Perception, 2003, volume 32, pages 211 – 233

DOI:10.1068/p5012

The effect of texture relief on perception of slant from texture

Jeffrey A Saunders¶
Department of Cognitive and Linguistic Sciences, Brown University, Providence, RI 02912, USA
Received 1 August 2001, in revised form 14 August 2002, published online 5 February 2003

Abstract. Texture can be an effective source of information for perception of slant and curvature. A computational assumption required for some texture cues is that texture must be flat along a surface. There are many textures which violate this assumption, and have some sort of *texture relief*: variations perpendicular to the surface. Some examples include grass, which has vertical elements, or scattered rocks, which are volumetric elements with 3-D shapes. Previous studies of perception of slant from texture have not addressed the case of textures with relief. The experiments reported here test judgments of slant for textures with various types of relief, including textures composed of bumps, columns, and oriented elements. The presence of texture relief was found to affect judgments, indicating that perception of slant from texture is not robust to violations of the flat-texture assumption. For bumps and oriented elements, slant was underestimated relative to matching flat textures, while for columns textures, which had visible flat top faces, perceived slant was equal or greater than for flat textures. The differences can be explained by the way different types of texture relief affect the amount of optical compression in the projected image, which would be consistent with results from previous experiments using cue conflicts in flat textures. These results provide further evidence that compression contributes to perception of slant from texture.

1 Introduction

When a surface with a regular texture is viewed from an angle, the image texture is distorted owing to perspective, forming a *texture gradient*. Texture gradients can evoke a strong impression of 3-D slant even in the absence of other cues, as demonstrated in figure 1. The concept of a texture gradient was introduced by Gibson (1950), who was the first to point out the potential role of texture in perception of 3-D surface orientation and shape. Texture has since been shown to be an effective source of information for perception of both 3-D slant and curvature (Beck 1960; Cutting and Millard 1984; Todd and Akerstrom 1987; Johnston et al 1993; O'Brien and Johnston 2000). Recovering shape from texture has also been a focus of study in computer vision

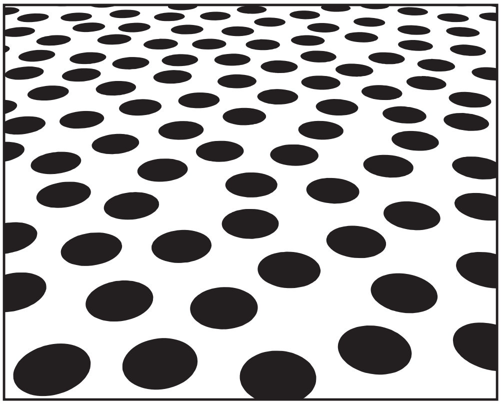

Figure 1. An example of a texture gradient. The image shows a perspective view of a slanted surface covered with circular elements. The surface is oriented 60° away from frontoparallel (slant = 60°) and recedes in the vertical direction (tilt = 0°). Following the convention of Stevens (1983), I define *slant* to be the angle between the line of sight and the surface normal, and *tilt* to be the angle of the projected surface normal relative to vertical (measured in the image plane).

¶ Address for correspondence: Center for Visual Science, University of Rochester, Meliora Hall, Rochester, NY 14627-0270, USA; e-mail: saunders@cvs.rochester.edu

(Stevens 1981a, 1981b; Witkin 1981; Ikeuchi 1984; Aloimonos and Swain 1988; Turner et al 1991; Garding 1992, 1993; Super and Bovik 1995; Malik and Rosenholtz 1997).

A basic problem in understanding how we use texture gradients is that they contain a number of regularities that could potentially be used to derive surface slant. One class of regularities are the gradients of texture properties: the rates of change of a texture property across an image. In figure 1, for example, the projected texture elements near the top of the image are smaller, closer together, and have more compressed shape than the elements at the bottom of the image. These three regularities have been described as *size*, *density*, and *compression gradients* (Cutting and Millard 1984). Since these cues correspond to distinct dimensions which could vary independently, each potentially provides an independent source of information about the slant and tilt of the surface.

Another source of information for surface slant is the amount of local *foreshortening* or *compression* of texture. For small regions of texture, perspective projection has the effect of compressing image texture in the tilt direction by an amount proportional to the cosine of local slant. Therefore, if the amount of compression can be measured in a region of the image, local slant can be recovered. For the texture shown in figure 1, texture elements are circular, so local compression is given by the aspect ratios of the elliptical texture elements in the projected image. For more natural textures, which often are not composed of such distinct and easily identified elements, local compression could be measured in other ways, such as from the distribution of edge orientations (Witkin 1981), or from asymmetries in the local power spectra (Garding 1993).

A number of researchers have attempted to determine which texture cues are used by the visual system, and how much each contributes to perception of slant (Stevens 1981a; Cutting and Millard 1984; Blake et al 1993; Passmore and Johnston 1995; Buckley et al 1996; Rosenholtz and Malik 1997; Knill 1998b, 1998c). Recently it has been found that compression tends to dominate perception of slant. Evidence comes primarily from cue-conflict experiments, in which stimuli have small discrepancies between the slants specified by different texture cues. Judgments are better predicted by compression than by size or density (Buckley et al 1996; Rosenholtz and Malik 1997; Knill 1998b). The results of Passmore and Johnston (1995) can also be interpreted as evidence that compression contributes to perceived slant. Observers were able to make reliable judgments for fractal textures which do not have a consistent notion of local size or density.

Compression can be distinguished from other texture cues based on the basis of the necessary assumptions underlying its use as depth information (Blake et al 1993; Rosenholtz and Malik 1997; Knill 1998a, 1998c). Size and density gradients provide valid sources of information as long as texture is *homogeneous*, which is defined as being statistically uniform across different regions of the surface. Using local compression as a cue does not require that a texture be homogeneous, but does depend on a local constraint of *isotropy*: texture must be statistically the same in all directions (ie circularly symmetric). The problem how we use texture information can be recast in terms of the underlying computational constraints (isotropy versus homogeneity) rather than specific texture cues. This approach is taken by Rosenholtz and Malik (1997) and Knill (1998b). In both studies, results can be interpreted as evidence that isotropy is used in perceptual processing of texture.[(1)]

[(1)] There is some evidence that textures with extreme violations of isotropy are perceived in a different way. When a texture is very elongated or has a dominant orientation (eg, wood grain), it provides contour information similar to that of a ruled surface. Knill (2001) found that, for curved surfaces with highly anisotropic textures, observers do not show the biases expected on the basis of an erroneous isotropy assumption, and appear to instead be using the texture as contour information.

Isotropy is one constraint that must be satisfied in order for texture compression to provide accurate information, but it is not sufficient. An additional constraint that has received less attention is that texture must also be flat and coplanar with a surface. This assumption does not always hold, since many surfaces have some depth structure perpendicular to a surface. Some examples of what I will term *texture relief* are bumps on a smooth surface, scattered objects like rocks which have volumes, coarsely woven fabrics, or blades of grass sticking up from the ground. The problem posed by texture relief for recovery of slant was pointed out by Stevens (1981a), and has recently been analyzed in more detail by Leung and Malik (1997). Depending on how texture is analyzed by the visual system, and also on the particular type of texture, the presence of texture relief might either interfere with use of texture as a slant cue, or provide additional information.

Previous experimental work has focused exclusively on textures that are flat and coplanar, so it is not known whether human observers are perceptually sensitive to violations of the flat-texture assumption. If texture gradients produced by surfaces with relief can still effectively convey slant, then models of human perception of structure from texture would have to account for this ability. In this case, the specific cases in which perception is robust might suggest solutions to how computational models can accurately compute slant and shape with these sorts of textures. On the other hand, if violations of the flat-texture constraint interfere with perception of slant from texture, it would lend support to models which incorporate such assumptions.

The experiments reported here use a novel class of stimuli, textures with relief, to explore how perception of slant from texture is affected by violations of the flatness assumption. An additional goal is to further test the relative contribution of compression and other texture cues. As I will describe in the next section, texture relief selectively affects the information from compression, which allows compression to be dissociated from other cues without explicitly introducing cue conflicts.

2 Texture relief

2.1 *Effect of texture relief on compression*

For a flat surface, perspective projection can be characterized in a simple way: given a top–down (frontal) view of a flat planar texture, the projected view from any other viewpoint can be generated by distorting the base texture by a 2-D projective transformation. This distortion depends only on viewpoint, and is independent of base texture. Consequently, there is direct mapping from local compression in a projected image to local slant (assuming isotropy), and from the gradient of compression to global slant (assuming homogeneity). When a textured surface has relief, this simplified characterization of perspective projection no longer holds, since its projected image is determined by its 3-D relief structure as well as by viewpoint. Compression measures like the aspect ratio of projected shapes, or asymmetry in edge orientations, would no longer depend solely on local surface orientation, and could be dramatically different than for a similar view of a flat texture.(2) Use of compression to estimate slant in such cases would be highly susceptible to error. In this section, I will present some examples of textures with relief, and describe how these different types of relief alter the information provided by compression.

For a texture composed of local elements, one can identify two components of the flat-texture constraint: texture elements must be planar, and their orientations must be

(2) The presence of relief can also affect other aspects of the projected appearance of a texture, which are not related to use of texture as a slant or shape cue. For example, the mean projected luminance of a texture with relief can vary depending on viewing angle, owing to local variations in shading or cast shadows. There has been some recent work on measuring and modeling such effects for real-world textures (Dana et al 1999).

aligned with the overall surface they form. An extreme example of texture which violates the planar part of this constraint would be a surface formed by many spheres (figure 2). The projected shapes of spheres remain circular regardless of viewing orientation, so perspective views of a texture formed by spheres would produce gradients of size and spacing, but no foreshortening or compression of individual spheres. In the absence of any compression, the compression gradient is also eliminated as a source of information.

