

Multiple texture cues are integrated for perception of 3D slant from texture

Zhongting Chen

Shanghai Key Laboratory of Brain Functional Genomics,
Shanghai Changning-ECNU Mental Health Center,
School of Psychology and Cognitive Science,
East China Normal University, Shanghai, China



Jeffrey A. Saunders

Department of Psychology, The University of Hong Kong,
Hong Kong, Hong Kong



The projected image of a textured surface contains multiple texture cues to three-dimensional (3D) surface orientation. Previous studies have reported conflicting findings about the roles of various texture cues. We tested the influence of texture compression relative to other texture cues using a cue conflict paradigm. Observers viewed images of textured planar surfaces with varied slants (0° – 70°) and estimated 3D slant by aligning their hand with the virtual surface. Conflicts between texture cues were created by stretching or compression the texture along the surface, which selectively changes the slant specified by texture compression. The texture distortions were relatively small ($\pm 10\%$ or $\pm 20\%$) to limit the size of the cue conflicts. Across three experiments, we varied the field of view (10° vs. 20°), texture regularity (circles vs. Voronoi), and availability of binocular cues. In monocular conditions, slant estimates were strongly affected by texture distortions. Analyses of cue weighting found that texture compression had more influence on slant settings than other texture cues and the relative influence of texture compression decreased with larger field of view and less regular textures. In binocular conditions, we also observed effects of texture distortion, and the influence of texture compression relative to information from stereo and other texture cues increased with slant. Our results provide evidence that texture compression contributes to perceived slant, in addition to other texture cues such as texture scaling. The observed effects of simulated slant, field of view, and texture regularity on cue weighting were all consistent with a model that integrates multiple sources of information according to their reliability.

Introduction

Texture gradients provide monocular information about three-dimensional (3D) structure that is utilized by the visual system. When a surface with a regular texture is viewed in perspective, the visual texture in the projected image is systematically distorted and scaled as a function of the 3D orientation and distance of surface regions relative to the observer, thereby providing multiple cues for the 3D surface structure. Many studies have investigated the use of monocular texture information to perceive the 3D slant of planar surfaces (Braunstein & Payne, 1969; Buckley, Frisby, & Blake, 1996; Chen & Saunders, 2019; Cutting & Millard, 1984; Knill, 1998a; Knill, 1998b; Norman, Crabtree, Bartholomew, & Ferrell, 2009; Rosas, Wichmann, & Wagemans, 2004; Saunders, 2003; Saunders & Chen, 2015; Tam, Shin, & Li, 2013; Tibau, Willems, van den Bergh, & Wagemans, 2001; Todd, Christensen, & Guckes, 2010; Todd, Thaler, & Dijkstra, 2005). Slant perception from texture and stereo cues has also been used as a test case for investigating 3D cue integration (Hillis et al., 2004; Knill & Saunders, 2003; Saunders & Chen, 2015). Results from these studies demonstrate that texture contributes to perceive 3D slant in both monocular and binocular conditions.

Although it is clear that texture information can be used to perceive 3D slant, there remains some debate about how texture is used. The texture gradient in a perspective view of a planar textured surface provides multiple cues that could be used to infer the 3D slant of the surface. Figure 1 illustrates two potential cues:

Citation: Chen, Z., & Saunders, J. A. (2020). Multiple texture cues are integrated for perception of 3D slant from texture. *Journal of Vision*, 20(7):14, 1–23, <https://doi.org/10.1167/jov.20.7.14>.

<https://doi.org/10.1167/jov.20.7.14>

Received June 3, 2019; published July 17, 2020

ISSN 1534-7362 Copyright 2020 The Authors






Figure 1. Illustration of texture scaling and compression cues for 3D slant. The left image shows a simulated perspective view of a slanted planar surface with texture composed of scattered dots. The middle column illustrates the texture compression cue. Foreshortening compresses the projected texture in the direction of surface tilt, which increases the spatial frequencies of the image texture in the tilt direction relative to the perpendicular direction. If a surface texture is isotropic, the ratio of spatial frequencies in a region of the image provides a cue for local surface slant. The right column illustrates the texture scaling cue. The small circular regions show texture from the top and bottom of the image, and the adjacent rectangles plot their spatial frequency spectra. Distance to the surface increases across the image in the tilt direction (vertical), causing an increase in the spatial frequencies of the image texture from the bottom to the top. The proportional rate of change in spatial frequency is a function of surface slant.

(1) compression of the image texture in the vertical direction relative to the horizontal direction due to foreshortening (texture compression), and (2) change in the size and spacing of projected texture across the image due to differences in relative depth (texture scaling). As discussed later, there is previous evidence that both texture compression and texture scaling can contribute to slant perception, but there have been conflicting findings about their relative contributions.

Texture compression and scaling cues

One texture cue for slant perception is local compression of image texture due to foreshortening. When a surface patch with isotropic texture is viewed from an angle, the projected texture becomes foreshortened in the tilt direction, causing the spatial frequencies in the tilt direction to increase by a factor of $1/\cos(S)$. In the example shown in Figure 1, the spatial frequencies in the vertical direction are higher than in the horizontal direction (middle panel). If texture is assumed to be isotropic, the anisotropy in the projected texture can be used to compute the local 3D orientation relative to the viewer. This computation could be applied at multiple locations in the visual field to derive a map of local 3D orientations.

Analysis of texture compression based on an assumption of isotropy has both advantages and disadvantages. Because texture compression can be analyzed in local image patches, this approach could be used to derive a map of local 3D orientations across the

visual field. Such a map would provide rich information that could be used to determine the slant of planar surfaces or the shape of curved surfaces. In the case of a planar surface, the local 3D orientations vary across an image in a predictable way, so the information from multiple regions could be integrated to better estimate geographic slant. A disadvantage of texture analysis based on an assumption of isotropy is that many real-world surface textures are not isotropic (e.g., bricks). In such cases, the isotropic interpretation of texture compression would result in biased estimates of local 3D orientation.

Another texture cue is the change in scaling of image texture across an image due to variations in distance relative to the observer. If texture is uniform across a surface (homogeneous), the scaling of the projected texture in the image varies as an inverse function of distance. In the example shown in Figure 1, the spatial frequencies are higher at the top of the image, which is farther from the viewer, than at the bottom of the image, which is closer. If texture is assumed to be homogeneous, this gradient of texture scaling can be used to recover the slant or shape of a 3D surface.

The gradient of texture compression provides an additional independent cue to 3D slant and shape under an assumption of homogeneity, but this cue is ambiguous by itself (Gårding, 1992). For a slanted planar surface with homogeneous texture, the local compression changes across an image in a systematic way, which could potentially be used as additional information to recover slant. However, Gårding (1992) showed that a compression gradient generated by a

planar surface can also be generated by a family of curved surfaces. Although a compression gradient could contribute to estimation of 3D slant, additional assumptions or other information would be required to resolve its ambiguity.

Field of view (FOV) has different effects on the information provided by texture cues based on homogeneity and isotropy assumptions. For texture cues based on an assumption of homogeneity, scaling, and compression gradients, the reliability of slant information would be highly dependent on the FOV. A larger FOV results in larger changes in scaling or compression across the visible region, which would increase the signal-to-noise ratio. For texture compression, which can be analyzed locally based on an assumption of isotropy, a larger FOV could improve reliability in a less direct manner. For a planar surface, the local slant estimates based on texture compression could be integrated to get a better estimate of the overall slant; however, this would require an assumption that the surface is planar, which is not always true. Furthermore, ideal observer analysis suggests that the benefit from spatial integration of local compression would be limited (Knill, 1998a). Although FOV would strongly affect the reliability of 3D slant from texture scaling or other gradient cues, the effect of FOV on the reliability of slant from local compression would depend on how the local information is integrated and would likely be smaller.

Scaling contrast model

Todd and colleagues have proposed that perceived slant from texture is based on an image-based measure related to the scaling gradient, which they term scaling contrast (Todd, Thaler, Dijkstra, Koenderink, & Kappers, 2007). Scaling contrast is the proportional change in texture scale across a region of the image, measured in the orthogonal direction. For the example in Figure 1, texture scaling would be the difference in horizontal size of texture at the top and bottom of the image divided by the average horizontal size. Based on empirical observations, Todd et al. (2007) proposed that perceived slant from texture is a direct function of the magnitude of scaling contrast across an image.

The scaling contrast model could potentially explain some systematic biases that have been observed in human estimates of slant and shape from texture. Estimates of slant from texture for planar surface are generally nonlinear functions of the depicted slant, with proportionally more underestimation of slant for surfaces close to the frontal plane. Slant estimates also tend to increase when a surface with a given slant is viewed with a larger FOV. Consistent with these observations, scaling contrast is a nonlinear function of slant and increases with FOV. Todd and colleagues

(Todd et al., 2005; Todd et al., 2007; Todd et al., 2010) observed that scaling contrast was strongly correlated with slant and shape estimates across a wide range of texture and FOV conditions and that it was a better predictor than other image-based measures.

The scaling contrast model, however, does not include a role for texture compression, which has been observed to affect slant perception in other studies (see next section). The experiments reported here provide an additional empirical test of the predictions of the scaling contrast model.

Relative contribution of compression and scaling

The relative contributions of texture compression and scaling have been investigated using stimuli that provide conflicting information from these two texture cues. The information from local texture compression under an assumption of isotropy can be manipulated independently from the information from gradient cues, based on an assumption of homogeneity, by adding various amounts of anisotropy to surface textures. If a texture were compressed or stretched in the tilt direction, this would change the slant specified by local compression but have no effect on the scaling gradient or compression gradient. An equivalent manipulation is to present a perspective image of a textured surface that was rendered with an incorrect FOV or center of projection. For example, Todd et al. (2010) tested conditions in which images presented with a 15° FOV were rendered with FOVs that varied between 15° and 60°. When the rendered FOV is larger than the presented FOV, the resulting image is an accurate perspective view of a surface with larger slant and texture that is stretched in the tilt direction. Conversely, when the rendered FOV is smaller than the presented FOV, it corresponds to a view of a surface with lower slant and texture that is compressed in the tilt direction. In cases where images correspond to a slanted surface with texture that is stretched or compressed, the slant specified by local compression is different than the slant specified by gradient cues.

Studies that have tested slant perception with conflicting texture cues have found varied results. Some have found that texture compression had a large influence on perceived slant from texture conditions with small cue conflicts and that texture scaling also contributed but had a smaller influence (Buckley et al., 1996; Knill, 1998a; Rosenholtz & Malik, 1997). Saunders (2003) observed effects of texture relief on slant estimates that were also consistent with the use of texture compression. However, other studies have observed little or no influence of texture compression in conditions with conflicting texture cues (Braunstein & Payne, 1969; Cutting & Millard, 1984;

Todd et al., 2005; Todd et al., 2007; Todd et al., 2010). Thus, there appears to be conflicting findings concerning the relative contributions of texture compression and texture scaling in slant perception.

Todd et al. (2010) suggested that the conflicting findings might be a consequence of differences in experimental tasks. Specifically, they suggested that the two-alternative forced choice discrimination task in Knill (1998a) and Knill and Saunders (2003) might have encouraged observers to directly compare the two-dimensional (2D) properties of image texture rather than comparing the perceived slant of surfaces. The adjustment tasks used by Rosenholtz and Malik (1997) and Saunders (2003) could have a similar problem. In these studies, subjects adjusted the slant of a circular gauge figure to match the perceived slant of texture surfaces, which might have encouraged them to match the 2D compression of the gauge figure and the image texture. In contrast, the tasks used in studies by Todd and colleagues (Todd et al., 2005; Todd et al., 2007; Todd et al., 2010) involved matching perceived 3D slant or shape to a side-view depiction of the surface orientation or shape. The differences in methods might explain why the earlier studies observed a strong influence of compression, whereas the studies by Todd and colleagues using estimation tasks did not.

