence of second-order features such as line crossings and intersections increased with decreasing texture coherence and reached the highest levels in the uncoherent background region. While differences in the distribution of second-order features may lead to texture segmentation (for an overview, see Julesz & Bergen, 1983), this was not a confound in our experimental design as the occurrence of second-order features is not correlated with our parameter of interest, the difference in orientation between texture and axes.

In total, there were 5 offset conditions, 6 coherence levels, and 4 sizes of both fat and thin peanuts, with all factors fully crossed. Five unique stimulus and mask textures were pregenerated for each combination of parameters. This resulted in a total of 1200 pairs of stimulus and mask, which were each presented once per testing session in random order.

Each trial commenced with the presentation of a fixation cross, which remained visible until the observer initiated the trial with a key press. A stimulus image was then presented for 200 ms. After a 100-ms blank interval, the mask image was presented for 100 ms. The screen then remained blank until the observer indicated with a key press whether the presented shape was "thin" or "fat". Feedback was given in the form of a quiet tone for an incorrect response. Subjects were instructed by means of an on-screen message to take a break after every 240 trials; they could take additional breaks whenever needed simply by delaying initiation of the next trial.

Subjects

Fourteen naive subjects took part in Experiment 1a. Each completed two testing sessions, generating a total of 480 responses for each texture–axis offset condition, spread over the six coherence levels.

Experiment 1a: Results

Texture coherence thresholds, at the level of 75% correct performance, are shown in Figure 3. The data clearly show an increase in thresholds (a drop in performance) as texture orientation is moved away from 0 degrees texture–axis offset, i.e., away from parallel to the shape axis, up to an offset of 45 degrees. An improvement in performance is then seen again as offset approaches 90 degrees, i.e., perpendicular to the axis orientation, giving the data a distinct "U-shaped" form overall. Notwithstanding the clear differences between subjects, a one-way within-subjects ANOVA confirmed that the variation of thresholds with texture–axis offset was significant ($F(4, 52) = 9.398$; $p < 0.001$) and that the form of the relationship between the two variables had a significant quadratic ("U-shaped") component ($F(1, 13) = 10.993$; $p = 0.006$). However, the fourth-order component of the relationship between the variables was also strongly significant ($F(1, 13) = 16.748$; $p = 0.001$), due to the "W-shaped" pattern of performance shown by some subjects, where a drop in thresholds was seen at texture–axis offsets of 22.5 and 67.5 degrees.

The different patterns of performance warrant different types of explanation: The quadratic pattern suggests a relationship between global shape segmentation and texture orientation that benefits segmentation at the parallel and perpendicular texture–axis offsets—an "axis alignment" effect. The possible reasons for this benefit will be discussed at greater length later. The fourth-order relationship may be related to a local edge saliency effect
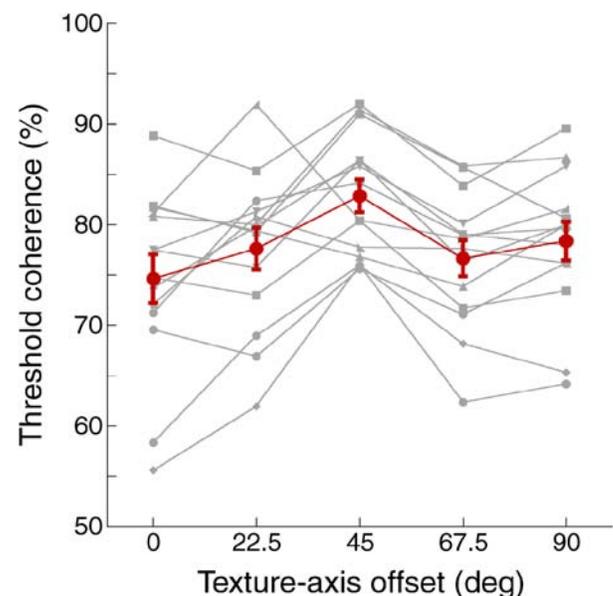
Figure 3. Results from Experiment 1a. Texture coherence thresholds (75% correct) in the shape discrimination task for all fifteen subjects (gray) and mean values (red, with error bars showing ±*SEM*) for the five texture–axis offset conditions.

mentioned earlier (Wolfson & Landy, 1995): At texture–axis offsets of 22.5 and 67.5 degrees, texture was approximately parallel to the steepest sections of the informative "dips" or "bulges" in the peanut shape, heightening the saliency of these contour portions—an "edge alignment" effect. This could perhaps facilitate use of these small portions of the boundary to guide performance. This would not be a very efficient strategy, due to variation in shape size, orientation, and location. Accordingly, we would suggest that this is why subjects who showed the fourth-order pattern of performance had higher thresholds on average than those who showed the quadratic pattern of performance. The large variation between subjects may partially have been caused by adoption of these two different strategies (and intermediate combinations). Given the range of absolute performance levels, it is all the more striking that the pattern of threshold variation when considered over all subjects was in fact strongly significant.

When summed over both patterns of performance, and as summarized by the mean thresholds in Figure 3, subjects were more able to distinguish between the two shapes when texture was parallel, and to a lesser extent perpendicular, to the shape axis. This was borne out by planned paired comparisons (with Bonferonni correction), which showed that thresholds at an offset of 0 degrees were significantly lower both than at 45 degrees and at 90 degrees ($\Delta threshold_{(45-0)} = 8.24\%$ coherence, $p < 0.001$; $\Delta threshold_{(90-0)} = 3.75\%$ coherence, $p < 0.012$) and that thresholds at 90 degrees were lower than at 45 degrees ($\Delta threshold_{(45-90)} = 4.49\%$ coherence, $p < 0.012$). Note that the "peanut" discrimination task required high spatial resolution of a specific portion of the shape edges and could conceivably be performed without processing the entire shape. Nevertheless, subjects did appear to benefit from processing the stimulus in a global manner, that is, in accord with the orientation of the axis of the global shape.

To further elucidate the nature of the relationship between texture segmentation and shape axes, and to consolidate our evidence that performance was based on the integrated global shape percept rather than on localized portions of boundary, we used a different shape discrimination task for the second part of this experiment.

Experiment 1b

We next explored shape discrimination in multi-limbed shapes. The task used here necessitated a much more holistic assessment of shape than in Experiment 1a: Shapes had either three or four limbs, and subjects' task was to discriminate between these two possibilities. Heightened spatial resolution of small portions of contour would not suffice for correct performance—large spatial extents of contour would have to be integrated into a shape percept in order to distinguish how many limbs, or

"parts", the shape has. In essence, this task required a judgment of the perceived underlying shape structure.

Methods

Stimuli

As in Experiment 1a, shapes were generated from skeletal axes. Ten three-limbed and ten four-limbed skeletons (Figure 4A) were created as follows: The angular direction of each limb was drawn randomly from the range 0–60, 120–180, and 240–300 degrees (three-limbed shapes) or 0–45, 90–135, 180–225, and 270–315 degrees (four-limbed shapes). The constraints on the geometry of
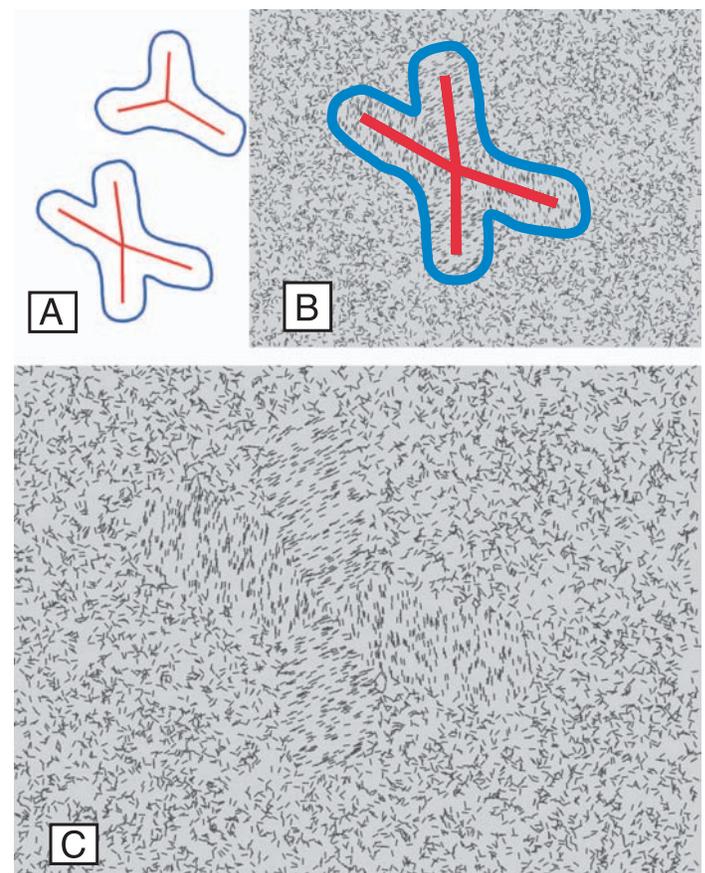

Figure 4. Examples of texture stimuli in Experiment 1b. (A) Examples of three-limbed and four-limbed shapes; shape boundaries are depicted in blue and shape axes are in red. (B) Allocation of texture elements in a four-limbed shape stimulus; orientation of elements within the shape boundaries is constrained by the texture–axis offset condition (here 67.5 degrees) with relation to the nearest part axis, and coherence level (here 100%). Elements in the surround have random orientation. (C) The texture stimulus in (B) as it appeared on a trial. Total stimulus size was 30 × 40 degrees of visual angle. Subjects' task was to discriminate between the two possible shapes.

these shapes meant that there was a greater chance of two limbs appearing as a single originating axis in the 4-limbed shape, a situation that the shapes depicted in Figure 4 exemplify. However, this potential influence, as well as in fact supposing an importance of axes, would not be expected to vary with texture–axis offset. Shapes were later randomly rotated before creating the textured stimuli, so these initial limb orientations were ultimately only relative orientations. The length of each limb was drawn from a normal distribution with a mean of 200 pixels and a standard deviation of 30 pixels; if the chosen value fell outside the range of 150–250 pixels (6–10 degrees), the limb length was assigned the limiting value. Shape boundaries were then calculated from these multi-limbed skeletons using a root width of 70 pixels, i.e., boundaries were located at a perpendicular distance of 70 pixels from the closest axis (or root). The location of the mid-point of each shape was drawn randomly from a central 100 × 100 pixel (4.0 × 4.0 degree) window.