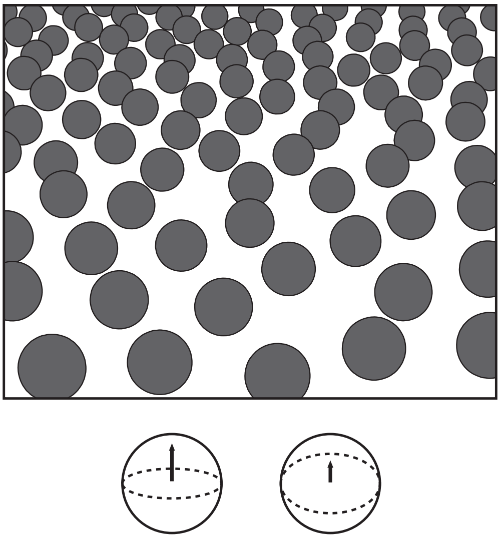

Figure 2. A perspective view of a texture composed of spheres (slant = 60°, tilt = 0°). Size and spacing gradients are present, but there is no compression of texture elements. The projections of spherical texture elements remain circular in the image regardless of the orientation of the underling plane, whereas a flat-texture element aligned with the same global surface would be foreshortened (dotted lines).

Other types of volumetric elements present similar problems for use of compression. Figure 3 shows an example of a texture composed of bumps, in which flat-texture elements have been replaced by the top halves of ellipsoids. The upper boundaries of projected texture elements in this example are due to self-occlusion, so their overall projected contours are larger and more circular than the projections of their planar bases. The result is an intermediate amount of compression, between the extreme cases of flat coplanar texture elements and spherical texture elements; some image compression remains, but it is attenuated, especially at high slants. This would be typical for textures composed of volumetric elements. In the general case, one can compare the projection of a planar slice through a volumetric texture element, which represents an equivalent flat-texture element, to the projected contour of the overall element. For smooth, compact texture elements, the difference between these contours would be greatest in the tilt direction, and would increase with viewing angle. The exact amount of compression would depend on the particular shape of texture elements,[3] but in general the presence of texture relief would reduce both the amount of local compression and the magnitude of the compression gradient. Using either of these cues would therefore lead to underestimation of slant.

[3] In the case of an isotropic texture composed of ellipsoids, the amount of image compression can be described analytically. On the assumption that texture elements are small relative to viewing distance (projection is locally affine) and have perpendicular height equal to h times the base diameter (h contained in [0,1]), the projected aspect ratio given is given by: $a(s) = \cos(s)[(1 + h^2 \tan^2(s)]^{-2}$, where s is the local surface slant. The case of $h = 0$ corresponds to flat coplanar circles, which compress by a factor $\cos(s)$, while the case of $h = 1$ corresponds to spheres, which have constant aspect ratios equal to 1. For the bumps in figure 3, the top halves of projected contours are equivalent to that of ellipsoids, and the bottom halves are equivalent to that of flat elements, so the overall projected aspect ratio is approximately halfway between these cases.

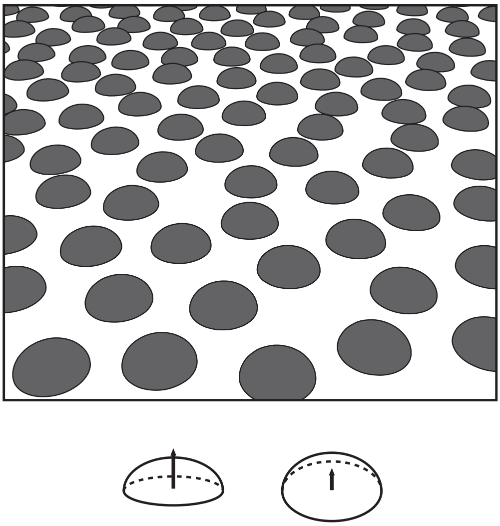

Figure 3. A perspective view of a texture composed of bumps (slant $= 60°$, tilt $= 0°$), each of which is the top half of an isotropic ellipsoid, with height equal to 65% of the base radius. Size and spacing gradients are the same as for a flat texture, but the amount of compression is reduced. Below the texture are two sample elements, taken from near the top of the projected image, where viewing angle is high (left), and near the bottom, where viewing angle is low (right). For both samples, the projected outline is closer to circular than a flat-texture element with the same orientation (dotted lines). The difference is much larger at high slant (left) than at low slant (right).

Surfaces composed of planar elements can have texture relief if the orientations of the elements are not aligned with the orientation of the surface. This class of textures would also present difficulties for methods which use compression. The projected shape of planar elements is a function of the orientation of elements relative to the observer's line of sight, not the orientation of the overall surface that the elements form. This is not a problem if individual elements are aligned with a surface, since the local orientation of the elements and the orientation of the overall surface agree. However, if texture elements are not aligned, using the local compression of the elements as information for the orientation of the overall surface would result in noise or biases.

One case would be if local orientations have random deviations from the overall slant. For example, if planar elements were distributed over a surface with small, random variations in topography, their orientations would randomly vary around the orientation of the surface (figure 4). If local compression was integrated over a large region, this information would remain accurate, since the mean orientation specified by the elements is equal to the overall orientation of the surface. However, the local differences in orientation contribute noise, so the variability of estimates of orientation

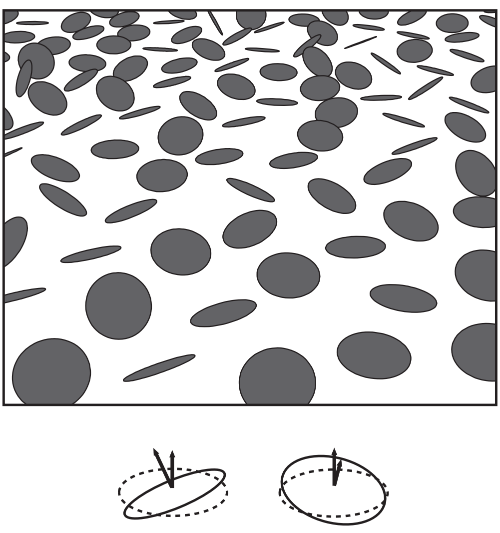

Figure 4. A texture composed of planar elements which are not all aligned with the global surface (slant $= 60°$, tilt $= 0°$). If the compression of these elements were used to determine local slant, the variations in local orientations would contribute noise.

would be increased. Orientations of local elements lying on a surface could also differ from the surface orientation in a correlated way. An example is the frontally oriented planar elements ('tombstones') shown in figure 5. For these elements, local compression is essentially eliminated as a potential cue. A compression gradient is present if one considers the change in projected shape of the tombstones, but the direction of the gradient is opposite that of flat planar texture elements.

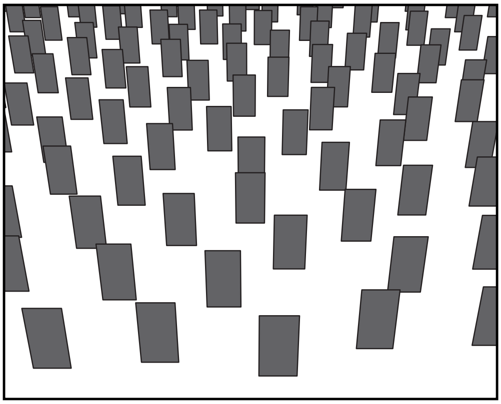

Figure 5. A texture composed of rectangular elements oriented perpendicular to the surface they form (slant $= 60°$, tilt $= 0°$). The projected shapes are foreshortened by an amount depending on the orientation of the texture elements. Because the elements are not aligned with the overall surface, compression does not provide reliable information for the global orientation.

I have thus far considered only how the presence of texture relief alters the projected shapes of individual texture elements. If texture elements are closely packed, there is also the potential problem of texture elements occluding one another, which Stevens (1981a) has termed *successive occlusion*. Successive occlusion of varying amounts can be observed in all of the previous examples of textures with relief. If a texture element is partially occluded, the projected shape of its visible portion no longer depends solely on its 3-D shape and viewing angle. The part of a projected boundary contour due to occlusion by nearer texture elements depends on the shape and positions of the occluding elements, which could dramatically alter its overall shape (ie aspect ratio and orientation) and shape properties (ie distribution of edge orientations) relative to an equivalent unoccluded texture element. Thus, if compression were measured simply on the basis of the visible portion of an occluded texture element, the resulting slant estimate would generally be inaccurate. The specific effect of successive occlusion is difficult to predict, since biases could be in either direction, depending on what portion of an element were occluded.

These examples demonstrate that there are qualitatively different ways that texture relief could interfere with use of compression to derive surface orientation. Volumetric, or bump-like, elements produce less compression when projected, which would lead to underestimation of slant. Variations in the orientation of texture elements can add noise to compression information, in the case of small random variations; or could eliminate compression entirely, in the case of large correlated differences between the orientation of texture elements and the overall surface. Finally, for any texture with relief that is densely packed (ie relief is large relative to spacing), successive occlusion could obscure the amount of compression, contributing noise or bias.

2.2 *Effect of texture relief on size and spacing*

In contrast to compression, the information provided by texture size and spacing is largely unaffected by the presence of texture relief. A distinguishing difference is that size and spacing describe the distribution of texture elements, and consequently are less affected by factors like shape and orientation of texture elements that strongly affect local compression.

