In: **Perception**, 2001.

The Perception of Texture on Folded Surfaces

MARY J. BRAVO [1], HANY FARID [2]

Do judgments of texture similarity reflect surface texture or image texture? To find out, we had observers view a rectangular surface that was folded into three panels much like a brochure. Each panel was textured with an oriented noise pattern and the observers' task was to determine which side panel matched the center panel in surface texture. Information about surface geometry was conveyed by binocular disparity and by the boundaries of the rectangular surface. We found that observers were often consistently wrong, selecting the texture that differed in the image and not on the surface. In sharp contrast, when observers judged the texture orientation on each panel individually, their judgments were accurate reflections of the surface texture. So even when observers can recover surface texture, their judgments of texture similarity may still be based on image texture.

Texture   Similarity Grouping   Perceptual Constancy

[1]Department of Psychology, Rutgers University, Camden NJ 08102.
[2]Department of Computer Science, Dartmouth College, Hanover NH 03755.

1 Introduction

In his laws of perceptual organization, Wertheimer included the tendency for similar stimulus elements to group together [22]. This law of similarity grouping has intrigued many vision scientists because the idea seems so simple, and yet a precise formulation has proven elusive.

The basic phenomenon is shown in Figure 1a. Each side of this stimulus is textured with a pattern of repeating lines. The lines differ in orientation, and observers report seeing them as forming two distinct groups separated by a perceptual boundary. In other words, the textures appear to segment. Early experiments [1, 15] showed that differences in orientation are especially potent for producing this effect. But textures that differ in other ways may also segment; for example, patterns of "+"s and "L"s segment even though they have the same component orientations.

Early attempts to explain texture segmentation were couched in terms of abstract features [10, 3, 20]. Thus, the segmentation of "+"s and "L"s was attributed to a difference in a line-crossing feature. However, a simpler explanation has since been proposed: the spatial frequency filters that are thought to be involved in early vision respond differently to these textures. This realization led to computational models of texture segmentation that were based on well-established models of threshold vision. According to these models, textures are analyzed by arrays of spatial frequency filters that are spatially local and selective for orientation. The outputs of these filter arrays are then analyzed by a second set of filters which detects spatial discontinuities [8, 6, 21, 5, 19, 11, 13].

These models signify an important advance in our understanding of texture segmentation because they are elegant and rigorous and yet still have great explanatory power. However, they cannot explain all aspects of texture segmentation. Rock et al. used cast shadows to change the retinal luminance, but not the perceived reflectance, of their texture stimuli [18]. Contrary to the predictions of filter-based models, they found that grouping was determined by surface reflectance. These authors

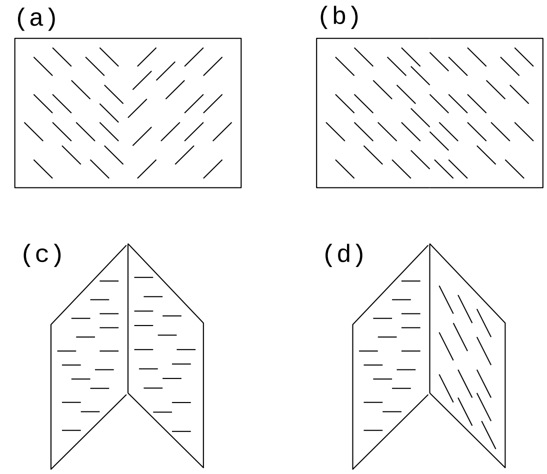

Figure 1: When a surface is folded, dissimilar surface textures (a) may appear similar in the image (b), and similar surface textures (c) may appear dissimilar in the image (d).

concluded that similarity grouping is not a product of early vision, but instead involves a surface representation. Similarly, He and Nakayama exploited the phenomenon of amodal surface completion to alter the perceived, but not the retinal, shape of their texture stimuli. Their stimulus manipulation should not have affected filter activity, and yet it had a significant effect on grouping ([9], see also, [16]). Finally, Palmer and Nelson used illusory contours to dissociate the retinal and the perceived structure of their stimuli, and they too found that grouping was determined by the perceived surface structure [17]. Thus, a number of researchers have provided evidence that filter-based models alone cannot account for human texture segmentation. They propose that at least some texture segmentation processes act on a surface representation in which such properties as reflectance and occlusion are made explicit.

Given the possibility that texture segmentation is based on a surface representation, it seems appropriate to reconsider even simple texture stimuli such as that shown in Figure 1a. What, exactly, are observers grouping in this display? Is grouping based on the retinal image textures, as the computational models suggest, or is it based on surface texture as recent findings suggest? With a fronto-parallel surface it is impossible to differentiate between these alternatives because the textures on the stimulus surface correspond closely to the textures in the retinal image. But the similarity of surface and image textures can be dissociated by folding this stimulus and viewing it obliquely. Dissimilar surface textures (Figure 1a) may then produce similar image textures (Figure 1c). And conversely, similar surface textures (Figure 1b) may produce dissimilar image textures (Figure 1d).

In this study, we used a folded stimulus to determine whether judgments of texture similarity are based on image texture or on surface texture. The stimuli were designed such that judgments based on image texture would produce qualitatively different from results judgments based on surface texture.

2 Methods

Our stimulus consisted of a computer generated surface that was folded into three panels much like a brochure, Figure 2(c). An oriented texture was mapped onto each panel and the observers' task was to tell us which of the side panels most closely resembled the center panel in surface texture. In two control conditions, observers performed the same task on modified versions of the stimulus. The basic stimulus and the experimental and control conditions are described below. First, however, we discuss our choice of method.