Texture stimuli were generated as in Experiment 1a, by random placement of 7000 line elements. Elements that fell within the shape boundary were assigned an orientation according to the texture–axis offset condition of 0, ±22.5, ±45, ±67.5, or 90 degrees; elements located outside the shape boundary were assigned random orientations. For these multi-limbed shapes, texture orientation within each limb was assigned relative to the orientation of that limb's axis, therefore texture had a different absolute orientation in each limb of a shape (Figures 4B and 4C). We also tested an additional condition in this experiment, in which a single, absolute, texture orientation, chosen at random for each occurrence, was applied across the entire shape. Note that in this condition, which we will refer to as "constant" (to differentiate it from the other conditions where each limb has texture of a different orientation), the axis orientations become irrelevant. Texture elements located outside the shape boundary were assigned random orientations. Images were pregenerated, along with mask screens consisting of 7000 randomly oriented line elements.

Texture shapes were presented at six coherence levels (50%, 60%, 70%, 80%, 90%, and 100%). When crossed with the six texture–axis offset conditions and ten versions each of three-limbed and four-limbed shapes, this resulted in a total of 1440 unique pairs of stimulus and mask, which were each presented once per testing session in random order.

The timing of the presentation sequence was identical to that shown previously: A central fixation cross was displayed until the subject initiated the trial with a key press. The stimulus was then presented for 200 ms, followed by a 100-ms blank screen, then a 100-ms mask. The screen then remained blank until the subject indicated with a key press whether a 3-limbed or 4-limbed shape had been presented. Feedback was given as before, in the form of a quiet tone to indicate an incorrect response. Subjects were instructed by means of an on-screen message to take a break after every 240 trials; they could take

additional breaks whenever needed simply by delaying initiation of the next trial.

Subjects

Four subjects took part in Experiment 1b. Each completed three testing sessions, generating a total of 720 responses for each texture–axis offset condition spread over the six coherence levels. Subjects were drawn from the pool of subjects who completed Experiment 1a, on the basis of availability and good task comprehension. Of the four subjects, only one had shown a clear axis alignment effect in Experiment 1a, and all four had shown edge alignment effects, to varying degrees. All subjects remained naive to the purpose of the experiment.

Experiment 1b: Results

Subjects found this type of shape discrimination task, which probed their perception of underlying shape structure, considerably easier than the previous task, which had required high spatial resolution of shape contours. Therefore, we extracted thresholds at the 85% correct performance level (Figure 5). Performance was even more strongly affected by texture–axis offset than in the previous task (ANOVA over all six offset conditions: $F(5, 15) = 36.728$, $p < 0.001$). The form of this dependence (excluding the constant texture condition) was very strongly quadratic ($F(1, 3) = 45.869$, $p = 0.007$), mirroring the previously seen performance pattern of higher thresholds at 45 degree texture–axis offset compared to 0 and 90 degrees. However, the data also showed a strong linear component ($F(1, 3) = 62.554$, $p = 0.004$), as thresholds did not decrease strongly between texture–axis offsets of 45 and 90 degrees.
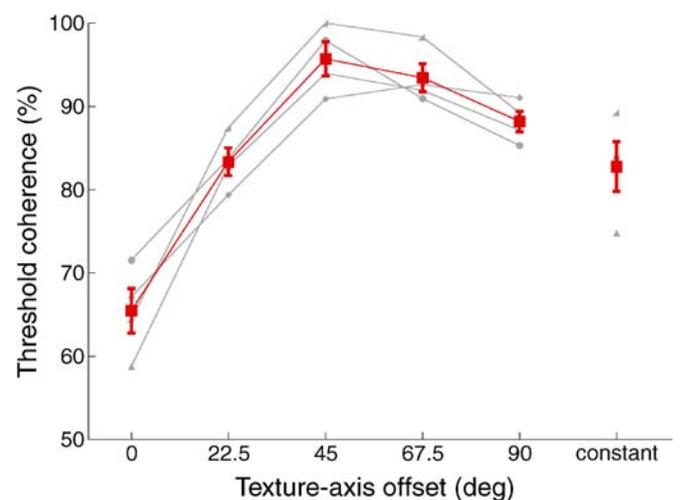

Figure 5. Results from Experiment 1b. Texture coherence thresholds (85% correct) in the shape discrimination task for all four subjects (gray) and mean values (red, with error bars showing ±SEM) for the six texture–axis offset conditions.

These patterns in performance were borne out by planned paired comparisons (with Bonferonni correction), which showed that thresholds at an offset of 0 degrees were significantly lower both than at 45 degrees and at 90 degrees ($\Delta threshold_{(45-0)}$ = 30.24% coherence, $p < 0.013$; $\Delta threshold_{(90-0)}$ = 22.73% coherence, $p < 0.032$), but that thresholds at 90 degrees were not significantly lower than at 45 degrees ($\Delta threshold_{(45-90)}$ = 7.51% coherence, $p < 0.458$). Notably, thresholds at 0 degrees offset, where texture was aligned with the axis within each limb of the multi-limbed shapes, were lower than for the constant texture orientation condition, where texture within the entire multi-limbed shape region had just one, randomly chosen, orientation ($\Delta threshold_{(constant-0)}$ = 17.33, $p = 0.043$).

The average texture orientation difference between adjacent limbs, as seen most notably at the limb junctures, was 90 degrees for 4-limbed shapes and 60 degrees for 3-limbed shapes, both highly suprathreshold values for texture segmentation. Additionally, varying the texture–axis offset does not change these values, and so this cannot be the cause of the variation in discrimination performance seen here. However, could it be that performance variation is driven by an effect of local texture alignment with inter-limb borders at shapes' centers? The number of borders infers the number of distinct shape parts, and so increased saliency of the borders could drive performance without requiring awareness of the full shape region. To the contrary, the geometry of our shapes means that texture–axis offsets intermediate to 0 and 90 degrees would most often result in texture alignment with these central inter-limb borders (45 deg texture–axis offset for 4-limbed shapes and 60 deg offset for 3-limbed shapes). Hence, if subjects were simply "counting" the number of regions of different texture orientation, texture–axis offsets of 22.5, 45, and 67.5 would be predicted to promote more strongly separable regions and lead to lower thresholds, the opposite of the pattern seen here.

Experiment 1: Discussion

Our results strongly support a role for skeletal axis representation in the segmentation of texture-defined shapes; alignment of texture structure with shape axes led to shape segmentation at lower texture coherence levels. In summary, Experiment 1a demonstrated that discrimination of shape contour benefits from texture alignment with the shape axis, and that to a slightly lesser extent the perpendicular relation facilitates performance also. This is despite the fact that the underlying shape structure, or axis, is unchanged by fine alterations in shape, such as those that differentiate "fat" and "thin" peanuts. Experiment 1b demonstrated that discrimination of the underlying structure of multi-axis shapes strongly benefits from texture alignment with the axis within each limb. Performance with parallel texture–axis structure was better even than for shapes with constant texture

orientation throughout the shape. While the straight edges of our multi-limbed shapes mean that alignment between texture and shape edges is an equally valid explanation for the results of Experiment 1b, in conjunction with Experiment 1a our results point to an effect of texture–axis alignment that is both separate from, but possibly complimentary to, that of texture-edge alignment.

Together, Experiments 1a and 1b fulfill our aim of looking at processes by which elements group into entire coherent objects, not just local aggregation of elements. This assertion is evidenced by our initial observation of the effect in shapes where local contours are not parallel to the axis, and our further observation of the effect in a task that required a holistic shape representation that could not be attained solely by local aggregation and contour extraction. Additionally, subjects reported that they perceived regularity in the stimulus texture at low coherence thresholds but could not necessarily discerning the underlying structure or shape of the organized region. We surmise that it was the perception of overall region shape that was affected by our manipulation of the texture, not the detection of regularity or of portions of boundary. The perception of discriminable regions, and of limbs within the shape, appears to benefit from alignment of texture with the shape axes. The strength of this influence is corroborated by the fact that it is seen in two shape discrimination tasks that could conceivably have been performed on the basis of shape size, or total spatial extent, alone.

Facilitation of segmentation by texture–axis alignment could be explained in at least two ways: Firstly, enhanced texture grouping could lead directly to segmentation, with grouping strongest for parallel, and to a lesser extent perpendicular, texture–axis arrangements. Recall that "parallel" and "perpendicular" are here defined relative to the global shape axis, hence, their influence on segmentation implies a global-to-local influence. The concept of global structure guiding local grouping operations, which themselves define the global structure, is contrary to standard hierarchical models, which would posit that the global structure is derived consequent to any local processing. A second explanation is that initial processing extracts shape axes, for instance by low-level lateral interactions as proposed by Zhaoping (2003): Recognition of axes could then cue subjects as to likely locations of informative boundary information (Experiment 1a) or in fact be sufficient for correct task performance without further processing (Experiment 1b). It is plausible that both of these mechanisms operate depending on the task demands.

An alternative to perceiving discriminable texture regions would be to segment the boundary of the entire region in a piecemeal manner, as suggested by a hierarchical segmentation scheme, regardless of the region's global form. "Shape" or "structure" would then be a later construct, derived from the extracted boundaries. Under this mechanism, it would not be hypothesized that texture

structure would impact on segmentation; clearly, this is not what we observed. In addition, we found that thresholds were higher in our "constant" condition (Experiment 1b) than when texture was parallel to the axes, reinforcing the notion that it is the distinctness of the separable shape regions that was enabling performance in this task, and that the strength of grouping within these separable regions is highest when texture is aligned with the region axes. While traditionally, performance might be expected to be best for constant texture orientation, we find that this is not the case with our more complex, multi-axis shapes. It appears that the internal coherence of the texture *relative to the shape axis* is more influential than simple uniform orientation "coherence."