The scaling component of perspective projection depends only on the distance of texture elements from the observer, so gradients of size produced by textures with relief would be the same as those produced by flat textures. All that would be required to use this cue is a measure of the projected size of a texture element size that is independent of its viewing angle. If tilt direction can be identified, or if texture elements have an identifiable planar 'base', then the width of texture elements would provide a sufficient measure for isotropic textures. Another possible measure is the length of a projected texture element measured in the maximal direction. For textures with shallow relief, this would generally be the width, but for other textures it might also be the height (eg tombstones, grass), or size along variable axes (eg oriented texture elements). These measures of object size would generally be unbiased for isotropic textures, with or without texture relief.[(4)] Successive occlusion could still potentially interfere with extraction of a local size measure. However, one would expect a size measure to be more robust to occlusion than compression, since it need not be based on the entire projected shape of a texture element. For example, in almost all instances of partially occluded texture elements in the previous figures, the visible projected shapes have the same maximal length as they would without occlusion.

The information provided by spacing of elements is also relatively unaffected by texture relief. The positions of texture elements within a projected image might be harder to define for textures like grass which occlude one another, but if a consistent measure of element positions can be identified (eg maximum point), the spacing of elements remains a valid source of information for surface orientation.

Thus, both size and spacing gradients would generally remain reliable sources of information when relief is present, regardless of how the relief structure affects the projected shapes of texture elements. Because texture relief selectively alters the availability and accuracy of compression as a cue for slant, using textures with relief provides a way to dissociate texture cues without explicitly introducing conflicts between cues.

2.3 *Additional information provided by texture relief*

Texture relief also has the potential to provide additional information for surface orientation. One case, discussed by Leung and Malik (1997), is when textures have perpendicular components that are uniform across a surface, such as the heights of the 'tombstones' shown in figure 5, or the vertical columns shown in figure 6. Compared to flat elements, the compression of perpendicular components scales in a different way as a function of surface orientation (figure 7). Flat elements are compressed by an amount equal to the cosine of angle between the surface and the line of sight of an observer, so compression is maximal at high slants (ie nearly parallel to the line of sight), while the compression factor of perpendicular elements is the sine of the angle between the surface and the line of sight, so compression is maximal at low slants (ie nearly frontal surfaces). Although the scaling function for perpendicular components is different than for flat textures, there is still a systematic relationship with the orientation of the global surface. Consequently, if perpendicular components can be segregated from flat components, and are roughly uniform in height, then the change in compression for these elements can also be used to determine surface orientation. The information provided by compression of perpendicular and flat elements is complementary. Flat textures are most informative at high viewing angles (Knill 1998a).

[(4)] Anisotropic textures with relief could still be problematic. For example, a texture composed of elongated texture elements, oriented away from the tilt direction, would produce misleading size information by a variety of measures. However, this difficulty is due to anisotropy, not texture relief, and would be shared to a lesser extent by flat anisotropic textures. Moreover, for some types of textures, the height of relief could provide an alternate unbiased measure of size even in the case of anisotropic texture.

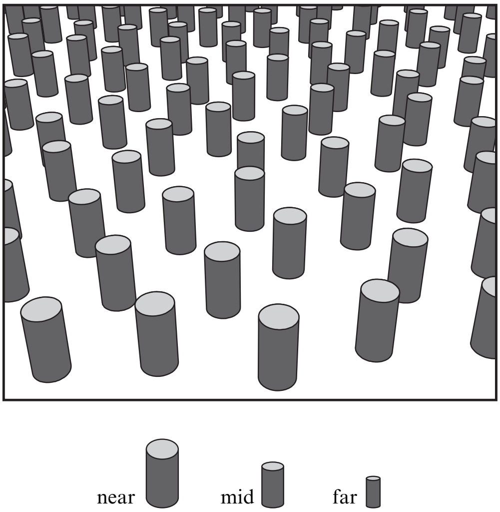

Figure 6. A texture composed of columns, which have planar top faces aligned with the surface and sides perpendicular to the surface (slant = 60°, tilt = 0°). The top faces and sides scale differently with distance, as illustrated by the sample columns taken from different parts of the image. The top faces change size and aspect ratio in the same way as a flat texture: elements get smaller, and the ratio of height to width of the faces in the projected image gets smaller. The projected heights of poles also change, but at a slower rate. The ratio of height to width gets bigger with distance and increased slant, opposite the change in aspect ratio for the top faces.

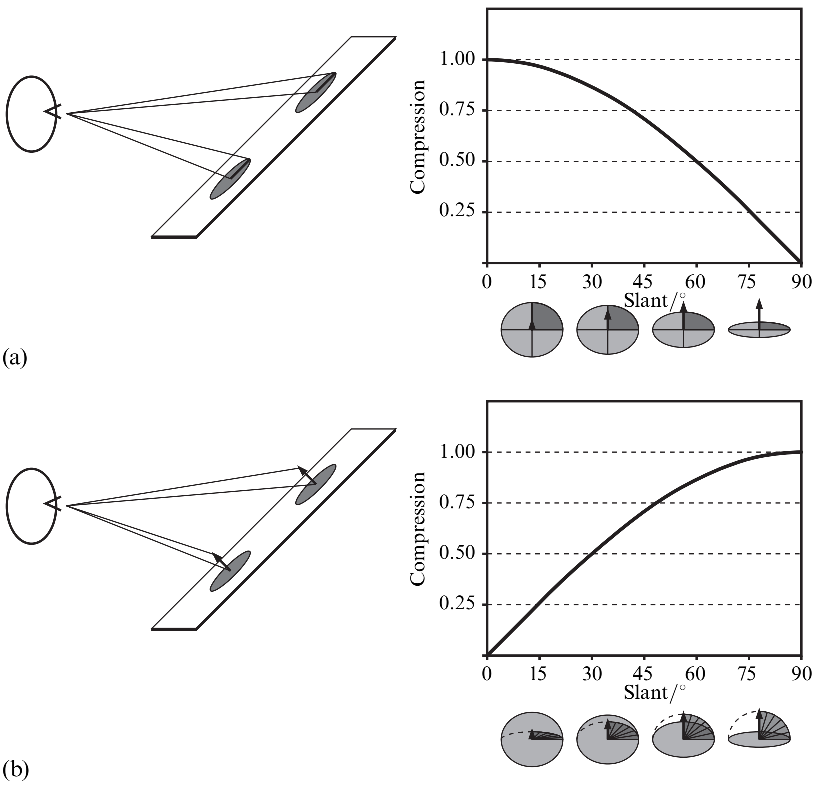

Figure 7. Illustration of the amount of compression due to perspective foreshortening as a function of slant for lengths along the surface in depth (a) and for perpendicular extents normal to the surface (b). The plots represent the ratios of visual angles of the foreshortened extents relative to the horizontal angular sizes, for equal lengths in 3-D. The amount of compression changes as viewing angle approaches the horizon and slant increases. For lengths in depth along the surface (a), there is little foreshortening at low slants, but as viewing angle increases the projected vertical sizes decrease at a progressively faster rate than the projected horizontal sizes, so the compression ratios decrease toward zero. For extents oriented perpendicular to a surface (b), the projected heights are highly foreshortened at low slants, but approach the projected horizontal sizes as viewing angle increases toward the horizon, so the compression ratios approach 1.

Intuitively, this is because compression changes quickly as a function of viewing angle at large slants, and slowly at low slants. The compression of perpendicular components is a mirror opposite function of viewing angle, with maximum rate of change of compression at 0° slant (frontal) rather than at 90° slant in the case of flat textures. Thus, the scaling gradient of perpendicular components would be most informative at small slants.

Another potential source of information that is specific to textures with relief is the occlusion of far texture elements by nearer ones. As discussed previously, successive occlusion could obscure the projected shape of texture elements, thereby interfering with the use of local compression or compression gradients. On the other hand, if occlusion can be identified, it could also provide information. For example, the projected shapes in figure 3, considered by themselves, could be mistaken as flat-texture elements. The resulting interpretation would underestimate slant, since the texture relief in this example reduces compression. However, the visible occlusion between elements is inconsistent with a flat-texture interpretation. From this, one could infer that the compression in the image underestimates the amount of foreshortening expected from a flat texture. To compensate, compression could possibly be given less weight or ignored, or could be reinterpreted on the basis of an estimate of the amount of relief present.

Occlusion could also be used as gradient cue to estimate slant directly, rather than as a constraint on interpreting other texture properties. The amount of successive occlusion depends on the difference between the projected height of texture elements and their projected spacing. Since both of these attributes form gradients that depend on slant, the amount of occlusion also provides a gradient cue. To provide a measurable gradient, successive occlusion would have to be sufficiently dense, which limits the use of this cue to textures with high relief relative to spacing, and large viewing angles (for example, the upper portions of figures 5 and 6). In addition, because noise in both height and spacing contributes to noise in the amount of occlusion, direct use of occlusion as a gradient cue would be highly sensitive to noise. In general, the informativeness of an occlusion gradient would strongly depend on viewing angle, height of texture relief, and the regularity of texture height and spacing.

The examples in this section illustrate that the presence of texture relief can introduce additional regularities in the view of a slanted textured surface, beyond those available in the image of a flat coplanar texture. If the height of texture relief is homogeneous and the perpendicular extent of projected texture elements can be measured, then the gradient of the perpendicular components provides a cue to slant. Occlusion between texture elements also provides information, which could be quantified by the gradient of the amount of successive occlusion, or as a constraint on the interpretation of other texture information.

3 Experiment 1: Types of texture relief

Experiment 1 tested perception of slant for four classes of textures that varied in the presence and type of texture relief: Flat elements, which served as a baseline or control; Oriented elements, which were flat but varied in orientation relative to the surface; and Bumps and Columns, which were raised relative to the plane of the surface. These texture types were chosen to sample different qualitative ways that texture relief could inhibit or facilitate use of texture to perceive slant. In general, differences in judgments across texture types would be expected only to the extent that perception of slant relies on compression.

