

Slant from Texture and Disparity Cues: Optimal Cue Combination

James M. Hillis

Department of Psychology, University of Pennsylvania,
Philadelphia, PA, USA



Simon J. Watt

Department of Psychology, University of Wales
Bangor, Wales, UK



Michael S. Landy

Department of Psychology & Center for Neural Science,
New York University
New York, NY, USA



Martin S. Banks

Vision Science Program, Department of Psychology, &
Wills Neuroscience Institute, University of California,
Berkeley, CA, USA



How does the visual system combine information from different depth cues to estimate 3D scene parameters? We tested a maximum-likelihood estimation (MLE) model of cue combination for perspective (texture) and binocular disparity cues to surface slant. By factoring the reliability of each cue into the combination process, MLE provides more reliable estimates of slant than would be available from either cue alone. We measured the reliability of each cue in isolation across a range of slants and distances using a slant-discrimination task. The reliability of the texture cue increases as $|\text{slant}|$ increases and does not change with distance. The reliability of the disparity cue decreases as distance increases and varies with slant in a way that also depends on viewing distance. The trends in the single-cue data can be understood in terms of the information available in the retinal images and issues related to solving the binocular correspondence problem. To test the MLE model, we measured perceived slant of two-cue stimuli when disparity and texture were in conflict and the reliability of slant estimation when both cues were available. Results from the two-cue study indicate, consistent with the MLE model, that observers weight each cue according to its relative reliability: disparity weight decreased as distance and $|\text{slant}|$ increased. We also observed the expected improvement in slant estimation when both cues were available. With few discrepancies, our data indicate that observers combine cues in a statistically optimal fashion and thereby reduce the variance of slant estimates below that which could be achieved from either cue alone. These results are consistent with other studies that quantitatively examined the MLE model of cue combination. Thus, there is a growing empirical consensus that MLE provides a good quantitative account of cue combination and that sensory information is used in a manner that maximizes the precision of perceptual estimates.

Keywords: depth perception, cue combination, stereopsis, Bayesian perception, texture gradient

Introduction

The fundamental problem in depth perception is due to the geometry of perspective projection, which reduces the 3-dimensional coordinates of the visual scene to the 2-dimensional coordinates of the retinal images. The third dimension of space has to be inferred from the 2D images. The visual system uses several sources of information—“depth cues” such as disparity, perspective, and motion parallax—to estimate the layout of the 3D scene. Estimates based on each individual cue are subject to error. By combining information from several depth cues, the visual system could estimate 3D layout with greater precision across a wider variety of viewing situations than it could by relying on any one cue alone. To realize this advantage, the reliability of each depth cue must be factored into the combination rule. Factoring in reliability is complicated

because the reliability of individual depth cues depends on scene parameters in different ways. For example, the reliability of binocular disparity cues to depth decreases as viewing distance increases, whereas the reliability of perspective cues, such as the texture gradient, does not depend on distance. Are depth cues combined in a manner that is appropriate for the varying cue reliabilities that occur with different scene geometries? To examine this question, we compared human slant discrimination ability based on disparity and texture cues to a model of statistically optimal cue combination. Slant estimation from texture and disparity is an interesting case to examine because the reliabilities of disparity and texture cues vary in different ways with slant and viewing distance. We first measured the reliability of slant estimates from each cue in isolation across a range of slants and distances. We used an optimal cue-combination rule to predict the appearance of two-cue stimuli and the precision of slant estimation with

two-cue stimuli. We then compared these predictions to the results of two-cue slant discrimination experiments.

Optimal cue combination

Visual estimates of slant from any depth cue are subject to error. For example, perceived slant from a given texture gradient will vary from one instance to another due to the statistical nature of slant information from texture and errors in the measurement of the gradient (Blake, Bülthoff, & Sheinberg, 1993; Cutting & Millard, 1984; Knill, 1998a). When more than one depth cue is available and informative, one can in principle reduce the uncertainty associated with any one of the cues by combining across cues (for a review and derivation of the following results, see Oruç, Maloney, & Landy, 2003).

One approach to optimizing cue combination is statistical: what cue-combination rule results in an estimator that is unbiased and has minimum variance? Assume that the observer has unbiased estimates \hat{S}_d and \hat{S}_t of the slant of a surface based on disparity and texture cues, respectively. Assume further that these estimators are uncorrelated and have variances \mathbf{s}_d^2 and \mathbf{s}_t^2 . If we combine the two estimates linearly, the rule that yields the minimum-variance, unbiased estimate is a weighted average that satisfies (Cochran, 1937)

$$\hat{S} = w_d \hat{S}_d + w_t \hat{S}_t, \quad (1)$$

where

$$w_d = \frac{r_d}{r_d + r_t} \quad \text{and} \quad w_t = \frac{r_t}{r_d + r_t}, \quad (2)$$

and r_d and r_t are the *reliabilities* of the two cues (e.g., $r_d = 1/\mathbf{s}_d^2$). Furthermore, if errors associated with the individual estimators are Gaussian, no other (non-linear) rule has lower variance.

An alternative approach is to apply Bayesian methods (for a review, see Mamassian, Landy, & Maloney, 2002). In the absence of any immediate consequences to an observer's actions (payoffs and penalties), the maximum *a posteriori* (MAP) estimate is typically employed. That is, the observer chooses a slant estimate \hat{S} that is most probable given the image data. We assume the image data can be segregated into those data I_d used to estimate slant from disparity and I_t used to estimate slant from texture. Thus, we choose the value of \hat{S} that maximizes $p(\hat{S} | I_d, I_t)$. Applying Bayes' rule, and assuming that the two cues are conditionally independent, we derive

$$p(\hat{S} | I_d, I_t) \propto p(I_d | \hat{S}) p(I_t | \hat{S}) p(\hat{S}). \quad (3)$$

The first two terms on the right side of the equation are the likelihood functions for each cue characterizing the probability of observing the image data if \hat{S} is the actual slant. The last term is the prior distribution, which is the probability of observing \hat{S} in the scene, independent of

the image data. If the likelihoods and prior are Gaussian, the MAP estimate has the same form as the minimum variance, linear combination estimate

$$\hat{S} = w_d \hat{S}_d + w_t \hat{S}_t + w_p \hat{S}_p, \quad (4)$$

where

$$w_d = \frac{r_d}{r_d + r_t + r_p}, \quad w_t = \frac{r_t}{r_d + r_t + r_p} \quad \text{and} \quad w_p = \frac{r_p}{r_d + r_t + r_p}. \quad (5)$$

Here, \hat{S}_d and \hat{S}_t are the maximum-likelihood estimates the observer would have made from each cue in isolation (the mean of the respective Gaussian distributions), and \hat{S}_p is the mean of the prior. The r_i are the reliabilities of the respective distributions (likelihoods and prior). If the prior has large variance relative to the individual cue likelihoods, Eqs. 4 and 5 reduce to Eqs. 1 and 2, which also yields the most likely slant to have caused the current sensory data (i.e., it is the maximum-likelihood estimate or MLE). For our conditions, the variance of the individual cues is much smaller than the prior's variance (see *Ideal observer models* in the Discussion), so we will use Eqs. 1 and 2 throughout the paper.

By following the strategy described by Eq. 2, the variance of the weighted average \hat{S} is

$$\mathbf{s}^2 = \frac{\mathbf{s}_d^2 \mathbf{s}_t^2}{\mathbf{s}_d^2 + \mathbf{s}_t^2} \quad \text{or, equivalently, } r = r_d + r_t, \quad (6)$$

which is lower than the variance of either single-cue estimate.