To conclude Experiment 1, both segmentation tasks benefited from alignment between the shape axes and texture within the shape, suggesting that the "lower" or "locally driven" process of OBTS is influenced by the "higher" process of global shape extraction. The way in which texture influenced performance in our tasks strongly supports the idea that the visual system holds a skeletal representation of shapes. Our results therefore provide evidence for the psychological reality of such representations.

Experiment 2

We next considered another aspect of texture perception, to see if the effect that we observed in Experiment 1 generalizes or is in fact specific to segmentation. The task under question in Experiment 2 is the ability to report the orientation of a texture. Experiment 1 showed that subjects' performance in two very different shape discrimination tasks was highest when texture was aligned with the shape axes. However, the phenomenon of texture segmentation does not automatically imply that subjects can report the composite texture orientations; boundaries can be perceived without knowledge of the surface properties that underlie the boundary percept (see, for instance, Ben-Shahar, 2006, Experiment 2; see also Rogers-Ramachandran & Ramachandran, 1998). In our experiments though, increased visibility of texture–axis aligned shapes at low texture coherence levels suggests that the property of the texture that mediates segmentation, i.e., its orientation, is more strongly represented, at some level, when texture and axes are aligned. Accordingly, Experiment 2 directly investigates this aspect of texture perception, in a series of texture orientation discrimination experiments.

The issue of orientation discrimination is particularly interesting because it potentially allows us to distinguish between two possible roles that alignment of texture with shape axes could be playing in our previous segmentation experiments: One possibility is that it causes the entire ensemble of oriented elements to be more strongly represented and grouped, leading to greater region saliency. An alternative is that it affects segmentation indirectly, by improving subjects' ability to locate the shape axis. Knowledge of the shape axis would then refine further judgment of shape properties. As mentioned previously, these possibilities are not mutually exclusive.

Using an equivalent experiment design to Experiment 1, we investigated the relationship between texture patch shape and texture orientation discrimination by varying the offset between the shape axis orientation and the texture orientation. Our rationale, as before, was that simply altering the shape itself, while conceivably producing changes in orientation discrimination performance, would not further inform us as to what property of the shape was influential. Texture orientation discrimination thresholds were established by sequentially presenting two stimuli. Subjects reported whether the texture in the second stimulus appeared to be rotated clockwise or counterclockwise compared to the first stimulus. The orientation of the texture shape itself (i.e., its axis) was unchanged in each pair of stimuli; only the orientation of the composite texture was changed. Segmentation of the texture shapes was not an issue in this series of experiments, as the textures were presented against an untextured (mid-gray) background.

In the series of experiments that follow, we ask whether the ability to discriminate texture orientation, which is a first-order orientation signal contained in a collective group of local orientation signals, is affected by the global shape of the region that the group of signals occupies. This should not be confused with studies such as that of Regan (1995), who investigated orientation discrimination of the structures *defined* by texture, not of the texture itself. Additionally, we do not seek to quantify any misperception of orientation, as has been shown to occur between incongruent local and global orientation signals (Morgan & Baldassi, 1997; Morgan, Mason, & Baldassi, 2000). In the following experiments, two orientation signals are being compared rather than one orientation signal compared to an internal reference such as vertical.

Experiment 2a

The first experiment simply sought to establish whether there was an effect of texture–axis offset on texture orientation discrimination for simple oblong-shaped texture patches (shown in Figure 6C). Oblongs had a total length of 16.0 degrees and a width of 4.0 degrees (generated from an axis of length 300 pixels and with root width of 50 pixels). Orientation discrimination thresholds are very sensitive to absolute orientation, with thresholds substantially lower at vertical and horizontal compared to 45 degrees (Appelle, 1972; Orban, Vandenbussche, & Vogels, 1984). So as not to confound texture–axis offset condition with absolute orientation of the texture, and also not introduce wide
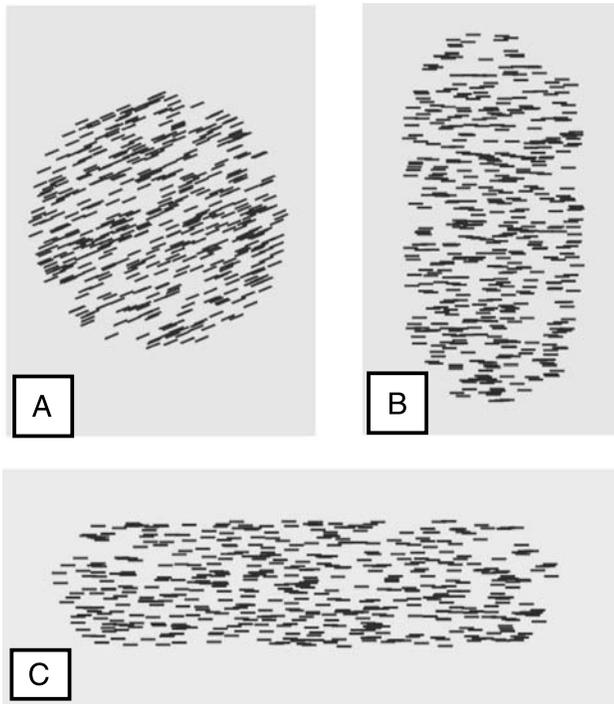
Figure 6. Sample stimuli used in Experiment 2. The most elongated shape, (C) Shape 3, was used in Experiment 2a. Experiment 2b used (C) Shape 3 together with (A) Shape 1 and (B) Shape 2.

variability in performance within any given offset condition, the possible values of absolute texture orientation were restricted. This was achieved by presenting shapes at orientations of ±22.5 and ±67.5 degrees and investigating the offset conditions of 0, ±45, and 90 degrees only. These parameters result in absolute texture orientation for each offset condition being balanced across the orientations of ±22.5 and ±67.5 degrees, and so any effect of absolute orientation on discrimination thresholds was both minimized and equalized across offset conditions.

Stimuli

Texture stimuli were generated similarly to those in Experiment 1, with the exception that texture patches were presented on an untextured background: The first stimulus of a pair was created by randomly placing 500 line elements within the shape boundary, with orientation as prescribed by the texture–axis offset condition (0, 45, or 90 degrees). The second stimulus was also created by random placement of line elements but with an increment change to the texture orientation. Textures were 100% coherent and had no local orientation jitter. However, overlap of randomly placed line elements often occurred, and so spurious orientation signals could arise within a local region due to the formation of longer "false" line elements. While likely the cause of the slightly higher-than-usual orientation discrimination thresholds that follow,

this detail of our stimuli also encouraged holistic processing of the texture region, rather than narrow processing of a single element or small group of elements, so as to achieve accurate comparison of the two sequentially presented textures.

There were five equally spaced levels of orientation change increment that were applied between the first and second stimuli, with the exact magnitude of the increments for each subject decided after practice trials. A total of 240 unique stimulus pairs was pregenerated, consisting of 8 repeats of each combination of shape orientation, texture–axis offset, and texture orientation increment. Each stimulus pair was presented once, in random order, per testing block.

Each trial commenced with presentation of a fixation cross. When the subject initiated the trial with a key press, the first stimulus was presented for 200 ms, followed by a blank screen for 1000 ms, and then the second stimulus for 200 ms. No mask was used. The subject then indicated with a key press whether the texture in the second stimulus was rotated clockwise or counterclockwise relative to the first stimulus (reporting the direction of smallest orientation change). Feedback was given for an incorrect response.

Subjects

Thirteen naive subjects took part in Experiment 2a. Twelve subjects completed 3 testing blocks, each generating a total of 240 trials for each texture–axis offset condition spread over all levels of orientation increment. One subject completed 5 testing blocks, generating 400 trials for each texture–axis offset condition.

Experiment 2a: Results

Orientation discrimination thresholds, extracted at the level of 85% correct, are presented in Figure 7. Mirroring the performance pattern that we saw in Experiment 1, highest performance is seen when texture is parallel or perpendicular to the shape axis. The variation in thresholds with offset was significant ($F(2, 24) = 11.943$, $p < 0.001$), and the form of this dependency was significantly quadratic ($F(1, 12) = 13.691$, $p = 0.003$). Paired comparisons (with Bonferonni correction) showed that the threshold at 45 degrees offset was significantly greater than that at 0 degrees or 90 degrees ($\Delta threshold_{(45-0)} = 4.082$ degrees, $p = 0.018$; $\Delta threshold_{(45-90)} = 4.740$ degrees, $p = 0.007$) and that thresholds at 0 and 90 degrees were not significantly different.

Experiment 2b

Having established that there was indeed an effect of texture–axis offset on orientation discrimination, we next wanted to chart the development of the axial influence
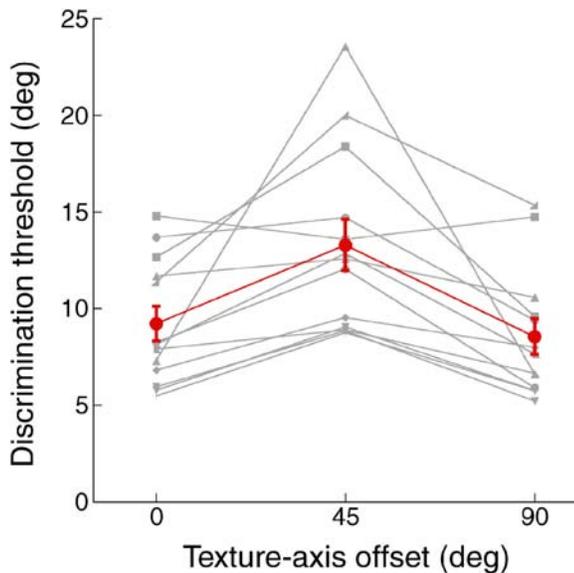
Figure 7. Results from Experiment 2a. Orientation discrimination thresholds (85% correct) for texture in the shape shown in Figure 6C. Thresholds are shown for 13 subjects (gray) and mean values (red, with error bars showing ±SEM) for the three texture–axis offset conditions.

as elongation increased. A circular patch by definition does not have an elongated axis. However, would an "axis effect" emerge gradually with axis length, or plateau for all non-circular shapes? In a Bayesian framework (Feldman & Singh, 2006), the length of an oblong shape relative to its width modulates the likelihood ratio of the axial interpretation relative to a neutral circular interpretation, and thus the posterior strength of "belief" in the shape axis. Varying the aspect ratio thus allowed us to investigate the strength of the axis alignment effect as a function of the theoretical perceived strength of the axis.