As reviewed in [4], researchers have used a variety of approaches to study texture segmentation. The most stringent method requires observers to make judgments about the shape of a texture defined region. We could not use this method here because, in our stimuli, the texture discontinuity necessarily coincides with highly salient depth discontinuity. Another approach to studying texture perception involves presenting observers with pairs of textures and asking them to judge the similarity (or conversely, the discriminability) of the textures. We are essentially using a two-alternative forced choice variant of this approach. Although this second measure does not directly assess grouping, it is thought to be closely correlated with grouping. That is, textures that are judged to be similar will group together, and textures that are judged to be dissimilar will segment [3]. For this relationship to hold, however, it is essential that the similarity judgments be based on an impression of the global texture pattern and not on scrutiny of the individual elements. To discourage such scrutiny we used continuous, stochastic textures in which local measurements of orientation were less reliable than global measurements. So although this experiment measures perceived similarity, we believe the results bear on similarity grouping as well.

2.1 Stimuli

The stimuli were generated in OpenGL on an SGI 02. Perspective projections of the right and left eye views of this stimulus were presented on alternating frames of the computer display. Observers viewed the stimuli with liquid crystal goggles (CrystalEyes by Stereographics) which synchronized the view of each eye with the alternating frames, yielding a stereoscopic display. A headrest was used to maintain the viewing distance of 70 cm.

The rendered stimulus consisted of three panels presented side-by-side, Figure 2. The two outer panels were 10cm squares, and the center panel was a 5cm wide and 10cm tall rectangle. Each panel was texture-mapped with white noise that had been convolved with an oriented band-pass filter. Because the filter was band-pass, contrast was spread across a small range of orientations. When considered globally, however, each texture had a clear dominant orientation. Fifteen random texture samples were generated for each of the texture orientations used in the experiment.

A -30, 0 or 30 degree texture was mapped onto the center panel of each stimulus. A different sample of the same texture orientation was mapped onto one of the side panels, Figure 2a. Note that the textures on these two panels were not continuous; there was always a seam between the panels. [3] The other side panel was mapped with a tex-

[3] This seam is a necessary artifact of our method of stimulus generation. One way to create seamless stimuli is to use

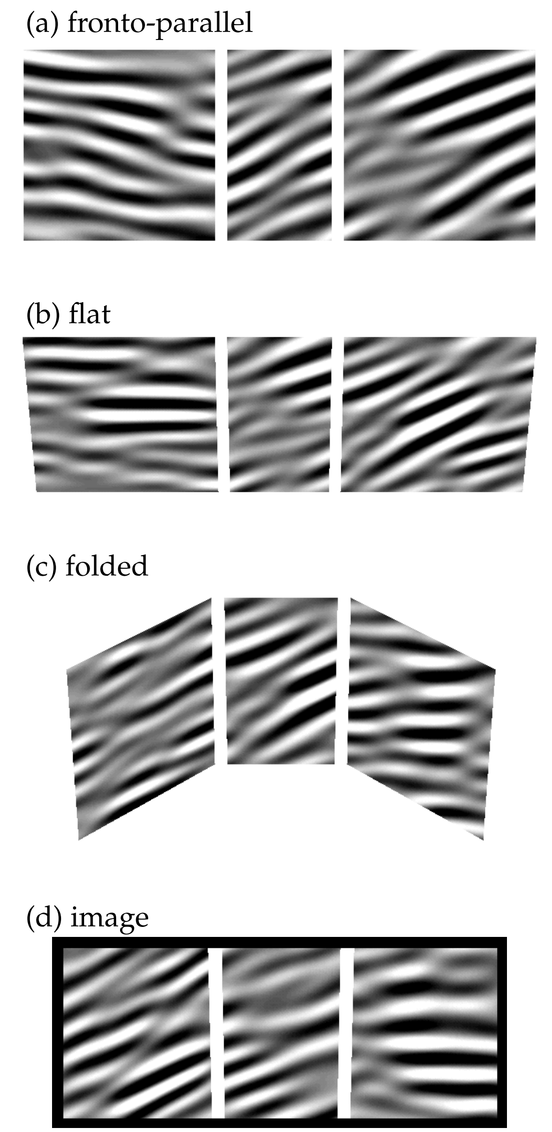

Figure 2: (a) The basic stimulus. Here, as in b-d, the right and center panels have the same surface texture orientation. Thus the right panel is the target, and the left panel is the non-target. (b) The flat condition: the panels have been rotated about the horizontal axis by 35 degrees. (c) The folded condition: same as b, but the side panels have been folded forward. (d) The image condition: same as c, but viewed behind an occluder. Conditions b and c were viewed stereoscopically.

ture orientation that differed from the center texture by 10, 20, 30 or 40 degrees. The observers' task was to select the side panel that had the same texture as the center panel. We will refer to this side panel as the "target" and to the dissimilar panel as the "non-target".

2.1.1 The folded condition

The three panels were rotated by 35 degrees about the horizontal axis so that the top of the stimulus appeared closer to the viewer. The left and right side-panels were then rotated forward by 50 degrees, Figure 2c. Under these viewing conditions, the similarity of the image textures can be readily dissociated from the similarity of the surface textures.