Figure 8a shows an example of a texture from the Flat condition. Textures were composed of planar shapes with random size and spacing, which were approximately isotropic and aligned with the overall surface they formed. The textures with relief in the other three conditions shared the same random distributions of size and spacing, but differed in either the shape or orientation of the individual texture elements.

In the Oriented condition, texture elements were planar but randomly varied in their orientation relative to the global surface. Figure 8b shows an example of a texture from the Oriented condition. One effect of the orientation perturbations is to add noise to information provided by compression. The local orientations indicated by the projected shapes of individual texture elements are not globally consistent, and deviate randomly around the overall surface slant. To the extent that compression contributes to perception of slant, judgments of slant in the Oriented condition would therefore be expected to show more variability or biases than in the Flat condition.

Information from size and spacing of texture elements would be less affected by orientation perturbations. The projected size of a texture element does change with rotation, but the resulting variability would be much smaller than the variability in the direction and aspect ratios of projected shapes. Thus, orientation perturbations would be expected to have less effect on use of the size gradient than on use of local compression. The distribution of positions of elements is not affected by the perturbations of orientation, so, to the extent that spacing information is used, no difference between the Oriented and Flat textures would be expected.

Figure 8c shows a texture from the Bumps condition. The base of these texture elements is the same as in the Flat condition, but the centers of the texture elements are raised in the direction perpendicular to the surface. The relief present in the Bumps texture reduces the foreshortening of individual elements, so use of either compression or the compression gradient would be expected to decrease perceived slant, biasing judgments toward the frontal plane. Information provided by the size and spacing of

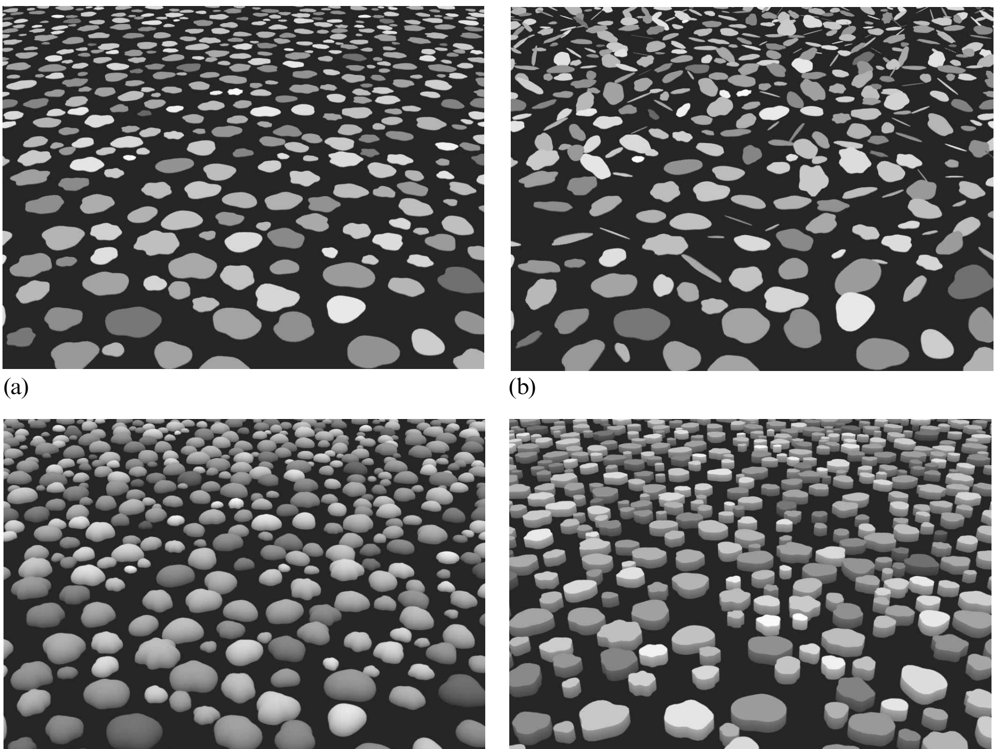

Figure 8. Examples of (a) a Flat texture (slant = 60°, tilt = 0°), (b) an Oriented texture (slant = 60°, tilt = 0°), (c) a Bumps texture (slant = 60°, tilt = 0°), and (d) a Columns texture (slant = 60°, tilt = 0°).

Bumps is unaffected by the relief, so, to the extent that perception of slant relies on these texture cues, no differences between Bumps and Flat elements would be expected.

Columns textures were generated by extruding texture elements in the Flat condition in the direction perpendicular to the surface (figure 8d). The projected shape of the top faces of Columns elements provides the same information as Flat elements, so, unlike the Bumps and Oriented textures, the texture relief in Columns does not degrade the information provided by compression. The sides of Columns could potentially interfere with analysis of compression, but they could also be used as additional information, which would be most advantageous for small slants. As in the Bumps condition, the information provided by size and spacing is unaffected.

In addition to varying the type of texture relief, I varied the degree of variability in the size and spacing of individual texture elements. Variability in the size and spacing of texture elements selectively decreases the reliability of information from size and spacing gradients in an image, which are the texture cues that remain unaffected by the presence of texture relief. Studies of the interaction of multiple depth cues have found that the relative weighting of different cues can be adaptive, with the same cues being weighted more in conditions in which they are more reliable (Johnston et al 1993; Landy et al 1995; Jacobs 1999; Jacobs and Fine 1999). If texture cues were integrated adaptively, the amount of weight given to compression would depend on the reliability of size and spacing information, and would be expected to increase if noise were added to these other cues. On the other hand, neither size nor spacing of elements directly affects the information provided by compression, so, to the extent this information contributes independently, no difference would be expected.

3.1 *Methods*

3.1.1 *Subjects.* Twelve subjects participated in experiment 1. Subjects were recruited by signs posted around Brown University, and were paid for their participation. Most were Brown University students. All were naïve to the purposes of the experiment, with the exception of two graduate students, who had a general idea of the research.

The data from two of the naïve subjects showed signs that they either did not understand the task or were unmotivated. Slant judgments for these two subjects showed larger variability than for the other subjects, and did not vary reliably as a function of stimulus slant, even for the control condition. The data for these two subjects were excluded from analysis.

3.1.2 *Apparatus.* Stimuli were rendered with an SGI Crimson Reality Engine computer with OpenGL and presented on an SGI computer monitor with a resolution of 1024×768 pixels and refresh rate of 60 Hz. The display was viewed monocularly through a masking box that occluded the edges of the monitor. The displays subtended visual angles of 40 deg horizontally and 32 deg vertically. Images were simulated to be accurate perspective projections when viewed at the opening of the box, which was 45 cm from the screen.

3.1.3 *Procedure.* Perceived slant was measured with a matching task. Subjects viewed images of textured surfaces presented on a monitor, and adjusted a mouse-controlled gauge figure (figure 9) to match the apparent slant of the texture surface. The gauge figure was 24 deg when viewed frontally, filling roughly 60% of the screen.

The test image and gauge were viewed sequentially, such that they were not visible at the same time, and subjects could toggle back and forth between the two. Subjects were encouraged to view the gauge and test image only a few times for each trial, but they were free to toggle as often as they needed to be confident of their judgment, and no time limit was enforced.

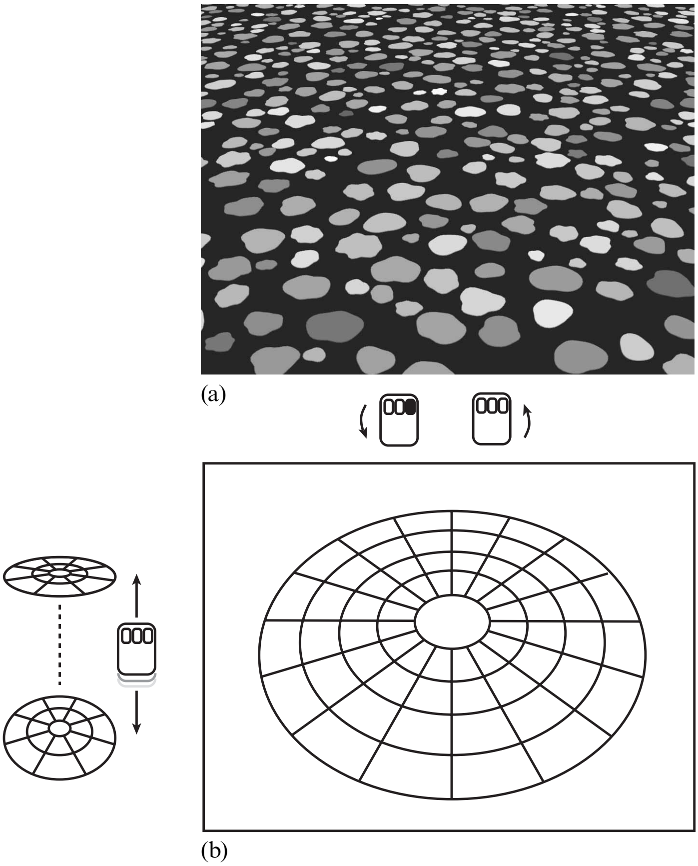

Figure 9. The slant judgment task. Subjects toggled back and forth between the test image (a) and a gauge figure (b), which was adjusted with the mouse.

Each subject was presented with images from all 24 of the Slant × Type × Randomness conditions (see below for conditions). The test images were presented in random order in two blocks of 192 trials, which yielded 8 trials per condition and block, for a total of 16 trials per condition. The experiment was self-paced, and generally took between 40 and 70 min.