Another possible reason for the different findings across previous studies is the size of the conflicts between texture cues. The previous studies that found contributions of both texture compression and scaling used relatively small cue conflicts between these cues (Buckley et al., 1996; Knill, 1998a; Rosenholtz & Malik, 1997), and the studies that found minimal effects of texture compression used much larger conflicts (Cutting and Millard, 1984; Todd et al., 2005; Todd et al., 2007; Todd et al., 2010). In Todd et al. (2010), for example, the difference between the slants specified by texture was between 10° and 50°, and in some cases the local texture compression specified an orthogonal tilt direction. When different cues have large conflicts, the visual system might use a strategy of relying on only one cue and “vetoing” the other rather than combining the information from the cues (Girshick & Banks, 2009; Landy, Maloney, Johnston, & Young, 1995). Previous studies of perceived slant from stereo and texture cues have found that these cues are integrated in a close to optimal manner when cue conflicts are small (Hillis et al., 2004; Knill & Saunders, 2003; Saunders & Chen, 2015), but cue vetoing can occur with large conflicts between stereo and texture (Girshick & Banks, 2009). Similarly, large conflicts between the slants specified by texture scaling and texture compression might encourage the visual system to rely solely on texture scaling. Texture compression may have more influence in situations when there is less discrepancy between texture cues.

Present study

The current study investigated the relative contributions of texture compression and scaling to perception of slant from texture using methodology designed to deal with these potential issues. We measured perceived slant in conditions with consistent and conflicting texture cues using a cross-modal slant estimation task, and we limited the size of the conflicts between cues. The slant specified by texture compression was independently varied by stretching or compressing the surface texture in the tilt direction. The amount of compressing or stretching was $\pm 10\%$ or $\pm 20\%$, which produces smaller conflicts between cues than in the conditions of Todd and colleagues. The smaller cue conflicts would be less likely to produce cue vetoing, and the adjustment task would encourage responses based on 3D perception.

We also manipulated factors that affect the reliability of texture compression and scaling cues: texture regularity and FOV. Figure 2 shows examples of stimuli. Both the circles and Voronoi textures are homogeneous and isotropic, but the texture elements in the Voronoi textures have variable sizes ($\pm 17\%$ *SD*) and aspect ratios ($\pm 26\%$ *SD*). The variability in aspect ratios might especially interfere with use of texture compression, which involves comparing spatial frequencies at different orientations within local image regions. FOV also affects the reliability of texture cues but would be expected to primarily affect the slant information provided by texture scaling or other gradient cues, as discussed previously. Thus, texture regularity and FOV differentially affect the slant information from texture scaling and texture compression. If both texture cues were used for slant perception and integrated in an approximately optimal manner, one would expect less use of texture compression for the Voronoi textures than for the circles textures, and less use of texture compression with a larger FOV.

Manipulation of FOV also provides a way to test the predictions of the scaling contrast model proposed by Todd et al. (2007). Scaling contrast is highly dependent on FOV. For the case of a planar surface, scaling contrast is equal to $\tan(\text{slant}) \times \tan(\text{FOV}/2)$. If perceived slant is proportional to scaling contrast, then increasing the FOV from 10° to 20° should approximately double the perceived slant. This is a larger effect than would be expected if FOV influenced perceived slant from texture in other ways, such as by increasing the reliability of texture information.

We conducted three experiments for the current study. In Experiments 1 and 2, stimuli were monocular images of textured surfaces with and without conflicts between texture compression and texture scaling, and



Exp 1 - Circle textures**Exp 2 - Voronoi textures**

Figure 2. Sample stimuli with circles textures (Experiment 1) and Voronoi textures (Experiment 2). The images are perspective views of planar surfaces mapped with homogeneous textures. The rows correspond to different combinations of texture type and FOV and two different simulated slants. In each row, the simulated slant is the same, but the surface texture has been stretched or compressed in the vertical direction before applying the 3D rotation and perspective projection. This alters the local texture compression but does not affect the change in texture scaling across the image or other gradient cues.

we measured the relative contributions of these cues to perceived slant for different texture types and FOVs. Texture compression had a strong influence in all conditions. In [Experiment 3](#), we manipulated texture compression in binocular images of textured surfaces to test whether texture compression still affects slant perception when richer 3D information from stereo is available.

Experiment 1

[Experiment 1](#) manipulated information from texture compression and scaling by presenting slanted surfaces with textures that were stretched or compressed along the surface in the tilt direction. This manipulation changes the slant specified by the texture compression cue but does not affect the gradient of texture scaling or the overall scaling contrast. Observers estimated slant in these cue conflict conditions, as well as consistent cue conditions with isotropic textures, and the combined data were used to estimate the relative influence of texture cues. To the extent that texture compression is used to perceive slant, stretching and compressing the surface texture would be expected to systematically bias estimates of slant, whereas use of other texture cues predicts no effect. We also varied FOV (10° vs. 20°), which greatly changes the amount of scaling contrast across a projected image but has limited effect on the information provided by texture compression.

Methods

Participants

Twelve adults from the University of Hong Kong (five males and seven females; mean age, 22.0 ± 3.1 years) were paid to participate in [Experiment 1](#). Subjects were required to be right-handed (for our hand tracking apparatus), have normal or corrected-to-normal visual acuity, and show no impairment in a screening test for stereo acuity. All subjects were naïve as to the purpose of the study and were paid for their participation. The procedures were approved by and conformed to the standards of the Human Research Ethics Committee for Non-Clinical Faculties.

Apparatus and stimuli

[Figure 3](#) illustrates the apparatus and task. Subjects viewed computer-generated perspective images of slanted planar surfaces presented on a liquid crystal display that was visible through a circular aperture on a black masking board. Two masking boards with different size apertures were used to constrain the FOV to be 10° or 20° . The display was viewed from a chin rest at a distance of 80 cm, with the masking board

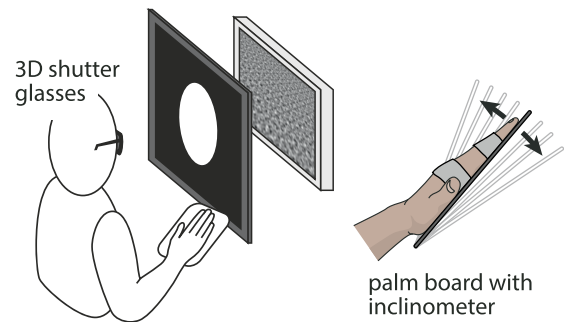


Figure 3. Illustration of the slant estimation task. Subjects viewed simulated planar surfaces with textures and adjusted their hand to align with the surface. A palm board with an inclinometer was attached to the subject's right hand, which they could otherwise move freely. The display was viewed through a round aperture that constrained the field of view to 10° or 20° . Stereo shutter glasses were used to present different images to the left and right eyes in binocular conditions. In monocular conditions, subjects wore an eye patch over their non-dominant eye.

placed 20 cm in front of the monitor. The VG278H monitor (ASUS, Taipei, Taiwan) had a 59.2-cm \times 33.6-cm viewable region, 1920×1080 pixel resolution, and 120-Hz refresh rate in stereo mode. Images were rendered with OpenGL using a Quadro 600 graphics card (NVIDIA, Santa Clara, CA) and were antialiased with subpixel resolution. Subjects wore an eye patch over their non-dominant eye to create monocular viewing conditions. Although displays were monocular, subjects wore NVIDIA 3D shutter glasses over the eye patch so that stimulus presentation would be compatible with the binocular conditions of [Experiment 3](#).

For the consistent cue stimuli of [Experiment 1](#), the simulated surface textures were scattered white circles of constant size on a dark gray background. The positions of the circles were determined by sequentially generating random positions and rejecting any circles that were less than a minimal distance away from any previously generated circle (10% of circle diameter). We randomly varied the overall size of surface texture across trials to discourage subjects from judging slants based on the projected size of the circles at the bottom or top of the images. In the unscaled version of the textures, the diameters of the circles were 2.4 cm, corresponding to 1.7° in a frontal view. On each trial, the texture was uniformly scaled by a random amount between 0.875 and 1.125.

To create cue conflict stimuli with inconsistent slant information from texture compression and scaling, we stretched or compressed textures by 10% or 20% in the vertical direction and then rendered accurate perspective images of the distorted texture at the simulated base slant. Examples are shown in [Figure 2](#). The image compression at the center of the stimuli in cue conflict conditions specified a surface slant of

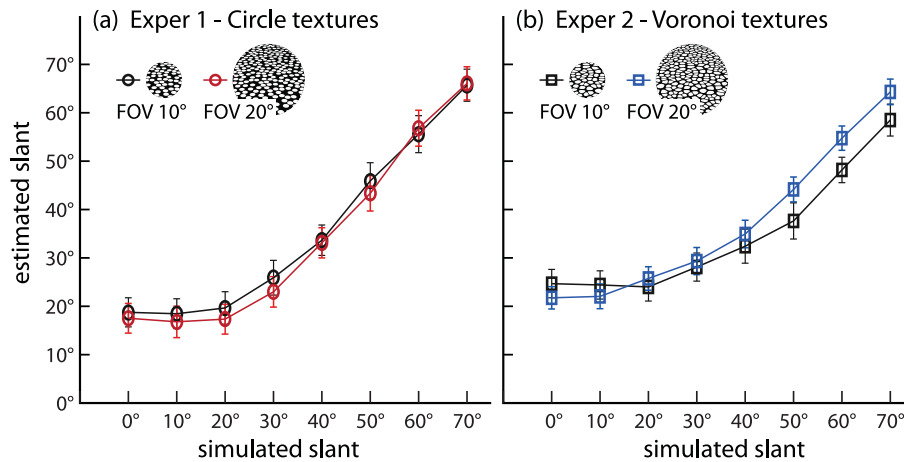


Figure 4. Mean slant estimates as a function of simulated slants in the consistent cue conditions of [Experiment 1](#) with circles textures (left) and [Experiment 2](#) with Voronoi textures (right). The two sets of data on each graph correspond to the 10° and 20° FOV conditions. Error bars depict ± 1 SE.

$S_{\text{comp}} = \arccos(\alpha \cos(S))$, where S is the simulated slant and α is the stretching ratio of the surface texture. For example, when simulated slant was 40° and the texture was compressed by 10% ($\alpha = 0.9$), local compression at the center of the stimulus was consistent with a slant of 46.4°.

Simulated surfaces were slanted around a horizontal axis (i.e., receding in the vertical direction) by 0° to 70° relative to the frontal plane. For consistent cue conditions, the simulated slants were sampled in 10° steps. For cue conflict conditions, we used a subset of base slants: 0°, 20°, 40°, and 60°.

To record the orientation of a subject's hand during responses, we used a DOG2 MEMS Series Inclinometer (TE Connectivity, Schaffhausen, Switzerland) attached to a rigid board that was worn on the palm of a subject's right hand. The inclinometer recorded 3D slant and tilt at a rate of 100 Hz. The palm board was lightweight and detached from any other surfaces, so the subject could move their hand freely.