To explore this issue, we used oblongs with three different aspect ratios. The exact parameters of the oblongs were chosen so as to keep maximum width × maximum height equal across all three aspect ratios.

Stimuli

Shapes had relative maximum width to maximum length ratios of 2:2, 1.33:3, and 1:4 (corresponding to absolute aspect ratios of 1:1, 1:2.25, and 1:4; Figures 6A–6C, respectively). The shape dimensions were 8.0 × 8.2 degrees (Shape 1, virtually circular), 5.36 × 12.0 degrees (Shape 2, medium elongation), and 4.0 ×16.0 pixels (Shape 3, the most elongated shape, identical to that used in Experiment 2a). These shapes were generated from single straight axes of lengths 5, 167, and 300 pixels, with root widths of 100, 67, and 50 pixels, respectively.

All other parameters and stimulus details were as for Experiment 2a. Note that the shape orientation and texture–axis offset parameters have no perceptual meaning

in the case of the near-circular shape, as there was no discernable shape axis. A total of 270 unique stimulus pairs was pregenerated, consisting of three repeats of all combinations of shape aspect ratio, shape orientation, texture–axis offset, and level of orientation increment. Each stimulus pair was presented once, in random order, per testing block.

Subjects

In this and all further parts of Experiment 2, subjects were chosen from the pool of subjects who had completed Experiment 2a, on the basis of availability and good task comprehension. The overall pattern of performance did not change despite the possibility that learning at the harder texture–axis offsets in particular could have occurred, weakening any potential axis effect. Subjects remained naive to the purpose of the experiments. Five subjects took part in Experiment 2b. The results of one subject were discarded, as performance did not reach threshold in some conditions despite showing high performance in Experiment 2a. The remaining four subjects completed five testing blocks, each generating 150 responses per texture–axis offset condition within each of the three shapes, distributed across texture orientation increment levels.

Experiment 2b: Results

Thresholds are presented in Figure 8. It can be seen that while the circular shape (Shape 1) did not show an "axis effect," both further elongations of the shape showed the previously observed pattern of low thresholds for texture–axis offsets of 0 and 90 degrees with a drop in performance (higher thresholds) at a texture–axis offset of 45 degrees. A two-way ANOVA (shape (3) × texture–axis offset (3)) confirmed that texture–axis offset was a significant factor in its own right $(F(2, 6) = 11.325, p = 0.009)$, due to the strong effect shown for Shapes 2 and 3, and that the interaction of shape and texture–axis offset was significant $(F(4, 12) = 8.027, p = 0.002)$. However, post-hoc analysis of the interaction of shape and texture–axis offset within pairs of shapes (two-way ANOVAs; shape (2) × texture–axis offset (3)) revealed that while the effect of texture–axis offset on segmentation was different for Shape 1 compared to Shape 2 $(F(2, 6) = 15.799, p = 0.004)$ and for Shape 1 compared to Shape 3 $(F(2, 6) = 9.118, p = 0.015)$, it was not different between Shapes 2 and 3 $(F(2, 6) = 0.24, p = 0.794)$. Apparently the influence of the axis on texture organization saturates by a 1:2.25 aspect ratio (Shape 2), perhaps because the axial posterior probability has by this point nearly reached a ceiling.

Shape elongation was not a significant main factor; threshold values did not differ significantly across the three shapes because the effect of texture–axis offset in
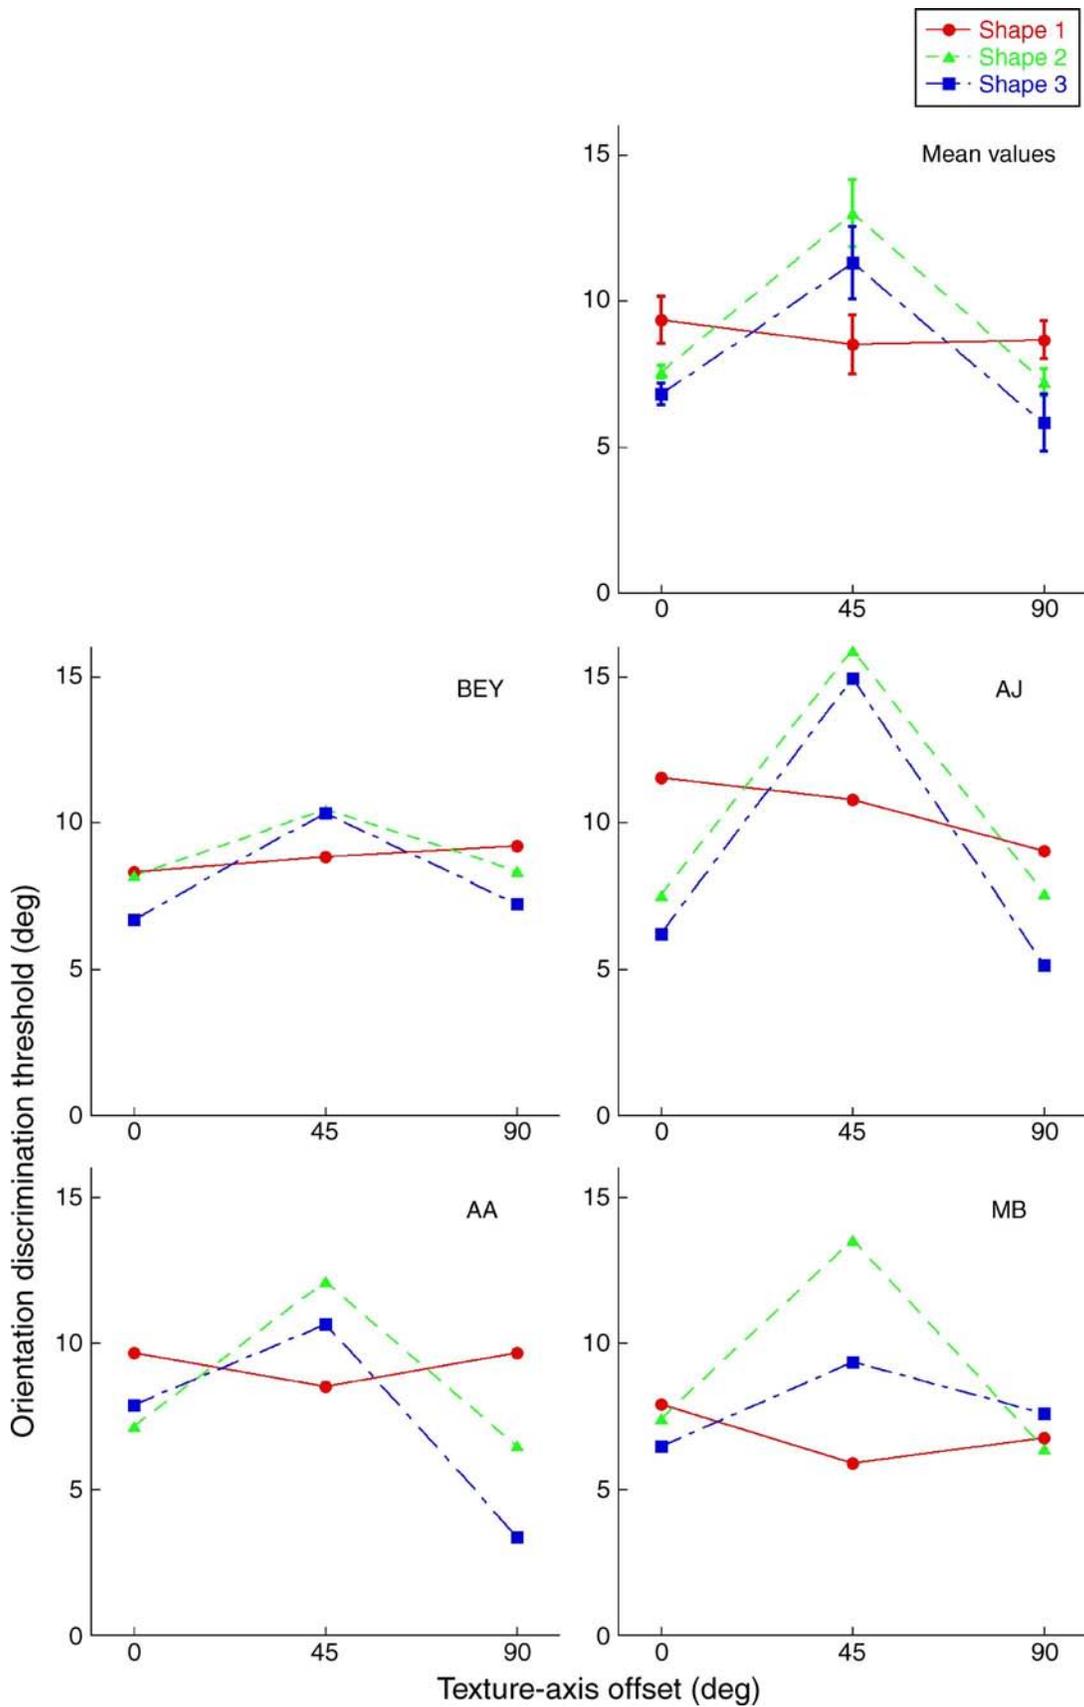
Figure 8. Results from Experiment 2b. Orientation discrimination thresholds (85% correct) for texture in the three shapes depicted in Figure 6. Thresholds are shown for 4 observers and mean values with error bars showing $\pm SEM$ (top right).

Shapes 2 and 3 was to improve performance at offsets of 0 and 90 degrees but to decrease performance at an offset of 45 degrees. We surmise that the axis effect is a differentiation of performance across all texture–axis offset levels rather than a simple facilitation or inhibition of performance at specific points in the texture–axis relationship.