3.1.4 *Stimuli.* Textures were composed of individual texture elements of varying size, scattered on a flat plane. Two size and position distributions were used across conditions, one regular and one irregular, and each distribution was combined with all types of texture elements. Stimuli with texture relief were constructed by replacing the flat planar elements with equally sized elements with relief.

The size and position distributions of texture elements were determined by first choosing a random set of sizes, and then randomly positioning elements with the constraint that elements cannot overlap. The sizes were chosen from either a uniform probability distribution or linearly decreasing probability distribution, with ranges chosen to equate means. The spacing was derived by first assigning all texture elements to random positions, and then relocating elements that overlapped. In the case of overlaps, the smaller texture element was moved. This method of generating random positions for variable-sized objects has been used to model natural distributions such as the locations of neighboring trees, which are produced by a competitive growth process (Stoyan and Stoyan 1994). The amount of randomness in spacing was not directly manipulated, but this procedure applied to objects with different distributions of sizes produces different degrees of homogeneity in spacing. Figure 10 shows examples of the two types of distributions. Distributions generated with sizes sampled from a uniform distribution were less variable in size and more evenly spaced (Regular condition) than distributions with sizes sampled from a linear distribution (Irregular condition).

The random shapes of Flat-texture elements were generated by adding perturbations to a circular contour. For each element, the radius of the shape perturbation as a function of angle was the weighted sum of five cosine functions with different frequencies. Weights for each of these functions were chosen randomly and scaled by a constant factor inversely proportional to frequency, so that high-frequency perturbations tended to be smaller in magnitude than low-frequency perturbations. After adding a perturbation to a base circle, the resulting contour function was then re-normalized so that the maximum radius was equal to the desired size of the element (as determined by the size and spacing distribution).

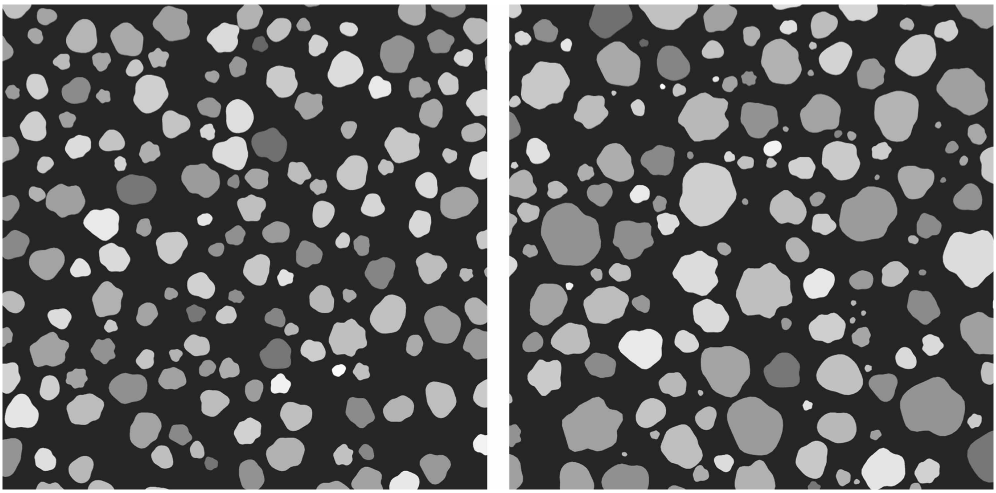

Figure 10. Examples of the two distributions of texture size and spacing used to generate the stimuli. The left distribution was generated by selecting texture element sizes from a uniform distribution, while the right distribution was generated by selecting sizes from a linearly decreasing distribution. Positioning was random, but with the constraint that elements cannot overlap. The uniform method produces more regular sizes and more even spacing (Regular condition) than the linear method (Irregular condition).

Oriented textures were derived from Flat textures by adding random 3-D perturbations to the orientation of individual texture elements. The perturbations can each be described by a pair of parameters: a direction of rotation and a magnitude of rotation. The magnitude of rotation for a texture element corresponds to the angular difference between its normal direction and the normal direction of the global surface. This angle was chosen from a Gaussian distribution with standard deviation of 30°. The maximum rotation was limited to 90°; larger rotations were replaced by new random values. The direction of rotation for each texture element was chosen from a uniform distribution, so that all possible directions were equally likely. Because the distribution of orientation perturbations was circularly symmetric, the expected value of the mean of local-surface normal directions would be the same as the normal direction of the global surface. The actual mean local orientations for the stimuli differed slightly owing to sampling noise, and because the subset of visible elements varied depending on slant, but were within 2° of the global surface orientation.

The Bumps and Columns textures were generated by using the shapes of elements in the Flat condition as a lower cross section and adding a volumetric component.

The basic shape of the Bumps elements was half of an ellipsoid, with height equal to 35% of the horizontal size. Shape perturbations for the corresponding Flat elements were applied to cross sections of the ellipsoid to produce a smooth but randomly shaped bump. The shading of these elements simulated a matte surface with a combination of ambient and overhead directed lighting. The Columns were generated by raising the faces of Flat elements by a constant amount equal to 35% of the average width of the elements, and adding sides along the outer contour. The same surface and lighting conditions as those used for the Bumps were applied to these textures, causing the sides of the Columns to be darker than the top faces.

The final step in generating the stimuli was to render the textures in accurate perspective at different orientations relative to the observer. The simulated surfaces in the displays were all slanted away from the observer in the vertical direction (surface tilt is vertical). I will use the convention that slant is the angle between the surface normal and a line of sight from the observer, so 0° corresponds to a frontal surface and 90° corresponds to a surface which is parallel to the line of sight. Three slants were used to generate the textured test surfaces: 30°, 45°, and 60°. Subjects were asked at the end of the experiment about the apparent range of test slants, and none of the subjects reported noticing that only three slants were used.

3.2 *Results*

In the two graphs in figure 11, the mean slant of gauge settings across subjects is plotted as a function of texture type and test slant. The top graph corresponds to the textures with the more regular size and spacing distributions, and the bottom graph corresponds to textures with the less regular and more variable size and spacing distributions. Differences between texture types varied depending on the test slant, but a general ranking of texture types is evident. Overall, Oriented and Bumps textures were judged to be less slanted than the baseline Flat condition, while Columns textures were judged to be more slanted. The same pattern is present for both of the size and spacing distributions.

An overall ANOVA of the slant judgments found significant main effects of texture type ($F_{3,207} = 52.09$, $p < 0.001$) and test slant ($F_{2,207} = 322.1$, $p < 0.001$), and an interaction between texture type and slant ($F_{6,207} = 3.337$, $p < 0.004$). The regularity of size and spacing had no effect ($F_{1,207} = 0.3806$, $p = 0.538$, ns), and did not interact with texture type ($F_{3,207} = 0.1884$, $p = 0.904$, ns), or with slant ($F_{2,207} = 0.4927$, $p = 0.612$, ns). The main effects and interactions for texture type and slant were further explored by independent comparisons between pairs of texture types for each test slant, with a Bonferroni adjustment applied to the critical p-value to compensate for multiple tests ($p_{\text{critical}} = 0.0057$).

Relative to the Flat texture, which serves as a baseline, the Oriented textures were judged to be less slanted at 45° and 60° test slants, but not at the lowest test slant of 30° (Oriented 30°: $F_{1,27} = 2.594$, $p = 0.119$, ns; Oriented 45°: $F_{1,27} = 55.7$, $p < 0.001$; Oriented 60°: $F_{1,27} = 84.36$, $p < 0.001$). The Bumps textures were judged to be less slanted than Flat textures at 60° test slant, but not at 30° or 45° test slants (Bumps 60°: $F_{1,27} = 56.11$, $p < 0.001$; Bumps 45°: $F_{1,27} = 7.689$, $p = 0.01$, ns; Bumps 30°: $F_{1,27} = 2.142$, $p = 0.155$, ns).

In contrast to Oriented and Bumps conditions, judgments of slant in the Columns condition were not underestimated relative to the Flat condition. At the largest test slant, for which Bumps and Oriented textures both showed biases, there was no significant difference between the Columns and Flat textures (Columns 60°: $F_{1,27} = 1.172$, $p = 0.289$, ns). For the 45° test slant there was also no difference between Columns and Flat textures (Columns 45°: $F_{1,27} = 6.448$, $p = 0.017$, ns), while for the