Procedure

The task of a subject was to align the palm of their right hand with a simulated slanted surface while wearing a flat board with an inclinometer attached to their palm. Trials began with the subject's hand resting on the table with fingers pointed forward. Subjects pressed the spacebar on a keyboard with their left hand to trigger the presentation of a simulated surface, then lifted their right hand and aligned it with the perceived orientation of the surface. They were instructed to keep their right hand at their side with fingers pointing forward and to adjust their palm orientation only in the y - z plane (i.e., vertical tilt direction). When finished with the alignment, subjects held their hand still and pressed the spacebar again with their left hand to conclude the

trial. If the tilt of the palm was more than 5° away from vertical, a warning beep would cue observers to readjust their hand and confirm again. Trials were self-paced and typically took 1 to 2 seconds to complete.

The experiment was conducted over two sessions on separate days, each with three blocks of 144 trials. The two FOV conditions (10° and 20°) were tested in separate sessions, with order counterbalanced across subjects. The other experimental factors—simulated slant and stretching/compression—were fully randomized across trials in a session. Each subject performed a total of 864 trials, which included 18 repetitions of each condition.

Results

Consistent texture cues

The slant estimates in consistent cue conditions (no stretching or compression) varied as a function of simulated slant in a nonlinear manner, as expected from previous findings. [Figure 4a](#) plots mean slant estimates, averaged across subjects, as a function of simulated slant for the 10° and 20° FOV conditions. When simulated slants were low ($\leq 20^\circ$), slant estimates showed a large bias toward frontal, whereas at higher simulated slants the slant estimates became closer to veridical. An ANOVA found a main effect of slant: $F(7, 77) = 127.19$, $p < 0.001$, partial $\eta^2 = 0.920$. Additional tests of polynomial contrasts found that there was a significant linear trend: $F(1, 11) = 166.22$, $p < 0.001$, partial $\eta^2 = 0.938$. There were also significant quadratic and cubic trends: quadratic, $F(1, 11) = 86.48$, $p < 0.001$, partial $\eta^2 = 0.887$; cubic, $F(1, 11) = 15.76$, $p = 0.002$, partial $\eta^2 = 0.589$. These results confirm that the psychometric functions were nonlinear.

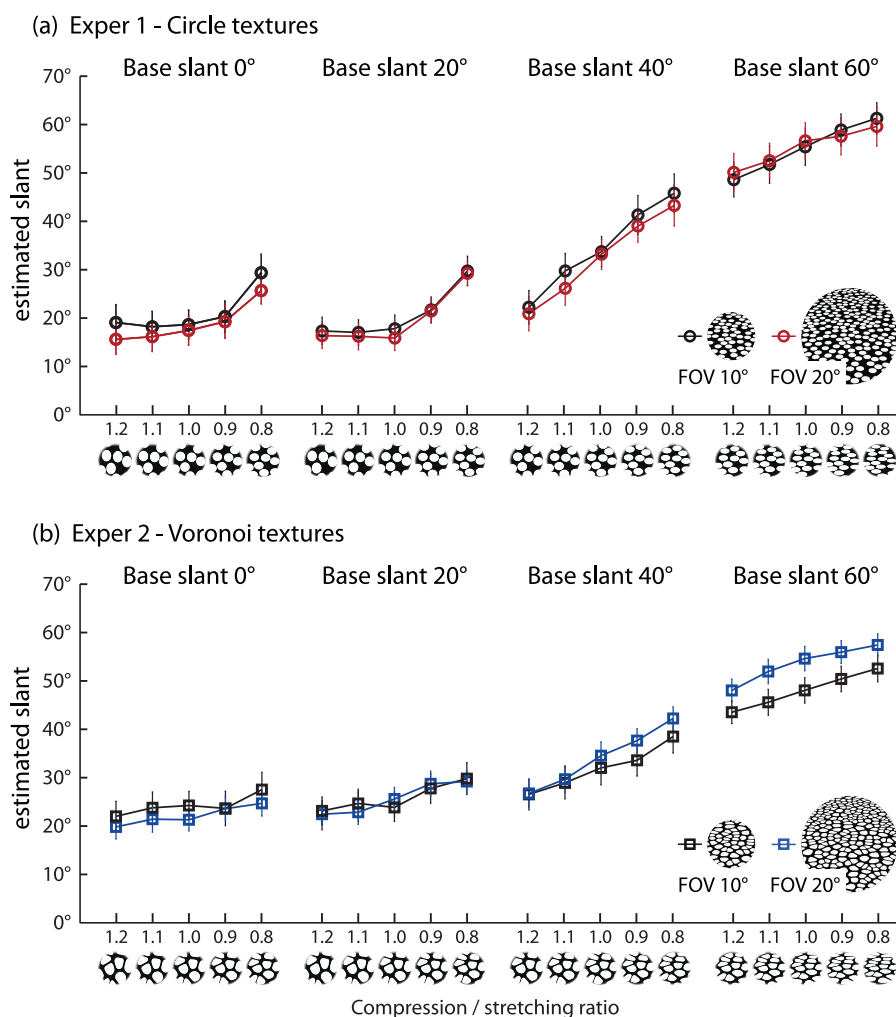


Figure 5. Mean slant estimates as a function of texture distortion in the cue conflict conditions of [Experiment 1](#) (upper) and [Experiment 2](#) (lower). Each set of connected points represents conditions with the same simulated slant and FOV but different amounts of stretching or compression in the tilt direction. Error bars depict ± 1 SE. Higher estimated slant for compressed textures compared to stretched textures is consistent with an influence of texture compression.

The consistent cue results were highly similar for the 10° and 20° FOV conditions of [Experiment 1](#), and we did not detect any reliable differences between slant estimates in these conditions. An ANOVA found no main effect of FOV, $F(1, 11) = 0.288$, $p = 0.602$, partial $\eta^2 = 0.025$, and no significant interaction between FOV and slant, $F(7, 77) = 0.785$, $p = 0.602$, partial $\eta^2 = 0.067$.

The results indicate that there was some response bias in the slant matching task. Because subjects received no feedback about their responses, they had no information to correct for any constant bias or scaling bias in their hand alignment. Simulated surfaces with low slant subjectively appear frontal (zero slant), but the mean orientation of the hand for these conditions deviated from frontal by about 20°. This is likely due to biases in the hand alignment task rather than perceptual bias. In a previous study using the same task ([Saunders](#)

& [Chen, 2015](#)), we similarly observed psychometric functions with a non-zero y -intercept and found that the intercepts from individual subjects were correlated across display conditions. This suggests that the non-zero intercepts are due to general response biases, which varied across individuals. There may also be bias in the overall scale of responses, but our data do not provide a way to evaluate the accuracy of scaling. Thus, the reported slant estimates should be interpreted as measures of relative perceived slant, which might differ from perceived slant by some scaling and constant bias.

Effect of texture compression/stretching

[Figure 5a](#) shows results from conditions with stretched or compressed textures that provide conflicting texture scaling and compression information. The graph plots mean slant estimates as a

function of texture distortion for the different base slant and FOV conditions. Each set of points corresponds to conditions with the same simulated slant but different amounts of stretching or compression applied to the surface texture. Use of texture compression predicts that perceived slant would be lower for stretched textures and higher for compressed textures, whereas the use of texture scaling and other gradient cues predicts that perceived slant could be constant.

At all simulated slants, the slant estimates varied as a function of texture distortion in the direction predicted by use of texture compression. We performed separate 5×2 ANOVAs for each base slant and confirmed that there was a significant main effect of the compression/stretching ratio in all cases: $F(4, 44) > 11.48$, $p < 0.001$, partial $\eta^2 > 0.51$. The conditions show a consistent general pattern of higher slant estimates for compressed textures than for stretched textures.

For the 0° and 20° base slant conditions, the effect of texture distortion was asymmetric: compression increased slant estimates but stretching had little effect. One important difference compared to other conditions is that stretched textures viewed from a low slant projected to vertically elongated ellipses. In this situation, texture compression specifies a different tilt direction than other texture cues, rather than a lower slant with the same tilt. The asymmetric effects of compression and stretching at low slants are likely due to this qualitative difference in cue conflicts. At higher slants, the stretched and compressed textures all project to vertically compressed ellipses, and there was no pronounced asymmetry between the effects of stretching and compression on slant estimates.

Cue weighting of texture compression and scaling

To further analyze the relative influence of texture compression and texture scaling cues on slant estimates, we computed cue weights for each base slant and FOV condition. The cue weights represent the influence of texture compression relative to the combined effect of all texture cues. Our approach is analogous to that used in some previous studies to infer cue weights from discrimination trials (e.g., Knill & Saunders, 2003; Hillis et al., 2004). Cue weights were based on a comparison between the change in slant estimates caused by compressing or stretching the surface texture, which varies only texture compression, and the change in slant estimates caused by varying the simulated slant of isotropic surfaces, which alters both texture compression and scaling. For example, in the condition with surface texture compressed by a factor of 0.9 and simulated slant of 40° , the texture compression was consistent with a slant of 46.4° , but texture scaling remained consistent with the simulated slant of 40° . If subjects relied entirely on texture compression, slant estimates in this condition would be the same as for

a consistent cue texture with slant of 46.4° , whereas relying entirely on texture scaling predicts that slant estimates would be the same as for a consistent cue texture with slant of 40° . If this cue conflict condition were perceived to have the same slant as a consistent cue texture with an intermediate slant of 45° , it would imply that subjects relied primarily on texture compression, but not entirely. The relative influence can be represented by the ratio of the effect of changing just compression to the effect of changing both texture cues. For the example above, the inferred cue weight would be $w = (45^\circ - 40^\circ)/(46.4^\circ - 40^\circ)$, which is approximately 0.78.

Our computation of cue weights assumes that the effects of cue perturbations are approximately linear but does not require that the psychometric functions in the consistent cue conditions be unbiased or globally linear. A cue weight is inferred from the consistent cue condition that is perceived to have the same slant as a cue conflict condition, with no assumptions that the slant estimates are veridical. The estimated weights would be invariant to linear transformations of the slant estimates, so response biases in the hand alignment task would not affect the cue weight results.

We fit the combined data from consistent cue and conflicting cue responses with a hierarchical Bayesian model. The slant estimates in the consistent cue conditions were assumed to vary around some nonlinear psychometric function $f(S)$. For the conflicting cue conditions, we assumed that slant estimates varied around $f(S + w\Delta S)$, where ΔS is the difference between the slant specified by texture compression and texture scaling ($S_{\text{comp}} - S$), and w is a linear weight representing the relative influence of texture compression for a given base slant and FOV condition. We excluded conditions where projected texture is vertically elongated (stretched textures with 0° or 20° slant) because texture compression does not specify a slant in the vertical direction in these cases. Other details of the cue weight analysis are provided in the [Appendix](#).