Many investigations of sensory cue combination have shown that cue reliability is taken into account in the estimation process (Backus & Banks, 1999; Banks, Hooge & Backus, 2001; Buckley & Frisby, 1993; Frisby, Buckley, & Horsman, 1995; Rogers & Bradshaw, 1995; van Beers, Sittig, & Denier van der Gon, 1998; van Beers, Wolpert, & Haggard, 2002; Young, Landy, & Maloney, 1993). Four of these—Ernst and Banks (2002), Gepshtein and Banks (2003), Knill and Saunders (2003), and Landy and Kojima (2001)—have examined whether the combination is the quantitatively optimal one. All four reported that the combination is quite close to the one predicted by Eqs. 1-2. In these four experiments, the variances of estimates derived from single cues were measured by conducting two-interval, forced-choice (2-IFC) discrimination experiments when only one cue was informative. For example, Ernst and Banks (2002) conducted size-discrimination experiments for vision alone and haptics alone and then fit cumulative Gaussians to the two psychometric functions. The variance parameter of the Gaussians provided estimates of the variances of the underlying visual and haptic estimators. Eqs. 1, 2 and 6 were then successfully used to predict the results of two-cue (visual-haptic) experiments.

Preview

The experiments presented here used the strategy of Ernst and Banks (2002) to ask whether texture and disparity cues to slant are combined in a statistically optimal fashion. The reliability of texture and disparity cues to slant vary with viewing geometry in different ways. First, the reliability of texture should increase with increasing slant because at large slants the image changes associated with a given change in slant increase (Blake et al., 1993; Knill, 1998a). This relationship between reliability and slant is reflected in human performance (Knill, 1998b). Theoretical and empirical analyses of the reliability of disparity as a function of slant have not been conducted, but it is unlikely that it changes significantly (Banks et al., 2001). Second, because the magnitude of binocular disparities for a given depth difference decreases as viewing distance increases, the reliability of slant and curvature estimated from binocular disparity should decrease as viewing distance is increased. Experiments confirm that it does (Howard & Rogers, 2002; Ogle, 1950). On theoretical grounds, the reliability of texture-specified slant, for a fixed retinal-image density, should not change with distance. If a given textured surface is doubled in size and viewed from twice the distance, the retinal image is unchanged. Thus, optimal combination of disparity and texture cues to slant should involve complex changes in the weights given to the two cues depending on base slant and viewing distance.

We looked for evidence that the visual system weights the two cues appropriately across a range of slants and distances. Because the reliability of the texture cue increases with slant, we expect the texture weight to increase as slant increases. Because the reliability of disparity decreases as viewing distance increases, we expect the texture weight to increase as distance increases. As in the previous studies, we determined the reliability of the individual cues with 2-IFC discrimination experiments. In our case, this was done for various viewing distances and base slants. Then, we measured the apparent slant and slant discrimination performance for two-cue stimuli. As we shall see, the MLE cue combination predictions based on the single-cue experiments (Eqs. 1-2 and 6) were largely in accord with the data from the two-cue experiment.

Methods

Subjects

Two observers participated. One was not aware of the experimental hypotheses. Both observers had normal stereopsis and did not manifest eye misalignment in normal viewing situations.

Apparatus

All stimuli were displayed on a custom-designed stereoscope with two mirrors and two CRTs (one for each

eye; see Backus, Banks, van Ee, & Crowell, 1999). Each mirror and CRT was attached to an arm that rotated about a vertical axis passing through the eye's center of rotation. With this arrangement, the eye and stereoscope arm rotate on a common axis, so when we change the vergence distance, the mapping between the stimulus array and the retina is unaltered (for fixed accommodation).

We used anti-aliasing to specify dot position to sub-pixel accuracy. To ensure accurate reproduction of visual direction, we spatially calibrated each CRT to eliminate distortions in the images on the CRTs (for details, see Backus et al., 1999).

The observer's head position was stabilized using a bite bar fastened to an adjustable mount. Each observer had a personal mount so that the vertical axes of rotation of left and right eyes were collinear with the rotation axes of the two stereoscope arms (for details, Hillis & Banks, 2001). The optical distance between the center of rotation of each eye and the face of the CRT was 40 cm.

Stimuli

Stimuli were virtual planes slanted about a vertical axis (i.e., tilt = 0 deg). We independently manipulated two cues to slant: binocular disparity and texture. In single-cue measurements, we isolated one or the other of the two cues. In two-cue measurements, both cues were informative, but could have different slant values. Viewing distance was 19.1, 57.3, or 171.9 cm. Example stimuli are shown in Fig. 1.

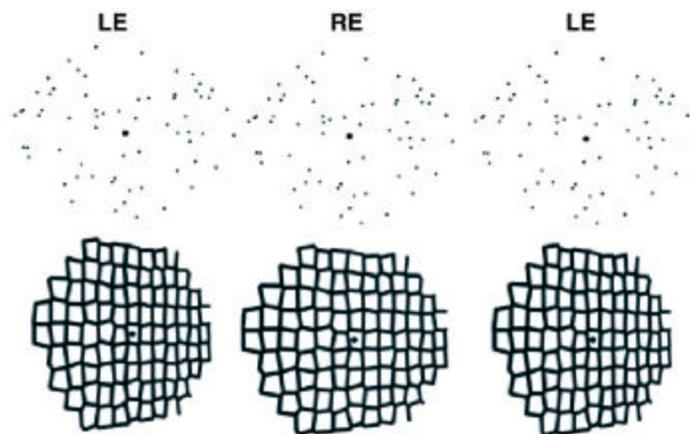


Figure 1. Examples of the stimuli. Cross-fuse or divergently fuse to see the appropriate slants. The upper stimulus is an example of the disparity-alone stimulus. It has a negative slant (right side near). The lower row provides examples of the texture-alone stimulus when viewed monocularly and the disparity-texture stimulus when viewed binocularly. The disparity- and texture-specified slants are positive (right side far).

The texture cue was the perspective projection of planar patches textured with Voronoi patterns (de Berg, van Kreveld, Overmars, & Schwarzkopf, 2000; Fig. 1, bottom panel). These patterns were generated from a

jittered grid of dots. On a frontoparallel plane, a regular grid of points was defined. Then, each point on the dot grid was perturbed horizontally and vertically (uniform distribution from -0.3 to 0.3 deg). The Voronoi pattern defined by these points was then computed. Finally, the resulting textured plane was rotated by an amount equal to the texture-defined slant. To isolate the texture cue, the stimuli were viewed monocularly. The visible portion of the plane was elliptical with a height of 15 deg. The width on each presentation was randomly chosen from a uniform distribution from 15 to 20 deg when the stimulus was frontoparallel. The stimulus was then rotated to the appropriate slant. Thus, the retinal shape of the stimulus outline was an unreliable cue to slant.

The disparity cue to surface slant was the difference between left- and right-eye projections (calculated for each observer's inter-pupillary distance). To isolate the disparity cue, the stimulus was defined by sparse random dots (Fig. 1, top panel). Each stimulus consisted of 64 dots, with positions randomly drawn from a uniform distribution (note that on average the texture gradient specified by the dots was therefore consistent with a fronto-parallel plane). Dot density was ~ 0.3 dots/deg².

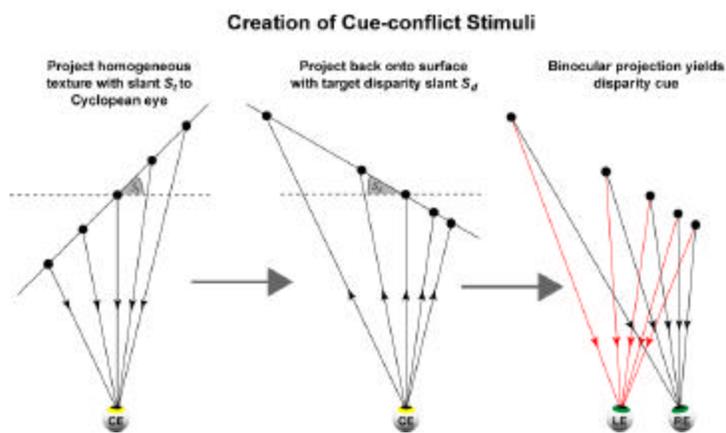


Figure 2. Creation of the cue-conflict stimuli. Left panel: Perspective projection of a homogeneously textured surface with the Cyclopean eye as the center of projection. This projection creates the texture-specified slant, S_t . The rays from the surface toward the eye are used in the next step. Middle panel: A virtual surface with the disparity-specified slant, S_d , is created. The rays from the first step are back-projected from the Cyclopean eye to find their intersections with the disparity-defined surface. They are marked in the diagram with black points. Right panel: Viewing the black points binocularly yields the cue-conflict stimulus containing the texture-specified slant in the left panel and the disparity-specified slant in the middle panel.