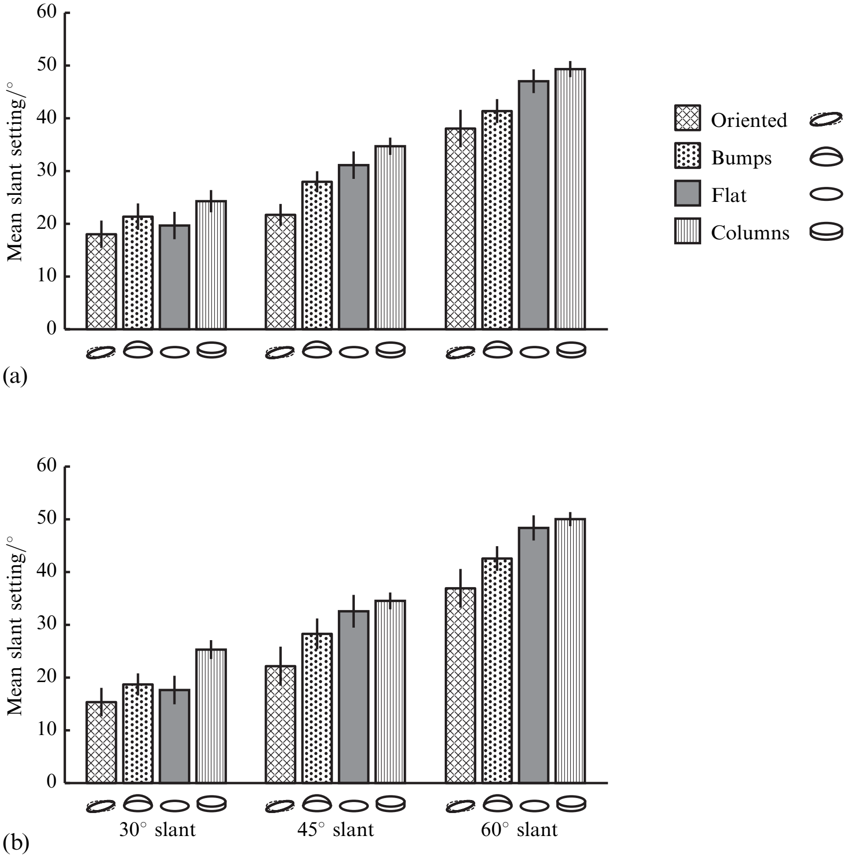

Figure 11. Results of experiment 1. The top graph (a) shows results for the more regular distributions of size and spacing (uniform distribution of sizes), and the bottom graph (b) shows results for irregular distributions of size and spacing (linear distribution of sizes). In each graph, the bars represent mean slant settings, averaged across subjects, for each of combination of texture type and test slant. The error bars show standard errors of the means.

30° test slant the mean judged slant for the Columns was significantly larger than for the Flat textures (Columns 30°: $F_{1,27} = 19.62$, $p < 0.001$).

Across texture conditions, the mean absolute slant settings of gauge figures were significantly less than the simulated slants of the texture stimuli ($t_{23} = -10.79$, $p < 0.001$). This implies that the textured test surfaces were perceived to be less slanted than gauge figures having the same simulated slant.

3.3 *Discussion*

The data in experiment 1 show that texture relief can influence perception of slant when other factors are equated, and that the effect of texture relief depends on the type of texture. Texture relief produced by raised bumps or variations in orientation both tended to produce underestimation of slant relative to Flat textures, while the texture relief present in Columns produced equal or larger slant estimates.

The relative underestimation of slant observed in the Bumps and Oriented textures can be explained by the effect of texture relief on the information provided by compression. In the case of Bumps textures, the main effect of relief is to reduce the amount of compression in an image. As described earlier, the aspect ratio of the projected contours of Bumps texture elements is larger (ie closer to circular) than that of flat elements, so use of compression for these textures would lead to underestimation of slant. Because the discrepancy between flat and raised textures increases with viewing angle, one would expect a larger difference in bias at higher slants, consistent with the data. Compression would also be affected, to a lesser extent, by the presence of successive occlusion in the Bumps textures. Occlusion is most prevalent at higher slants, so its effect on compression could contribute to explaining the biases observed for Bumps textures.

Variations in orientation, as in the Oriented textures, also alter the information provided by compression, but most directly by reducing the reliability of local compression rather than its accuracy. The observed effect of orientation variations was to reduce slant judgments. Considering only the information provided by texture, one might expect the orientation noise to correspondingly increase noise in slant judgments, rather than bias judgments. However, in the context of cue integration, a bias toward underestimation could be easily explained. In addition to texture cues, the stimuli provide conflicting information from accommodation, and lack normally occurring information from stereo and motion. The visual system may also have an a priori bias toward interpreting surfaces as frontal (Gogel and Tietz 1973). For any method of cue integration that takes reliability into account (eg Bayesian optimal estimation; see also Landy et al 1995), the noisier texture information in the Oriented textures would be less effective in counteracting the conflicting cues and/or prior assumptions specifying frontal orientation, resulting in greater underestimation of slant. Thus, the effect of texture relief on the reliability of compression can also explain the bias observed in the Oriented condition.

In contrast to the Bumps and Oriented conditions, textures from the Columns condition produced slant settings as large as in the Flat condition. The presence of texture relief in this case did not appear to interfere with perception of slant from texture. The Columns are the only one of the three textures with relief that retains reliable local compression information, because of the top faces of the texture elements are flat and coplanar with the surface, and are not occluded by nearer elements. Thus, the lack of interference is also consistent with contention that texture compression is a significant perceptual cue.

For 30° slant conditions, the judged slants for Columns textures were actually larger than that of Flat textures. One possible explanation is that observers use the scaling of perpendicular components (the sides of the columns) as additional information, which is not available for Flat textures. Because the informativeness of perpendicular elements is greatest at low slants, this interaction between texture type and slant is consistent with the hypothesis that the perpendicular components contribute to perception of slant. Another possibility is that, instead of being used as a direct slant cue, the perpendicular sides of columns somehow modulate the interpretation of the flat tops. For example, if compression were measured on the basis of anisotropy in the local power spectrum (as in Garding 1993), the resulting slant estimate would be higher for Columns textures than for matched Flat textures, because the darker ridges below the projections of the planar top faces would increase the asymmetry in the power distribution.[5] The present data cannot distinguish these qualitatively different ways that the column sides could influence perceived slant.

Another potential source of information available for textures with relief, but not flat textures, is successive occlusion. The results provide no evidence that this additional information is used, and some aspects of the observed effects suggest that it is not.

[5] Specifically, the dark ridges would selectively reduce power at low frequencies for horizontal components, but would reduce power uniformly across frequencies for vertical components. The distribution of power for horizontal components would therefore be shifted toward higher frequencies, while the distribution of power for vertical components would be unchanged, resulting in greater anisotropy in the power spectrum, which is consistent with higher slant.

First, if successive occlusion contributed to perceived slant, the expected beneficial effect would be greatest at high slants, since both the amount of occlusion and the proportion of overlapping texture elements increase with viewing angle. This prediction is contrary to the data, since the interaction between texture type and slant is in the opposite direction for both Bumps and Columns. Second, because the reliability of information from successive occlusion is highly dependent on the regularity of size and spacing of texture elements, one would expect the contribution of occlusion to be modulated by the amount of regularity. No such interaction was observed in the data: the effects of texture relief were unaffected by regularity for both Bumps and Columns. These findings suggest that the contribution of occlusion, if any, is weak compared to other factors influenced by the presence of relief.

Although the data provide no indication that successive occlusion is used as additional information, the results are consistent with successive occlusion having a negative effect. Occlusion would interfere primarily with extraction of slant from compression, and the expected adverse effects are similar to those caused by the projected shape of individual elements. For the Bumps and Oriented textures, in which slants were underestimated, successive occlusion obscures the projected shapes of some texture elements, potentially interfering with compression. For the Columns textures, which did not show biases toward underestimation, successive occlusion does not affect the projected shapes of the top faces, so accurate compression remains available. Moreover, in the case of Bumps textures, the amount of successive occlusion increases with slant, consistent with the increasing biases observed in the data. It is unlikely that successive occlusion is the sole cause of the effects, since the amount of occlusion present in the stimuli was limited. However, to the extent that the observed perceptual biases can be attributed to inaccurate compression, some part of these effects could be due to successive occlusion.

In addition to differences between texture conditions, there was also an overall bias in absolute slant settings. This indicates that when the perceived slant of texture stimuli and gauge figure are matched, the actual simulated slant of the gauge figure is on average lower than the simulated slant of the textured surface. This difference is likely due to the different depth information available for gauge and test surfaces. The gauge figure moved continuously during adjustment, so it provided additional depth information from motion that was not available for the stationary test surfaces. Differences in the monocular slant information available for test surfaces and gauge figures (eg texture versus contour, noise in texture) could also have contributed to the overall bias. Although there are reasons to expect that perceived slant of the test surfaces would be underestimated—the artificially generated stimuli lack normally occurring information from stereo, and contain conflicting information from accommodation—one does not know the extent that such factors also affect the perceived slant of the gauge figures. Because matching judgments involve the perceived slant of the gauge figures as well as the test surfaces, the absolute slant settings do not provide a direct measure of the accuracy of perceived slant from texture, and the overall bias cannot be interpreted as an underestimation relative to the simulated slant.

4 Experiment 2: Height of texture relief

The purpose of experiment 2 was to test the effect of texture relief for Bumps and Columns when the height of texture elements is varied. If compression contributes to perceived slant, one general prediction is that Bumps textures will appear less slanted than Flat textures. This effect was observed in experiment 1. An additional prediction is that the amount of underestimation should increase as a function of the height of the bumps. This prediction is tested in experiment 2.

Two variants of the Bumps textures were tested, which differed only in the height of texture relief: the Tall Bumps and Short Bumps conditions (figure 12). As a comparison, Columns textures with the same distributions of texture element heights were also tested: Tall Columns and Short Columns (figure 13).(6)

In experiment 1, the height of texture elements in the Bumps condition co-varied with the size of the planar cross section (ie the equivalent Flat-texture element), while in the Columns condition the height was constant. To more closely match the height distributions for the Bumps and Columns textures, the Columns textures in experiment 2 had elements with heights proportional to their widths.