The results from the cue weight analysis support contributions from both texture compression and other texture cues. [Figure 6a](#) plots the estimates of the mean cue weights for texture compression as a function of base slant for the small and larger FOV conditions, and [Figure 6d](#) plots the results collapsed across slant conditions. The error bars depict the 95% highest density intervals (HDIs) of the posterior samples for each condition, which represent the credible range of mean cue weights given the data and assumptions. In all slant and FOV conditions, the HDIs of the mean cue weights are entirely positive. This is evidence that texture compression had a reliable influence on responses for all slant and FOV conditions. The results also provide evidence that subjects did not rely entirely on compression. When collapsed across slant conditions, the HDIs of the mean cue weights

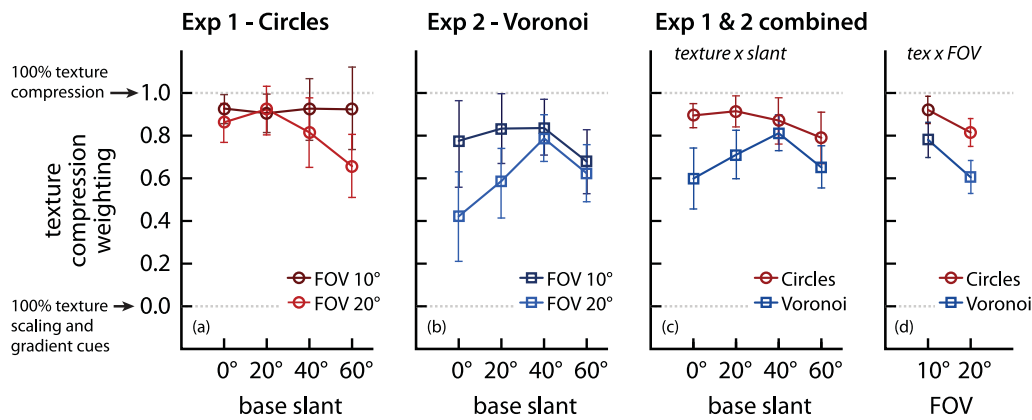


Figure 6. Estimated cue weights representing the influence of texture compression relative to other texture cues in Experiments 1 and 2. The graphs plot means and 95% HDIs of the estimated posteriors for the mean cue weights as a function of slant and/or FOV. The left two panels show results from Experiments 1 and 2 separately. The right two panels plot the results of Experiments 1 and 2 together, after averaging across FOV or base slant conditions.

were entirely below 1 for both the 10° and 20° FOV conditions. The HDIs were below 1 for some slant and FOV conditions when considered separately, whereas in other cases the data do not rule out the possibility that subjects relied entirely on compression. Taken together, these results show that texture compression had a large influence on slant estimates in all conditions, and that other texture information also had an influence in at least some of the slant conditions for both FOVs.

The cue weight results also provide evidence that texture compression had more overall influence when the FOV was smaller. To test whether FOV had a reliable effect on cue weights, we used the 95% HDIs of the differences between mean cue weights in the 10° and 20° FOV conditions. The 95% HDI of a difference measure can be used in a manner analogous to a statistical hypothesis test: If the HDI is entirely positive or negative, then there is a 95% probability of a positive or negative difference given the data and assumptions, so the difference can be interpreted as a reliable effect. When slant conditions were collapsed, the mean cue weights from the 10° and 20° FOV conditions were reliably different by this criterion. When slant conditions are analyzed separately, however, only the 60° slant conditions showed conclusive evidence for an FOV effect. From these results, we can conclude that FOV did have some effect on the relative influence of texture compression, but we cannot determine whether FOV had an effect for all slant conditions or only a subset.

Discussion

Perceptual biases

In the consistent cue conditions, we found that slant estimates were a nonlinear function of the simulated slant, with a large bias toward the frontal plane when

simulated slant was low. This finding is consistent with a number of previous studies (e.g., Norman et al., 2009; Saunders & Chen, 2015; Todd et al., 2005; Todd et al., 2010). One proposed explanation for perceptual biases is that the visual system relies on the scaling contrast across an image to perceive slant (Todd et al., 2007); however, the scaling contrast model predicts a large effect of FOV on perceived slant from texture, which was not observed. We found no detectable difference in perceptual bias for the 10° and 20° FOV conditions, for which scaling contrast would differ by a factor of two.

Another proposed explanation for perceptual biases involves the influence of a Bayesian prior or additional frontal cues (Saunders & Chen, 2015). By this explanation, the amount of underestimation would be a function of the overall reliability of slant information. This predicts that there should be proportionally more frontal bias for surfaces with low slant than high slant, which is consistent with the observed nonlinearity of slant estimates. It also predicts an effect of FOV, which was not observed. A larger FOV would generally provide more reliable information about slant, so less frontal bias would be expected with larger FOV; however, the effect of FOV on the reliability of slant information depends on what texture information was used. Slant estimation using texture scaling involves comparisons across the image, so the reliability would strongly depend on FOV. On the other hand, texture compression could be estimated from a local region of the image, particularly for the circular texture elements used in Experiment 1, so a large FOV is not necessarily required to reliably estimate slant from compression. If observers relied primarily on texture compression, the overall reliability of slant information may have been similar in the 10° and 20° FOV conditions. This could explain why FOV had little effect on perceptual biases in Experiment 1. In Experiment 2, we tested Voronoi textures that have less regular projected shapes.

If the irregular shapes reduce the reliability of the texture compression cue, FOV might have more effect on perceptual biases.

Relative influence of texture compression and scaling

The results from cue conflict conditions revealed a large influence of texture compression on slant estimates, as well as a significant influence of texture scaling. Compressing or stretching a texture along the surface, which selectively varies the texture compression, caused slant estimates to be biased toward the slant specified by compression. The effect of conflicting texture compression information was large but less than would be expected if observers relied entirely on texture compression. Our cue weight analysis found that the proportional influence was greater than zero in all conditions, and also less than one in most cases. These results are consistent with some previous findings that both texture scaling and compression are used to perceive slant from texture (Knill, 1998a; Rosenholtz & Malik, 1997).

The relative influence of texture compression and texture scaling depended on FOV, and the direction of the observed FOV effect is consistent with the relative reliability of these cues. The cue weight analysis found that overall weighting of texture compression was lower for the 20° FOV stimuli than for the 10° FOV stimuli. This difference was clearly present for stimuli with 60° slant, and there were trends in the same direction for other slants. The size of the visible region of a surface would be expected to have a large effect on the reliability of texture scaling, which involves changes across the visible region, and comparatively less effect on texture compression, which can be measured in local regions. Our finding that texture compression had less influence with the larger FOV is consistent with the different expected effects of FOV on reliability of texture compression and texture scaling cues.

For simulated surfaces with low slant, the effects of stretching and compression were asymmetric. This asymmetry is likely due to the fact that the stretched textures viewed at low slants projected to vertically elongated ellipses, which creates a qualitatively different conflict between texture compression and other texture cues. Vertically elongated ellipses could be the projection of circles that are slanted leftward or rightward but could not be the projection of a circle slanted in the vertical direction. Thus, the stretched textures viewed from a low slant present a conflict between the tilt direction implied by texture compression and the tilt direction specified by other texture cues, rather than just conflicting information about slant. Our results suggest that the texture compression cue was either ignored or interpreted as specifying zero slant in these cases. At higher slants of 40° and 60°, there was no inconsistency between the tilt direction specified by texture cues,

and we did not observe a qualitative asymmetry in the effects of compressing and stretching texture.

The fact that the textures in [Experiment 1](#) were composed of perfect circles or ellipses might have contributed to the large influence of texture compression on slant estimates. The foreshortening of texture elements might be especially salient when they are perfect ellipses, potentially encouraging observers to rely on this information more than they would for other classes of texture. Previous studies by [Knill \(1998b\)](#) and [Rosenholtz & Malik \(1997\)](#) also observed an influence of both texture compression and scaling, but the influence of compression was somewhat smaller than observed here. Slant estimates in cue conflict conditions with low slant also suggest more influence of texture compression than might be expected. For the cue conflict conditions with 0° slant, the images consisted of identical ellipses with no variation in size or shape across the image. Such images would be expected to be perceived as frontal surfaces, but the mean slant estimates in the conditions with compressed ellipses were greater than zero ([Figure 5a](#), left). This could indicate that slant estimates were not based solely on perceived 3D slant. In [Experiment 2](#), we tested whether texture compression also has a large influence for less regular textures.

Experiment 2

[Experiment 2](#) was the same as [Experiment 1](#) except that we used Voronoi textures with less regular texture elements. One motivation is to test whether the large influence of texture compression observed in [Experiment 1](#) generalizes to less regular textures. Textures composed of Voronoi cells have been used in a number of previous studies of slant from texture (e.g., [Knill, 1998b](#); [Knill & Saunders, 2004](#); [Rosenholtz & Malik, 1997](#); [Saunders & Backus, 2006](#); [Saunders & Chen, 2015](#)) and have naturally occurring variations in size and shape. Texture regularity also affects the reliability of texture cues. According to a Bayesian model, reduced reliability would be expected to increase perceptual biases, and changes to relative reliability of texture compression and texture scaling would affect the relative weighting of cues. These predictions were tested in [Experiment 2](#), as well.

Methods

Participants

Twelve right-handed adults from the University of Hong Kong (seven males and five females; mean age, 21.8 ± 3.7 years) were paid to participate in

Experiment 2. All of the subjects had normal or corrected-to-normal visual acuity and passed a stereo acuity screening test as in [Experiment 1](#). None of the subjects participated in [Experiment 1](#), nor were they aware of the purpose of the study. Each subject finished two 1-hour sessions in two separate days. The procedures were approved by and conformed to the standards of the Human Research Ethics Committee for Non-Clinical Faculties.

Apparatus and stimuli

The stimuli were presented with the same apparatus as in [Experiment 1](#), and hand orientation was measured using the same palm board and inclinometer. As in the previous experiment, subjects viewed the images monocularly through an aperture, and we varied the size of the aperture to manipulate FOV.

The simulated textures in [Experiment 2](#) were composed of Voronoi cells with curved corners (see [Figure 2](#)). The base textures were statistically isotropic, like the circle texture used in [Experiment 1](#), but the Voronoi cells had variations in shape that would add noise to the compression cue and variations in size that would noise to the scaling cue, as well. The center points of the Voronoi cells were chosen randomly in a tileable square region with the constraint that the distance between each pair of dot was greater than 3.5% of the region width. We generated the Voronoi diagram from these points and shrank the cells by 10% of their sizes to create space between elements and smoothed the contours of polygons. The texture was scaled so that the average size of the Voronoi cells was matched to the diameters of the circles in the texture used in [Experiment 1](#). At the base size of the Voronoi texture, the mean diameters of the cells varied between 1.5 and 3.5 cm, with a mean of 2.4 cm and *SD* of 0.4 cm. The size of the cells measured in the vertical and horizontal directions had ratios between 0.5 and 2.1, with a mean of 1.03 and *SD* of 0.26. As in the previous experiment, we added random variations to the overall size of the texture on each trial to discourage the use of projected size as a direct cue to slant. On each trial, the base texture was uniformly scaled by a random amount between 0.875 and 1.125.

Procedure

The task, design, and procedure of [Experiment 2](#) were the same as those of [Experiment 1](#). Subjects reported their perceived slant by aligning the palm with the virtual surfaces. Each subject performed a total of 864 trials. The 10° and 20° FOV conditions were tested in separate sessions on different days, with order counterbalanced across subjects. The slant and stretch ratio conditions were randomized within sessions.

Results

Consistent texture cues

The mean slant estimates in the consistent cue conditions of [Experiment 2](#) are plotted in [Figure 4b](#). As in [Experiment 1](#), we observed a nonlinear relationship between simulated slants and slant estimates. An ANOVA found a significant main effect of simulated slant: $F(7, 77) = 87.27, p < 0.001$, partial $\eta^2 = 0.888$. Tests of polynomial contrasts found a significant quadratic trend, $F(1, 11) = 58.31, p < 0.001$, partial $\eta^2 = 0.841$, confirming that the psychometric functions were nonlinear, in addition to a significant linear trend: $F(1, 11) = 114.17, p < 0.001$, partial $\eta^2 = 0.912$.