To see how folding the stimulus dissociates the image and surface textures, compare the fronto-parallel and folded stimuli depicted in Figure 2 (a and c). In the fronto-parallel stimulus, the right and center textures have a 30 degree orientation, while the left texture is -10 degrees. When this same stimulus is folded, the left and center image textures now appear more similar, while the right image texture differs. Note that if the surface texture on the left panel had differed by $+40$ rather than -40 degrees (i.e., if we had used a 70 degree rather than a -10 degree texture), then the left panel would differ from the center panel in both surface and image texture. Since our goal was to determine whether similarity judgments are based on surface texture or on image texture, we always selected non-target textures that would dissociate the two. Thus, whenever the left panel was the non-target, we rotated its texture in a clockwise direction to compensate for the counterclockwise rotation caused by the image projection. And conversely, whenever the right panel was the non-target, we rotated its texture in the counterclockwise direction.

The plot in Figure 3 shows, for each of the three

discrete textures and position the fold between columns of texture elements as in Figure 1. We made a few such seamless stimuli out of cardboard and observed the same perceptual effects reported here.

panels, the relationship between the texture orientation on the surface and the texture orientation in the image. [4] The filled circles correspond to a 30 degree texture mapped onto each of the panels. The open circles correspond to the non-target textures used in the experiment. Thus if the target was on the right, the non-targets would correspond to the four open circles on the left panel. Note that on the left panel, the more the non-targets differed from the center in surface texture, the more they resembled it in image texture.

2.1.2 The flat condition

In the flat control condition, the three panels were rotated by 35 degrees, but the side panels were not folded forward (Figure 2b). Because the panels were coplanar, the similarity of the textures in the image closely corresponded to the similarity of the textures on the surface. As such, this stimulus resembles stimuli used in a standard texture discrimination experiments. It also provides a rough estimate of how observers would perform in the folded condition if they based their judgments on surface texture. [5]

2.1.3 The image condition

In this condition, the folded stimulus described above was presented behind a cardboard mask which obscured the edges of the panels (Figure 2d). The vertical extent of the stimulus was made larger for this control condition so that the visible stimulus area would be unchanged. The stimulus was also presented non-stereoscopically. Without depth cues, the panels appeared coplanar and fronto-parallel.

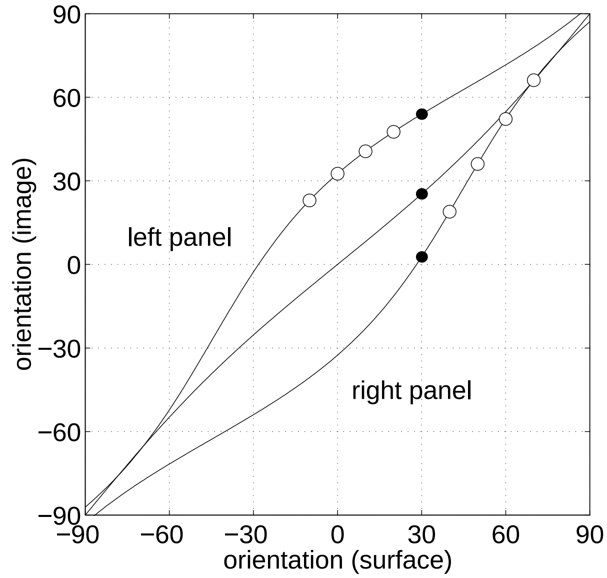

Figure 3: The orientation transformation caused by image projection. On the horizontal axis is the orientation of the surface texture, on the vertical axis is the orientation of the image texture. The three curves correspond to the three stimulus panels in the folded condition (Figure 2c).

This image control condition shows how observers would perform in the folded condition if they based their judgments on image texture.

2.2 Procedure

We used a physical model of the stimulus to explain the task to the observers. The side panels of the model could be adjusted so that they were either flat (coplanar with the center panel) or folded. The panels of the model had plastic covers for mounting pieces of cardboard that had been printed with texture patterns. To explain the task, we first mounted several texture triplets (e.g., 0, 0, 30 and 0, 30, 30) on the flat model. We then asked the observers to indicate which side panel had the same texture orientation as the center panel. Next, we folded the model, mounted new textures, and again asked the observers to judge the similarity of the textures. To ensure that the observers understood that we were asking about the textures on the surface and

[4] Note that Figure 3 shows only part of the texture transformation that is caused by the image projection. In addition to changing the texture orientation, the perspective projection also changes spatial frequency. So while this graph gives some indication of how observers might perform if they base their judgments on image texture, the best indicator of this is the direct empirical measure described in Section 2.1.3.

[5] This is only a rough estimate because it does not take into account any biases or imprecision in the observers' assessment of surface geometry. Nor does it take into account the nonlinearity of the curves in Figure 3 and the effect this nonlinearity would have on texture discrimination thresholds.

not in the retinal image, we explained that their response should be the same whether the model was folded or unfolded.

This same model was then used to asses how accurately the observers perceived the geometry of the folded surface. Each observer adjusted the physical model to match the folded surface displayed on the computer screen. This task was repeated three times and after each time we measured the angles between the center and side panels and between the center panel and vertical. Overall, there was a slight tendency for observers to underestimate both angles, but their average judgments were accurate within 6 degrees of rotation.

After these preliminaries, observers ran the experiment. The flat, folded and image conditions were run in 30 separate but interleaved blocks of 96 trials each. Within a block, the different texture orientations were randomly intermixed. Each stimulus was displayed for two seconds. The observers responded by pressing a key on the computer keyboard and the next stimulus was presented after a two second delay. No feedback was given. Each observer ran five, one-hour sessions over a period of $2 - 3$ weeks.

2.3 Observers

Three undergraduates from Rutgers University (females, ages $18 - 38$) were paid to participate in the study. The observers had never before participated in a psychophysical experiment and were unaware of the purpose of the study.