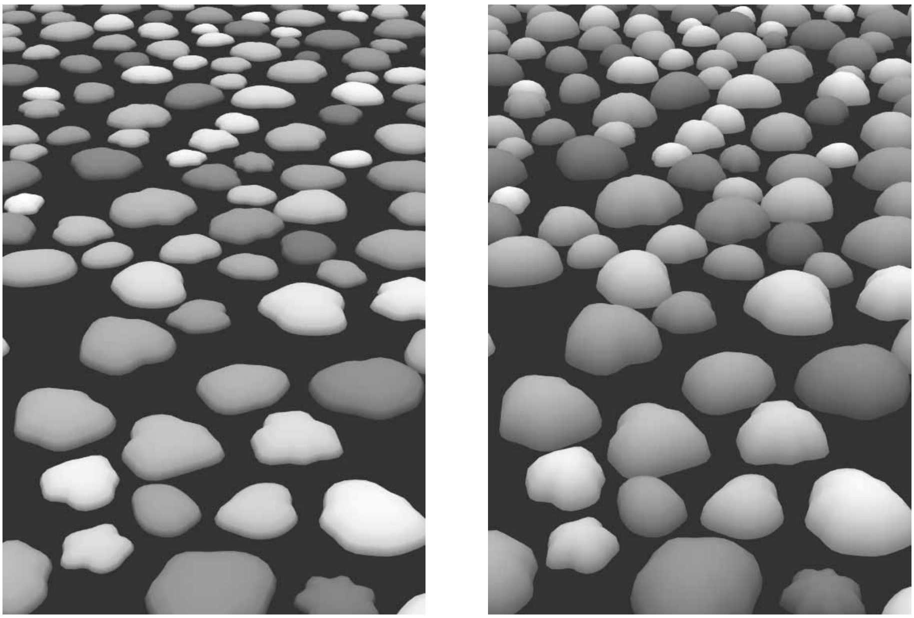

Figure 12. Short Bumps (left) and Tall Bumps (right) textures.

If the difference between Bumps and Columns textures observed in experiment 1 is due to their differing compression information, changing the height of texture should have different effects for the two types. For the Bumps textures, decreasing the height of relief increases the amount of compression (closer to that of Flat textures), which would be expected to increase perceived slant in the Short Bumps condition relative to the Tall Bumps condition. For the Columns textures, compression is due to the flat top faces, so on the basis of this cue one would expect no difference between the Short Columns and Tall Columns.

4.1 *Methods*

4.1.1 *Subjects.* Ten subjects participated in experiment 2. Subjects were recruited by signs posted around Brown University, and were paid for their participation. All were Brown University students, and were naïve to the purposes of the experiment. The data from one subject were found to have large random fluctuations, indicating that the subject was not successfully performing the slant-matching task. The data from this subject were excluded from the analysis.

(6) The term 'Tall' is somewhat of a misnomer, since the average height of elements in this condition is no bigger than in the previous experiment. However, this label is useful to clearly distinguish this condition from the Short version.

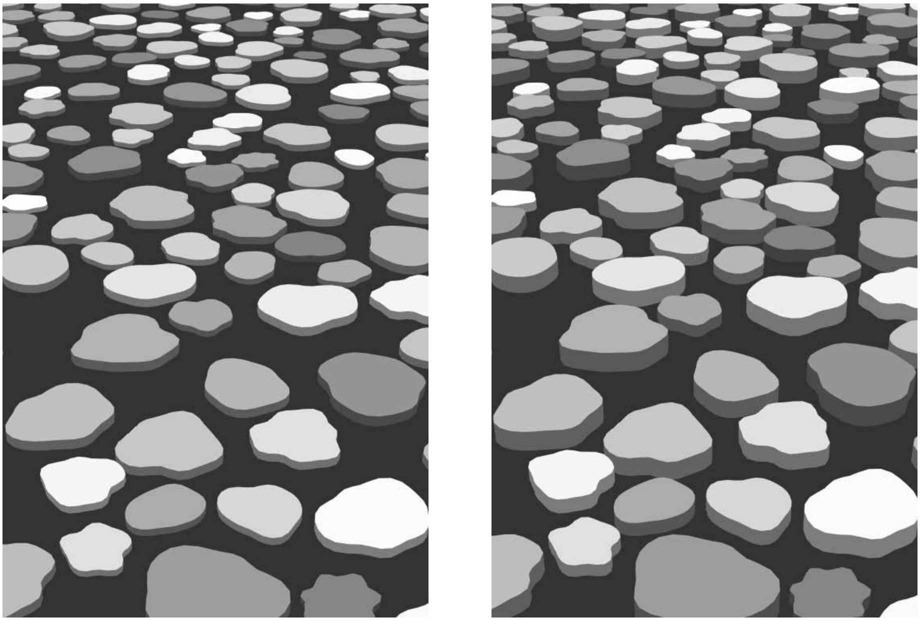

Figure 13. Short Columns (left) and Tall Columns (right) textures.

4.1.2 *Procedure.* Subjects performed slant judgments in two randomized blocks of 96 trials, which each contained images from all 12 of the Slant $\times$ Type $\times$ Height conditions, yielding 16 trials per condition. The experimental task and procedure were the same as in experiment 1.

4.1.3 *Stimuli.* The two Bumps textures were generated in the same way as in experiment 1: shape perturbations were added to ellipsoids, which varied in size but had height proportional to width. The heights of the base ellipsoids distinguished the Short Bumps and Tall Bumps. For the Tall Bumps textures, the height of an ellipsoid was equal to 35% of its width, which is the same as in the Bumps textures in experiment 1. For the Short Bumps, the height of a generating ellipsoid was half that of the Tall Bumps, equal to 17.5% of its width.

The two Columns textures were generated as before, by raising the top faces and adding perpendicular sides. However, for experiment 2 the heights of the Columns texture elements were proportional to their size. For the Tall Columns, the column heights were equal to 35% of their widths; while for Short Columns, column heights were half as large, equal to 17.5% of their widths.

The regularity of the distributions of positions and sizes was not manipulated, since this factor produced no effect in the previous experiment. The distributions of sizes and positions for all the stimuli in experiment 2 were identical to the more regular distributions used in experiment 1.

The slants used for the test surfaces were 45°, 52.5°, and 60°, which are larger than the slants used in experiment 1. Larger test slants were chosen because judgments in the 30° slant condition were much more noisy, and because the differences between texture types were more pronounced for 60° slant stimuli.

4.2 *Results*

In figure 14 mean slant settings are plotted for the data from experiment 2. As in the previous experiment, Columns textures were judged to be more slanted than Bumps textures. In addition, the slant settings depended on the height of the texture. An effect

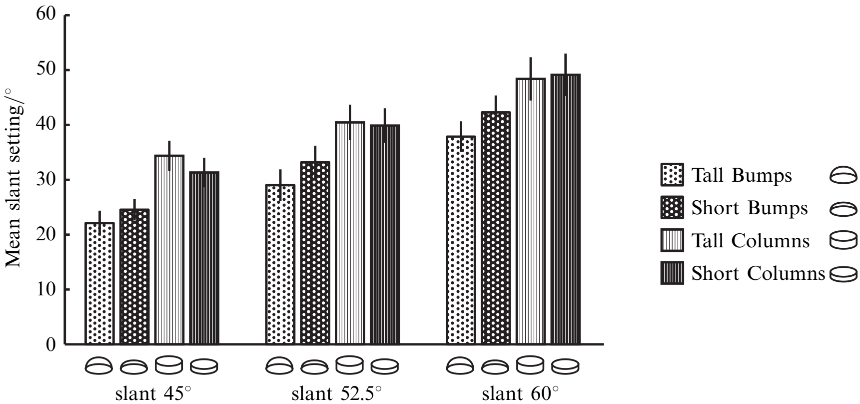

Figure 14. Results of experiment 2. The bars show mean slant settings, averaged across subjects, for each of combination of texture type and test slant. The error bars show standard errors of the means.

of texture height is most evident for the Bumps textures: slant judgments in the Short Bumps condition were higher than those in the Tall Bumps condition.

An ANOVA revealed main effects of texture type ($F_{1,88} = 116.8$, $p < 0.001$), and test slant ($F_{2,88} = 125.6$, $p < 0.001$). There was no main effect of texture height ($F_{1,88} = 2.594$, $p = 0.111$, ns), but there was a significant interaction between texture type and height ($F_{1,88} = 7.591$, $p < 0.007$). To explore the interaction between texture type and height, additional tests were performed with the Bumps and Columns conditions treated separately.

For the Bumps textures, there was a significant effect of texture height ($F_{1,40} = 13.44$, $p < 0.001$), but no interaction between height and test slant ($F_{2,40} = 0.381$, $p = 0.686$, ns), confirming that the Short Bumps produced an overall increase in slant judgments relative to the Tall Bumps. For the Columns textures, there was no main effect of texture height ($F_{1,40} = 0.8442$, $p = 0.3637$, ns), nor an interaction between height and test slant ($F_{2,40} = 1.132$, $p = 0.3328$, ns).

4.3 *Discussion*

The results of experiment 2 demonstrate that the effect of texture relief on perceived slant can vary depending on the height of texture relief, and that the effect of texture height depends on the type of texture relief. For the Bumps textures, reducing the height of texture relief led to an increase in slant judgments, while for Columns textures there was no effect. There was also an overall difference between texture types, with Columns textures judged to be more slanted than Bumps textures, which replicates one part of the findings from experiment 1.

The effect of texture height in the Bumps conditions is in the direction that would be expected on the basis of the amount of compression in the projected shapes. Short Bumps produce more compression than Tall Bumps, since they are closer to being flat, so, if compression were used, one would expect larger perceived slant for the Short Bumps, as was observed in the data. Thus, the difference between judgments in the Tall Bumps and Short Bumps conditions provides further support for the hypothesis that compression contributes to perceived slant from texture.

Texture height had no effect on judgments for the Columns textures, which is also consistent with use of compression. The top faces of the Columns provide the same compression information regardless of height, so on the basis of this cue no difference would be expected.