Unlike in the previous experiment, the slant estimates in consistent cue conditions of [Experiment 2](#) showed some effects of FOV. An ANOVA found no main effect of FOV, $F(1, 11) = 0.859, p = 0.374$, partial $\eta^2 = 0.072$, but there was a significant interaction between simulated slant and FOV: $F(7, 77) = 5.19, p < 0.001$, partial $\eta^2 = 0.320$. At low slants, the slant estimates in the two FOV conditions were similar, but, at high slants, the slant estimates were lower with the smaller FOV. Pairwise tests for individual slant conditions found significant differences between the FOV conditions at slants of 60°, $t(11) = 3.274, p = 0.007$, Cohen's $d = 0.945$, and 70°, $t(11) = 2.744, p = 0.019$, Cohen's $d = 0.792$, as well as a trend toward a difference at 50°: $t(11) = 2.016, p = 0.069$, Cohen's $d = 0.582$. At other slants, there was no detectable effect of FOV: $t(11) < 0.798, p > 0.442$, Cohen's $d < 0.230$. These results indicate that the Voronoi textures appeared less slanted with the smaller FOV, but only at high slant conditions.

As in [Experiment 1](#), the y -intercepts of the psychometric functions were not zero, which can be attributed to response bias. The simulated surfaces with low slants clearly appear frontal, so the deviation from zero in the slant estimates would be due to errors in matching the orientation of the hand to the perceived slant, rather than indicating a non-zero perceived slant. There may also be some bias in the overall scaling of responses, which cannot be assessed from our data.

Effect of texture compression and stretching

As in [Experiment 1](#), slant estimates in conditions with stretched and compressed textures varied as a function of texture distortion, consistent with an influence of texture compression. [Figure 5b](#) plots the mean slant estimates as a function of the amount of compression or stretching for different base slants. There were significant main effects of texture distortion at all of the base slants: $F(4, 44) > 9.08, p < 0.001$, partial $\eta^2 > 0.45$. As in [Experiment 1](#), we also observed asymmetric effects of compression and stretching for slants of 0° and 20° but not for higher slants.

Cue weighting of texture compression and scaling

The cue weight analysis revealed a large influence of texture compression on slant settings in the cue conflict conditions of [Experiment 2](#) and also a smaller influence of other texture cues. The cue weights were estimated in the same way as in [Experiment 1](#). [Figure 6b](#) plots the means and 95% HDIs of the posterior samples of mean cue weights for each slant and FOV condition. In all conditions, the HDIs were entirely positive and lower than 1, consistent with nonzero contributions from both texture compression and other texture cues.

The cue weights from [Experiment 2](#) also showed evidence of an effect of FOV on the weighting of texture compression. To test for a main effect of FOV, we collapsed the mean cue weights from all slant conditions. The 95% HDI for the difference between 10° and 20° FOV conditions was entirely negative, indicating that the overall weighting of texture compression was lower for the larger FOV. When slant conditions are considered separately, the evidence for FOV effects was less conclusive. For slants of 0° and 20°, the 95% HDIs of the differences between FOV conditions were entirely negative, whereas for the other slants the mean differences were in the same direction but the HDIs overlapped zero. From these results, we can conclude that the relative influence of texture compression was lower on average for the larger FOV, but this difference might only be present for some of the slant conditions.

Cue weighting for circles versus voronoi textures

We compared the compression weights from [Experiments 1](#) and [2](#) to test the effect of texture type. [Figure 6c](#) plots results from the two experiments together as a function of slant after averaging over FOV conditions, and [Figure 6d](#) plots results from the two experiments as a function of FOV after collapsing across base slant conditions. We computed the 95% HDIs for the differences between mean cue weights in the circles and Voronoi texture conditions to evaluate the effects of texture type. When collapsed across slant conditions, the HDIs for texture type effects were entirely negative for both the 10° and 20° FOV conditions, providing evidence that cue weights were lower overall for the Voronoi textures than for the circles textures. For the separate slant conditions, the HDIs were entirely negative for slants of 0° and 20°, but overlapped zero for the higher slant conditions. All slant conditions showed a trend toward lower compression weight for Voronoi textures, but there was only strong evidence for a difference at low slants. We also evaluated the difference between the FOV effects for the two textures to test for an interaction between FOV and texture type. There was a trend toward more effect of FOV for the Voronoi textures, but this interaction

was only present for 82% of the posterior samples, so the evidence was not conclusive. Taken together, the results from the two experiments demonstrate that the influence of texture compression depends on both texture type and FOV, but the effects of these factors might not be consistent across all slants.

Discussion

The results of [Experiment 2](#) show that texture compression still has a large influence on slant estimates when the surface texture is less regular. In all of the slant and FOV conditions, stretching or compressing the surface texture caused slant estimates to be biased toward the slant specified by texture compression. The cue weight analysis found that perceived slant in the cue conflict conditions was determined primarily by texture compression and that there was also a contribution from texture scaling or other gradient cues. This finding is consistent with the results from the cue conflict conditions of [Experiment 1](#) and results of some previous studies that used Voronoi textures but different slant judgment tasks ([Knill, 1998a](#); [Rosenholtz & Malik, 1997](#)).

We also replicated our finding that relative influence of texture compression depended on FOV, with comparatively more influence of texture compression when the FOV was smaller. As discussed previously, this effect could be explained in terms of the relative reliability of texture compression and texture scaling cues. Increasing FOV would be expected to improve the reliability of texture scaling more than the reliability of texture compression, so optimal integration would predict less relative influence of texture compression. When slant conditions are considered separately, the pattern appears somewhat different for the two experiments ([Figure 6a](#) vs. [Figure 6b](#)). We suspect that these variations are primarily noise. The cue weight estimates for individual slant and FOV conditions are not very precise, and there is no reason to expect a qualitative difference in the effect of FOV for the two types of texture. When averaged across slant conditions ([Figure 6d](#)), the estimates of the overall cue weighting are more precise and show a clear reduction in the influence of texture compression with larger FOV.

In [Experiment 2](#), slant estimates in consistent cue conditions were also affected by FOV, which was not observed in [Experiment 1](#). For slants above 40°, the slant estimates were lower for the smaller FOV. This effect of FOV could be explained in terms of the reliability of texture information. If perceptual underestimation of slant from texture is due to the influence of a frontal prior or frontal cues, as we proposed in [Saunders and Chen \(2015\)](#), then decreasing the reliability of texture information would be expected to result in more bias toward frontal. Decreasing

FOV would reduce the overall reliability of texture information, so more underestimation of slant would be expected. The lower slant estimates observed in the 10° FOV conditions compared to the 20° conditions in [Experiment 2](#) are consistent with this explanation.

Comparing across experiments, we found that influence of texture compression relative to other texture cues was lower overall for the Voronoi textures than for the circles textures. This difference was reliably observed at low slants, and there was a similar trend at higher slants ([Figure 6c](#)). A possible explanation for the overall difference is that the irregularity of the Voronoi textures degraded the information from texture compression more than the information from texture scaling, causing the visual system to rely more on texture scaling. This interpretation would be consistent with the different effects of FOV observed in the consistent cue conditions of [Experiments 1 and 2](#). If perceived slant were more dependent on texture scaling for the Voronoi textures, then FOV would be expected to have more effect on perceptual biases. The same reasoning would apply to the effect of FOV on cue weighting. The mean cue weights did not show conclusive evidence for an interaction, but there was a trend toward a larger effect of FOV for the Voronoi textures ([Figure 6d](#)). The combined results from the two experiments demonstrate that texture regularity modulates the relative influence of texture compression and other texture cues and suggest that this modulation is a function of how texture regularity affects the information provided by different texture cues.

In summary, the results of [Experiment 2](#) replicate the main findings of [Experiment 1](#) and also demonstrate that texture regularity affects the use of texture cues for slant perception. Texture compression strongly influenced slant settings, and the influence of texture compression relative to other texture cues was modulated by both FOV and texture regularity.

Experiment 3

In the previous experiments, we observed a strong effect of texture compression on slant estimates in conditions with conflicting texture cues. One possible concern is that the slant estimates were not actually based on perceived slant. Observers might have instead used some heuristic strategy based on 2D cues. If so, the cue weightings might not reflect the contributions to slant perception. The fact that we used an estimation task rather than slant discrimination would encourage judgments based on perceived slant, but it remains possible that observers used some 2D strategy.

[Experiment 3](#) tested the effect of manipulating texture compression for surfaces viewed binocularly rather than monocularly. Under binocular viewing, the

simulated surfaces convey a strong percept of a surface slanted in depth, which would encourage observers to base their estimate on perceived 3D slant. As in the previous experiments, we created cue conflict conditions by stretching or compressing the texture along the surface, which selectively affects texture compression. In [Experiment 3](#), however, the actual simulated slant was specified by stereo information, as well as other texture cues such as texture scaling. Effects of stretching or compressing the textures in binocular conditions would be strong evidence that texture compression is used to perceive 3D slant. We used both the circles and Voronoi textures, allowing us to further test whether cue weights depend on the texture regularity. For this experiment, we used only the larger FOV (20°) from the previous experiments. Texture scaling information is better for a larger FOV, so an influence of compression would not be due to the lack of other texture information.

Methods

Participants

Twelve right-handed adults from the University of Hong Kong (nine males and three females; mean age, 21.6 ± 2.7 years) were paid to participate in [Experiment 3](#). All subjects had normal or corrected-to-normal visual acuity and passed a stereo acuity screening test, as in the two previous experiments. None of the subjects participated in the other experiments, and all were naïve to the purpose of the study. The procedures were approved by and conformed to the standards of the Human Research Ethics Committee for Non-Clinical Faculties.

Procedure

The task in [Experiment 3](#) was the same as of the previous experiments. We tested both circles and Voronoi textures, but only in the 20° FOV condition. Slant, compression, and texture type conditions were fully randomized. Subjects completed a total of 864 trials separated into two sessions on separate days.

Apparatus and stimuli

The stimuli were presented with the same apparatus as in the previous experiment. The only difference was that subjects did not wear an eye patch and viewed binocular images of the simulated surfaces. Shutter glasses (NVIDIA 3D Vision 2) were used to present left and right stereo images to the two eyes separately. The stimuli were accurate perspective renderings of textured planar surfaces based on each subject's measured interocular distance. The binocular region when viewed through the aperture was about 17.5° wide, and

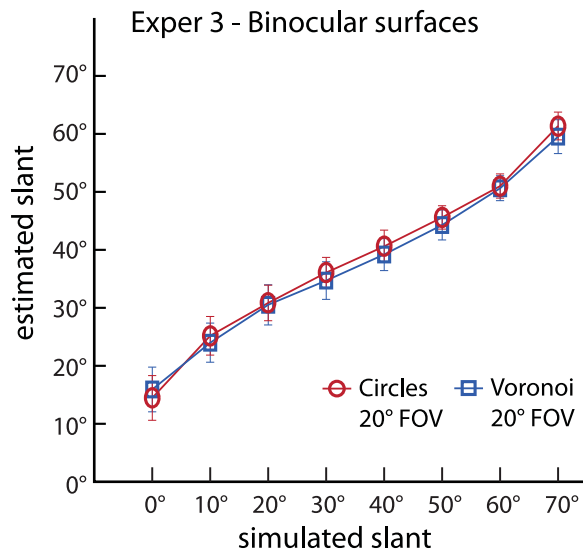


Figure 7. Mean slant estimates as a function of simulated slants in the consistent cue conditions of [Experiment 3](#), in which stimuli were viewed binocularly. The two sets of data correspond to the conditions with circles and Voronoi textures. Error bars depict ± 1 SE.

subjects reported no difficulty fusing the images. Hand orientation was measured and recorded using the same equipment as in the previous experiments.