When both cues were present, disparity and texture could be consistent ($S_d = S_t$) or they could be in conflict. In the no-conflict case, homogeneous Voronoi-textured surfaces were projected directly to the two eyes. In cue-conflict cases, we first calculated a perspective projection of

the texture with slant S_t at the Cyclopean eye (Fig. 2, left panel). We then found the intersections of rays through this Cyclopean projection with a surface patch at the disparity-specified slant S_d (Fig. 2, middle panel). The markings on this latter surface were then projected to the left and right eyes to form the two monocular images (Fig. 2, right panel).

Control experiments and procedures to validate single-cue measurements

We went to some lengths to ensure that the single-cue experiments measured the variances of the disparity and texture estimators in a fashion appropriate for making two-cue predictions. In this section, we describe control experiments and methodological procedures we used to that end.

To make sure that only binocular information determined slant discrimination in the disparity-alone case, we conducted two control experiments.

First, we wanted to make sure that the stimulus did not provide a monocular cue to slant. To check this, we measured monocular thresholds as a function of slant for the 64 -dot stimulus. Observers could not reliably discriminate anything but large slant changes, and those changes were much larger than the thresholds in the single-cue experiment. We conclude that there is no useful monocular slant information in the 64 -dot random-dot stimulus.

Second, we wanted to make sure that we presented enough dots in the display for disparity-based thresholds to be as low as possible while still isolating the disparity estimator. For this purpose, we measured slant-discrimination thresholds as a function of the number of dots. We ran three conditions: 1) texture-specified slant = 0 deg and disparity-specified slant = 0 deg + d (where d is the increment or decrement given to the base slant in order to obtain a threshold), 2) texture slant = 0 and disparity-specified slant = 45 + d , and 3) texture slant = 45 and disparity-specified slant = 45 + d . The results are plotted in Fig. 3 which shows the just-discriminable change in slant as a function of dot number for the three conditions. For observer JMH, discrimination thresholds decreased as dot number was increased from 2 to 32 and then thresholds reached an asymptote by 32 - 64 dots. With 64 dots, disparity-based thresholds were essentially as low as they could be. The results were more complicated for observer ACD. When the disparity-specified slant was 0 deg, her data were quite similar to JMH's. However, when the disparity slant was 45 deg, thresholds decreased from 4 - 64 dots and then increased with more than 64 dots. One would observe such an effect if the uninformative texture signal was given increasing weight with increasing dot number. Thus, ACD may have given some weight to the uninformative texture signal when shown the random-dot stimulus at base slants different from 0 deg. We will return

to this point when we discuss her two-cue data (in *Summary of results* in the Discussion).

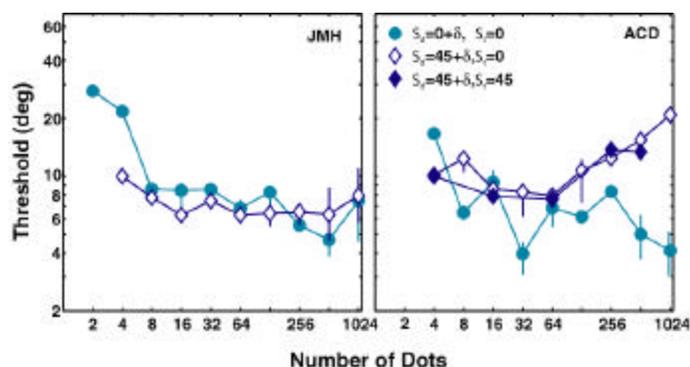


Figure 3. Slant discrimination thresholds for the disparity-alone condition as a function of the number of dots on the surface for the 57.3-cm viewing distance. Filled blue circles represent discrimination thresholds for a base slant of 0 deg ($S_d = 0 + d$, $S_t = 0$). Unfilled diamonds represent thresholds for a base slant of 45 deg when the texture specifies a slant of 0 deg ($S_d = 45 + d$, $S_t = 0$). Filled diamonds represent thresholds for a base slant of 45 deg when the texture specifies a fixed slant of 45 deg ($S_d = 45 + d$, $S_t = 45$). Error bars are 95% confidence intervals.

To combine two cues for slant, the cues must be promoted to the same units. For example, disparity signals alone do not provide a slant estimate because they must be scaled or normalized for distance (Gårding, Porrill, Mayhew, & Frisby, 1995). We were concerned that observers might perform the slant-discrimination task in the single-cue, disparity-alone case by comparing only the disparity gradients in the two stimulus intervals. Said another way, they could in principle perform the task without normalizing the disparity signals into slant estimates. Because estimates from the two cues must be in the same units, the raw disparity gradient without normalization would not work for combining disparity and texture signals in the two-cue experiment. For this reason, we looked for evidence that observers normalize the disparity signal in a discrimination task with our disparity-alone stimuli. We did so by having observers perform the slant-discrimination task with the disparity-alone stimulus with the comparison stimuli appearing at different distances relative to the standard stimulus.

Five observers participated. Two were unaware of the experimental hypotheses. All had normal stereopsis and did not manifest eye misalignment in normal viewing situations.

A different apparatus was used than the one in the main experiment. Stimuli were displayed on a CRT at a distance of 57 cm. Dichoptic presentation of the left and right eye's images was achieved using CrystalEyes™ liquid-crystal shutter glasses. Left and right eye images were rendered on alternate frames so each eye's image was drawn only when the corresponding shutter was open. The monitor refresh rate was 100 Hz, so each eye's image was

redrawn at 50 Hz. The stimuli were drawn using the red phosphor only because this minimized cross-talk through the shutter glasses. The room was otherwise dark. Precise reproduction of visual directions was achieved using the same anti-aliasing and spatial calibration techniques as in the main experiment. The same bitebar set up was used to position and stabilize the observer's head.

The stimuli were virtual planes slanted about a vertical axis. They were very similar to the stimuli used in the disparity-alone condition in the main experiment (Fig 1, top): random-dot stereograms with the same dot density and size as in the main experiment. Stimulus size for a given presentation was drawn randomly from a uniform distribution from 12.5-17.5 deg. The stimuli were generated taking each observer's inter-pupillary distance into account. Changes in simulated distance were achieved by shifting the two eyes' images laterally on the CRT; this technique produces the correct eye vergence and horizontal gradient of vertical disparity for the simulated distance. The size of the stimulus, dot density, and dot size were the same in angular terms at each simulated distance.

As in the main experiment, two stimuli were presented sequentially and observers indicated the one containing the apparently greater slant. No feedback was provided. The standard stimulus was always presented at a simulated distance of 57 cm, and the comparison stimulus was presented at one of several simulated distances (selected randomly before each trial). Thus, observers had to judge relative slant even when the standard and comparison were presented at different simulated distances. There were three comparison distances (45.2, 57.0, and 71.8 cm) and two base slants (± 30 deg). Each stimulus was presented for 1 sec with an inter-stimulus interval of 1.5 sec. To facilitate fusion, a fixation marker appeared 500 msec prior to each stimulus presentation at the distance of the upcoming stimulus. Observers were all able to make a vergence eye movement and fuse the fixation marker before the stimulus appeared. 1-down/2-up and 2-down/1-up staircases were used to vary the slant of the comparison stimulus. For each base slant/distance combination, both reversal rules were used to sample points on the psychometric function either side of the 50% point. At least four staircases were employed for each psychometric function, corresponding to approximately 200 trials per function. We fit the resulting psychometric data with a cumulative Gaussian using a maximum-likelihood criterion.