Experiment 2c

Following on from our investigation of how the axis effect related to shape aspect ratio, we conducted further trials to investigate whether aspect ratio or oblong width was the more influential factor in determining the strength of the axis effect. This issue was raised by Zhaoping's (2003) model of interactions within V1, which proposes that previous observations of heightened firing rate in response to elements along a shape's axis, and corresponding increased perceptual saliency, varies with the width of the texture region: Would the axis effect observed in our experiments be similarly affected by shape width? Alternatively, the axis effect could remain constant for all shapes of the same aspect ratio, regardless of their size—this would mean that the axis effect was scale-invariant and was related to the saliency of the shape axis.

Stimuli

Stimuli were three oblong-shaped texture regions: One shape corresponded to the medium elongation used in Experiment 2b (Figure 9A, Shape 2). The next shape had the same width as Shape 2 but was reduced in length by 25%, thus diminishing the aspect ratio (Figure 9B, Shape 4). The last shape (Figure 9C, Shape 5) was reduced in both length and width by 25%; this meant that it has the same aspect ratio as Shape 2 and the same length as Shape 4. The shape dimensions were 5.36 × 12 degrees, 5.36 × 9.0 degrees, and 4.0 × 9.0 degrees. These shapes were generated from single straight axes of lengths 167, 92, and 126 pixels, with root widths of 67, 67, and 50 pixels, respectively. Texture stimuli were created as previously, but to maintain constant density of line elements across all three shapes, the number of line elements placed was 500, 362, and 289, respectively.

A total of 270 unique stimulus pairs was pregenerated, consisting of three repeats of all combinations of shape aspect ratio, shape orientation, texture–axis offset, and levels of orientation increment. Each pair was presented once, in random order, per testing block.

Subjects

Five subjects took part in Experiment 2c. Each subject completed five testing blocks, each generating 150 responses per texture–axis offset condition within each of the three shapes, distributed across texture orientation increment levels.
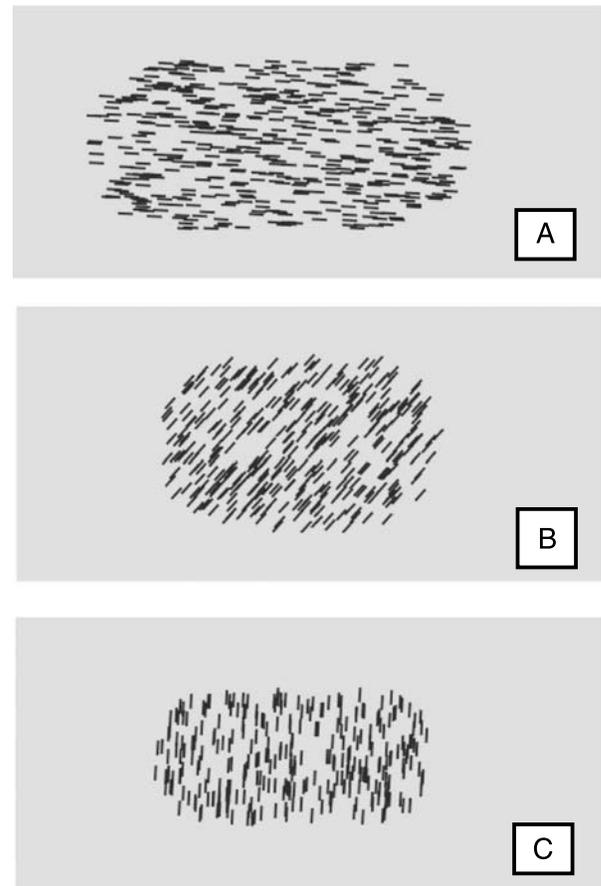

Figure 9. Texture shapes used in Experiment 2c. (A) Shape 2 was identical to that used in Experiment 2b. (B) Shape 4 had the same width but was reduced in length by 25% compared to Shape 2. (C) Shape 5 was reduced by 25% in both width and length compared to Shape 2.

Experiment 2c: Results

Thresholds are presented in Figure 10. It can be seen that thresholds for the three shapes are not significantly different overall and that axis effects may be seen for all three shapes. However, thresholds are far more variable than for the shapes used in Experiment 2b. In particular, two subjects failed to show an axis effect for Shape 2, a shape that had shown a robust axis effect in Experiment 2b (Figure 8), including for one of the same subjects (BEY). A 2-way ANOVA (shape (3) × texture–axis offset (3)) confirmed that texture–axis offset was a significant factor overall ($F(2, 8) = 7.914$, $p = 0.013$), and the interaction between shape and texture–axis offset also reached significance ($F(4, 12) = 3.093$, $p = 0.046$).

Post-hoc analysis of the interaction of shape and texture–axis offset within pairs of shapes (two-way ANOVAs; shape (2) × texture–axis offset (3)) revealed that the effect of texture–axis offset was more pronounced for Shape 5 compared to the less-elongated Shape 4 ($F(2, 8) = 11.681$, $p = 0.004$) and was not significantly
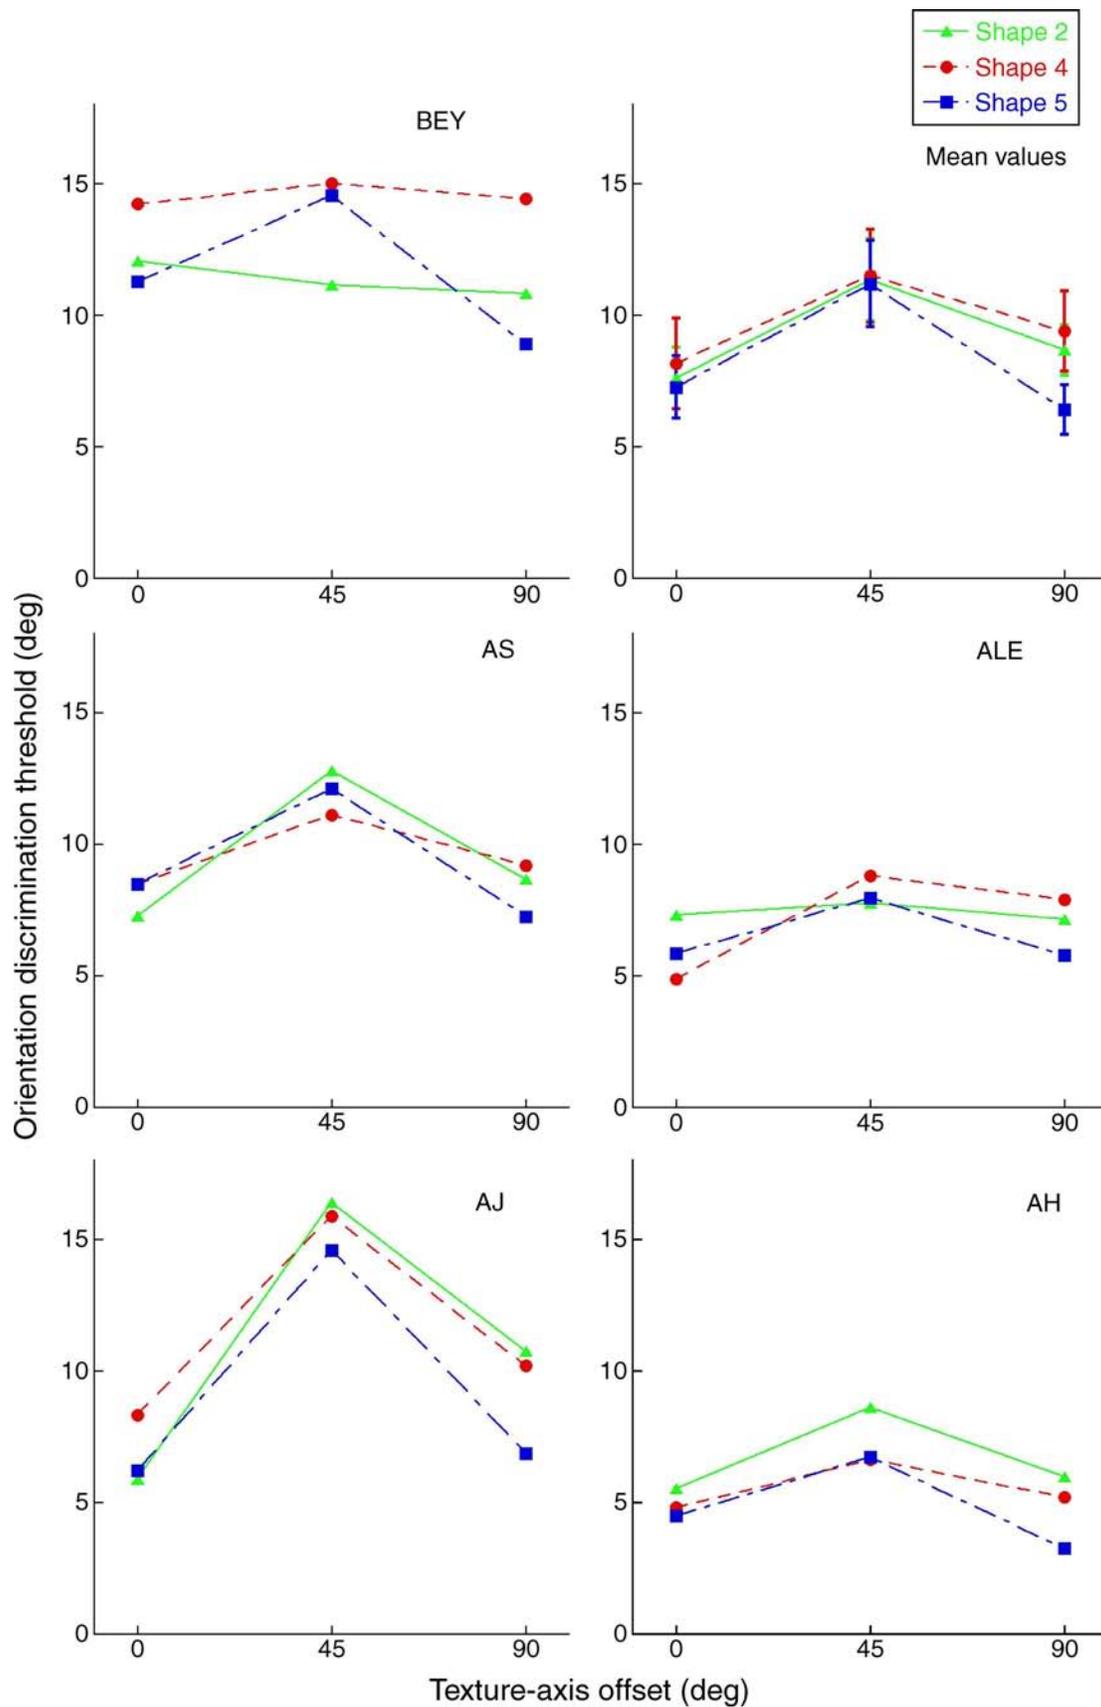
Figure 10. Results from Experiment 2c. Orientation discrimination thresholds (85% correct) for texture in the three shapes depicted in Figure 9. Thresholds are shown for 5 observers and mean values with error bars showing $\pm SEM$ (top right).

different between the equally elongated Shapes 2 and 5 ($F(2, 8) = 3.042$, $p = 0.104$). However, the variability between subjects for Shape 2 meant that the overall pattern of performance was not significantly different to that for the less elongated Shape 4 ($F(2, 8) = 0.117$, $p = 0.841$), despite showing a numerically slightly greater axis effect overall (a greater difference between thresholds at texture–axis offsets of 0 and 90 degrees compared to 45 degrees).