Results and discussion

Consistent texture cues

[Figure 7](#) plots the mean slant estimates as a function of simulated slants for the consistent cue conditions with isotropic textures. The psychometric functions were much more linear than in the monocular conditions of the previous experiments. Tests of polynomial contrasts revealed a significant linear trend, $F(1, 11) = 154.74$, $p < 0.001$, partial $\eta^2 = 0.934$; no quadratic trend, $F(1, 11) = 0.005$, $p = 0.946$, partial $\eta^2 < 0.001$; and a small but significant cubic trend, $F(1, 11) = 10.98$, $p = 0.007$, partial $\eta^2 = 0.500$.

The results were similar for consistent cue conditions with circles and Voronoi textures, but there was evidence for a small difference in overall slant estimates. An ANOVA found a marginally significant main effect of texture type, $F(1, 11) = 3.59$, $p = 0.085$, partial $\eta^2 = 0.246$, and no interaction between texture type and slant, $F(7, 77) = 1.20$, $p = 0.313$, partial $\eta^2 = 0.098$. Slant estimates were higher for the circles textures than for the Voronoi textures, although the magnitude of this difference was small ($0.79^\circ \pm 1.45^\circ$).

The non-zero intercepts of the psychometric functions indicate that there was some response bias in the binocular conditions. The constant bias was smaller than in the previous experiments. For surfaces

that appear frontal, the mean hand orientation in [Experiment 3](#) was 15.2° , whereas in the previous two experiments it was 18.1° and 23.0° . The different constant biases suggest that the mapping from perceived slant to hand orientation was not exactly the same across experiments and conditions. It is therefore possible that the overall scaling of responses may have been adjusted according to the range of perceived slants across the stimuli. Across the range of simulated slants (0° – 70°), the range of responses was similar for the monocular and binocular conditions. At the high slant (70°), the slant of the stimuli appears similar but not identical, so there was likely some difference in response scaling. The response scaling would not affect our measures of the relative influence of texture cues but would affect comparisons in the amount of overall bias in responses.

Effect of texture compression/stretching

In the cue conflict conditions of [Experiment 3](#), varying texture compression affected slant estimates in all cases except when slant was zero. [Figure 8](#) plots mean estimates as a function of texture distortion for each base slant, and [Figure 9](#) plots cue weights representing the influence of texture compression relative to stereo and other texture cues. Cue weights were estimated in the same way as in the previous experiments. For frontal surfaces, the 95% HDIs of cue weights included zero, but for all other cases the HDIs were entirely positive. This was true for both the circles and Voronoi textures. The cue weights in [Experiment 3](#) were lower than in the monocular conditions of the previous experiments, as would be expected given that the binocular stimuli provided additional stereo information that specified the base slant. Despite the conflicting stereo information, texture compression had an observable influence on slant estimates in most conditions.

The weighting of texture compression relative to stereo and other texture cues increased with slant. To test for effects of slant on the cue weights, we computed 95% HDIs of pairwise differences between slant conditions. The results showed evidence that cue weights were larger for surfaces with 60° slants compared to the other conditions and that cue weights were lower for surfaces with 0° slant compared to the other conditions, but there was no evidence for a difference between the 20° and 40° slant conditions.

There was a trend toward more weighting of texture compression for the circles textures compared to the Voronoi textures, but the evidence was not conclusive. This trend was present for all of the slant conditions, but the 95% HDIs of the differences between cue weights in the circles and Voronoi texture conditions included zero in all cases. When collapsed across slant conditions, the mean cue weights were higher for the

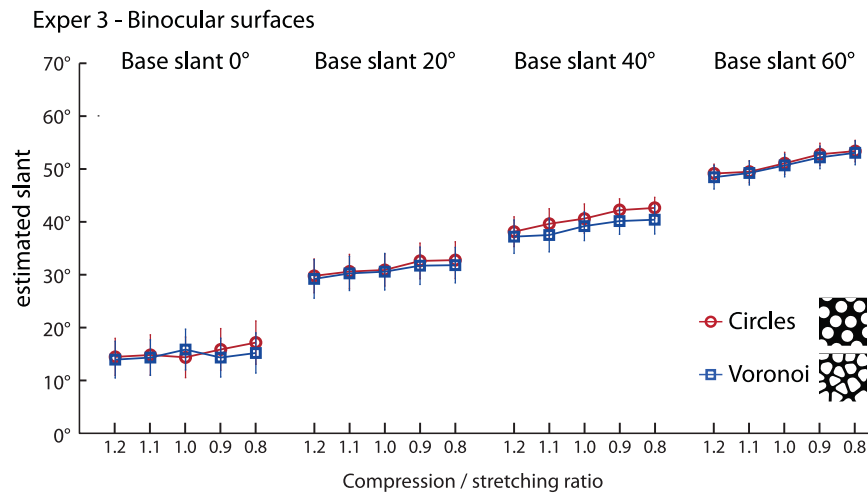


Figure 8. Mean slant estimates as a function of texture distortion in the cue conflict conditions of Experiment 3. Error bars depict ± 1 SE.

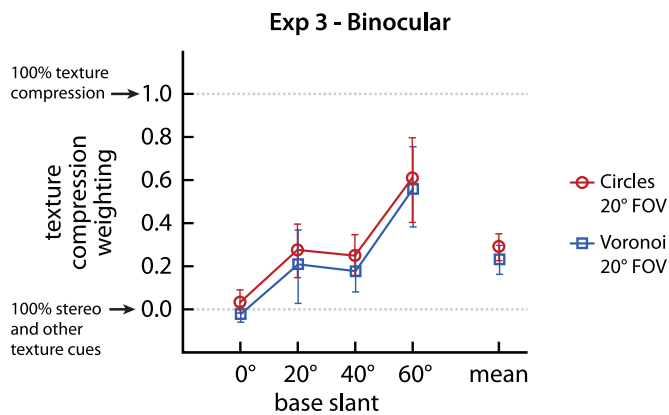


Figure 9. The mean cue weights representing the influence of texture compression relative to stereo information and other texture cues in Experiment 3, plotted as a function of the simulated slant. The mean cue weight averaged across slants is also shown. Error bars show the 95% HDIs of the estimated posterior distributions for each variable.

circles textures in 91% of the posterior samples, which is suggestive but not conclusive.

Varying texture compression had less effect in binocular conditions of Experiment 3 than in the monocular conditions of the previous experiments, as would be expected given that the binocular conditions provide additional stereo information that is consistent with simulated slant. The cue weights from Experiment 3 are comparable to the cue weights observed in previous studies that measured the overall influence of texture cues relative to stereo cues (Hillis et al., 2004; Knill & Saunders, 2003; Saunders & Chen, 2015). This is consistent with our finding that slant from texture is primarily determined by texture compression in the monocular conditions with the same textures and FOVs.

General discussion

Role of texture compression in perceived slant

The results from the three experiments demonstrate a strong influence of texture compression on perceived slant. Stretching or compressing the simulated texture on a virtual planar surface produced systematic biases in slant estimates toward the slant specified by texture compression. This was observed in both monocular conditions and binocular conditions. In the latter case, binocular slant cues were available and specified the same slant as texture scaling and gradient texture cues, and there was still an influence of texture compression at most slants.

In monocular conditions with conflicting texture cues, texture compression was a stronger determinant of slant estimates than texture scaling or other texture cues. This agrees with the findings from some previous studies that tested perception of slant from textures with conflicting slant cues (Buckley et al., 1996; Knill, 1998a; Rosenholtz & Malik, 1997) but conflicts with the findings from other studies (Todd et al., 2005; Todd et al., 2007; Todd et al., 2010). Todd et al. (2010) suggested that subjects might have relied on 2D cues rather than perceived 3D slant in the earlier studies that observed an influence of texture compression, but this is unlikely for our experiments. We used a cross-modal slant matching task, which would encourage subjects to rely on perceived 3D slant. Furthermore, we observed an influence of texture compression even for binocular stimuli. The binocular stimuli create a strong percept of a slanted 3D surface, so there would be little motivation for subjects to rely on artificial strategies for the slant matching task. Our results suggest that the influence of texture compression on slant judgments observed in

earlier studies was not due to some artificial strategy and that the conflicting findings of Todd and colleagues have some other explanation.

We suspect that the size of cue conflicts is a main reason for the conflicting findings about the role of texture compression in perceived slant. The conflicts between texture cues were often quite large in previous studies that did not observe an influence of texture compression (Todd et al., 2005; Todd et al., 2007; Todd et al., 2010) and were relatively small in our study and previous studies that observed a strong influence of texture compression (Buckley et al., 1996; Knill, 1998b; Rosenholtz & Malik, 1997). Large conflicts can cause the visual system to rely on single cues rather than integrating cues, as has been demonstrated for the case of texture and stereo slant cues (Girshick & Banks, 2009). Texture compression depends on an assumption that a texture is isotropic, which is violated by many real-world textures (e.g., bricks), whereas texture scaling depends on a weaker assumption of homogeneity. When there is a large conflict between the slant specified by texture compression and the slant specified by other texture cues, the visual system might give less weight to or veto the compression cue.

The lack of influence from texture compression observed in some studies could also be due to the availability of stronger information from other texture cues. Todd et al. (2010) used textures with rectangular structure. The converging parallel lines in the projected texture provided an additional slant cue, which has been shown to contribute to perceived slant from texture (e.g., Chen & Saunders, 2019; Saunders & Backus, 2006; Tam et al., 2013). Perspective convergence is not affected by stretching or compressing texture in the tilt direction, so this manipulation might have less effect if perspective convergence is available and used for slant perception. Another factor is field of view. Todd et al. (2005) and Todd et al. (2007) included conditions with a 60° FOV, which is larger than in our conditions. We found that increasing FOV from 10° to 20° produced a detectable reduction in the relative weighting of texture compression. For larger FOVs, the relative influence of texture compression might reduce further and become difficult to detect.

Effects of FOV and texture regularity

The experiments also tested how perception of slant from texture depends on FOV and texture regularity. Both of these factors affected the relative weighting of texture compression in cue conflict conditions. Texture compression had more influence for the circles textures than for the Voronoi textures with less regular shape, and the influence of texture compression was reduced with larger FOV. There was also a small but detectable

effect of FOV on perceptual biases for the Voronoi textures.

The effects of FOV and texture regularity on the relative weighting of texture compression are qualitatively consistent a Bayesian model that integrates multiple slant cues according to their reliability. In both Experiments 1 and 2, texture compression had less influence on slant estimates with larger FOV. The reliability of texture scaling is strongly dependent on FOV, whereas the reliability of texture compression is less dependent (Knill, 1998c). This predicts that texture compression should have relatively less influence for stimuli with large FOV, as we observed. The effect of texture regularity on the relative weighting of texture cues could also be explained. We observed less weighting of texture compression for the Voronoi textures in monocular conditions and a trend in the same direction in binocular conditions. If the irregular shapes in the Voronoi texture degrade the information from texture compression more than the information from texture scaling, then the relative weighting of texture cues would be expected to change in the observed direction. Thus, the effects of both FOV and texture regularity are qualitatively consistent with optimal integration of texture cues.