In Fig. 4, we plot the slants of the comparison stimulus that had the same perceived slant on average as the standard stimulus as a function of the simulated distance of the comparison. There were no significant differences between base slants of -30 and $+30$ deg, so we pooled the data from the two slants. If observers performed the task by comparing slants (as they were instructed), they would have to take distance into account (Eq. 11). If they did so veridically, the PSEs would all have the same slant as the standard stimulus (horizontal dashed line). If, on the other hand, observers performed the task by comparing disparity

gradients, they would not need to take distance into account. In this case, the PSEs would have the same disparity gradients ($\#ISRs$), but different slants (diagonal dashed line).

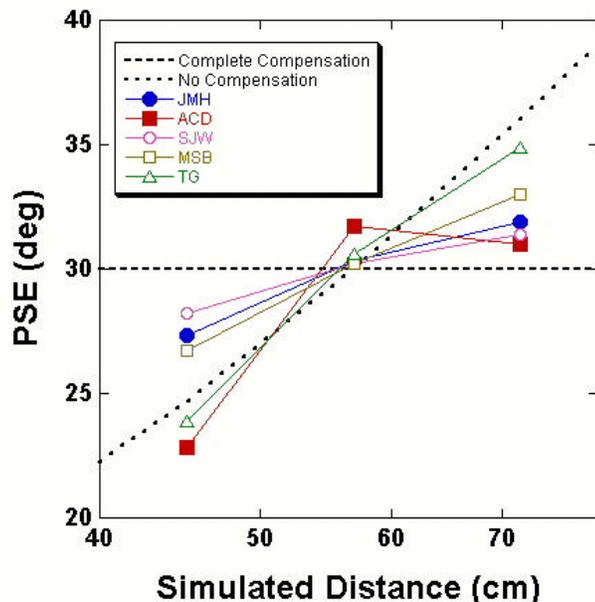


Figure 4. PSEs as a function of distance. The slant of the comparison stimulus that was on average perceived as the same slant as a standard at ± 30 deg slant and a distance of 57 cm is plotted as a function of the distance of the comparison stimulus. The dashed horizontal line represents the predicted PSEs if observers compensated for the change in distance veridically. The dotted diagonal curve represent the predicted PSEs if observers made their judgments from the disparity gradients of the two stimuli only; that is, if they failed to compensate for distance. The filled symbols are the data from the two observers whose data are presented in the main experiment. The unfilled symbols are the data from three who did not participate in the main experiment.

The PSEs for four of the five observers were not consistent with the no-compensation predictions; thus, they *did* take the change in stimulus distance into account. However, none exhibited complete compensation, which means that the distance compensation was not veridical. The data from observers JMH and ACD are represented by filled symbols. They are the observers in the main experiments reported below. The results of this control experiment show that most observers (importantly including observers JMH and ACD) do not perform the slant-discrimination task by only comparing the disparity gradient in the two stimulus intervals.

As stated earlier, the single-cue data were used to specify the model's parameters (single-cue variances) and the model was then used to predict the two-cue data. An important assumption is that the appropriate variances are being measured in the single-cue experiments. One might question this assumption for the measurements of the

disparity estimator's reliability because a different type of stimulus was used in the single-cue, disparity-alone condition (random-dot stimuli) than in the two-cue condition (Voronoi stimuli). This concern cannot be addressed by using Voronoi stimuli in the disparity-alone condition because such stimuli provide salient texture cues to slant. We can, however, check the validity of using random-dot stimuli by comparing single-cue thresholds with those stimuli to two-cue thresholds when the texture weight is expected to be approximately zero. This check, which is described in the Results (*Just-noticeable differences*), confirmed the validity of our assumption for observer JMH (the only observer for whom the required data were available).

A similar concern can be raised about using monocular stimuli in the texture-alone condition to measure texture reliability in the two-cue experiment. The stimulus in the two-cue experiment is binocular, so the visual system receives two samples while there is only one sample in the texture-alone experiment. The two samples will not be the same because the texture-specified slant at the left eye necessarily differs from the texture slant at the right eye (Appendix). Thus, the visual system must integrate two texture-gradient signals into one binocular estimate before combining with the slant estimated from disparity. The presence of two samples of the texture gradient in the two-cue experiment might reduce the uncertainty associated the texture cue in a fashion similar to the reduction in contrast threshold with binocular viewing (Legge, 1984). We can check this by comparing the monocular single-cue thresholds to binocular two-cue thresholds when the texture weight is expected to be approximately one. This check, described in the Results (*Just-noticeable differences*), confirmed the validity of our assumption that the monocular measurements were a valid estimate of the variance of the texture estimator in the two-cue experiment.

We wanted to isolate the slant cues of disparity and texture in our experiments, so we had to consider whether other slant cues might be present in the display. Three cues—the blur gradient, accommodation, and the phosphor grid of the CRTs—always signaled a slant of 0 deg. If the observer failed to ignore those uninformative cues, the variances we measured would be higher than the true variances associated with disparity and texture. To reduce the salience of all three cues, we placed diffusers on the faces of the CRTs to blur the stimuli slightly. Blurring the stimuli decreases the probability that the observers used the blur gradient because the blur gradient is a less reliable depth cue with blurred as opposed to sharply focused stimuli (Mather & Smith, 2002). Blurring should also decrease the probability that accommodation was used as a depth cue because humans accommodate inaccurately if at all to blurred stimuli (Heath, 1956). The diffusers also made the phosphor grid invisible.

Procedure: Single-cue conditions

To estimate the reliabilities of the texture- and disparity-based slant estimates, we obtained psychometric functions for texture and disparity presented in isolation at several base slants (± 70 , ± 60 , ± 45 , ± 30 , ± 15 , and 0 deg) and distances (19.1, 57.3, and 171.9 cm). We used a 2-IFC task with no feedback. On each trial, the observer indicated which of two stimuli—one at the base slant and the other with an increment of $\pm DS$ —had the greater apparent slant. The stimuli were displayed for 1.5 sec with a 0.3-sec inter-stimulus interval. We used staircases to control the value of DS and four reversal rules—3-down/1-up, 1-down/3-up, 2-down/1-up, and 1-down/2-up—to sample points along the entire psychometric function. At least eight staircases were employed for each psychometric function, which corresponds to approximately 350-450 trials per function. In each session, at least four interleaved staircases were run: two base slants (one positive and one negative to avoid adaptation) with two staircases each. Viewing distance was fixed in each session.

Procedure: Two-cue conditions

The procedure in the two-cue conditions was the same as in the single-cue conditions except that a no-conflict stimulus (disparity- and texture-specified slants equal to one another) and a conflict stimulus (disparity and texture slants not necessarily equal) were presented on each trial. Fig. 5 depicts the disparity- and texture-defined slants of the no-conflict and conflict stimuli. In both panels, the slant specified by the disparity cue is plotted on the abscissa and the slant specified by the texture cue on the ordinate. The conflict stimulus had one cue set to a base slant ($S_b = \pm 60$, ± 30 , or 0 deg) and the other cue was perturbed. The left panel depicts the conflict stimuli when disparity was perturbed and the right panel the stimuli when texture was perturbed. The perturbed cue had incremental slants of ± 10 , ± 5 , or 0 deg relative to the unperturbed cue, so the conflict was always small. In previous work with quite similar stimuli, a difference of 10 deg between the disparity- and texture-specified slants was generally not detectable (Hillis, Ernst, Banks, & Landy, 2002). The five possible perturbed cue values are represented along the abscissa and ordinate in the left and right panels, respectively. On each trial, a conflict stimulus and a no-conflict stimulus were presented and the observer indicated the one containing the apparently greater slant. No feedback was given. The value of the no-conflict stimulus was varied according to staircase procedures (with different reversal rules as described above) to map out the psychometric function. At least four staircases were run per experimental session: two conflict conditions for each of two base slants.

