



Discrimination of Orientation-defined Texture Edges

S. SABINA WOLFSON,* MICHAEL S. LANDY*

Received 30 June 1994; in revised form 31 August 1994

Preattentive texture segregation was examined using textures composed of randomly placed, oriented line segments. A difference in texture element orientation produced an illusory, or orientation-defined, texture edge. Subjects discriminated between two textures, one with a straight texture edge and one with a “wavy” texture edge. Across conditions the orientation of the texture elements and the orientation of the texture edge varied. Although the orientation difference across the texture edge (the “texture gradient”) is an important determinant of texture segregation performance, it is not the only one. Evidence from several experiments suggests that configural effects are also important. That is, orientation-defined texture edges are strongest when the texture elements (on one side of the edge) are parallel to the edge. This result is not consistent with a number of texture segregation models including feature- and filter-based models. One possible explanation is that the second-order channel used to detect a texture edge of a particular orientation gives greater weight to first-order input channels of that same orientation.

Texture Orientation

INTRODUCTION

The visual perception of texture plays a role in depth perception, distinguishing figure from ground, estimating surface orientation, defining the shape of objects, and other everyday visual tasks. Although the visual world is rarely, if ever, defined solely in terms of texture, for experimental purposes one often isolates a single cue, such as texture, to gauge its independent effects and efficacy. Texture segregation may be based on any number of image cues: size, orientation, and so on (Julesz, 1981). In this paper we examine instantaneous texture segregation of textures composed solely of oriented line segments (also called texture elements or *texels*). Such textures have been studied extensively (e.g. Bergen & Julesz, 1983; Bergen & Landy, 1991; Grossberg & Mingolla, 1985; Landy & Bergen, 1991; Malik & Perona, 1990; Nothdurft, 1985a, b, 1992; Sagi & Julesz, 1985; Wolfe, 1992; for an excellent review of texture segregation, not specifically focusing on oriented line segment textures, see Bergen, 1991).

Instantaneous texture segregation (also called effortless or preattentive texture segregation or segmentation) is said to occur when two different regions are “instantaneously” seen in an image composed of two differently textured areas (as opposed to regions which can only be distinguished with scrutiny). An example of this can be

seen in Fig. 1(A), where the texture on the left and the texture on the right are effortlessly segregated. In our experiments, subjects must see the regions as being different, and they must see the boundary formed by the two textures. In Fig. 1(A), one can instantly see that the two textured regions abut to form a vertical “edge”. Of course, not all textures instantly segregate when placed next to one another, nor do all textures which segregate form a clear, distinct boundary. Numerous theories and models have been proposed to predict how well textures segregate, but no theory accounts for all the results in the field. However, both feature- and filter-based models of texture segregation suggest that performance with textures like those illustrated in Fig. 1 is a function of line segment orientation difference across the texture border.

Feature-based models are those based on specific texture element attributes. For example, Julesz (1981) suggested that at least two classes of “textons” are important: terminators and elongated “blobs”. The blobs are characterized by their length, orientation, width and other properties. Beck (1982) suggested that color, brightness, size, and the slopes of component lines are important. Nothdurft (1985b) suggested that line segment spacing, length, and orientation difference are important. The feature-based models suggest that when the gradient of any such textural property (the “texture gradient” across space) is sufficiently large then a texture edge will be perceived.

Filter-based models have also been suggested to account for texture segregation results. These models generally have three stages: (1) convolution with

*Psychology Department and Center for Neural Science, New York University, 6 Washington Place, 8th Floor, New York, NY 10003, U.S.A. [Email sabina@cns.nyu.edu].




FIGURE 1. Sample stimuli. Notice that the edge in (A) appears more salient than that in (B) even though both edges are signaled by a 90° orientation difference across the edge. The subject's task in all experiments was to discriminate a straight edge (A) from a "wavy" edge (C). In the experiments, stimuli consisted of blurred white line segments on a gray background.

orientationally-tuned bandpass linear spatial filters (similar to processing carried out by simple cells in cortical area V1); (2) some nonlinearity such as computation of texture energy, possibly followed by spatial pooling (similar to processing carried out by cortical complex cells); and (3) segmentation based on the texture gradient. Bergen and Adelson (1988) used isotropic rather than orientation-tuned linear filters followed by

full-wave rectification. Malik and Perona (1990) based the segmentation stage on the maximum of the responses from a set of odd-symmetric filters. Landy and Bergen (1991) and Bergen and Landy (1991) used the Sobel edge detector to segment the textures. Caelli (1985) used correlation and grouping across different second-stage channel outputs to segment the images. Bovik, Clark and Geisler (1990) based segregation on differences in texture channel amplitude or phase. Other similar models include those of Sagi (Fogel & Sagi, 1989; Rubenstein & Sagi, 1990), Graham, Beck and Sutter (1992; Sutter, Beck & Graham, 1989; Graham, Sutter & Venkatesan, 1993), and Knutsson and Granlund (1983); similar models have also been used in other spatial vision tasks (e.g. Morrone & Burr, 1988; Morrone & Owens, 1987). The popularity of such models led Chubb and Landy (1991) to refer to them as the "back pocket models" of texture segregation, as researchers routinely pull these models out of their back pockets to explain new results in texture segregation.