3 Results and Discussion

The observers' task was to determine which of the side-panels matched the center panel in surface texture. The texture orientation on the center panel was either horizontal or diagonal. Because these orientations produced very different results, we consider them separately.

The results for the horizontal textures are plotted in Figure 4 which shows the individual and average data for the three observers. The vertical axis indicates the percentage of trials in which

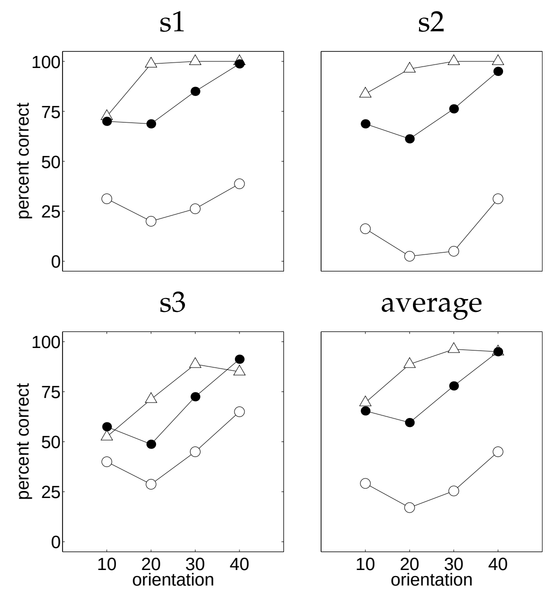

Figure 4: Results for stimuli that had a horizontal texture (0 degrees) on the center and target panels. The x-axis indicates the texture orientation on the non-target panel . The y-axis shows the percentage of times the observer selected the target as similar to the center. Triangles correspond to the flat condition, filled circles to the folded condition and open circles to the image condition.

observers correctly selected the target panel as being similar to the center panel. The horizontal axis shows the texture orientation of the dissimilar, non-target panel. Each point is the average of 80 trials from two mirror symmetric stimuli. Thus, the data points for the 40 degree orientation reflect the average results for the $0, 0, 40$ and $-40, 0, 0$ stimuli.

The triangles in Figure 4 correspond to the flat condition in which the three panels of the stimulus were coplanar. As expected, the results resemble those of a typical texture discrimination experiment. Performance improves as the orientation difference between the non-target and center panels increases, reaching an asymptote by 30 degrees. The filled circles show the data for the same stimulus but with the side panels folded forward. In this condition, the observers must factor-in surface geometry to judge the similarity of the sur-

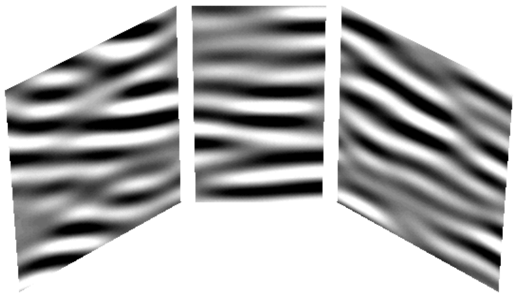

Figure 5: A stimulus with 0 degree surface texture on the center and right panels and -30 degree texture on the left panel. Note that the horizontal surface textures run parallel to the panel boundaries.

face textures. The open circles show the data for the image condition. These data indicate how observers would perform in the folded condition if they based their judgments solely on image texture. As expected, performance here is consistently below chance. This is because, in the image, the texture of the center panel is more similar to the non-target panel than it is to the target panel.

For our purposes, the important feature of these data is the separation between the image and folded conditions (i.e., the open and filled circles). The displacement of the folded data away from the image data and towards the flat data shows that observers did not rely solely on image texture in making their similarity judgments. Instead, they based their judgments on surface texture.

Although the results suggest that texture discrimination involves a 3D surface representation, these results could instead reflect a purely image-based process. This process cannot be based directly on image texture, of course, but it could be based on the parallelism between the image texture and the surface edge, since this is invariant to image projection, Figure 5. Because observers may have used the image cue of parallelism to determine which textures were horizontal, these findings may not generalize to other texture orientations.

Figure 6 shows the data for the diagonal textures. Again, the vertical axis shows the percentage of trials in which the observers correctly judged the target texture as similar to the center texture. The horizontal axis corresponds to the orientation of the non-target texture. Also, as in Figure 4, we have collapsed the data for mirror symmetric stimuli. This time, however, the symmetry is between stimuli with different textures on the center pane. So, for example, 0, 30, 30 is not the mirror image of 30, 30, 0, but it is the mirror image of -30, -30, 0. Thus, in graphing the results for diagonal textures, the data for 30 and -30 degree stimuli were combined. The labels on the graph correspond to the 30 degree stimuli.

The open triangles correspond to the flat condition. As before, these data look fairly typical for a texture discrimination experiment: performance increases as the orientation difference between the non-target and center increases up to about 30 degrees. The filled circles correspond to the data for the folded condition, while the open circles show the data for the image condition. Unlike the data in Figure 4, here the folded and image data are very similar. Thus instead of selecting the target, which resembled the center panel in surface texture, observers were consistently selecting the non-target, which resembled the center in image texture. This effect is especially striking on the left panel. Here, the dissociation between the similarity of image and surface textures is especially large: the more the textures differ on the surface, the more similar they are in the image. As a result, observers were remarkably bad at this task when the left panel was the non-target.

To summarize the results for this first experiment, when the center texture was horizontal, similarity judgments accurately reflected surface texture. In contrast, when the center texture was diagonal, these judgments were consistently wrong; instead of judging surface texture, observers judged image texture. How can we account for this discrepancy between the results for horizontal and diagonal textures?