Stimulus Variation in Two-cue Experiment

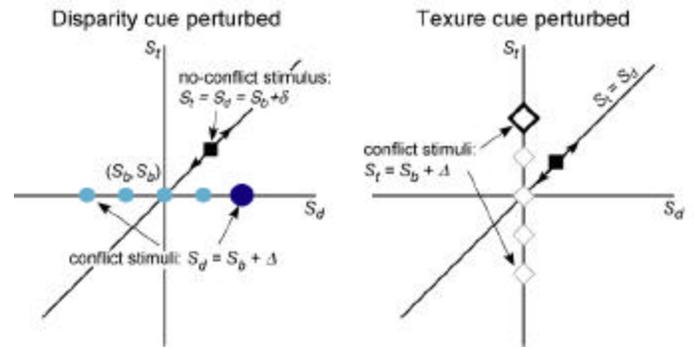


Figure 5. Depiction of the stimulus values and their manipulation in the two-cue experiment. Both graphs plot the disparity-specified slant (S_d) and the texture-specified slant (S_t) for the conflict and no-conflict stimuli. The base slant is represented by the origins (S_b, S_b). The conflict stimulus either had the disparity-specified slant perturbed from the base slant by D (depicted in the left panel) or the texture-specified slant perturbed by D (right panel). Five different conflicts were presented for each base slant and those are represented by the blue circles in the left panel and gray diamonds in the right panel. In the no-conflict stimulus, the disparity- and texture-specified slants were equal to one another. The staircase procedure varied the increments added to the base slant: $S_d = S_t = S_b + d$. For each conflict stimulus, there will be some value of the no-conflict stimulus that on average has the same perceived slant. This is the PSE.

Fig. 5 also shows how we determined the *point of subjective equality* (PSE), the value of the no-conflict stimulus with the same average perceived slant as the conflict stimulus. The enlarged bold symbols—the filled dark-blue circle in the left panel and open black diamond in the right—represent two particular conflict stimuli. D represents the incremental slant of the perturbed cue in the conflict stimulus and d represents the increment given to the no-conflict stimulus as the staircase procedure varies its slant. As d is increased and thereby the slant of the no-conflict stimulus is increased (represented in the figure by displacement up along the main diagonal where $S_t = S_d$), the observer will be increasingly likely to report that it had greater slant than the conflict stimulus. At some value of d , the no-conflict stimulus will on average have the same apparent slant as the conflict stimulus; this is the PSE. If the cue weights are constant across small variations in slant, we can determine the weights from this value of d .

Consider first the conflict stimulus. From Eqs. 1-2 and the fact that disparity-defined stimulus slant $S_d = S_t + \Delta = S_b + \Delta$ (where S_b is the base slant), the expected value of the estimated slant of the conflict stimulus is

$$\begin{aligned} S_c &= w_d (S_b + \Delta) + (1 - w_d) S_b \\ &= S_b + w_d \Delta. \end{aligned} \quad (7)$$

Now consider the no-conflict stimulus. From Eqs. 1-2 and the fact that $S_d = S_t = S_b + \mathbf{d}$, the expected value of the estimated slant of the no-conflict stimulus is

$$S_{nc} = w_d (S_b + \mathbf{d}) + (1 - w_d)(S_b + \mathbf{d}) = S_b + \mathbf{d}. \tag{8}$$

The conflict and no-conflict stimuli will have the same perceived slants when $S_c = S_{nc}$ and from Eqs. 7 and 8, we have

$$w_d = \mathbf{d} / \Delta. \tag{9}$$

Thus, the two-cue experiment yields an estimate of the PSE from which we can determine the weights given to disparity and texture. The assumption that the weights are constant for even small variations is inconsistent with statistically optimal slant estimation, in which the weights vary as a function of slant. However, given the precision of our measurements and the rate of change of cue reliability, the fixed local weight assumption provides a reasonable approximation.

We can also plot the percentage of judgments for which the no-conflict stimulus appeared to have greater slant as a function of its slant. The slopes of such psychometric functions index the discriminability of the stimuli (discussed below under *Just-noticeable differences*).

Results

Specifying the predictions

To quantify the predictions of the MLE model, we need estimates of the variances of the single-cue estimators (\mathbf{s}_d^2 and \mathbf{s}_t^2 , or equivalently, the reliabilities r_d and r_t in Eqs. 2 and 6). To estimate these variances, we fit the psychometric data with a cumulative Gaussian using a maximum-likelihood criterion. The standard deviations of the resulting functions were divided by $\sqrt{2}$ (because the psychophysical procedure was 2-IFC) to yield estimates of the standard deviations of the underlying slant estimators (Green & Swets, 1974). We call these just-noticeable differences (JNDs) because they represent the slant difference that is correctly discriminated ~76% of the time.

Fig. 6 shows the JND estimates. Each row of panels represents data from one observer. The left column shows the texture-alone data: JNDs in units of slant are plotted as a function of the absolute value of base slant (there was no apparent difference in the results for positive and negative slants). Different symbols represent data from different viewing distances. As expected, texture JNDs did not vary systematically with distance. Also as expected (Knill, 1998a), texture JNDs decreased as the absolute value of slant increased.

The middle column of Fig. 6 shows the data for the disparity-alone condition: slant JNDs are plotted as a function of the absolute value of base slant (there were no obvious differences for positive and negative slants). As expected from the viewing geometry (Eqs. 1 & 2 in Backus et al., 1999), disparity slant JNDs increased systematically with an increase in viewing distance. JNDs also tended to decrease with base slant at the medium and far viewing distances (see also Knill & Saunders, 2003). At the near viewing distance, JNDs tended to increase with base slant. In fact, as indicated by the symbol with the yellow star, JMH's thresholds were infinite at base slants of $\pm 70^\circ$ and a viewing distance of 19.1 cm because the binocular images could not be fused in this condition. This difference in the trend between near and far distances can be understood in terms of the retinal signal to slant.

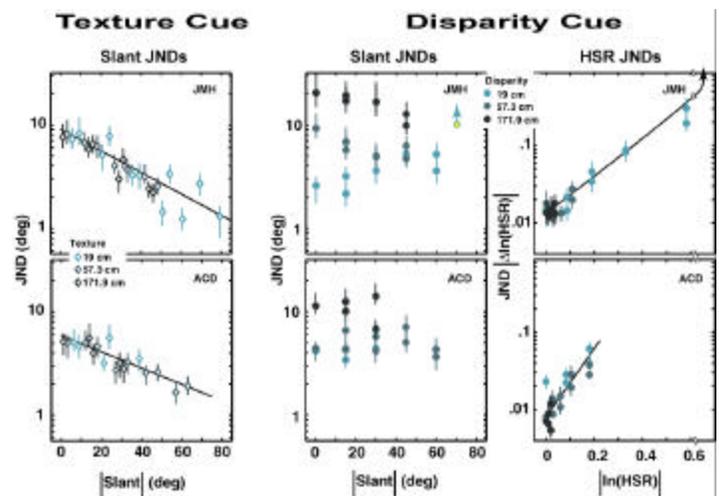


Figure 6. Just-noticeable slant differences (JNDs) for the single-cue experiment. JNDs are plotted as a function of slant or horizontal disparity. Different symbol colors represent different viewing distances: 19.1, 57.3, and 171.9 cm. Left and middle panels: JNDs for texture-alone and disparity-alone, respectively, as a function of the absolute value of the slant. Error bars are 95% confidence intervals. Right panel: JNDs for disparity alone plotted in terms of *HSR*. The ordinate is the difference between the absolute values of the natural log of *HSR* for the base slant and the just-noticeably different slant. The abscissa is the absolute value of the natural log of *HSR* of the base slant. Error bars are 95% confidence intervals. The lines in the left panel (texture-alone) and right (disparity-alone) panel are maximum-likelihood curve fits to the data ($JND = a e^{bx}$). The break in the curve and the upward pointing arrow in JMH's disparity fit indicates that JNDs go to infinity somewhere between $\ln(HSR)$ of 0.58 and 0.97 (slants of 60 and 70 deg at 19.1 cm).