The perceptual biases in slant estimates across FOV, texture type, and slant conditions are also generally consistent with a Bayesian model. In Saunders and Chen (2015), we proposed that underestimation of slant from texture is a consequence of combining weak texture information with either a frontal prior or conflicting frontal cues (e.g., accommodation). If so, any factors that decrease the reliability of texture information would result in more bias toward frontal. As in our previous study, the bias toward frontal in monocular conditions was proportionally greater at low slants than high slants, consistent with the fact that texture information is less informative at low slants than high slants (Knill, 1998c). The reliability of texture information also depends on FOV and texture regularity. We found that slant estimates were lower for Voronoi textures with the small FOV than in the other conditions, as expected if less reliable texture information leads to more frontal bias. We did not find a detectable effect of FOV on perceptual biases for the circles textures or a difference between perceptual biases in circles and Voronoi texture conditions with the larger FOV. However, the reliability of texture information in these conditions might be similar. If the visual system relies on texture compression more for the circles textures than the Voronoi textures, as suggested by our results, then increasing the FOV would have less effect on the overall reliability of texture information. Similarly, if the visual system relies less on texture compression with a larger FOV, as suggested by our results, then the variability in the Voronoi textures might have less effect on the overall reliability of texture




Figure 10. Perspective images of surfaces rendered with 60° slant and FOV of 40° or 20°. The stimulated textures are (a) circular dots, (b) ellipses with an aspect ratio of 0.8, (c) square bricks, and (d) rectangular bricks with an aspect ratio of 0.8. If texture cues are optimally integrated, compressing the surface texture (a vs. b, c vs. d) would be expected to have less effect on perceived slant with larger FOV and even less effect when additional information from perspective convergence is available (bricks). If perceptual biases are a function of the reliability of texture information, slant estimates would be expected to be more accurate with larger FOV and with additional perspective convergence information.

information. Thus, the expected effects of FOV and texture regularity on perceptual biases would likely be small in cases where they were not detected, and the condition that did show a detectable increase in perceptual bias had the least reliable texture information (Voronoi, 10° FOV). Overall, the pattern of results is consistent with a Bayesian model that integrates texture information with frontal prior or other frontal cues.

The demonstrations in Figure 10 suggest that perceived slant varies in a way consistent with a Bayesian model when FOV is larger and additional texture cues are added. For a texture composed of circles, we observed little effect of increasing the FOV from 10° to 20°, but increasing the FOV to 40° appears to increase the perceived slant (Figure 10a), consistent with reduced frontal bias. Conflicting information from texture compression also appears to have less effect with larger FOV (Figure 10b). Both of these observations could be explained by the improved quality of texture scaling information with a larger FOV. Figures 10c and 10d show images of surfaces with brick-like textures that provide an additional perspective convergence cue to slant, which are similar to the textures used in

Todd et al. (2010). Adding this cue appears to further reduce the relative influence of texture compression. These apparent effects are qualitatively consistent with optimal integration of multiple texture cues.

Implications for the scaling contrast model

The scaling contrast model of Todd et al. (2007) predicts that perceived slant from texture would be highly dependent on FOV, but we observed only limited effects in our monocular conditions. The scaling contrast across a planar surface patch is equal to $\tan(\text{slant}) \times \tan(\text{FOV}/2)$ (Saunders & Chen, 2015), so the scaling contrast for a 20° FOV would be approximately double the scaling contrast for a 10° FOV. This is much larger than the observed differences in our experiments. The mean slant estimates showed no significant differences between the two FOV conditions in Experiment 1, and the differences were only 17% to 39% in Experiment 2. Varying FOV had much less effect on slant estimates than would be predicted if perceived slant were a direct function of scaling contrast.

One issue with interpreting the effect of FOV on slant estimates in our experiments is that the 10° and 20° FOV conditions were tested in separate blocks, so subjects might have had different response biases in the two FOV conditions. In particular, subjects might have adjusted the mapping from perceived slant to hand orientations according to the range of perceived slants experienced during a block of trials. If subjects used an exaggerated range of hand orientations to indicate variations in perceived slant when the range of perceived slant was small, and vice versa, the effect would be to reduce the apparent differences between FOV conditions. However, it is doubtful that this could account for the discrepancy between our results and the predictions from scaling contrast. For the largest slant tested in our experiments, the apparent slant of surfaces with the circles texture is subjectively similar for the 10° and 20° FOV conditions (Figure 11). The perceived slant of the smaller stimuli would have to be half of the perceived slant of the larger stimuli to be consistent with the difference in scaling contrast. For the Voronoi textures, we did observe a detectable difference between perceived slant in the 10° and 20° FOV conditions. If a large difference in response scaling counteracted a difference in perceived slant for the circles texture, the same thing would have been expected for the Voronoi textures. Based on the subjective appearance, we believe that any difference between the response scaling in the two FOV conditions was relatively small and could not fully account for the fact that the difference between FOV conditions was smaller than predicted by scaling contrast.

The observed effects of stretching and compressing surface texture are also inconsistent with the scaling contrast model. This manipulation does not affect




Figure 11. Examples of consistent cue stimuli with circles texture and 70° slant, rendered with FOVs of 20° and 10°. The mean slant settings were approximately the same for these two conditions.

scaling contrast but strongly affected slant estimates. This discrepancy cannot be explained by differences in response bias. Our analysis of the relative weighting of texture compression and scaling cues was based entirely on comparisons between stimuli tested within blocks, and the estimated cue weights would be unaffected by an overall linear transformation. The cue conflict results suggest that the information from texture compression under an assumption of isotropy was a primary determinant of perceived slant.

Figure 12 re-plots the mean slant settings from all of our monocular conditions as functions of three possible predictors: scaling contrast, simulated slant, and slant from texture compression. The graphs include

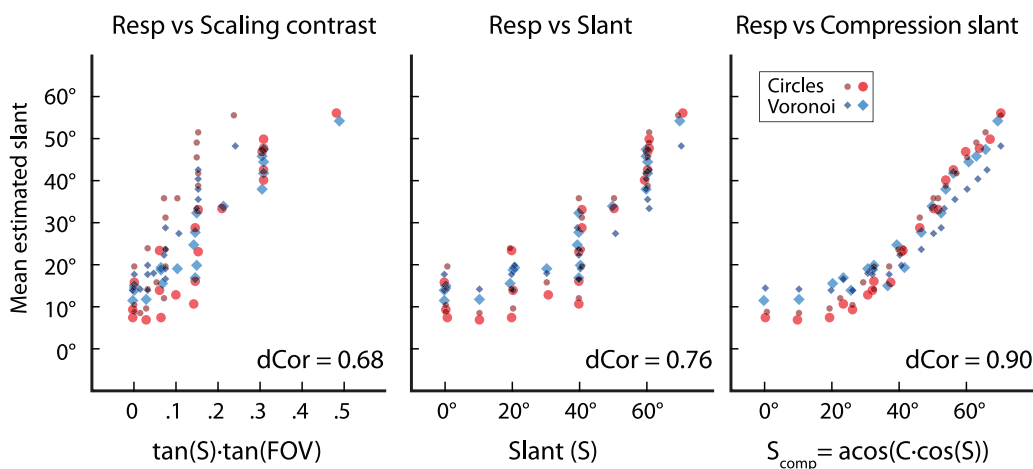


Figure 12. Mean slant estimates from all monocular conditions replotted as a function of scaling contrast (left), the simulated slant (middle), or the slant indicated by texture compression (right). The graphs plot the mean slant estimates averaged across subjects for each condition in Experiments 1 and 2, including both consistent cue and cue conflict conditions. Red circles and blue diamonds are results from Experiments 1 and 2, respectively. The smaller symbols correspond to the smaller FOV conditions. A small amount of random jitter was added to the horizontal positions to make nearby points more discernible. The distance correlation ($dCor$) between each predictor and the response is shown on the graphs.

the consistent cue and conflicting cue conditions from both [Experiments 1](#) and [2](#), and both 10° and 20° FOV conditions. To measure the strength of the relationship between the slant estimates and predictors, we computed the distance correlations between each predictor and the mean slant estimates across the set of conditions, as shown in the graphs. Distance correlation ($dCor$) is a normalized measure of the codependence of two predictors that can detect nonlinear as well as linear relationships ([Székely, Rizzo, & Bakirov, 2007](#)). Overall, there is a strong relation between scaling contrast and the mean slant estimates ($dCor = 0.68$), but there are conditions with approximately the same scaling contrast that have widely varying slant estimates. For example, the scaling contrast for conditions with 60° slant and 10° FOV is almost the same as for conditions with 40° slant and 20° FOV (0.152 vs. 0.148), but mean slant estimates in these conditions ranged from 10° to 70°. For our conditions, the distance correlation between slant estimates and scaling contrast was lower than the distance correlation between slant estimates and actual slant ($dCor = 0.76$) or slant from compression ($dCor = 0.90$). Although scaling contrast can account for a large portion of the overall variance in slant estimates, it does not accurately predict the observed effects of varying FOV and cannot account for the large observed effects of stretching and compressing the surface textures.

Another feature of scaling contrast is that it is a nonlinear function of slant, which [Todd et al. \(2007\)](#) suggest could account for nonlinear psychometric functions observed in studies of perceived slant from texture. To evaluate this possibility, we compared the shapes of psychometric functions observed in [Experiments 1](#) and [2](#) to the predictions of the scaling contrast model. The left graphs in [Figure 13](#) plot the mean slant estimates from monocular conditions with consistent cues as a function of slant and the slopes of these functions between successive slants. The right graphs plot predicted psychometric functions if perceived slant were a linear function of scaling contrast and the corresponding slopes of these functions. Although both slant estimates and scaling contrast show upward curvature as a function of slant, the qualitative shapes are different. The observed psychometric functions are highly nonlinear at low slants and become more linear at the highest slants, whereas the scaling contrast functions are linear at low slants and then become increasingly nonlinear at higher slants. For the monocular conditions tested here, scaling contrast does not accurately account for the nonlinear relation between simulated slant and perceived slant.

Integration of texture and stereo slant cues

The results from binocular conditions suggest that texture compression and stereo slant cues are also

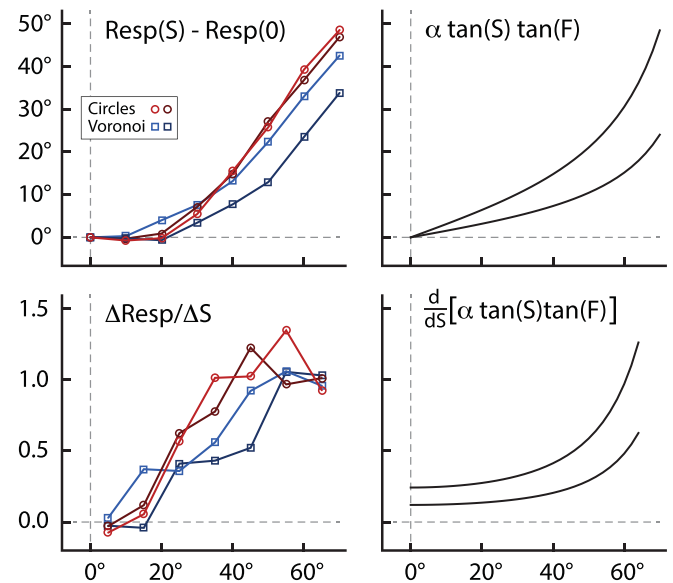


Figure 13. Comparison between the nonlinearity of the psychometric functions from monocular conditions and the nonlinearity predicted by the scaling contrast model. The upper left graph re-plots the mean slant settings as a function of slant from monocular conditions with consistent cues, and the lower left graph plots the slopes between successive slants. The mean response at zero was subtracted from the psychometric functions to remove constant bias. The upper right graph plots predicted psychometric functions if perceived slant were proportional to scaling contrast, and the lower right graph plots the slopes of these functions. The scaling contrast for a planar surface is $\tan(S) \cdot \tan(\text{FOV}/2)$. This was scaled by a factor of $a = 100$ to approximately match the overall range of responses in the 20° FOV conditions. The observed psychometric functions are highly nonlinear at low slants, with slopes systematically increasing from 0° to 50°, and then become more linear at the higher slants. In contrast, the scaling contrast model predicts that psychometric functions would be approximately linear for low slants and become increasingly nonlinear at higher slants.

integrated according to the relative reliability. The reliability of texture compression is highly dependent on slant ([Knill, 1998c](#)), so the weighting of texture compression relative to stereo would be expected to increase with slant, as observed in our results. Some previous studies have demonstrated optimal integration of texture and stereo cues ([Hillis et al., 2004](#); [Knill & Saunders, 2003](#)). These studies varied all texture cues in a consistent way, whereas our experiment selectively varied texture compression. If texture compression is a strong determinant of perceived slant from texture, as indicated by the results from the monocular conditions, then the effect of varying only texture compression should be close to the effect of varying all texture cues. Consistent with this prediction, the weighting of texture compression in our binocular conditions was similar to

the weighting of texture information in previous studies that tested conflicting texture and stereo cues (Hillis et al., 2004; Knill & Saunders, 2003; Saunders & Chen, 2015).