As noted earlier, the parallelism between the horizontal textures and the surface edges is preserved in the image projection, and so our observers could have used this image cue to determine which tex-

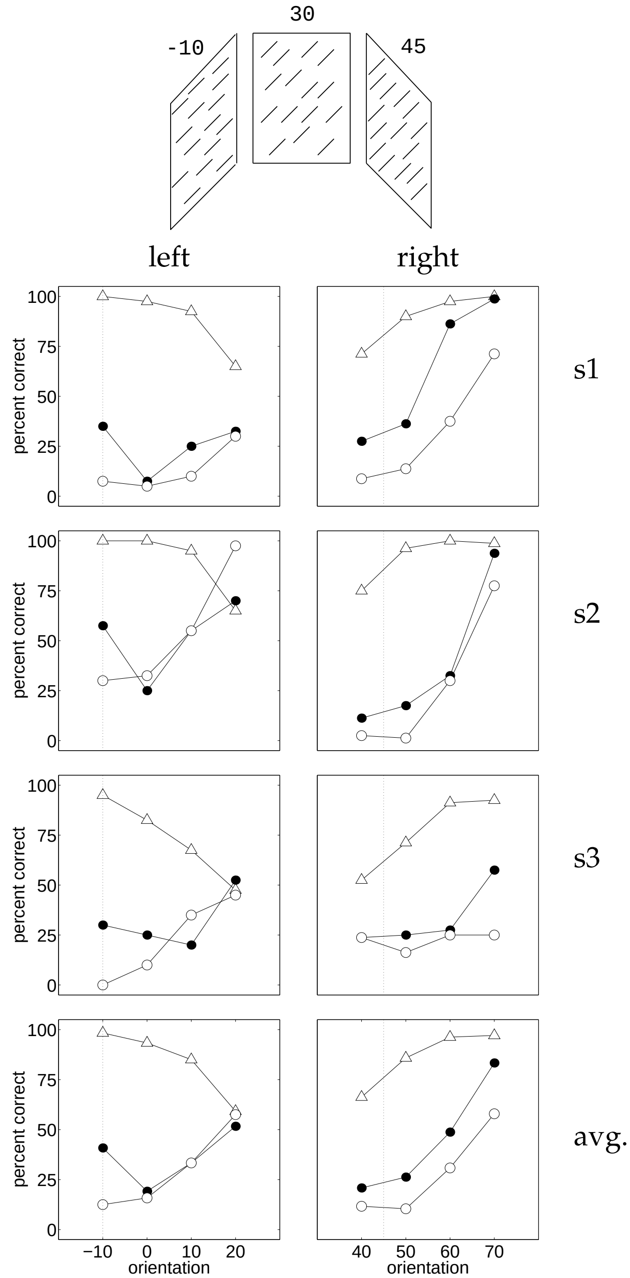

Figure 6: Results for stimuli that had a diagonal texture (30 degrees) on the center and target panels. The x-axis shows the texture orientation on the non-target panel. (These stimuli do not have the same bilateral symmetry as those in Figure 4, and so left and right panels are plotted separately.) The y-axis shows the percentage of times the observers selected the target as similar to the center. Triangles correspond to the flat condition, filled circles to the folded condition and open circles to the image condition. The dashed lines indicate the non-target textures that matched the center texture in the image. These textures are also depicted in the schematic at top.

tures were horizontal on the surface. [6] Additionally, horizontal is the most common texture orientation for man-made surfaces (consider bricks, tiles and aluminum siding), and so it is likely that our observers had more experience with these textures. And, as discussed below, there is recent evidence that observers can use surface texture to recover surface shape only for textures that are oriented perpendicular to the lines of maximal curvature on a surface (e.g., our horizontal textures) [12]. For all these reasons, the horizontal textures may be a special case.

Even if we can account for the superior performance for horizontal textures, we are still left with the surprisingly poor performance for diagonal textures. If observers base their similarity judgments on image texture, then these results reveal an unusually pronounced failure of perceptual constancy. In the next experiment we look more closely at this failure of perceptual constancy.

4 Experiment 2

In the previous experiment, observers judged texture similarity on a folded surface. When the textures ran diagonally along the surface, their judgments were dominated by image texture. Two conclusions can be drawn from this result. The first is that observers do not have orientation constancy. [7] That is, lines on a surface may appear to run one way from one vantage point, and a different way from a different vantage point. Alternatively, observers may have orientation constancy, but they may not use it when making similarity judgments. Or, in other words, orientation constancy may be calculated after, or in parallel with, similarity judgments. To distinguish between these alternatives, we measured orientation constancy directly by asking observers to make absolute judgments about

[6] But we should note that the data in Figure 6 show that observers did not take advantage of this cue when the center texture was diagonal and the non-target panel was horizontal.

[7] We are using this term in a somewhat non-standard way. Traditionally, orientation constancy has referred to the stability of perceived shape under image rotation.

the surface texture on each of the three stimulus panels.

4.1 Methods

Three new observers were recruited for this experiment. All three were Rutgers students who were paid for their time and who were naive as to the purpose of the study.

To measure orientation constancy, we first had the observers learn a standard texture orientation on the center panel. This center panel had the same dimensions as the center panel in the previous experiment, but, rather than being rotated about the horizontal axis, it was fronto-parallel. Initially, observers passively viewed 10 samples of a 30 or -30 texture mapped onto this panel. We then measured the accuracy of the observers' internal representation of the standard by having them perform an orientation discrimination task. In four blocks of 40 trials each, the observers were shown textures that deviated from the standard orientation by $\pm$ 10, 20, 30 or 40 degrees. Their task was to judge whether these deviations were in a clockwise or counterclockwise direction. Each stimulus was presented for 1 second. The observer responded by pressing a key on the keyboard, feedback was provided, and the next stimulus was presented 2 seconds following a 2 second delay.