Both feature- and filter-based models suggest that peak performance should result from an orientation difference (across a texture border) of 90°, regardless of the particular orientations on either side of the edge. However, Olson and Attneave (1970) and Nothdurft (1992) have noted that performance improves when texture elements align with the edge being formed. Olson and Attneave (1970), using arrays of texture elements, examined the effectiveness of configural effects in producing similarity groupings. They measured the reaction time for locating a target region (in the texture array), and found that differences in the slopes of the elements enhanced grouping. In addition, they found that the target region was located "more quickly when its boundaries were in a parallel-perpendicular relationship to elements ... than when the boundaries and elements were in a 45° relationship". Nothdurft (1992), using arrays of line segments at various orientations, examined the effects of *differences between regions* and *similarity within regions*. He found that texture segmentation occurs when the orientation contrast between regions is greater than the orientation contrast within regions. In addition, he notes that there is "some additional contribution [to texture segmentation] from alignment effects between line elements and the texture border", although "[how] the alignment of texture elements with [the target region] can account for the better visibility ... remains an open question". In both of these studies, the effect of alignment was not a primary focus of the investigation and was not systematically explored. In this paper, we investigate this effect by determining the aspects of orientation-defined texture edges which characterize good performance in an edge-discrimination task. For an example of this alignment effect, consider the two edges shown in Fig. 1(A, B). In both cases, there is a 90° orientation difference across the edge, but the edge involving the vertical and horizontal line segments appears more salient.

Why is the texture-defined border in Fig. 1(A) more salient than that in Fig. 1(B)? One possible explanation

would invoke the oblique effect. In some threshold detection tasks, oblique patterns are more difficult to detect than vertically or horizontally oriented patterns (Lennie, 1974). The pattern in Fig. 1(B) may yield a less salient edge due to an oblique effect for the orientation of the line segments which define the texture. Alternatively, good performance may relate to the configuration of the texture elements relative to the edge they define (see, e.g. Olson & Attneave, 1970; Nothdurft, 1992; Field, Hayes & Hess, 1993; Palmer & Bucher, 1982, on similar effects of configuration). Thus, any of the following might contribute to good texture segregation performance using textures composed of oriented line segments:

- orthogonal texels on either side of the texture-defined edge (due to a largest-possible texture gradient);
- vertical texels on one side of the texture-defined edge (due to an oblique effect);
- horizontal texels on one side of the texture-defined edge (due to an oblique effect);
- texels parallel to the texture-defined edge on one side of the edge (due to a configural effect);
- texels perpendicular to the texture-defined edge on one side of the edge (due to a configural effect).

In this paper, these possible determinants are compared by examining texture segregation performance across sets of stimuli for which the above mentioned theories provide different predictions. This is done by varying texel orientation on either side of a texture edge and by varying the orientation of the edge itself. The results indicate that, while the orientation difference across the edge is important, there is a strong configural effect. Specifically, texture elements parallel to the texture-defined edge lead to improved segregation performance. We attempt to account for these results using a number of recent models of texture segregation and variants of them. One model that accounts for our

results is a back pocket model (a filter-based model) with the added constraint that when a second-stage channel has a preference for texture edges of a given orientation, it gives greater weight to first-stage channels with peak sensitivity at that same orientation.

METHOD

Subjects

There were three subjects in each experiment. One was the first author. The other two were naive to the purposes of the experiment. All subjects had normal or corrected-to-normal vision.

Stimuli

The images consisted of randomly positioned texels. Figure 1 displays example stimuli, although in the experiments stimuli were white on a gray background. The texels were blurred line segments of length 0.3 deg and density 7.2 texels/deg², blurred with a Gaussian ($\sigma = 0.03$ deg).^{*} The display was viewed from a distance of 1 m, resulting in 33 pixels/deg. Each stimulus contained a texture edge. On each side of the edge all texels shared a common orientation. The edge was either *straight* [Fig. 1(A)] or *wavy* [Fig. 1(C)]. The wavy edge was sinusoidal in form (1 c/12.2 deg, 0.46 deg amplitude). We were concerned with configural effects and thus required a discrimination task which did not substantially vary the orientation of the edge between the two discriminanda. The parameters of the wavy edge imply that, where the wavy edge differs most in orientation from the straight edge (i.e. at the wavy edge's zero crossing), its orientation is 13.3° rotated from the straight edge, which is reasonably small, relative to the orientation shifts required for the effects we are investigating. At the same time, the amplitude of the wavy edge and texel density were chosen so as to provide some dynamic range in the results (away from performance ceiling and floor effects). We were also concerned with observer strategies involving scrutinizing particular parts of the display. To foil such strategies, the phase of the sine wave of the wavy edge was chosen randomly, and edge location was randomized as well (a shift from the center of the texture of as much as ± 2.3 deg). Background luminance was 26 cd/m². Texels had a nominal 100% white-on-gray Weber contrast, resulting in an incremental luminance of approx. 1.9×10^{-4} cd/texel.

To generate the stimuli, a fixed number of texel spatial positions were chosen at random. The texel position relative to the edge determined texel orientation. Then, a texel was painted centered on the given position. We were concerned that line segments could intersect and that this would be a cue to the task which varied with stimulus parameters. For large orientation differences across the edge, one might expect more such texel intersections at the edge itself. Thus, the stimulus generation procedure was modified to ensure a minimum distance was maintained between texture elements to eliminate any crossings.[†]

^{*}To avoid confusion, "deg" is used to denote measurements of stimulus size and "°" for angles within the stimulus such as edge and line segment orientation.

[†]This modification may provide another cue, density at the texture-defined edge, which varies with texture element orientation (with respect to the edge). Recall that the texture elements are randomly placed; a point is randomly selected and that point serves as the center of the element. Thus, in the case of a horizontal texture element (and a vertical edge), if the element is within 0.15 deg (half the length of a texture element) of the edge, it crosses the edge and might intersect with a differently oriented texel (from the other side of the edge). However, in the case of a vertical texel, the texel can abut the edge without crossing it (although a texel on the other side of the edge might cross the edge and intersect the vertical texel). In reality, few texels intersect at the edge and have to be deleted. However, we were still concerned that this might provide a useful cue, so we ran an additional experiment in which all texture elements were kept a minimum distance from the edge. Specifically, if the texel crossed the edge it was deleted even if it didn't intersect with another texel. The results for this experiment are not included here since they are not revealing. This rule produces a noticeable gap at the edge (when the texture elements are not parallel to the edge), and performance improves dramatically (over 95% correct in all but the most difficult conditions).