The pattern of results in three subjects (AJ, AS, AH), with a more pronounced axis effect for the two shapes with greater aspect ratio (Shapes 2 and 5) suggests that the effect is modulated by shape elongation regardless of absolute scale. This is consistent with a scale-invariant account of axial strength such as that given in Feldman and Singh (2006). However, the overall pattern of results does not rule out an influence of shape width, as the two shapes of equal width, Shapes 2 and 4 were not significantly different despite having different elongations. We prefer the first interpretation, as the second is heavily influenced by the results of one observer (BEY), who showed a very different pattern of performance to that shown previously in Experiment 2b (a point to which we return below). In any case, this result should be interpreted with considerable caution given both the small magnitude of the observed differences between shapes and the variability in performance across subjects.

Additionally, the previous lack of difference between the two elongated shapes used in Experiment 2b suggests that we may be looking for differences in axis effects relatively close to their asymptotic levels. However, another possible explanation for the pattern of performance, whereby Shape 2 did not show a significantly greater axis effect than did the less elongated Shape 4, is suggested by the loss of axis effect in Shape 2 when it had been observed previously in Experiment 2b (including for one observer, BEY, who took part in both experiments): The extent of spatial integration of orientation signals may in part be mediated by observers' expectation of shape size. While no longer a popular concept in the attention literature, the area encompassed by the "zoom lens" of attentional locus (Eriksen & St James, 1986) could be affected by subjects' expectations of the size of object that will appear in the brief presentations. This would mean that the full extent of the most elongated shape (Experiment 2b, Shape 3) or largest shape (Experiment 2c, Shape 2) would not be integrated, or distal portions would be given less weight in the integration process—a type of "set effect". Any such influence on performance would moderate the observed axis effect in these shapes as compared to smaller shapes interleaved in the same testing block.

Experiment 2d

The effect seen so far in this second series of experiments could conceivably be explained by the known effect
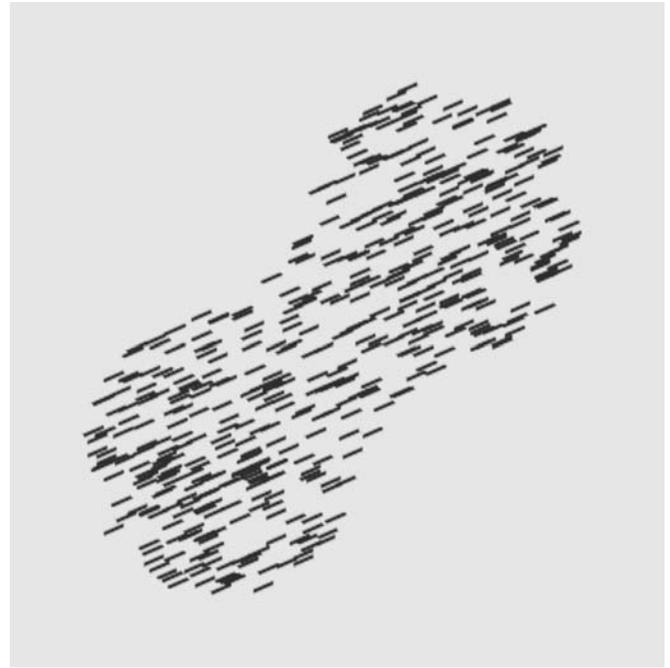

Figure 11. An example of the peanut-shaped texture stimuli used in Experiment 2d.

of enhanced edge saliency when texture is locally aligned with a region edge (Nothdurft, 1992; Wolfson & Landy, 1995). For instance, when the edge and the texture orientation are congruent in the first of a pair of presentations, perhaps subjects could simply assess texture orientation relative to the shape edge in the second presentation rather than having to compare with the first texture orientation. Hence, as in Experiment 1a, we once again sought to confirm that the pattern of performance was not due to local alignment of texture with edges. We tested texture orientation discrimination using "thin peanut" shapes, similar to those used in Experiment 1a, in which shape edges undulate relative to the shape axis (Figure 11). If the axis effect in fact relies on a local parallel relationship between texture and patch *edge,* then it should be disrupted in peanut shapes.

Stimuli

Peanuts were created, as before, from a single straight axis (or "root"). The shape contour was calculated by varying the perpendicular distance (or "root width") from the axis according to one cycle of a cosine function starting at 0 deg. Root width was held constant around the end of each shape, giving the shapes rounded endings. Axis length was 167 pixels, mean root width was 60 pixels, and cosine amplitude was 14 pixels. This results in shapes with a maximum length of 11.5 degrees and a central width of 3.7 degrees. Texture stimuli were created as previously, with 500 line elements randomly placed within the shape boundary.
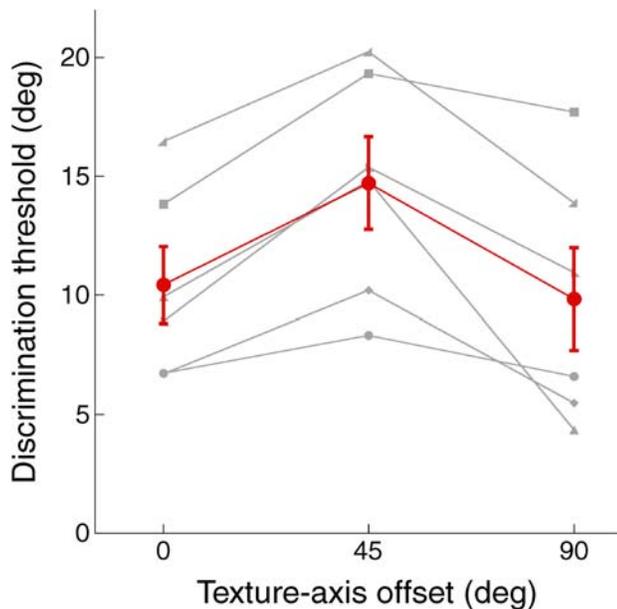
Figure 12. Results from Experiment 2d. Orientation discrimination thresholds (85% correct) for texture in the shape as depicted in Figure 11. Thresholds are shown for six subjects (gray) and mean values (red, with error bars showing ±*SEM*) for the three texture–axis offset conditions.

A total of 240 unique stimulus pairs was pregenerated, consisting of eight repeats of each combination of shape orientation, texture–axis offset, and texture orientation increment. Each stimulus pair was presented once, in random order, per testing block.

Subjects

Six subjects took part in Experiment 2d. Subjects completed three testing blocks, resulting in 240 trials per offset condition.

Experiment 2d: Results

Results for the "peanut" stimuli are shown in Figure 12. Performance is clearly degraded when the offset between texture orientation and shape axis is increased to 45 degrees. As seen in the peanut segmentation task (Experiment 1a), texture oriented both parallel and perpendicular to the shape axis facilitated performance. The variation of thresholds with offset was significant (1-way ANOVA; $F(2, 8) = 7.141$, $p = 0.017$), as was the quadratic relationship between performance and texture–axis offset ($F(1, 4) = 21.986$, $p = 0.009$). Paired comparisons (with Bonferonni correction) showed that the threshold at 45 degrees offset was significantly greater than that at 0 degrees ($\Delta threshold_{(45-0)} = 4.444$ degrees, $p = 0.018$), but the difference between 45 and 90 degrees, though of greater magnitude, was not significant ($\Delta threshold_{(45-90)} = 4.916$ degrees, $p = 0.120$). Thresholds at 0 and 90 degrees were not significantly different. Hence, the data clearly show that the axis effect persists in peanut shapes and therefore cannot be attributed to local alignment of texture with region edges.

Experiment 2: Discussion

Experiment 2 demonstrates that texture orientation discrimination is affected by the relationship between texture orientation itself and the axis orientation of the shape that the texture region composes. We find lower thresholds when texture is aligned with or perpendicular to the shape axis, compared to when texture and axis are offset by 45 degrees. This axis effect cannot be attributed to alignment of texture with region edges, as the effect is robust to the undulating edges of our peanut shapes.

It is not clear over what spatial extent axis information is integrated: While three of our subjects in Experiment 2b showed a slightly stronger axis effect for the most elongated shape compared to the shape with lesser elongation, one subject showed the reverse effect with much greater magnitude (subject MB, Figure 8). In addition, Experiment 2c showed an axis effect in a shape with even lesser elongation (Figure 10, green data points). Hence, although we can say that the axis effect is observed at lesser elongations, at least in this orientation discrimination task, we cannot say at exactly what aspect ratio it would first be observed.

The ability to accurately compare the orientations of two sequentially presented textures relies on the encoding of orientation within each of the textures. Our finding of lower orientation discrimination thresholds when texture is aligned with or perpendicular to the shape axis implies that in these conditions the texture orientation is more strongly encoded, enabling subjects to make a more accurate comparison of sequential orientation signals. These results add to those of Experiment 1, showing that the "locally driven" process of texture orientation representation is influence by the "higher" process of global shape extraction, specifically axis extraction, by the visual system.