The right column of Fig. 6 plots the same data as the middle column but in units of relative disparity; specifically, the horizontal-size ratio (*HSR*) (Backus et al., 1999). $HSR = \mathbf{a}_L / \mathbf{a}_R$, where \mathbf{a}_L and \mathbf{a}_R are the horizontal angles subtended by a surface patch in the left and right

eyes. Plotted in these units, JNDs do not vary systematically as a function of viewing distance. This implies that the increase in slant discrimination threshold is caused only by the geometric relationship between distance and disparity and not by greater error in the calculation of disparity nor by greater error in estimates used to scale for distance (such as vergence; Eq. 2 in Backus et al., 1999). JNDs plotted in these units increase with increasing base $|\ln(HSR)|$. This increase may reflect difficulties in solving the binocular-matching problem as the disparity gradient (which is linearly related to HSR) increases (Banks, Gepshtein, & Landy, 2004; Burt & Julesz, 1980). The increase may also reflect the fact that surfaces with large $|\ln(HSR)|$ contain fewer points near the Vieth-Müller Circle where stereoacuity is highest. $|\ln(HSR)|$ increases more rapidly as a function of slant at near distances (indicated by the fact that JMH's data at high base $HSRs$ all come from the near viewing distance). For example, a change in slant from 60 to 70 deg results in a change in $\ln(HSR)$ from 0.58 to 0.97 at 19.1 cm and from 0.06 to 0.1 at 171.9 cm. (We did not plot the point at 70 deg, $\ln(HSR) = 0.97$, because thresholds were infinite.) We will return to a discussion of the effects of distance and base slant in the Discussion (*Comparison of observed and expected effects of slant and distance on disparity- and texture-based JNDs*).

To make predictions for the two-cue conditions, we needed estimates of the variances of the disparity and texture estimators at slants between the ones for which we have measurements. For the interpolation, we fit smooth curves to the data ($JND = a \cdot e^{b|x|}$, where x is $|\text{slant}|$ and JND is in deg [texture], or x is $|\ln(HSR)|$ and JND is $|\Delta \ln(HSR)|$ [disparity], and a and b are free parameters). This was done by performing a maximum-likelihood fit to all of the raw psychometric data for a given condition (texture or disparity), varying a and b . The curves and the data are shown in the left and right columns of Fig. 6. The curve fits represent a fit to the data at all three viewing distances. Thus, they give us a way to estimate disparity and texture reliability between slants where we have measurements, and they also allow us to interpolate across distance. While the reliability of the disparity cue to slant, HSR , does not vary systematically with distance, the relationship between HSR and slant does vary with viewing distance. Movie 1 shows how the reliability of disparity slant estimates varies with slant and distance, based on the curve fits to JMH's data. The reliability of the disparity cue to slant decreases as distance increases and the reliability of the disparity cue varies with base slant in different ways at different distances. At near distances the disparity cue is more reliable than the texture cue (and hence, should be given more weight according to the MLE model).

Given a pair of JND values for texture and disparity, we can use Eq. 2 to calculate optimal weights. We did this as a function of base slant using the raw JND data from Fig. 6 as well as using the JNDs predicted by the fitted curves. Those predicted weights, which were determined solely from the standard deviations of the single-cue thresholds,

are shown in Fig. 7 as data points (based on the raw JND data) and curves (based on the fitted curves in Fig. 6). The filled circles and blue curves are the predicted disparity weights and the unfilled diamonds and gray curves are the predicted texture weights. Because JMH's thresholds were infinite at 70 deg in the disparity-alone condition at 19.1 cm, the MLE weight given to disparity in this condition is 0. The curve used to fit JMH's disparity data in Fig. 6 does not capture this fact. To incorporate this fact, we smoothly extrapolated the predicted weights curve so that the disparity weight reached 0 at 70 deg.

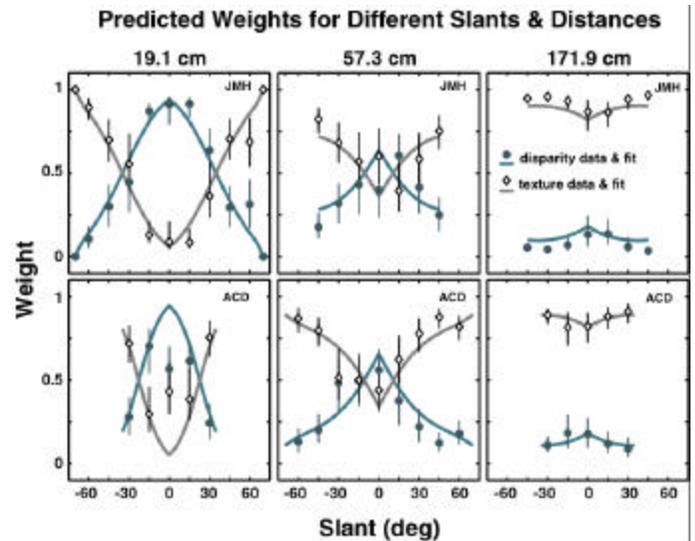


Figure 7. Predicted weights for disparity and texture cues. From left to right, the panels show data from viewing distances of 19.1, 57.3, and 171.9 cm. The weights were calculated using Eq. 2 and the single-cue discrimination data and curve fits shown in Fig. 6. Unfilled diamonds are predicted weights for the texture cue and filled circles are predicted weights for the disparity cue. The solid lines are predictions calculated from the curve fits in Fig. 6. Error bars are 95% confidence intervals.

The predicted weights exhibit two trends. First, with increasing slant, the texture weight increases and the disparity weight decreases. The reciprocal relationship between the texture and disparity weights occurs because the weights are constrained to sum to 1. The texture weight becomes relatively greater than the disparity weight with increasing slant because it becomes a relatively more reliable estimate (Fig. 6). Knill and Saunders (2003) observed a similar effect. Second, with increasing distance, disparity weight decreases (and texture weight increases). Although the reliability of the texture estimator does not change with distance, its relative reliability increases because the reliability of the disparity estimator decreases. In addition, differences in the individual observers' disparity and texture estimators (i.e., the single-cue data, Fig. 6) are manifest in their predicted weights, a point we will discuss later.

Points of subjective equality (PSEs)

From the two-cue data, we can derive the weights the observers actually gave the disparity- and texture-specified slants. Figs. 5 and 8 illustrate how this was done. The left panel of Fig. 8 shows one observer's psychometric data for a base slant of 0 deg and viewing distance of 57.3 cm. It plots the proportion of trials on which the observer indicated that the no-conflict stimulus appeared to have greater slant (right side farther away) than the conflict stimulus. Psychometric data from four cue-conflict conditions are shown. Unfilled diamonds represent data for which the disparity-specified slant was 0 deg and the texture-specified slant was -10 (gray) or +10 deg (black). Filled circles represent data when the texture slant was 0 deg and the disparity slant was -10 (light blue) or +10 deg (dark blue). It is readily apparent that the texture and disparity cues both affected perceived slant because perturbing the texture-specified slant affected judgments (shown by the separation between the gray and black diamonds) and perturbing the disparity-specified slant also affected judgments (the separation between the light and dark blue circles). The effect of disparity perturbation was greater than the effect of texture perturbation, so the weight given to disparity was larger in this condition. The PSEs are the values of the no-conflict stimulus that appeared on average to have the same slant as the conflict stimuli. The arrows indicate those values. The right panel of Fig. 8 illustrates how those PSEs were used to determine the empirical weights. If the perturbed cue (texture for the diamonds and disparity for the circles) were the sole determinant of perceived slant (meaning that its weight equaled 1; Eq. 1), the PSEs would lie along the diagonal line. If the non-perturbed cue were the sole determinant, the PSEs would fall on the horizontal line. Thus, the slope of the PSE data reflects the weight given to the perturbed cue.

We next examined the correspondence between the observed and predicted PSEs. Figs. 9 and 10 plot the empirical PSEs and the predictions of the cue-combination model in the same format as the right panel of Fig. 8. Fig. 9 shows the data from observer JMH and Fig. 10 the data from ACD. The columns of panels show data, from left to right, for viewing distances of 19.1, 57.3, and 171.9 cm. The rows of panels show data, from top to bottom, for base slants of +60, +30, 0, -30, and -60 deg. The abscissa in each panel is the value of the perturbed cue's slant in the conflict stimulus and the ordinate is the PSE.