TABLE 1. Summary of the experiments

Experiment	Edge orientation	Left-hand texel orientation	Right-hand texel orientation	Orientation difference	Results
1	Vertical	Horizontal	Varied	Varied	Figs 2, 12(A)
2	Vertical	Vertical	Varied	Varied	Figs 3, 12(B)
3	Vertical	Varied	Varied	90	Figs 4, 12(C)
4	Vertical	Left oblique	Varied	Varied	Figs 5, 12(D)
5	Right oblique	Left oblique	Varied	Varied	Figs 6, 12(E)
6	Right oblique	Right oblique	Varied	Varied	Figs 7, 12(F)
7	Right oblique	Varied	Varied	90	Figs 8, 12(G)
8	Right oblique	Vertical	Varied	Varied	Figs 9, 12(H)
9	Right oblique	Horizontal	Varied	Varied	Figs 10, 12(I)
10	Varied	Varied	Varied	90	Figs 11, 12(J)

Finally, we took care not to provide any external configural cues which might affect performance, such as visible vertical and horizontal environmental features which might frame the stimuli. We took three precautions: (1) all images were circular (12.2 deg dia); (2) the edges of the stimuli were smoothed so the textures appeared to gradually (over 0.3 deg) fade at the edges of the circular window; and (3) a large, circular black board (with a circular hole in the center) was placed directly in front of the monitor, and a floor-to-ceiling black curtain was placed behind the monitor. Since the images were viewed in a dark room, the edges between the black board and curtain were not visible, resulting in a featureless dark surround.

All displays were generated prior to the experimental sessions, and each image was unique. The images were computed using the HIPS image processing software (Landy, Cohen & Sperling, 1984a, b). Images were presented in random order (within an experiment) on a Barco Calibrator color monitor. The lookup tables were set so that the relationship between pixel value and display luminance was linear.

Procedure

The task was form discrimination (straight vs wavy edge) using a two-interval two-alternative forced-choice procedure and method of constant stimuli. In each block of trials there were 50 trials for each of seven conditions, resulting in 350 trials per block (except for Expts 3, 7 and 10, which had 50 trials for each of eight conditions). Each dataset represents at least two blocks of trials per subject. Each trial consisted of a 750 msec cue followed by a 250 msec blank; a 250 msec stimulus interval; a 250 msec blank; and then a second 250 msec stimulus interval. The screen remained blank (at the same mean luminance) after the second interval until the subject responded. The subject's response initiated

the subsequent trial. The straight edge stimulus was in either the first or second interval (chosen randomly), and the subject's task was to identify the interval containing the straight edge stimulus by pressing a response key. Auditory feedback was provided after each trial.

RESULTS

In all, 10 experiments were carried out and are summarized in Table 1. The results for each experiment are shown in Figs 2–11 (one experiment per figure, one subject per panel), and a summary of the results, averaged across subjects, is shown in Fig. 12 (one experiment per panel). Since the averaged data, on the whole, accurately reflect the individual subjects' data, we generally refer to the averaged graphs in this section. All experiments shared the same methods and procedure, and varied simply in the one-dimensional slice taken through the three-dimensional space of possible stimuli (determined by the orientation of the texture edge and the orientations of the texels on either side of the edge). We consider each experiment in turn.

The results for Expt 1 are shown in Fig. 12(A). The circular icon in the upper-right-hand corner of this graph, and all subsequent graphs, represents the experimental setup. For this experiment the icon indicates that:

- the texture edge was vertical (i.e. $\theta_{\text{edge}} = 0^\circ$);*
- the texels on the left side of the edge were horizontal ($\theta_{\text{left}} = 90^\circ$); and
- the texels on the right side of the edge varied across trials, i.e. θ_{right} was pseudo-randomly set to 0° , 22.5° , 45° , 67.5° , 112.5° , 135° , or 157.5° across trials, although it was the same for both intervals within a single trial.

The error bars indicate ± 2 SEMs.† The important data points in each graph are marked with symbols directly above the abscissa (see Table 2 for a key to the symbols). Each condition in Fig. 12(A) is marked with a symbol indicating that the texels (on the left side of the edge) are perpendicular to the edge. In addition, at a θ_{right} of 0°

*Angles are measured clockwise from vertical.

†The error bars indicate ± 2 SEMs in each direction. For interpreting the results, the rule of thumb for the error bars and significance is (without considering corrections for multiple tests): if the error bars do not overlap, then the difference is highly significant ($P < 0.005$). However, even when the error bars overlap considerably (the points being 2.8 SEMs apart), the difference is still significant ($P < 0.05$).




FIGURE 2. Results from Expt 1 for three subjects. In this and all subsequent figures a circular icon is used to summarize the experimental conditions. Here, the icon indicates that the texture edge was vertical ($\theta_{\text{edge}} = 0^\circ$), texture elements to the left of the edge were horizontal ($\theta_{\text{left}} = 90^\circ$), and the orientation of texture elements to the right of the edge (θ_{right}) was varied over trials. For all subjects, best performance results from a 90° orientation difference across the edge, as expected. Error bars indicate ± 2 SEMs. (The lower half of the error bar at 112.5° for RGP was removed for clarity.)




FIGURE 3. Results from Expt 2. Again, best performance results from a 90° orientation difference across the edge.




FIGURE 4. Results from Expt 3. These results confirm the observation that the edge in Fig. 1(A) is more salient than that of Fig. 1(B). For a constant 90° orientation difference across the edge, best performance results when texels are parallel and perpendicular to the edge.




FIGURE 5. Results from Expt 4. For a vertical edge and left oblique texels to the left of the edge, best performance does *not* result from right-oblique texels to the right of the edge (i.e. from a 90° orientation difference across the edge). For all three subjects, best performance is obtained with texels parallel to the edge, and for two subjects (SSW and JCE) good performance is also obtained with texels perpendicular to the edge.