General discussion

Our data converge on a broad argument that the shape of the enclosing region, and more specifically its axial structure, influences perception of the enclosed texture. First, our segmentation data demonstrate that texture is most strongly integrated into a coherent shape region when it is aligned with, and to a lesser extent perpendicular to, the nearby shape axis. This is suggestive of a local coordinate frame aligned with the axis, with respect

to which texture is organized. The observed axis effect is not due to alignment with the region's contour, as demonstrated by the persistence of the effect in the segmentation of our "peanut" shapes. Indeed, alignment with the skeletal axis facilitates texture integration to such an extent that axis-aligned textures are integrated into coherent shapes and segmented more strongly than textures with uniform orientation, in the case of multi-part shapes. Secondly, our orientation discrimination data show that the effect of axial alignment on segmentation is likely due, at least in part, to stronger encoding of texture orientation. While it is conceivable that congruency between texture and axes could assist in identifying or locating shape axes and so facilitate consequent shape discrimination performance, an end result of this congruency, however it comes about, is a stronger encoding of the ensemble of local orientation signals. That is, perception of the local elements themselves has been altered.

Both orientation-based texture segmentation and discrimination of texture orientation involve, at some stage, a representation of texture orientation. However, they are very different tasks that rely on different mechanisms and place different demands on the visual system. The similar pattern of results for both tasks and under a range of conditions, always showing elevated performance for textures aligned with the axial coordinate frame, argues for the broad conclusion we have drawn: texture seems to be organized relative to shape axes. This axis effect would presumably not be observable when shape axes are absent or non-salient, as for the circular or square shapes most commonly used in the OBTS literature. However, when the axes are strong, as is the case for the shapes used here, the effect is readily measurable.

To what extent can alignment of texture with region edges, rather than with region axes, explain our findings? As mentioned earlier, Wolfson and Landy (1995) showed that texture boundaries are most salient when texture elements are parallel or perpendicular to the boundary (an effect also noted by Nothdurft, 1992). Indeed, the texture-edge alignment effect may well underlie a subset of our data in the peanut shape discrimination task (Experiment 1a) that cannot be explained by texture–axis alignment, where some subjects had decreased thresholds at texture–axis offsets of 22.5 and 67.5 degrees. Note, however, that Wolfson and Landy's (1995) task required discrimination between a perfectly straight edge and an edge defined by one cycle of sinusoidal modulation (where the maximum deviation of edge orientation from the "straight" orientation was just over 13 degrees). Subjects may have performed this task by deciding whether the edge was "sharp" or "blurry" rather than by resolution of the boundary per se. In contrast, our peanut shape discrimination task required a choice between two equally curved boundaries. While it is straightforward to envisage that local texture alignment with an edge should improve edge resolution, as per Wolfson and Landy (1995), it is less intuitive that alignment of texture with

the global average orientation of our varying edges should do so also. Enhanced edge saliency by local texture-edge alignment appears to be a separate effect to the broader finding here, where texture–axis alignment increases not only resolution of global contours (Experiment 1a) and structural saliency (Experiment 1b) but also orientational resolution of the texture as a whole (Experiment 2). However, several of our conditions are as well explained by edge alignment as they are by axis alignment—it is by considering results from all experiments together that we have drawn our conclusion as to the importance of axes.

Wolfson and Landy's (1995) stimuli contained only edges, in the absence of unambiguous figural assignment, and therefore their conclusions with regard to edge saliency were, de facto, justified. Additionally, given that edges and axes are irrevocably linked, one might ask whether the effect of texture alignment with edges and axes can truly be disentangled. After all, the axis orientation will approximate the average orientation of the two nearest segments of the figure boundary. In consideration of such issues questioning the relevance of axial representations in visual processes, we would suggest that axis recognition is likely more pivotal to our comprehension of the visual world than precise resolution of edges. Additionally, in a world consisting of coherent objects and coherent regions, not of unrelated edges, the orientations of regions rather than of edges may be more likely to infer organizational structure on the textures that they enclose. In summary, we propose that while local texture alignment impacts local edge saliency, strength of region representation and consequently segmentation of global forms appears to be strongest when texture is aligned with the shape axis.

The effect that we have begun to detail in this series of experiments is particularly interesting because of the connections it suggests between global or holistic aspects of stimulus interpretation and purportedly locally driven percepts such as texture segmentation. Such effects are not unprecedented; as detailed in the Introduction section, strong effects of global interpretations on the saliency of local elements are well known (Navon, 1977; Palmer, 1977; Pomerantz et al., 1977). Additionally, the effects of global influences on texture perception itself have increasingly been a topic of debate.

Texture segmentation has, in the past, been regarded as stemming primarily from analysis of local differences rather than global similarities (e.g., Nothdurft, 1994), with the allowance that the salience of local orientation contrast was related to its magnitude in comparison to the surrounding contrasts (Landy & Bergen, 1991; Nothdurft, 1991). However, more recent takes on texture perception have shed light on the complex role of texture element orientation in segmentation, with some interesting parallels with our study. Ben-Shahar and Zucker (2004) demonstrated that globally salient geometric patterns within regions of local elements can be highly decisive in the resulting texture percept. Equal magnitudes of orientation contrast across a region border can have very

different perceptual effects depending on the *vector* quantity of the orientation variations, or "flow", in the surrounding texture region. The effect of "flow" becomes increasingly significant in textures with higher levels of orientational variance. As Ben-Shahar and Zucker (2004) point out, textures with orientational variance are common in the visual environment because a strictly uniform texture requires a coincidence of viewing angle and object angle. This influence of globally defined patterns on the significance of local elements is reminiscent of the axis effect we have observed here.

Further, even in textures that vary smoothly in orientation, not all local contrasts of the same magnitude are created equal (Ben-Shahar, 2006). While some variations in texture orientation are interpreted as texture "flow" occurring within a coherent region, others of equal magnitude are perceived as "breaks", or region boundaries. Again, Ben-Shahar notes that processes by which textures are formed rarely result in regions with uniformly oriented content. These findings draw attention to the possible relationship of texture segmentation to natural texture formation processes; it may be that texture formation within an object rarely results in "perceptual singularities" (Ben-Shahar, 2006) or "edge hallucinations" (Nothdurft, 1992), or that when such perceptual breaks in structure do occur, they have structural significance. Likewise, could it be that the link between structure and local salience observed in our study is due to the relationship between texture orientation and the processes that form texture?

In broad terms, the interaction between texture and object axes can be viewed from the perspective of the statistics of naturally occurring textures. Textures that are predominantly parallel or perpendicular to the extended regions that they occupy may occur more frequently due to natural growth processes, where skeletal axes denote the growth path of an object. Such an arrangement could result when shape-forming processes such as growth influenced, or at least were correlated with, the processes that determined surface patterns. Naturally occurring textures are, after all, the result of a complex of formative processes that relate to biological, physical, or even chemical processes involved in the creation of the objects on which the textures reside (Ball, 1999; Thompson, 1917). Kass and Witkin (1985) have argued for the utility of characterizing textures in terms of the formative processes from which they developed.

Additionally, even if not statistically dominant, axis-aligned texture patterns might be especially informative about underlying structure. Extraction of skeletal axes may be of functional significance in vision not only for recognition, memory, and assessment of shape similarity but also for comprehension of shape structure and form function. If parallel and perpendicular textures promote extraction of these important axes, the axes themselves may in turn further strengthen the texture representation and consequent grouping. Structural principles underlying man-made structures may also result in a preponderance of meaningful parallel and perpendicular texture–axis arrangements, although admittedly man-made patterns presumably played little role in the evolution of the visual system. Ultimately, our greater sensitivity to textures parallel or perpendicular to region axes may be related to the predominance of such structures in our visual world and to the significance of such structures when they do occur.

Theoretical implications

The "why?" of the axis effect does not however answer the "how?" of axis extraction. Zhaoping's (2003) model of lateral interactions in V1 presents one framework within which our results could be explained. The model does indeed predict stronger axial extraction when texture is parallel to the region axis (though it is not immediately clear whether the model would predict a benefit for perpendicular texture arrangements, particularly in our peanut shapes). Yet another account of axis extraction that does not rely on cortical processing of shape is that of Burbeck and Pizer (1995). In their model, "medialness detectors" in V2 are activated by simultaneous input from locally symmetric fragments detected in V1, regardless of their global significance. The conjunction of multiple activated detectors then results in a probabilistic description of a region axis. Such models, describing the biological plausibility of increased axial saliency due to interactions within V1, do not speak as to whether such axes are refined, or have perceptual consequence, at higher cortical levels. Indeed, if the visual system does make use of medial axes in later stages of shape processing and recognition, then a mechanism is needed to extract at the very least an approximation of the axes in the first instance.

There are also models of axis extraction that *do* make use of global configuration. For instance, the skeletal structure may be gradually refined by resonance between local input signals and global context, an approach used by Lee et al. in their model of texture segmentation (Lee & Mumford, 2003; Lee et al., 1998). This type of solution has also recently been implemented by Adluru et al. (2007), who coupled extraction of shape contours and shape axes, using particle filtering to maintain multiple hypotheses while both local and global evidences gradually converge.

The above models appear to extract axes on the basis of spatial relations in a 2-dimensional image. An outstanding question therefore is whether axis extraction would correspond to the true, structural, axes of a 3-dimensional object, or to the axes in the object's 2-dimensional retinal image. Axial structure is a type of local symmetry (Brady & Asada, 1984), meaning that each shape axis induces a local symmetry axis. The above question can therefore also be considered within the framework of the visual system's recognition of symmetry: A mirror-symmetric shape viewed from the normal of its surface has a corresponding symmetric retinal image. However, from many other viewpoints, the shape would exhibit *skewed* symmetry in

the retinal image. Skewed symmetry is known to be more difficult to detect than regular symmetry, but nonetheless it is recognized as a non-accidental property by the visual system and is influential in perceptual organization (Wagemans, 1992, 1993; Wagemans, Van Gool, & d'Ydewalle, 1991, 1992). Skewed symmetry would often occur not only in the retinal image projected by a symmetric 3-dimensional object but also in regular texture on objects' surfaces. However, detection of skewed symmetry in the texture may be more difficult (Sawada & Pizlo, 2008), and therefore less influential in perceptual organization of the texture as seen here, compared to detection of symmetry of the shape contour itself. It is likely that the skewed symmetry of a real 3D object under everyday viewing conditions would be easily recognized given cues to depth such as disparity.