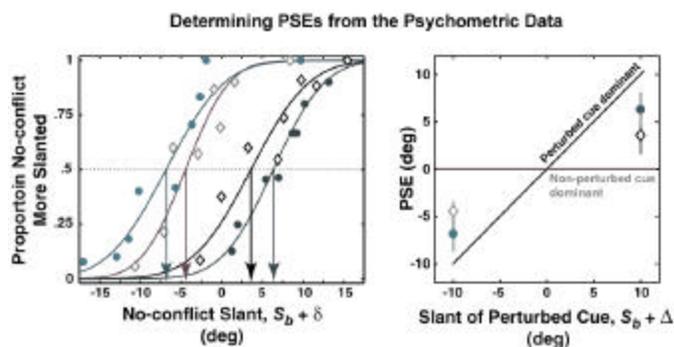


Figure 8. Determination of points of subjective equality (PSEs) from two-cue data. Left panel: one observer's results for four cue-conflict stimuli with $S_b = 0$ deg, $D = \pm 10$ deg, and distance = 57.3 cm. The conflict stimuli are: $S_d = 0$, $S_t = -10$ deg (unfilled gray diamonds), $S_d = 0$, $S_t = 10$ deg (unfilled black diamonds), $S_d = -10$, $S_t = 0$ deg (filled light-blue circles), $S_d = 10$, $S_t = 0$ deg (filled dark-blue circles). Data represent the proportion of times the observer indicated the no-conflict stimulus was more slanted than the conflict stimulus. Staircase data with fewer than four observations at a given value of the no-conflict stimulus have been removed for clarity. Curves are maximum-likelihood fits of cumulative Gaussians (which used all the points including the ones removed for the clarity). The mean of the fits are PSEs, the value of the no-conflict stimulus that on average had the same apparent slant as the conflict stimulus. The PSEs for each of the four conflict stimuli are indicated by the arrows. Right panel: PSEs for the four psychometric functions. The values of the no-conflict stimulus (indicated by the arrows in the left panel) are plotted as a function of the conflict D . If perceived slant were determined by one cue only (meaning its weight = 1), the data would lie on the diagonal line labeled "Perturbed cue dominant" when that cue was perturbed and would lie on the horizontal line labeled "Non-perturbed cue dominant" when the other cue was perturbed. Error bars are 95% confidence intervals.

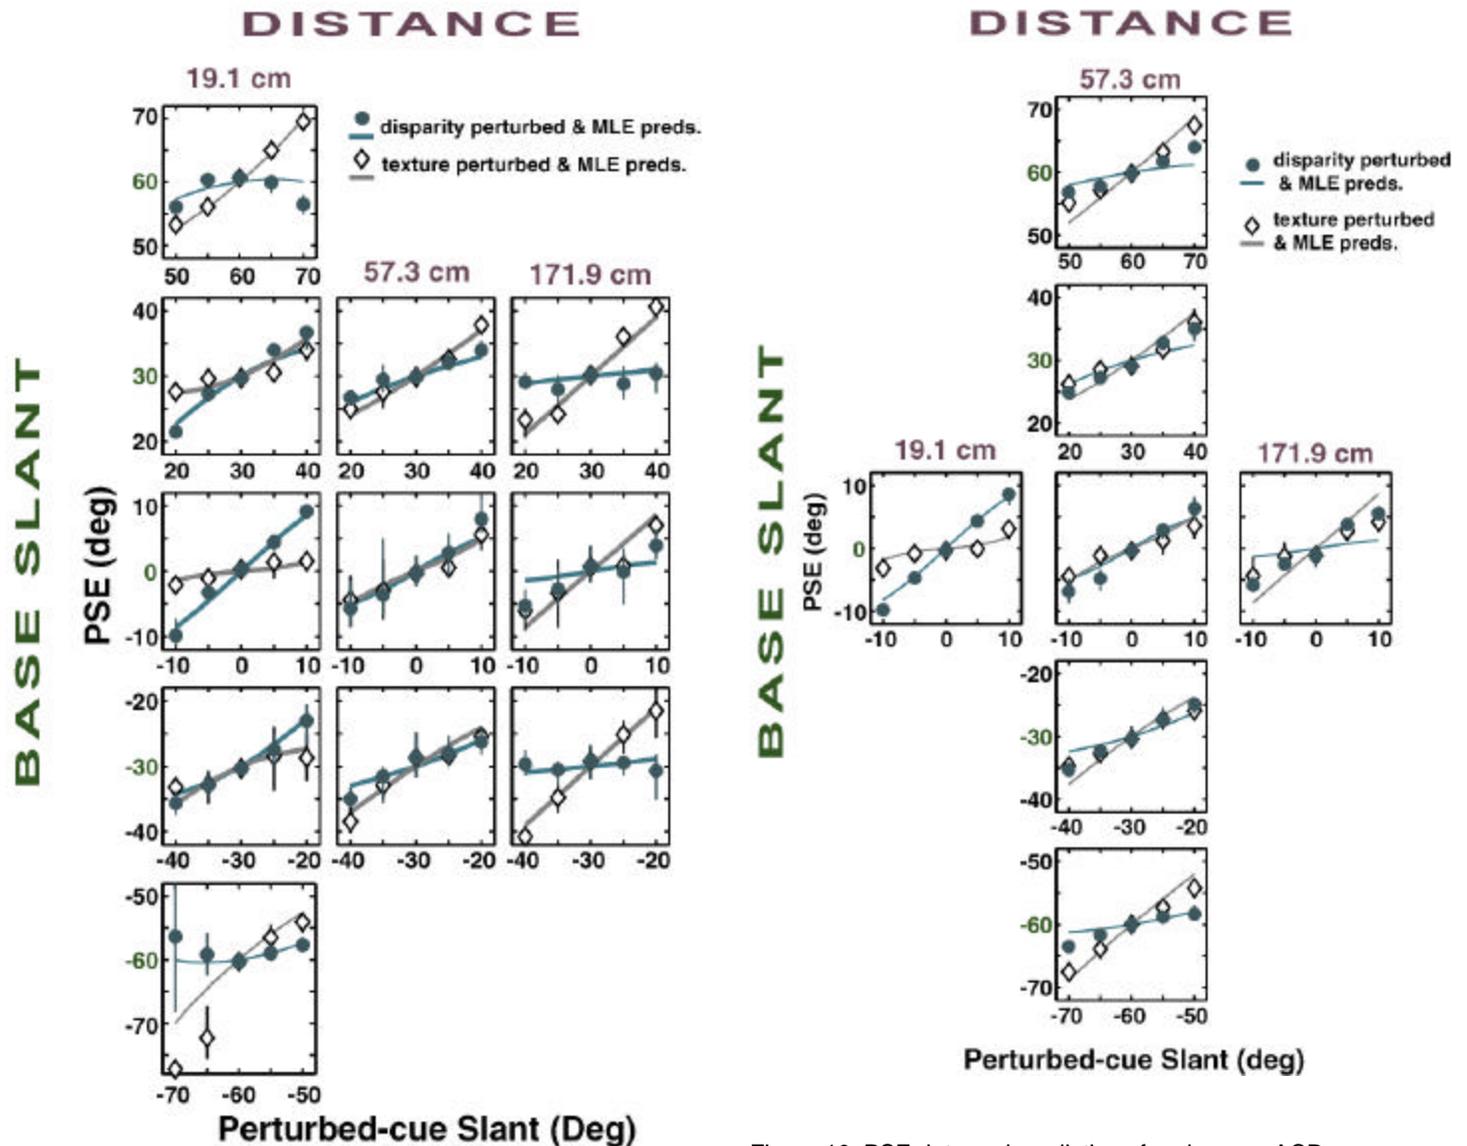


Figure 9. PSE data and predictions for observer JMH. PSE (slant of the no-conflict stimulus perceived on average as the same as conflict stimulus; $S_b + d$) is plotted as a function of the value of the perturbed cue in the conflict stimulus ($S_b + D$). The left, middle, and right columns are data from viewing distances of 19.1, 57.3, and 171.9 cm. The rows are for base slants (S_b) of -60, -30, 0, 30, and 60 deg. Those base slants are the middle abscissa value in each panel. Blue filled circles are PSEs when the disparity cue was perturbed and black unfilled diamonds are PSEs when the texture cue was perturbed. Blue and gray lines are the predictions based on Eq. 2 and the curve fits in Fig. 6. Error bars are 95% confidence intervals.