The results of experiment 1 suggested that the sides of the Columns could have a facilitatory effect: Columns textures had the same compression, size, and spacing information as Flat textures, but were judged to be more slanted than Flat textures at low test slant. If the scaling of perpendicular sides of Columns were used as information, one might have expected a similar difference between Short Columns and Tall Columns in experiment 2, because the Short Columns are closer to being flat and have less salient perpendicular extents. However, this effect was not observed in the data. There was a trend in this direction at 45° slant, so it is possible that a difference between Short Columns and Tall Columns would be observed if lower slants were tested. It is also possible that the scaling gradients produced by short and tall perpendicular extents are equally informative, in which case no difference would be observed even if perpendicular scaling did contribute. A different possible explanation is that the column sides modulate or influence use of other cues, rather than directly contributing as a gradient cue, in which case height might not be an important factor.

5 General discussion

The results of the two experiments provide evidence that texture relief affects perception of slant from texture, and the pattern of biases suggests that this is mainly due to the way in which texture relief alters the amount of compression in projected views. Greater underestimation of slant was observed for the types of relief that affect compression, Bumps and Oriented textures, but not in the case of the Columns textures, for which compression remains accurate. For Bumps textures, increasing the height of texture relief decreases the amount of compression, and the results of experiment 2 show that slant judgments for Bumps textures decrease with height. In the case of Columns textures, for which height does not affect compression, increasing height had no detrimental effect. These aspects of the data can all be explained in terms of the effect of texture relief on compression. Size and spacing, the other separable components of a texture gradient, cannot account for differences between texture types, since by design they were matched across stimuli.

In previous experiments with flat textures it was found that varying the information provided by compression produces biases in perceived slant (Rosenholtz and Malik 1997), and that compression tends to dominate when size and density information conflict (Buckley et al 1996; Knill 1998b). The results reported here are consistent with these previous findings. In contrast to methods used previously, the experiments reported here alter the information available from compression without explicitly introducing conflicting cues. The amount of compression was manipulated by the presence of texture relief, which can occur in many examples of real-world textures, and under generic viewing conditions. Another distinguishing factor is that all of the textures were isotropic. Information from compression was manipulated by varying relief rather than by varying texture anisotropy, whereas in previous experiments these factors were confounded. Thus, an additional implication of these results is that compression remains a dominant factor even when textures are isotropic and no explicit cue conflicts are present.

While these results provide further evidence that compression plays a role in perception of slant from texture, they do not rule out a contribution from other cues. The addition of texture relief did influence slant judgments in ways predicted by the effect of texture relief on compression. But even in 'bad' texture-relief conditions, a strong percept of slant remained, and judgments varied reliably as a function of stimulus slant. This residual ability to perceive slant from texture could be due to use of other cues, or from the degraded information from compression. Studies of textures without relief have found that, although compression is a dominant factor, other texture cues also contribute (Cumming et al 1993; Rosenholtz and Malik 1997; Knill 1998b).

It is likely that these other cues would similarly contribute for the textures with relief used here.

I have focused on the effect of texture relief on local compression, because compression has been implicated in previous studies of slant from texture and because it is selectively degraded by some types of texture relief. But texture relief has the potential to add information as well as interfere with texture cues. One aspect of the results suggests that texture relief can indeed contribute information: for small test slants, Columns textures were judged to be more slanted than Flat textures. If these textures were analyzed solely in terms of the local compression of the planar faces, there would be no difference between these texture types. Thus, the effect of texture relief in the Columns condition must be mediated by some other factor. The sides of the Columns could provide additional information, distinct from processing of compression of the top faces, or the sides might modulate the interpretation of the top faces (eg local spectral analysis). Successive occlusion is another factor, specific to textures with relief, that could potentially add information. The results provide no evidence that occlusion contributes to perception of slant, though it could have had a negative effect, by means of interfering with compression.

A motivation for studying textures with relief is that many natural textures have some 3-D structure, which is inconsistent with the traditional conception of texture as a 2-D pattern mapped onto a surface. Although the artificial stimuli used here did have relief, they had some other properties that are unrealistic. One simplification is that the textures were composed of individual elements, which had clearly identifiable position and shape in the projected images. I have also neglected effects of lighting such as cast shadows and specularities, which can interact with 3-D structure to produce complex effects (Dana et al 1999). So there remains a question whether the results reported here would generalize to more complex texture gradients produced by real-world textures with relief. In an unpublished study, I have tested slant judgments for stimuli based on photographs of natural textures with relief, and found similar results (Saunders 1999), but these findings remain preliminary.

In conclusion, perception of slant from texture is not robust to the presence of texture relief. When texture relief alters the availability of compression as an accurate cue for orientation, systematic biases are observed. In other cases, the vertical components of texture relief may contribute additional information.

References

Aloimonos J, Swain M J, 1988 "Shape from patterns: regularization" *International Journal of Computer Vision* **2** 171 – 187

Beck J, 1960 "Texture gradients and judgments of slant and recession" *American Journal of Psychology* **73** 411 – 416

Blake A, Bülthoff H H, Sheinberg D, 1993 "Shape from texture: ideal observers and human psychophysics" *Vision Research* **33** 1723 – 1737

Buckley D, Frisby J, Blake A, 1996 "Does the human visual system implement an ideal observer theory of slant from texture?" *Vision Research* **36** 1163 – 1176

Cumming B G, Johnston E B, Parker A J, 1993 "Effects of different texture cues on curved surfaces viewed stereoscopically" *Vision Research* **33** 827 – 838

Cutting J E, Millard R, 1984 "Three gradients and the perception of flat and curved surfaces" *Journal of Experimental Psychology: General* **113** 198 – 216

Dana K J, Ginneken B van, Nayar S K, Koenderink J J, 1999 "Reflectance and texture of real-world surfaces" *ACM Transactions on Graphics* **18**(1) 1 – 34

Garding J, 1992 "Shape from texture for smooth curved surfaces in perspective projection" *Journal of Mathematical Imaging and Vision* **2** 329 – 352

Garding J, 1993 "Shape from texture from texture and contour by weak isotropy" *Journal of Artificial Intelligence* **64** 243 – 297

Gibson J J, 1950 *The Perception of the Visual World* (Boston, MA: Houghton Mifflin)

Gogel W C, Tietz J D, 1973 "Absolute motion parallax and the specific distance tendency" *Perception & Psychophysics* **13** 284 – 292

Ikeuchi K, 1984 "Shape from regular patterns" *Journal of Artificial Intelligence* **22** 49 – 75
Jacobs R A, 1999 "Optimal integration of texture and motion cues to depth" *Vision Research* **39** 3621 – 3629
Jacobs R A, Fine I, 1999 "Experience-dependent integration of texture and motion cues to depth" *Vision Research* **39** 4062 – 4075
Johnston E B, Cumming B G, Parker A J, 1993 "Integration of depth modules: stereopsis and texture" *Vision Research* **33** 813 – 826
Knill D C, 1998a "Discrimination of planar surface slant from texture: human and ideal observers compared" *Vision Research* **38** 1683 – 1711
Knill D C, 1998b "Ideal observer perturbation analysis reveals human strategies for inferring surface orientation from texture" *Vision Research* **38** 2635 – 2656
Knill D C, 1998c "Surface orientation from texture: ideal observers, generic observers, and the information content of texture cues" *Vision Research* **38** 1655 – 1682
Knill D C, 2001 "Contour into texture: information content of surface contours and texture flow" *Journal of the Optical Society of America A: Optics, Image Science, and Vision* **18** 12 – 35
Landy M, Maloney L, Johnston E, Young M, 1995 "Measurement and modeling of depth cue combination: in defense of weak fusion" *Vision Research* **35** 389 – 412
Leung T, Malik J, 1997 "Perpendicular texture: Why do we see more flowers in the distance?" *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Puerto Rico* (New York: IEEE Press)
Malik J, Rosenholtz R, 1997 "Computing local surface orientation and shape from texture for curved surfaces" *International Journal of Computer Vision* **23** 149 – 168
O'Brien J, Johnston A, 2000 "When texture takes precedence over motion in depth perception" *Perception* **29** 437 – 452
Passmore P J, Johnston A, 1995 "Human discrimination of surface slant in fractal and related textured images" *Spatial Vision* **9** 151 – 161
Rosenholtz R, Malik J, 1997 "Surface orientation from texture: isotropy or homogeneity (or both)?" *Vision Research* **37** 2283 – 2293
Saunders J A, 1999 "The effect of texture relief on perceived slant" *Investigative Ophthalmology & Visual Science* **40**(4) S344
Stevens K A, 1981a "The information content of texture gradients" *Biological Cybernetics* **42** 95 – 105
Stevens K A, 1981b "The visual interpretation of surface contours" *Artificial Intelligence* **17** 47 – 73
Stevens K A, 1983 "Slant-tilt: The visual encoding of surface orientation" *Biological Cybernetics* **46** 183 – 195
Stoyan D, Stoyan H, 1994 *Fractals, Random Shapes and Point Fields* (Chichester, Sussex: John Wiley)
Super B J, Bovik A C, 1995 "Planar surface orientation from texture spatial frequencies" *Pattern Recognition* **28** 729 – 743
Todd J T, Akerstrom R A, 1987 "Perception of three-dimensional form from patterns of optical texture" *Journal of Experimental Psychology: Human Perception and Performance* **13** 242 – 255
Turner M R, Gerstein G L, Bajcsy R, 1991 "Underestimation of visual texture slant by human observers: a model" *Biological Cybernetics* **65** 215 – 226
Witkin A P, 1981 "Recovering surface shape and orientation from texture" *Journal of Artificial Intelligence* **17** 17 – 45

ISSN 0301-0066 (print) ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 32 2003

www.perceptionweb.com