FIGURE 6. Results from Expt 5. This is effectively Expt 1 with stimuli rotated clockwise 45°. Again, an orientation difference of (approx.) 90° results in best performance (cf. Fig. 2). The results indicate that there is not an oblique effect for the texels.




FIGURE 7. Results from Expt 6. This is effectively Expt 2 with stimuli rotated clockwise 45°. Again, an orientation difference of (approx.) 90° results in best performance (cf. Fig. 3).




FIGURE 8. Results from Expt 7. This is effectively Expt 3 with stimuli rotated clockwise 45° . The data are unclear but show hints of better performance with texels parallel and perpendicular to the edge (cf. Fig. 4). The data are clearer when averaged across subjects [see Fig. 12(G)].




FIGURE 9. Results from Expt 8. This is effectively Expt 4 with stimuli rotated clockwise 45° . Two subjects (SSW and JCE) show a slight peak at 67.5° suggesting a combined effect of orientation difference across the edge and texels parallel to the edge (cf. Fig. 5).




FIGURE 10. Results from Expt 9. This is similar to Expt 8. Two subjects (SSW and RGP) show a slight peak at 45° suggesting a configural effect for texels parallel to the edge.




FIGURE 11. Results from Expt 10. In this experiment the texels are kept parallel and perpendicular to the edge while the orientation of the edge varies. These results suggest a slight oblique effect for the texture edge. The data are clearer when averaged across subjects [see Fig. 12(J)].

(180°) there are two extra symbols indicating that, in this condition, the texels are perpendicular across the edge, and the texels are parallel to the edge. The results show a peak in performance around a θ_{right} of 0°. This is consistent with most of the suggested determinants of good performance: maximal orientation difference (90°) across the edge, vertical texels, and texels that are parallel to the edge. Similarly, the results shown in Fig. 12(B) (Expt 2) show best performance around a θ_{right} of 90° which is consistent with maximum orientation difference (90°) across the edge, horizontal texels, and texels which are perpendicular to the edge being determinants of good performance.

The results shown in Fig. 12(C) (Expt 3) confirm our observation that the edge in Fig. 1(A) appears more salient than the edge in Fig. 1(B), even though there is a 90° orientation difference across the edge in both cases. In this experiment θ_{left} and θ_{right} covaried so as to preserve a 90° orientation difference across the edge. The results agree with our observation, showing a clear performance peak at 0° and 90°. These results are *not* consistent with theories based solely on orientation difference across the edge. Such theories would predict constant performance in each condition shown in Fig. 12(C). Thus, a determinant of good performance must be involved other than simply orientation difference across the edge. However, the results of this experiment are consistent with both an oblique effect for the texels and a configural effect (for texels parallel and perpendicular to the edge).

Figure 12(D) (Expt 4) shows results which verify that orientation difference across the edge is not the sole determinant of good performance. In this experiment, θ_{left} was held constant at 135° while θ_{right} was varied. Best performance did *not* result from a 90° orientation difference across the edge. Subjects performed best for a θ_{right} at or near 0°. That is, best performance results either because the texels on the right-hand side of the edge are vertical or because they are oriented parallel to the edge. There is also a weak peak around a $\theta_{\text{right}} = 90°$. That is, good performance also may result either from horizontal texels or from texels which are perpendicular to the edge. In any case, these results are inconsistent with theories

based solely on orientation difference across the edge. Such theories would predict peak performance at the maximum orientation difference, $\theta_{\text{right}} = 45°$, which lies between the two performance peaks. However, these results do not distinguish between a configural effect (which favors texels that are parallel and/or perpendicular to the edge) and an oblique effect (which favors vertical and horizontal texels).

To distinguish between configural and oblique effects, we repeated the above experiments with the stimuli rotated 45°. The results for the experiments with an oblique edge [Fig. 12(E–I)] are noisier and the effects weaker than in the corresponding vertical edge experiments [Fig. 12(A–D)], but some trends are suggested by the data.

The results shown in Fig. 12(E–F) (Expts 5 and 6) are analogous to those shown in Fig. 12(A–B) (Expts 1 and 2) respectively. For both Expts 5 and 6, best performance resulted from an orientation difference of (approx.) 90° across the edge. This is consistent both with an orientation difference effect (as best response resulted from the maximum orientation difference) and with configural effects (as best response resulted from texels oriented perpendicular and parallel to the edge). However, best performance did *not* result from vertical and

TABLE 2. Summary of graph symbols

Symbol	Meaning
—	Vertical edge, perpendicular texels
⋮—	Vertical edge, texels perpendicular to edge
⋮	Vertical edge, texels parallel to edge
<—	Oblique edge, perpendicular texels
<⋮	Oblique edge, texels perpendicular to edge
<	Oblique edge, texels parallel to edge




FIGURE 12. Results averaged across subjects (see Figs 2-11 for the individual subject data). The circular icon in each panel summarizes the experimental conditions (see text for details). Error bars indicate ± 2 SEMs. (A) Results from Expt 1. Best performance results from a 90° orientation difference across the edge. (B) Results from Expt 2. Again, best performance results from a 90° orientation difference across the edge. (C) Results from Expt 3. Best performance occurs when texels are parallel and perpendicular to the edge. (D) Results from Expt 4. Best performance occurs when texels are parallel to the edge, *not* when there is a 90° orientation difference across the edge. (E) Results from Expt 5. This is Expt 1 rotated 45°. Again, an orientation difference of (approx.) 90° results in best performance. (F) Results from Expt 6. This is Expt 2 rotated 45°. Again, an orientation difference of (approx.) 90° results in best performance. (G) Results from Expt 7. This is Expt 3 rotated 45°. The results suggest that performance is best when texels are parallel and perpendicular to the edge. (H) Results from Expt 8. This is Expt 4 rotated 45°. The results hint that best performance is a combined effect of orientation difference across the edge and texels parallel to the edge. (I) Results from Expt 9. This is similar to Expt 8. The results hint that performance is best when texels are (approximately) parallel to the edge. (J) Results from Expt 10. The results show an oblique effect for the edge.

horizontal texels, suggesting that an oblique effect for the texels is not operative. The results also appear to indicate an oblique effect for the texture edge orientation since the overall performance levels (in Expts 5 and 6) dropped from those in Expts 1 and 2. This performance drop (between the vertical edge experiment and corresponding oblique edge experiment) was seen for five of the six datasets, although this difference was significant for only two of the datasets. Experiment 10 further examines this performance drop.