After discriminating textures on the center panel, the observers repeated the task for the same textures mapped onto the side panels. These side panels were identical to those in the previous experiment, and so were rotated about both the horizontal and vertical axes and offset from the midline (left and right panels of Figure 2c). Only one panel was shown at a time, and the observers' task was to judge whether the texture on the panel was rotated clockwise or counterclockwise relative to the standard. Trials involving the right panel were randomly intermixed with trials involving the left panel. In all, eight blocks of 45 trials each were run for this one-panel, in-depth condition.

Interleaved with blocks of the one-panel, in-depth condition were blocks of a one-panel, no-depth condition. In this condition, observers judged the textures on the side panels in the absence of depth information. The stimuli were presented non- stereoscopically behind an aperture which obscured the edges of the panel.

Each observer ran the constancy experiment twice, once with the 30 degree texture as the standard and a second time with the -30 degree texture as the standard. After completing both texture constancy experiments, the same observers repeated the texture similarity experiment described in Section 2.

4.2 Results

The results of the orientation constancy experiment for three observers are shown in Figure 7. The observers' task was to judge whether a texture pattern was rotated clockwise or counterclockwise relative to a learned standard. The standards were 30 and -30 degrees, and because these stimuli are mirror symmetric, we have averaged the data from the two conditions. As before, the labels on the graphs correspond to the 30 degree texture. The graph shows the percentage of counterclockwise responses plotted against the texture orientation for the left and right panels. Recall that observers first performed this task, with feedback, on the fronto-parallel center panel. These data are plotted as open triangles in the left and right graphs. The filled circles show the data for the in-depth condition in which the observers repeated the task on the slanted side panels. The observers were given a full stereoscopic view of the panels and so could potentially have used information about surface geometry to recover the surface texture. Finally, the open circles show the data for the no-depth condition in which observers were forced to base their judgments on the image texture of the slanted panels.

If, as in the previous experiment, observers base their orientation judgments on image texture, then the data for the in-depth condition (filled circles) should fall near the data for the no-depth condition (open circles). Although this trend is seen in one observer's data for one panel (s2, right panel), this is clearly not the general pattern. In most cases,

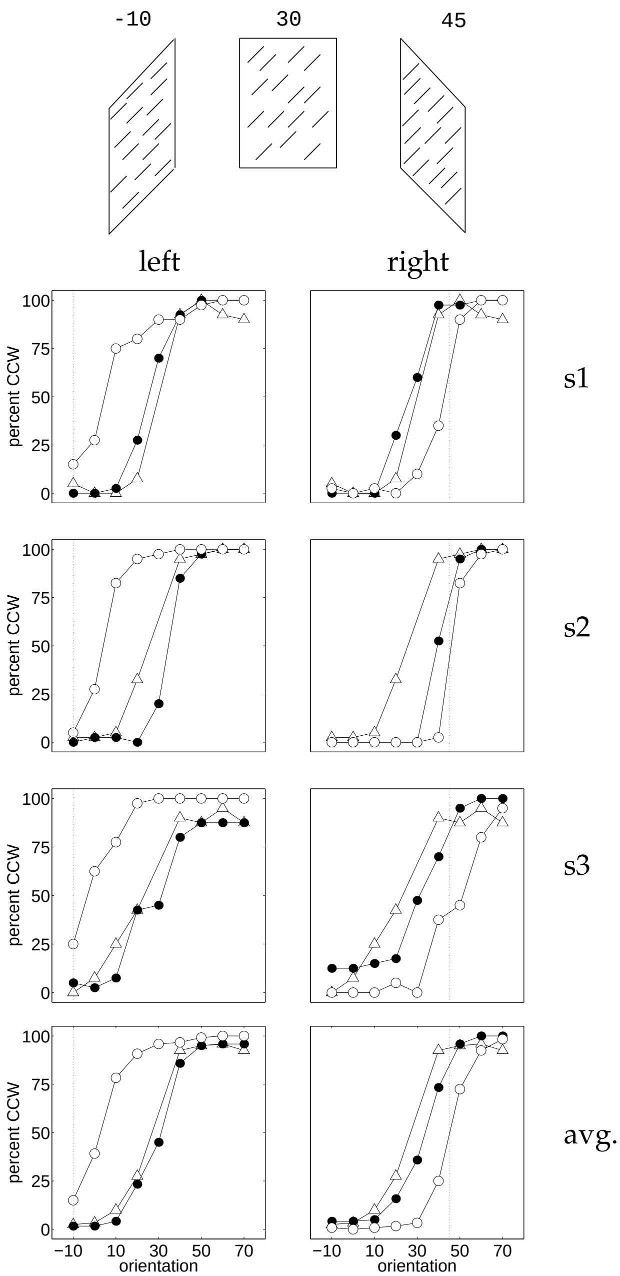

Figure 7: Perceived texture orientation. The x-axis shows the texture orientation on the panel surface. The y-axis shows the percentage of times the observers judged the texture to be rotated counter-clockwise from the 30 degree standard. Triangles correspond to the flat condition, filled circles to the one-panel, in-depth condition and open circles to the one-panel, no-depth. Dashed lines indicate which surface texture project to a 30 degree image texture. These texture orientations are also depicted in the schematic at top.