The question remains then, would recognition of the true object shape drive axis extraction, or would the 2D projection of the shape (when the two do not coincide) result in "incorrect" axes being inferred? Investigation of this issue may help clarify the separate roles of texture edge and axis alignments, as the two would be deconfounded for certain classes of shapes and viewpoints. A somewhat simpler manipulation might be to study the axis effect under conditions of shape occlusion: Would the effect survive when a substantial portion of the elongated shape is actually occluded? In this case, the shape axis still exists as a subjective global construct, provided the shape is perceptually completed behind the occluder. However, the effect of local alignment of texture with shape edges would not be predicted to have the same influence as in the unoccluded shape.

Conclusion

To conclude, the results presented here, in both shape segmentation and orientation discrimination tasks, converge on an axial organization of texture. The "axis effect" raises many questions that warrant further investigation. As well as furthering our understanding of the function of skeletal representations in visual perception, our results also add to the weight of evidence for global-to-local influences and recurrent interactions in vision in general and specifically in the realm of texture perception.

Acknowledgments

Corresponding author: Sarah Harrison.
Email: sharrison@sunyopt.edu.
Address: SUNY State College of Optometry, Vision Sciences, 33 West 42nd Street, New York, NY 10036, USA.

References

Adluru, N., Latecki, L. J., Lakaemper, R., Young, T., Bai, X., & Gross, A. (2007). Contour grouping based on local symmetry. In *11th IEEE International Conference on Computer Vision (Rio de Janeiro, Brazil)* (pp. 1–8). Institute of Electrical and Electronics Engineers (IEEE).

Appelle, S. (1972). Perception and discrimination as a function of stimulus orientation: The "oblique effect" in man and animals. *Psychological Bulletin, 78,* 266–278. [PubMed]

Ariely, D. (2001). Seeing sets: Representation by statistical properties. *Psychological Science, 12,* 157–162. [PubMed]

Ball, P. (1999). *The self-made tapestry: Pattern formation in nature*. Oxford, UK: Oxford University Press.

Ben-Shahar, O. (2006). Visual saliency and texture segregation without feature gradient. *Proceedings of the National Academy of Sciences of the United States of America, 103,* 15704–15709. [PubMed] [Article]

Ben-Shahar, O., & Zucker, S. W. (2004). Sensitivity to curvatures in orientation-based texture segmentation. *Vision Research, 44,* 257–277. [PubMed]

Bergen, J. R., & Adelson, E. H. (1988). Early vision and texture perception. *Nature, 333,* 363–364. [PubMed]

Blum, H. (1967). A transformation for extracting new descriptors of shape. In W. Wathen-Dunn (Ed.), *Models for the perception of speech and visual form* (pp. 362–380). Cambridge, MA: MIT Press.

Blum, H. (1973). Biological shape and visual science. I. *Journal of Theoretical Biology, 38,* 205–287. [PubMed]

Brady, M., & Asada, H. (1984). Smoothed local symmetries and their implementation. *International Journal of Robotics Research, 3,* 36–61.

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

Burbeck, C. A., & Pizer, S. M. (1995). Object representation by cores: Identifying and representing primitive spatial regions. *Vision Research, 35,* 1917–1930. [PubMed]

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

De Valois, R. L., Albrecht, D. G., & Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque

visual cortex. *Vision Research, 22,* 545–559. [PubMed]

De Valois, R. L., Yund, E. W., & Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research, 22,* 531–544. [PubMed]

Enns, J. (1986). Seeing textons in context. *Perception & Psychophysics, 39,* 143–147. [PubMed]

Eriksen, C. W., & St James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics, 40,* 225–240. [PubMed]

Feldman, J., & Singh, M. (2006). Bayesian estimation of the shape skeleton. *Proceedings of the National Academy of Sciences of the United States of America, 103,* 18014–18019. [PubMed] [Article]

Graham, N. (1994). Non-linearities in texture segregation. *Ciba Foundation Symposium, 184,* 309–322. [PubMed]

He, Z. J., & Nakayama, K. (1994). Perceiving textures: Beyond filtering. *Vision Research, 34,* 151–162. [PubMed]

Hegde, J. (2008). Time course of visual perception: Coarse-to-fine processing and beyond. *Progress in Neurobiology, 84,* 405–439. [PubMed]

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology, 160,* 106–154. [PubMed]

Julesz, B., & Bergen, J. R. (1983). Textons, the fundamental elements in preattentive vision and perception of textures. *Bell System Technical Journal, 62,* 1619–1645.

Kass, M., & Witkin, A. (1985). Analyzing oriented patterns. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence* (pp. 944–952). Los Angeles, CA.

Katz, R. A., & Pizer, S. M. (2003). Untangling the Blum medial axis transform. *International Journal of Computer Vision, 55,* 139–153.

Kimia, B. B. (2003). On the role of medial geometry in human vision. *The Journal of Physiology, 97,* 155–190. [PubMed]

Kovacs, I., Feher, A., & Julesz, B. (1998). Medial-point description of shape: A representation for action coding and its psychophysical correlates. *Vision Research, 38,* 2323–2333. [PubMed]

Kovacs, I., & Julesz, B. (1993). A closed curve is much more than an incomplete one: Effect of closure in figure-ground segmentation. *Proceedings of the National Academy of Sciences of the United States of America, 90,* 7495–7497. [PubMed] [Article]

Kovacs, I., & Julesz, B. (1994). Perceptual sensitivity maps within globally defined visual shapes. *Nature, 370,* 644–646. [PubMed]

Lamme, V. A. (1995). The neurophysiology of figure-ground segregation in primary visual cortex. *Journal of Neuroscience, 15,* 1605–1615. [PubMed] [Article]

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

Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 20,* 1434–1448. [PubMed]

Lee, T. S., Mumford, D., Romero, R., & Lamme, V. A. (1998). The role of the primary visual cortex in higher level vision. *Vision Research, 38,* 2429–2454. [PubMed]

Livne, T., & Sagi, D. (2007). Configuration influence on crowding. *Journal of Vision, 7*(2):4, 1–12, http://journalofvision.org/7/2/4/, doi:10.1167/7.2.4. [PubMed] [Article]

Malik, J., & Perona, P. (1990). Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 7,* 923–932. [PubMed]

Morgan, M. J., & Baldassi, S. (1997). How the human visual system encodes the orientation of a texture, and why it makes mistakes. *Current Biology, 7,* 999–1002. [PubMed]

Morgan, M. J., Mason, A. J., & Baldassi, S. (2000). Are there separate first-order and second-order mechanisms for orientation discrimination? *Vision Research, 40,* 1751–1763. [PubMed]

Motoyoshi, I., & Nishida, S. (2001). Visual response saturation to orientation contrast in the perception of texture boundary. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 18,* 2209–2219. [PubMed]

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology, 9,* 353–383.

Nothdurft, H. C. (1985a). Orientation sensitivity and texture segmentation in patterns with different line orientation. *Vision Research, 25,* 551–560. [PubMed]

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

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

Nothdurft, H. C. (1992). Feature analysis and the role of similarity in preattentive vision. *Perception & Psychophysics, 52,* 355–375. [PubMed]

Nothdurft, H. C. (1994). Common properties of visual segmentation. *Ciba Foundation Symposium, 184,* 245–259. [PubMed]

Orban, G. A., Vandenbussche, E., & Vogels, R. (1984). Human orientation discrimination tested with long stimuli. *Vision Research, 24,* 121–128. [PubMed]

Palmer, S. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology, 9,* 441–474.

Parkes, L., Lund, J., Angelucci, A., Solomon, J. A., & Morgan, M. (2001). Compulsory averaging of crowded orientation signals in human vision. *Nature Neuroscience, 4,* 739–744. [PubMed]

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

Pomerantz, J. R., Sager, L. C., & Stoever, R. J. (1977). Perception of wholes and of their component parts: Some configural superiority effects. *Journal of Experimental Psychology: Human Perception and Performance, 3,* 422–435. [PubMed]

Regan, D. (1995). Orientation discrimination for bars defined by orientation texture. *Perception, 24,* 1131–1138. [PubMed]

Rogers-Ramachandran, D. C., & Ramachandran, V. S. (1998). Psychophysical evidence for boundary and surface systems in human vision. *Vision Research, 38,* 71–77. [PubMed]

Sawada, T., & Pizlo, Z. (2008). Detection of skewed symmetry. *Journal of Vision, 8*(5):14, 1–18, http://journalofvision.org/8/5/14/, doi:10.1167/8.5.14. [PubMed] [Article]

Siddiqi, K., Shokoufandeh, A., Dickinson, S., & Zucker, S. (1999). Shock graphs and shape matching. *International Journal of Computer Vision, 30,* 1–24.

Thompson, D. W. (1917). *On growth and form.* Cambridge, UK: Cambridge University Press.

Wagemans, J. (1992). Perceptual use of nonaccidental properties. *Canadian Journal of Psychology, 46,* 236–279. [PubMed]

Wagemans, J. (1993). Skewed symmetry: A nonaccidental property used to perceive visual forms. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 364–380. [PubMed]

Wagemans, J., Van Gool, L., & d'Ydewalle, G. (1991). Detection of symmetry in tachistoscopically presented dot patterns: Effects of multiple axes and skewing. *Perception & Psychophysics, 50,* 413–427. [PubMed]

Wagemans, J., Van Gool, L., & d'Ydewalle, G. (1992). Orientational effects and component processes in symmetry detection. *Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 44,* 475–508.

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

Wolfson, S. S., & Landy, M. S. (1999). Long range interactions between oriented texture elements. *Vision Research, 39,* 933–945. [PubMed]

Zhaoping, L. (2003). V1 mechanisms and some figure-ground and border effects. *The Journal of Physiology, 97,* 503–515.

Zhu, S. C. (1999). Stochastic jump-diffusion processes for computing medial axes. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 21,* 1158–1169.