Figure 12(G) (Expt 7) displays results from a rotated version of Expt 3. In this experiment θ_{left} and θ_{right} covaried so as to maintain a 90° orientation difference across the edge. The results suggest that performance is best when the texture elements are parallel and perpendicular to the edge, but this configural effect is clearer and stronger when the edge is vertical [Fig. 12(C)] rather than oblique [Fig. 12(G)].

Two versions of Expt 4 [Fig. 12(D)] were carried out with rotated stimuli. The results shown in Fig. 12(H) (Expt 8) do not show a distinct peak performance, although peak performance appears to be around a θ_{right} of 67.5°, suggesting a combined configural (for texels parallel to the edge) and orientation difference (orthogonal texels) effect. There was a slight dip in performance for horizontal texels giving evidence against an oblique effect for the texels and again showing that orientation difference is not the sole determinant of good performance. Similarly, the results shown in Fig. 12(I) (Expt 9) do not show a distinct peak performance, although peak performance appears to be around a θ_{right} of 45°, suggesting a configural effect for texels parallel to the edge.

The previous experiments may indicate a small oblique effect for the orientation of the texture edge being formed, since performance generally drops when the edge is oblique. In the final experiment [Expt 10, Fig. 12(J)], the texture elements were kept parallel and perpendicular to the edge while the edge orientation was varied. Performance peaks when θ_{edge} is vertical and horizontal. However, this effect is clearly small (see the individual subject data in Fig. 11).

DISCUSSION

Overall, the results indicate that there are multiple determinants of texture segregation performance. There is a configural effect for texels parallel and perpendicular to the texture edge, and texels parallel to the edge appear to be more effective than texels perpendicular to the edge. There is an effect of orientation difference across the edge, as expected. And, there is small oblique effect for the orientation of the texture-defined edge.

While orientation difference across a border is clearly important, theories based solely on orientation difference cannot account for our results. Specifically, as orientation difference across the edge is held constant, but the specific orientations on each side of the edge

vary, the results indicate that peak performance occurs when the texels are parallel and perpendicular to the edge [Fig. 12(C)]. Further, in some experiments segregation performance does not consistently improve as the orientation difference across the edge increases. Instead, peak performance occurs when texels are parallel (and to a small extent when they are perpendicular) to the edge [e.g. Fig. 12(D)]. Lastly, the results support a small oblique effect for the edge orientation [Fig. 12(J)], and they do not support an oblique effect for the texel orientation.

MODELING

This section concerns our attempts to model these results. We begin by summarizing salient aspects of the data that a model needs to account for. Then, energy models are discussed; such models fail to account for our results. Next, a model is described based on figural completion; this model also fails to account for our data. Finally, a modified energy model is described which accounts for our data, although it is not particularly compelling.

What does the model need to do?

To account for the data presented here, a model would need to show:

- the influence of the orientation difference across the edge {seen best in Expts 1 and 2 [Fig. 12(A, B)]}; and
- the configural effect of texels parallel and perpendicular to the edge {seen best in Expts 3 and 4 [Fig. 12(C, D)]}.

The influence of the orientation difference (Expts 1 and 2) across the edge [the structure gradient of Nothdurft (1985b)] is a consequence of almost any texture segregation model one can imagine. In feature-based models (e.g. Julesz, 1981), a 90° orientation difference is simply the largest featural difference achievable across this set of patterns. In filter-based models which use orientationally-tuned linear filters, a 90° orientation difference provides the greatest signal-to-noise ratio from the filters (e.g. Landy & Bergen, 1991). And, the influence of the orientation difference across the edge would result from any model based on discriminability of the individual texels and a signal-detection viewpoint [such as used to model visual search with similar texels (Pavel, Econopouly & Landy, 1992)].

On the other hand, the configural effect (Expts 3 and 4) is not a consequence of most models. The data suggest that texels parallel to the texture edge (and, to a lesser extent, those perpendicular to it) are more successful at edge formation, at least for the purposes of discriminating a straight from a wavy edge. We have considered several explanations of this effect, and have tested them by constructing models. All of the models described below, except for the last one, failed qualitatively to produce data similar to those in Expts 3 and 4.




FIGURE 13. The model begins by applying a Gaussian pyramid reduction. Next, oriented filters are used at six different orientations followed by a square-law nonlinearity or texture "energy" computation. This converts a textural difference to a difference in overall activation. Next, an edge enhancement filter is used to look for a straight vertical edge (as this model was only applied to the experiments involving a vertical edge) and the maximum response is chosen. Finally, the maximal responses from each orientation channel are combined. The model is applied to both stimuli in a trial (a straight and a wavy edge stimulus) and the stimulus resulting in the largest overall response is chosen. With the combination rule of equation (2) and a proper choice of weights, this model account for our data, but with equal weights [i.e. equation (1)], it does not.

The standard energy model

The standard energy model (Fig. 13) is a version of the back pocket model of texture segregation. It begins with the addition of noise to the input,* followed by a bank of oriented linear filters, and then an energy computation. As in Bergen and Landy (1991), rather than using a fairly large kernel for computations, we reduce the computational load by performing one level of Gaussian pyramid reduction [an operation which combines blur and subsampling (Burt & Adelson, 1983)] followed by orientationally-tuned filtering. Without the pyramid reduction a filter

kernel of double the size (in both dimensions) would have been required for the same effective spatial frequency preference. In the bulk of our simulations, six filters were used (spaced in orientation at intervals of 30°). Gabor filters were used with parameters chosen so as to result in half-height bandwidths of one octave in spatial frequency and 30° in orientation (see Appendix for details). Next, a pointwise nonlinearity is applied: each pixel is squared. This set of operations converts a difference in orientation into a difference in average response [i.e. an intensity difference, see Fig. 14(A)].