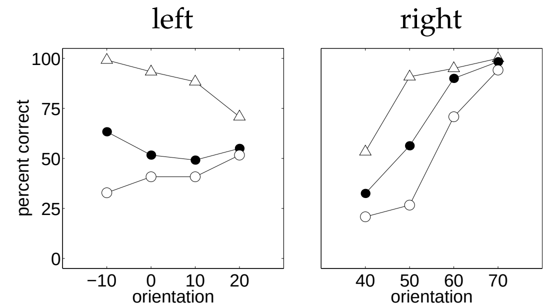

Figure 8: Diagonal texture: average data for three observers. Triangles correspond to the flat condition, filled circles to the folded condition and open circles to the image condition.

the orientation judgments for the in-depth stimulus are quite similar to those for the center panel (open triangles). Thus these data show that the observers can accurately judge surface texture.

The same three observers also ran in the texture similarity experiment described in Section 2. Their average results are shown in Figure 8. Previously we had found that observers relied on image texture in judging texture similarity, and this observation holds here as well. Overall, the data for the folded condition (filled circles) are similar to those of the image condition (open circles). These results contrast with those in (Figure 7) where the data of the one-panel in-depth condition, which is analogous to the folded condition, were more similar to the flat condition.

We should note that the effect in this second texture similarity experiment is not as strong as in the original experiment (Figure 6). One possible explanation for this reduced effect is that during the orientation constancy experiment, the observers may have learned to associate the image textures with their corresponding surface textures. This association might also explain why the second set of observers did noticeably better than the first set on the image condition (open circles).

To summarize, observers demonstrated good constancy when making absolute judgments of texture orientation. Thus they were able to judge whether the surface texture on a slanted panel was rotated clockwise or counterclockwise relative to a learned standard. Because they learned the standard orientation on a fronto-parallel panel, the observers must have compensated for the change in surface geometry when making these judgments. Nonetheless, these same observers showed poor constancy when making relative judgments of texture similarity. That is, when asked to compare the orientations of several textures presented simultaneously, their performance was strongly biased by the image textures.

At first glance, it may seem that these results do not add up. If observers could accurately judge the texture orientations on the left, right and center panels, then surely they should have been able to determine which panels had the same orientation. Presumably, if we had instructed our observers to analyze the stimulus one panel at a time, we would have found this result. However, we instructed our observers to tell us which panels appeared to have the same surface texture. Thus our results likely reflect not what observers can calculate, but what they perceive.

5 General Discussion

Our central finding is that observers fail to take surface geometry information into account when judging texture similarity for diagonal textures. Thus, their similarity judgments reflect not the textures on the surface, but the textures in the surfaces' projected image. Although this finding suggests that observers cannot recover surface texture, a control experiment showed that this is not the case. When presented with a single planar surface, observers were able to accurately judge surface texture even when the surface was rotated in depth. So although observers can recover surface texture, they do not use this information when judging texture similarity across a folded surface.

This surprising conclusion has a parallel in recent work by Li and Ziadi [12]. While we were interested in the recovery of surface texture from shape information, these authors were interested in the recovery of surface shape from texture information. Li and Ziadi mapped oriented textures onto a surface that was sinusoidally modulated in depth and then asked observers to make judgments about the shape of this surface. These judgments were accurate only when the surface textures were oriented along the lines of maximum and minimum curvature. When the textures ran diagonally along the surface, performance was poor. Thus, the visual system appears to be limited both in its ability to use image texture to recover surface shape and in its ability use surface shape to recover surface texture.

Our finding that the visual system might be unable to use surface shape to recover surface texture is also related to work by Cavanagh et al. [7]. They reported a study and demonstration which show that observers do not compensate for surface geometry when interpreting surface markings. Their demonstration is easy to recreate and so we will describe it here. It requires folding a dollar bill so that it appears as a "$\vee$" when viewed from the top and so that the crease vertically bisects George Washington's face. When the folded bill is viewed straight-on with the crease orthogonal to the viewing direction, Washington appears expressionless. But when the bill is rotated about the horizontal axis, Washington changes expression. Bringing the bottom of the bill closer makes him smile, bringing the top closer makes him frown. When the top of the bill is closer, the retinal image of Washington's mouth is a "$\wedge$", and observers attribute this "$\wedge$" to a frown and not to the fold. So just as in our experiment, the observers seem to ignore the fold when interpreting the markings on the surface.

Our study is also related to an experiment by Beck [2]. His stimulus was an array of intermixed vertical and diagonal lines drawn on cardboard, and his observers' task was to rate the extent to which the two orientations formed separate groups. When the stimulus was fronto-parallel, grouping was near threshold and observers gave very low ratings. When the stimulus was slanted towards the ground plane, however, these ratings increased considerably. Beck concluded that since the line orientation on the surface did not change, observers must base their judgments on image orientation.

However, as Beck himself acknowledged, there are difficulties with this interpretation (see also [18]). The most fundamental difficulty is that surface orientation is calculated from image orientation. Thus, any change in the reliability of the image information will necessarily affect judgments of surface orientation. So his finding is consistent both with the idea that similarity judgments are based on image orientation, and with the idea that these judgments are based on a property derived from image orientation, namely, surface orientation.

This same objection does not apply to our experiment, however, because we approached this question in a different way. Rather than present two textures on a single planar surface, we presented two textures on a folded surface and we asked whether observers take the fold into account when judging texture similarity. If they do, then even if two textures are similar in the image, they should not be judged as similar on the surface. But, of course, our experiment rests on the assumption that observers had sufficient depth information to recover the surface geometry. We think they did because all observers gave accurate verbal descriptions of the folded stimulus. In addition, the observers were able to adjust a physical model of the stimulus to match the computer generated surface. And finally, when we measured orientation constancy, observers were able to recover surface texture, suggesting that they could recover surface geometry.