Conclusions

Our results demonstrate that texture compression has a strong influence on perceived 3D slant from texture. Previous studies have observed some conflicting findings. We used a cross-modal slant estimation task and small cue conflicts to avoid some potential issues with previous methods, and we observed clear effects of varying texture compression on perceived slant.

The influence of texture compression on perceived slant was modulated by FOV and texture regularity. The effects of FOV and texture regularity could potentially be explained in terms of the reliability of different texture cues. Our findings support a model that integrates multiple sources of texture information to estimate 3D surface slant in a close to optimal manner, with a large contribution from texture compression when the field of view is 20° or less.

Keywords: slant, depth, texture, cue integration

Acknowledgments

Supported by grants from the Hong Kong Research Grants Council (GRF HKU-752010H), China Postdoctoral Science Foundation (2018M630410), Key Specialist Projects of Shanghai Municipal Commission of Health and Family Planning (ZK2015B01), and Programs Foundation of Shanghai Municipal Commission of Health and Family Planning (201540114).

Commercial relationships: none.

Corresponding author: Jeffrey A. Saunders.

Email: jsaun@hku.hk.

Address: Department of Psychology, The University of Hong Kong, Hong Kong.

References

- Braunstein, M. L., & Payne, J. W. (1969). Perspective and form ratio as determinants of relative slant judgments. *Journal of Experimental Psychology*, *81*, 584–590.
- Buckley, D., Frisby, J. P., & Blake, A. (1996). Does the human visual system implement an ideal observer theory of slant from texture?. *Vision Research*, *36*(8), 1163–1176.
- Chen, Z., & Saunders, J. A. (2019). Perception of 3D slant from textures with and without aligned spectral components. *Journal of vision*, *19*(4), 7–7.
- Cutting, J. E., & Millard, R. T. (1984). Three gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General*, *113*(2), 198–216.
- Fritsch, F. N., & Carlson, R. E. (1980). Monotone piecewise cubic interpolation. *SIAM Journal on Numerical Analysis*, *17*(2), 238–246, doi:10.1137/0717021.
- Gärding, J. (1992). Shape from texture for smooth curved surfaces in perspective projection. *Journal of Mathematical Imaging and Vision*, *2*(4), 327–350.
- Girshick, A. R., & Banks, M. S. (2009). Probabilistic combination of slant information: Weighted averaging and robustness as optimal percepts. *Journal of Vision*, *9*(9), 8.1–8.20.
- Hillis, J. M., Watt, S. J., Landy, M. S., & Banks, M. S. (2004). Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision*, *4*(12), 967–992.
- 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., & Saunders, J. A. (2003). Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision Research*, *43*, 2539–2558.
- Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. (1995). Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision Research*, *35*(3), 389–412.
- Norman, J. F., Crabtree, C. E., Bartholomew, A. N., & Ferrell, E. L. (2009). Aging and the perception of slant from optical texture, motion parallax, and binocular disparity. *Perception & Psychophysics*, *71*(1), 116–130.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, F. Leisch, & A. Zeileis (Eds.), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003)*

- (p. 10). Vienna, Austria: R Foundation for Statistical Computing.
- Rosas, P., Wichmann, F. A., & Wagemans, J. (2004). Some observations on the effects of slant and texture type on slant-from-texture. *Vision Research*, 44(13), 1511–1535.
- Rosenholtz, R., & Malik, J. (1997). Surface orientation from texture: isotropy or homogeneity (or both)? *Vision Research*, 37(16), 2283–2293.
- Saunders, J. A. (2003). The effect of texture relief on perception of slant from texture. *Perception*, 32(2), 211–233.
- Saunders, J. A., & Backus, B. T. (2006). Perception of surface slant from oriented textures. *Journal of Vision*, 6(9), 882–897.
- Saunders, J. A., & Chen, Z. (2015). Perceptual biases and cue weighting in perception of 3D slant from texture and stereo information. *Journal of Vision*, 15(2), 14.
- Székely, G. L., Rizzo, M. L., & Bakirov, N. K. (2007). Measuring and testing independence by correlation of distances. *Annals of Statistics*, 35(6), 2769–2794.
- Tam, D. M., Shin, J., & Li, A. (2013). Dominance of orientation over frequency in the perception of 3-D slant and shape. *PLoS One*, 8(5), e64958.
- Tibau, S., Willems, B., Van Den Bergh, E., & Wagemans, J. (2001). The role of the centre of projection in the estimation of slant from texture of planar surfaces. *Perception*, 30(2), 185–193.
- Todd, J. T., Christensen, J. C., & Guckes, K. M. (2010). Are discrimination thresholds a valid measure of variance for judgments of slant from texture? *Journal of Vision*, 10(2), 20.1–20.18.
- Todd, J. T., Thaler, L., & Dijkstra, T. M. H. (2005). The effects of field of view on the perception of 3D slant from texture. *Vision Research*, 45, 1501–1517.
- Todd, J. T., Thaler, L., Dijkstra, T. M. H., Koenderink, J. J., & Kappers, A. M. L. (2007). The effects of viewing angle, camera angle and sign of surface curvature on the perception of 3D shape from texture. *Journal of Vision*, 7(12), 1–16.

Appendix

This appendix describes the model and assumptions used to fit the relative weighting of texture compression and scaling cues. The general concept of the analysis was presented earlier in the Results section of [Experiment 1](#).

For each experiment, the combined data from all subjects and conditions were fit with a hierarchical Bayesian model that included parameters representing the psychometric functions for consistent cue conditions, as well as cue weight parameters representing the influence of texture compression relative to other slant cues. The posterior probability distributions of the unknown parameters were estimated from the data by using Gibbs sampling to generate a Monte Carlo Markov chain of posterior samples. The Gibbs sampling was implemented using the JAGS package (Plummer, 2003). For each experiment, we generated a chain of 100,000 samples and thinned them by a factor of 10 to decorrelate the samples, resulting in 10,000 final samples of the posterior distribution. We computed the 95% highest density intervals from the samples for the variables of interest and differences across conditions and used the HDIs as measures of the credible ranges for the values and differences.

We assumed that slant estimates on consistent cue trials for a given subject and condition could be described by some nonlinear psychometric function of slant, $f(S)$, plus noise due to trial-to-trial variability in perceived slant and responses. Slant estimates on conflicting cue trials were assumed to vary around $f(S + w\Delta S)$, where ΔS is the difference between the slant specified by texture compression and texture scaling, and w is an unknown cue weight representing the relative influence of texture compression.

The model for responses on individual trials was a scaled t -distribution around the expected mean response for a subject and condition: $R \sim dt(f(S + w\Delta S), \sigma_R^{-2}, \nu)$. The parameter σ_R^{-2} is the inverse variance of responses around the mean across trials, and ν is the degree-of-freedom parameter for the t -distribution representing deviation from normality. We used t -distributions to model trial-to-trial variability rather than normal distributions to be more robust to outliers. The distributions of σ_R^{-2} and ν across subjects in each condition were modeled as gamma distributions, with generic hyperpriors for the gamma parameters.

The cue weight parameters were the main variables of interest. The distribution of cue weights across subjects in a condition was modeled as a normal distribution with unknown mean and variance. Each FOV and texture type condition had separate group mean and variance parameters, and the hyperpriors for these parameters were generic and uninformative. We used the posterior samples of the mean cue weight parameters to evaluate the influence of texture compression across the various conditions.

We modeled the nonlinear psychometric functions for consistent cue trials as cubic splines with control points at slants of 0° , 35° , and 70° . The spline functions

can be expressed in terms of the values and derivatives at these points:

$$f(S) = \begin{cases} f_0 + f'_0 \left(\frac{S}{35} \right) + 3(f_{35} - f_0) \\ \quad - 35(f'_0 + 2f'_{35}) \left(\frac{S}{35} \right)^2 \\ \quad + (2(f_{35} - 2f_0) + 35(f'_0 + f'_{35})) \left(\frac{S}{35} \right)^3, \\ \quad S \in (0, 35) \\ f_{35} + f'_{35} \left(\frac{S - 35}{35} \right) + 3(f_{70} - f_{35}) \\ \quad - 35(f'_{35} + 2f'_{70}) \left(\frac{S - 35}{35} \right)^2 \\ \quad + (2(f_{70} - 2f_{35}) + 35(f'_{35} + f'_{70})) \\ \quad \left(\frac{S - 35}{35} \right)^3, \quad S \in (35, 70) \end{cases}$$

We constrained the free parameters of the splines so that the psychometric functions were strictly monotonic as a function of slant within the range. We required that $f_{70} > f_{35} > f_0$ and that the derivatives be positive and less than three times the slope between endpoints.

These constraints are sufficient to ensure monotonicity (Fritsch & Carlson, 1980). The range $[0^\circ, 70^\circ]$ spans the set of slants used in our experiments, but the perceptually matching slant for cue conflict conditions ($S + w\Delta S$) could potentially be outside this range if w is less than zero or much greater than one. For negative slants, we assumed symmetry around the zero point: $f(-S) - f_0 = -(f(S) - f_0)$. For slant values larger than 70° , we extended $f(S)$ by linearly extrapolating from f_{70} and f'_{70} .

For the Bayesian model, we used the following priors for parameters describing the psychometric functions of individual subjects and conditions:

$$\begin{aligned} f_0 &\sim \text{dnorm}(0, \sigma_{\text{zero}}^{-2}) \\ f_{70} &\sim \text{dgamma}(\alpha_{\text{max}}, \beta_{\text{max}}) \\ f_{35}/f_{70} &\sim \text{dbeta}(\alpha_{\text{mid}}, \beta_{\text{mid}}) \\ f'_0 &\sim \text{dunif}(0, 3 \cdot (f_{35} - f_0)/35) \\ f'_{35} &\sim \text{dunif}(0, 3 \cdot \min(f_{35} - f_0, f_{70} - f_{35})/35) \\ f'_{70} &\sim \text{dunif}(0, 3 \cdot (f_{35} - f_0)/35) \end{aligned}$$

where $\sigma_{\text{zero}}^{-2}$ is the variance of f_0 across subjects, α_{max} and β_{max} are parameters for a gamma distribution fit to the distribution of f_{70} across subjects, and $(\alpha_{\text{mid}}, \beta_{\text{mid}})$ are parameters for a beta distribution fit to the distribution of the ratio of f_{35} to f_{70} . We used separate distributional parameters for the different FOV and texture type conditions and uninformative hyperpriors on these parameters.