For a given trial, the output of these first two stages is a collection of images (each trial interval results in six images from each of the six oriented channels). Next, a simulated response is generated; the response is an answer to the question: "Which interval contained the straight edge?" The simulated response for each trial is generated by combining the information within each interval (across the six oriented channels) to produce a value indicating how consistent the outputs are with the presence of a straight vertical edge (we are only applying this model to Expts 1-4, all of which have vertical edges). To accomplish this we apply an edge enhancement filter† to each energy output [see Fig. 14(B)], and then find the best-fitting‡ straight, vertical edge [these latter stages of the model are task-specific, as were those of Bergen and Landy (1991)]. This results in a collection of 12 responses $R_{\theta,i}$, where θ ranges over the six oriented channels, and i represents the interval ($i = 1, 2$). The model then

*Without added noise, the model's performance is far more accurate than humans. In our simulations, to lower the model's performance, we added noise (independent, identically-distributed Gaussian distributed pixel noise) to the input stimuli before applying the model. The variance of the noise was raised until model performance was in the range of human performance. The signal-to-noise ratio was on the order of 1.0.

†The filter kernel used in the simulations presented here was 32×32 pixels in size. Each pixel on the left-hand side had a value of -1 , and each pixel on the right-hand side had a value of 1 . Other sizes were also explored (rows $\times$ columns): 16×16 , 32×8 , 32×16 , and 64×64 .

‡Specifically, to find the best-fitting straight, vertical edge, we sum the entries in each column, sum successive runs of four columns, and choose the largest absolute value of such a sum. That is, we find the best straight, vertical edge that is 4 pixels wide. Other widths of edge were also explored. Close to a width of 4, the width does not matter. But, a width of 1 yields noisier responses (since non-edges are more often chosen as the best-fitting edge) while a width of 8 blurs the difference between the straight and wavy edges.




FIGURE 14. Sample model outputs to the stimulus of Fig. 1(A). (A) Squared filter output for the vertically oriented channel. (B) Edge-enhanced output for the vertically oriented channel.

combines these responses using a Minkowski metric and responds that the straight edge appeared in interval 1 if:

$$\left(\sum_{\theta} R_{\theta,1}^{\beta} \right)^{1/\beta} > \left(\sum_{\theta} R_{\theta,2}^{\beta} \right)^{1/\beta}. \quad (1)$$

We have run this model and it fails to account for our data. It shows an effect of orientation difference [like Expts 1 and 2, Fig. 12(A, B)] as predicted. However, it shows no configural effect [unlike Expts 3 and 4, Fig. 12(C, D)], performing best at an orientation difference of 90° regardless of the texels' orientation relative to the edge. We thought that this model might account for our

results. The intuition (which is the basis for line-element filter models of vernier acuity) is that if we make the oriented filters elongated, the vertical texture elements will be *localized* better than the nonvertical elements relative to a vertical texture edge. That is, imagine a vertical texel with a vertical receptive field located for optimal response to the texel. A large vertical displacement of the receptive field is required to reduce the response substantially, whereas only a small horizontal displacement produces the same reduction. Thus, the filter smears the edge little for vertical texels (so, the vertical edge will be "crispest" when texels are parallel to it). However, this intuition is not borne out in simulations with filters having a half-height orientation bandwidth of 30° (nor in simulations with a half-height orientation bandwidth of 22.5°).

Completion and illusory contours

When looking at the textures in our experiments, the texture-defined edge appears continuous. By referring to the effect as configural, one is tempted to relate it to such Gestalt notions as "good continuation", and to boundary formation with illusory contours. Thus, we reasoned that if texels first "grew" toward one another (along the direction of the texel's orientation) then vertically oriented texels would likely produce a crisper edge (and thus better edge discrimination). We implemented the edge completion neural network of Grossberg and Mingolla (1985) to test this intuition. The relevant aspect of this algorithm is the following: a strongly responding oriented unit (similar to units in our energy computation) will cause a nearby unit to become more active if the two units are approximately aligned with each other and with their respective orientation preferences. Thus, for example, two vertically aligned vertically oriented texels will cooperate to generate responses in regions between them (thus completing a partially specified vertical line). We explored various definitions of neighborhood and orientation similarity and failed to account for all our results. A wide range of model parameters were tried. For any parts of the model parameter space we explored, either very little completion occurred or so much completion occurred that the difference between straight and wavy edges was destroyed. Thus, where completion was effective, the task was made more difficult rather than easier.

The work on orientation and tangent fields by Zucker and colleagues (Parent & Zucker, 1985; Zucker, 1985) and in particular their "Type II Processes" appear to be highly relevant. Zucker is interested in the definition of orientation in images with a variety of orientations present at each location, as well as measuring curvature when local orientation changes across an image. While it seems unlikely that past versions of the model would account for our results, Zucker (personal communication) suggests that new versions of their model, which have yet to appear in print, will be applicable to this problem and may well account for our results.

A working model

We have constructed a model which does account for our results. We began with the energy mode (Fig. 13), and considered what modifications to it would be required to reproduce the data of Expts 1–4 [Fig. 12(A–D)]. The data suggest that, for whatever reason, vertically oriented texels result in better performance when the task involves a vertically oriented edge. Thus, we modified the model in the simplest way possible to include this effect: the combination rule of equation (1) was changed to give a greater weight to the vertical channel. The model responded that the straight edge was in the first interval if:

$$\left(\sum_{\theta} W_{\theta} R_{\theta,1}^{\beta} \right)^{1/\beta} > \left(\sum_{\theta} W_{\theta} R_{\theta,2}^{\beta} \right)^{1/\beta}, \quad (2)$$

where W_{θ} was set to one for $\theta \neq 0$, and W_0 was increased to give the vertically oriented input channel greater weight.