But if observers could recover surface texture, why did they base their similarity judgments on image texture? One possibility is that judgments of texture similarity and texture segmentation are based on a process that occurs in early vision. Of course, to accommodate the results mentioned in the introduction, some segmentation processes must occur after occlusions and reflectance have been recovered. Thus, texture segmentation may involve an early, rudimentary surface representation that indicates depth ordering but not surface geometry. Surface geometry and surface texture may still be recovered by the visual system and used for orientation constancy, but this occurs only at a later stage in processing.

Alternatively, the discrepancy between texture similarity and orientation constancy may arise not because texture processing occurs in early vision but because the process underlying orientation constancy involves cognition rather than visual perception. That is, orientation constancy may not be perceived directly, but may involve some form of deliberate mental rotation. According to this second idea, judgments of texture similarity are based on a high-level surface representation. But even at the highest levels of visual processing, surface texture is not represented.

In any event, a great deal of research has shown that observers can readily detect the boundary between two textures with different orientations [1, 15, 14, 11, 23]. These texture boundaries may correspond to a number of situations in the world. That is, a texture boundary could reflect an occlusion, a change in surface texture or a surface fold. Our results indicate that observers cannot discriminate between texture discontinuities that are due to a change in surface texture and those that are due to a surface fold. Thus this finding extends earlier work by showing that even though observers can readily detect texture boundaries, they cannot reliably interpret these boundaries.

Acknowledgments

We are grateful for the support from NSF grant SBR-9729015 (Bravo) and NSF grant EIA-98-02068 (Farid) and NSF CAREER award IIS-99-83806 (Farid). We thank Stijn Oomes for providing the stereo viewing OpenGL code and Bob Hoskins for helping to collect the data.

References

[1] Jacob Beck. Perceptual grouping produced by changes in orientation and shape. *Science*, 154:538–540, 1966.

[2] Jacob Beck. The relation between similarity grouping and perceptual constancy. *American Journal of Psychology*, 88:397–409, 1975.

[3] Jacob Beck. *Textural segmentation*, chapter Organization and representation in perception. Erlbaum Associates, 1982.

[4] James R. Bergen. *Theories of visual texture perception*, chapter Spatial Vision, vol. 10 of Vision and Visual Dysfunction, pages 114–134. MacMillan Publishing, 1991.

[5] James R. Bergen and Edward H. Adelson. Early vision and texture perception. *Nature*, 333:363–364, 1988.

[6] Terry Caelli. On discriminating visual textures and images. *Perception and Psychophysics*, 31:149–159, 1982.

[7] Patrick Cavanagh, S. Peters, and M. von Grunau. Rigidity failure and its effect on the queen. *Perception*, 17:27a, 1988.

[8] L.O. Harvey and M.J. Gervais. Visual texture perception and fourier analysis. *Perception and Psychophysics*, 24:534–542, 1978.

[9] Zijiang He and Ken Nakayama. Perceiving textures: Beyond filtering. *Vision Research*, 34:151–162, 1994.

[10] Bela Julesz. Textons, the elements of texture perception and their interactions. *Nature*, 290:91–97, 1981.

[11] Michael S. Landy and James R. Bergen. Texture segregation and orientation gradient. *Vision Research*, 31:679–691, 1991.

[12] Andrea Li and Qasim Ziadi. Perception of three-dimensional shape from texture is based on patterns of oriented energy. *Vision Research*, 40:217–242, 2000.

[13] Jitandra Malik and Pietro Perona. Preattentive texture discrimination with early vision mechanisms. *Journal of the Optical Society of America (A)*, 7(5):923–932, 1990.

[14] H.C. Nothdurft. Orientation sensitivity and texture segmentation in patterns with different line oreintation. *Vision Research*, 25:551–560, 1985.

[15] Richard K. Olson and Fred Attneave. What variables produce similarity grouping? *American Journal of Psychology*, 83:1–21, 1970.

[16] Stephen Palmer, Jonathan Neff, and Diane Beck. Grouping and amodal completion. *Psychonomic Bulletin and Review*, 3:75–80, 1996.

[17] Stephen Palmer and Rolf Nelson. Late influences on perceptual grouping: illusory figures. *Perception and Psychophysics*, 62:1321–1331, 2000.

[18] Irvin Rock, Stephen E. Palmer Romi Nijhawan, and Leslie Tudor. Grouping based on phenomenal similarity of achromatic color. *Perception*, 21:779–789, 1992.

[19] Anne Sutter, Jacob Beck, and Norma Graham. Contrast and spatial variables in texture segregation: Testing a simple spatial frequency channels model. *Perception and Psychophysics*, 46:312–332, 1989.

[20] Anne Treisman. Features and objects in visual processing. *Scientific American*, 255:114–125, 1986.

[21] M.R. Turner. Texture discrimination by gabor functions. *Biological Cybernetics*, 55:71–82, 1986.

[22] Max Wertheimer. Untersuchungen zur lehre von der gestalt ii. *Psycologische Forschung*, 4:301–350, 1923. Translation published in Ellis, W. (1938). A source book of Gestalt psychology (pp. 71-88). London: Routledge & Kegan Paul.

[23] Sabina S. Wolfson and Michael S. Landy. Discrimination of orientation-defined texture edges. *Vision Research*, 20:2863–2877, 1995.