In the simulations we now report, W_0 was set to 9.0. The extra weight is given to the channel whose orientation corresponds to the texture-defined edge orientation, but since we have only modeled the experiments involving vertical texture edges, the extra weight is always given to the vertical channel. In a fully-defined model, texture edge channels would exist at a variety of orientations, and each would give greater weight to input channels at the same orientation.

Simulations were carried out for Expts 1–4 [Fig. 12(A–D)]. Figure 15(A–D) shows results for the simu-

lations corresponding to the human observer data shown in Fig. 12(A–D). The results for Expts 1 and 2 show that performance improves as the orientation difference across the edge increases for both the model and the subjects, as expected. The results for Expt 3 show that while the orientation difference across the edge is held constant at 90° , performance is best when the texture elements are parallel and perpendicular to the edge for both the model and the subjects. Lastly, the results for Expt 4 show that performance is best when the texture elements are parallel to the edge rather than when there is a maximum orientation difference across the edge for both the model and the subjects. These latter two results are not predicted by most current models of texture segregation.

Figure 16 shows results of model simulations for the same inputs (for Expt 3) using four different decision rules. When a standard decision rule is used, i.e. when all channels are weighted equally, the performance is inconsistent with that of human observers. This is true regardless of the value of the exponent β . The model shows no variation with the orientation of the texels relative to the texture edge. However, when the vertical channel is given extra weight (two different values of W_0 are shown) followed by the maximum rule (a large value of β), performance becomes consistent with our subjects' performance. The choice of decision rules is crucial: both the high value of β and the large value of W_0 were required to produce a qualitative fit to the data. This model does account for the data, but it is neither




FIGURE 15. Monte Carlo results for the model (described in A Working Model) for Expts 1–4. The human results for these experiments are shown in Fig. 12(A–D).




FIGURE 16. Monte Carlo results for the model for Expt 3 using a variety of decision rules to combine across orientation channels. —◇— A Minkowski metric with an exponent of 1. —□— A maximum rule. —×— A maximum rule which gives the vertical channel 4 times the weight of the other channels. —+— A maximum rule which gives the vertical channel 9 times the weight of the other channels. Note that the decision rule is critical in modeling the results of Expt 3 even qualitatively.

normative (since the added weight to the vertical channel effectively ignores useful data) nor compelling. It remains to be discovered why the visual system has chosen to make this particular compromise.

CONCLUSIONS

We have found at least three stimulus parameters which result in improved performance in a texture segregation task with orientation-defined texture edges:

- increase in orientation difference across the edge;
- orientation of texels parallel, and to some extent perpendicular, to the texture edge (a configural effect); and
- orientation of the texture-defined edge either vertically or horizontally (a small oblique effect for the edge).

An energy model can account for these results by giving extra weight to the oriented channel which is oriented similarly to the edge.

REFERENCES

- Beck, J. (1982). Texture segregation. In Beck, J. (Ed.), *Perceptual organization and representation*. Hillsdale, N.J.: Lawrence Erlbaum.
- Bergen, J. R. (1991). Theories of visual texture perception. In Regan, D. (Ed.), *Vision and visual dysfunction, 10B* (pp. 114–134). New York: Macmillan.
- Bergen, J. R. & Adelson, E. H. (1988). Early vision and texture perception. *Nature*, 333, 363–364.
- Bergen, J. R. & Julesz, B. (1983). Rapid discrimination of visual patterns. *IEEE Transactions on Systems, Man, and Cybernetics*, 13, 857–863.
- Bergen, J. R. & Landy, M. S. (1991). Computational modeling of visual texture segregation. In Landy, M. S. & Movshon, J. A. (Eds.), *Computational models of visual perception* (pp. 253–271). Cambridge, Mass.: MIT Press.
- Bovik, A. C., Clark, M. & Geisler, W. S. (1990). Multichannel texture analysis using localized spatial filters. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12, 55–73.
- Burt, P. J. & Adelson, E. H. (1983). The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31, 532–540.
- Caelli, T. (1985). Three processing characteristics of visual texture segmentation. *Spatial Vision*, 1, 19–30.
- Chubb, C. & Landy, M. S. (1991). Orthogonal distribution analysis: A new approach to the study of texture perception. In Landy, M. S. & Movshon, J. A. (Eds.), *Computational models of visual processing* (pp. 291–301). Cambridge, Mass.: MIT Press.
- Daugman, J. G. (1985). Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by two-dimensional visual cortical filters. *Journal of the Optical Society of America A*, 2, 1160–1169.
- DeValois, R. L., Albrecht, D. G. & Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research*, 22, 545–559.
- DeValois, R. L., Yund, E. W. & Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vision Research*, 22, 531–544.
- Field, D. J., Hayes, A. & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local “association field”. *Vision Research*, 33, 173–193.
- Fogel, I. & Sagi, D. (1989). Gabor filters as texture discriminators. *Biological Cybernetics*, 61, 103–113.
- Graham, N., Beck, J. & Sutter, A. (1992). Nonlinear processes in spatial-frequency channel models of perceived texture segregation: Effects of sign and amount of contrast. *Vision Research*, 32, 719–743.
- Graham, N., Sutter, A. & Venkatesan, C. (1993). Spatial-frequency and orientation-selectivity of simple and complex channels in region segregation. *Vision Research*, 33, 1893–1911.
- Grossberg, S. & Mingolla, E. (1985). Neural dynamics of perceptual grouping: Textures, boundaries, and emergent segmentations. *Perception & Psychophysics*, 38, 141–171.
- Julesz, B. (1981). Textons, the elements of texture perception, and their interactions. *Nature*, 290, 91–97.
- Knutsson, H. & Granlund, G. (1983). Texture analysis using two-dimensional quadrature filters. In *IEEE Computer Society Workshop on Computer Architecture for Pattern Analysis and Image Database Management* (pp. 206–213). Silver Spring, Md: IEEE Computer Society.
- Landy, M. S. & Bergen, J. R. (1991). Texture segregation and orientation gradient. *Vision Research*, 31, 679–691.
- Landy, M. S. & Wolfson, S. S. (1993). Discrimination of orientation-defined texture edges. *Investigative Ophthalmology & Visual Science*, 34, 819.
- Landy, M. S., Cohen, Y. & Sperling, G. (1984a). HIPS: A UNIX-based image processing system. *Computer Vision, Graphics and Image Processing*, 25, 331–347.
- Landy, M. S., Cohen, Y. & Sperling, G. (1984b). HIPS: Image processing under UNIX—Software and applications. *Behavior Research Methods Instruments and Computers*, 16, 199–216.
- Lennie, P. (1974). Head orientation and meridional variations in acuity. *Vision Research*, 14, 107–111.
- Malik, J. & Perona, P. (1990). Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America A*, 7, 923–932.
- Morrone, M. & Burr, D. (1988). Feature detection in human vision: A phase dependent energy model. *Proceedings of the Royal Society of London B*, 235, 221–245.
- Morrone, M. & Owens, R. (1987). Feature detection from local energy. *Pattern Recognition Letters*, 6, 303–313.
- Nothdurft, H. C. (1985a). Orientation sensitivity and texture segmentation in patterns with different line orientations. *Vision Research*, 25, 551–560.
- Nothdurft, H. C. (1985b). Sensitivity for structure gradient in texture discrimination tasks. *Vision Research*, 25, 1957–1968.
- Nothdurft, H. C. (1992). Feature analysis and the role of similarity in preattentive vision. *Perception & Psychophysics*, 52, 355–375.
- Olson, R. K. & Attneave, F. (1970). What variables produce similarity grouping? *American Journal of Psychology*, 83, 1–21.
- Palmer, S. E. & Bucher, N. M. (1982). Textural effects in perceived pointing of ambiguous triangles. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 693–708.

- Parent, P. & Zucker, S. W. (1985). Trace inference, curvature consistency, and curve detection. Technical Report CIM-86-3, Department of Electrical Engineering, Computer Vision and Robotics Laboratory, McGill University, Montréal, Canada.
- Pavel, M., Econopoulou, J. & Landy, M. S. (1992). Psychophysics of rapid visual search. *Investigative Ophthalmology & Visual Science*, 33, 1355.
- Rubenstein, B. S. & Sagi, D. (1990). Spatial variability as a limiting factor in texture-discrimination tasks: Implications for performance asymmetries. *Journal of the Optical Society of America A*, 7, 1632–1643.
- Sagi, D. & Julesz, B. (1985). "Where" and "what" in vision. *Science*, 228, 1217–1219.
- Sutter, A., Beck, J. & Graham, N. (1989). Contrast and spatial variables in texture segregation: Testing a simple spatial-frequency channels model. *Perception & Psychophysics*, 46, 312–332.
- Watson, A. B. (1983). Detection and recognition of simple spatial forms. In Braddick, O. J. & Sleigh, A. C. (Eds), *Physical and biological processing of images* (pp. 100–114). New York: Springer.
- Wolfe J. M. (1992). "Effortless" texture segmentation and "parallel" visual search are not the same thing. *Vision Research*, 32, 757–763.
- Wolfson, S. S. & Landy, M. S. (1994). Modeling the discrimination of orientation-defined texture edges. *Investigative Ophthalmology & Visual Science*, 35, 1667.
- Zucker, S. W. (1985). Early orientation selection: Tangent fields and the dimensionality of their support. *Computer Vision, Graphics, and Image Processing*, 32, 74–103.

Acknowledgements—The work described in this paper was supported in part by a grant from the National Eye Institute EY08266. We would like to thank John Econopoulou, Mark Young and two anonymous reviewers for their helpful comments; Larry Maloney for continuing encouragement and help; and Robert Picardi for technical assistance. We thank Rick Gurnsey for bringing the Olson and Attneave paper

to our attention. We thank our subjects, Russell Phillips and John Econopoulou. Portions of this work have been presented at the annual meeting of the Association for Research on Vision and Ophthalmology and published in abstract form (Landy & Wolfson, 1993; Wolfson & Landy, 1994).

APPENDIX

We use even Gabor functions as the filters in our model. A Gabor function is a cosine wave windowed by a two-dimensional Gaussian. For a filter located at the origin oriented vertically, the weighting function is

$$w(x, y) = \exp[-(x^2/a^2 + y^2/b^2)]\cos(2\pi fx), \quad (\text{A1})$$

which has a corresponding Fourier transform

$$W(u, v) = \exp[-\pi^2((u-f)^2a^2 + v^2b^2)] + \exp[-\pi^2((u+f)^2a^2 + v^2b^2)], \quad (\text{A2})$$

where f is frequency in c/deg, a and b determine the shape of the Gaussian window and thus control the spatial frequency and orientation bandwidth (see, e.g. Daugman, 1985; Watson, 1983). Other orientations are derived by rotation of the coordinates. To obtain a frequency bandwidth of 1 octave and an orientation bandwidth of 30° at half-height, we derive

$$a^2 = 9 \ln(2)/(\pi^2 f^2) \quad (\text{A3})$$

and

$$b^2 = \ln(2)/(\pi^2 f^2 \tan^2(\pi/12)). \quad (\text{A4})$$

The filters used in our simulations, unless noted otherwise, had an orientation bandwidth of 30° and spatial frequency bandwidth at half-height of 1 octave, which are biologically plausible values (DeValois, Albrecht & Thorell, 1982; DeValois, Yund & Hepler, 1982).