Figure 10. PSE data and predictions for observer ACD. Conventions the same as Fig. 9.

The blue and gray lines are MLE predictions for the disparity-perturbed and texture-perturbed conditions, respectively. For each conflict stimulus, the reliability for each cue was computed based on the fitted curves in Fig. 6. The optimal weights were then computed using Eq. 2. These weights, together with the displayed slants for each cue were combined using Eq. 1 to predict the PSE (i.e., the perceived slant for the conflict stimulus). The predictions are curved because the relative reliabilities (and hence the cue weights) change as the perturbation is changed (Hillis et al., 2002). We used a shortcut to generate the prediction curves. Specifically, we used the reliability based on the displayed slant to calculate the weight, rather than the reliability based on the observer's estimate of slant from each cue (which varies from trial to trial). Predictions based on a full Monte Carlo simulation in which weights were calculated separately for each simulated trial were, however, indistinguishable from these.

The agreement between the PSE data and predictions is generally excellent. The two main expected trends are observed in the data: the influence of disparity decreases with increasing distance and with increasing slant. We will discuss exceptions to the close agreement in the Discussion (*Summary of the results*).

Just-noticeable differences (JNDs)

The estimation model for cues with uncorrelated noise sources (Eqs. 12) produces the least-variable estimate of slant given the available cues. If human observers employ this cue-combination scheme, we should see improvements in JNDs when both cues are available as compared to when only one cue is available. Eq. 6 specifies the variance of the optimal cue-combined estimator which is lower than either of the single-cue estimators. We used the estimates of JNDs from the single-cue conditions (Fig. 6) and Eq. 6 to calculate the predictions for JNDs when both cues were available. Fig. 11 shows measured and predicted JNDs for both observers as a function of base slant for the three distances. The pale symbols represent the single-cue JNDs: diamonds for texture alone and circles for disparity alone. The filled red squares are the observed two-cue JNDs and the shaded red areas contain the 95% confidence intervals for the MLE predictions. With few exceptions (discussed in *Summary of results*), the two-cue data follow the MLE predictions very closely. Importantly, the two-cue JNDs are consistently lower than the single-cue JNDs, which shows that the visual system does indeed benefit from having both cues available.

Earlier we mentioned a test of the assumption that the reliability of the disparity estimator measured in the single-cue experiment with random-dot stimuli is a valid estimate of the estimator's reliability in the two-cue experiment with Voronoi stimuli. We tested the assumption by examining situations in the two-cue experiment in which the texture weight was nearly zero. The texture weight was less than 0.15 in three situations, all with observer JMH: distance = 19.1 cm and base slants of -15 , 0 , and $+15$ deg. His two-cue thresholds in those situations were respectively 2.9, 2.4, and 2.9 deg (Fig. 11). His single-cue, disparity-alone thresholds in the same situations were respectively 3.2, 2.6, and 2.1 deg (Fig. 6). The close correspondence supports our assumption that the disparity-alone thresholds provided an estimate of the appropriate reliability for the two-cue experiment.

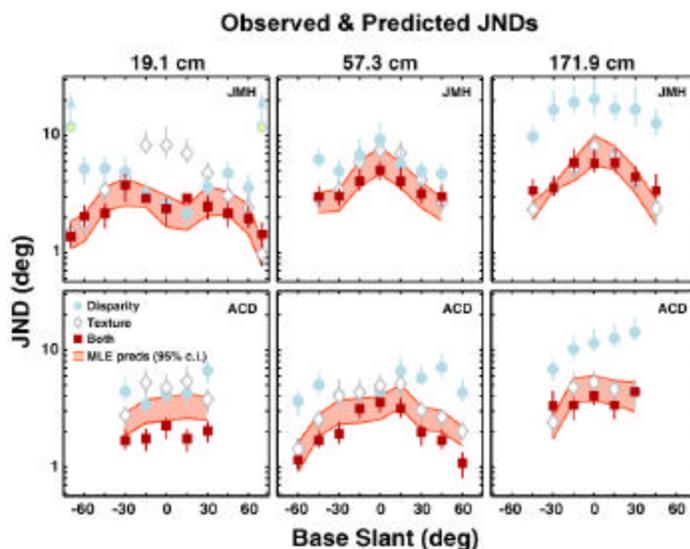


Figure 11. Predicted and observed JNDs. The just-noticeable difference in slant (JND) is plotted as a function of base slant S_b . Those values are the sigma parameters for the cumulative normal fits to the psychometric data divided by $\sqrt{2}$ and represent our estimates of the standard deviation of the slant estimators. Filled red squares are the observed JNDs when both texture and disparity were present. The faint gray diamonds are the observed JNDs for texture alone (Fig. 6, left panel) and the faint blue circles are the observed JNDs for disparity alone (Fig. 6, middle panel). Disparity JNDs for ± 70 deg base slant at 19.1 cm for JMh were infinite (indicated by pale blue symbols with inscribed yellow stars). Error bars represent 95% confidence intervals. The red curves represent 95% confidence intervals for the predicted JNDs (Eq. 6). The left, middle, and right panels represent the data from viewing distances of 19.1, 57.3, and 171.9 cm, respectively.

By similar reasoning, we can test the assumption that the reliability of the texture estimator measured in the single-cue experiment with monocular stimuli is a valid estimate of the estimator's reliability in the two-cue experiment with binocular stimuli. To generate one slant estimate from the texture-specified slants at the two eyes, the visual system should combine the monocular estimates in some fashion. The combination could occur in two ways. 1) The visual system might combine the two eyes' images before computing slant. This could be done in principle by averaging the visual directions for each corresponding point in the two images. Then slant would be computed from the combined Cyclopean image. 2) The visual system might estimate eye-centered slants before combining. Specifically, it could estimate the slants from the texture signals received by each eye and then average the two estimates. These two means of combining the monocular images are geometrically equivalent and yield the same slant as would be observed at the Cyclopean eye as long as the coordinate origin is on the Vieth-Müller Circle. At any rate, averaging the two eyes' inputs is a reasonable way to form a texture-

based slant estimate. If we assume that the two monocular inputs are equally informative and that their noises are uncorrelated (perhaps an implausible assumption), the variance of the combined estimate would be half the variance of either monocular estimate. In other words, discrimination thresholds based on the texture information alone should be lower in the binocular than in the monocular case by $\sqrt{2}$ (Legge, 1984). We tested this possibility by examining situations in the two-cue experiment in which the disparity weight was nearly zero. This occurred for both observers and all base slants when the viewing distance was 171.9 cm. It also occurred for observer JMH at 19.1 cm and base slant = $|70|$ deg. JMH's texture-alone JNDs at 171.9 cm for base slants of -45 to $+45$ deg (the range of tested slants) were 2.3–8.0 deg (the lowest values occurring at the greatest slants; Fig. 6). His two-cue JNDs at 171.9 cm for base slants of -45 to $+45$ deg ranged from 3.4–5.9 deg (again the lowest values occurring at the greatest slants; Fig. 11). JMH's texture-alone JNDs at 19.1 cm for base slants of -70 and $+70$ deg were 1.5 and 1.0 deg, respectively, and his corresponding two-cue JNDs were 1.4 and 1.4 deg. ACD's texture-alone JNDs at 171.9 cm ranged from 2.4–5.3 deg and her two-cue JNDs ranged from 3.3–4.1 deg. Thus, when the disparity weight was low, the texture-alone thresholds were generally similar to the corresponding two-cue thresholds. The good correspondence supports our assumption that the texture-alone thresholds provided an estimate of the appropriate reliability for the two-cue experiment. It also implies that the slant specified by texture is not made more reliable by averaging the two eyes' images, perhaps because the noises are highly correlated.

Discussion

Summary of the results

The generally excellent agreement between observed and predicted PSEs and JNDs indicates that humans use a statistically optimal strategy for combining slant information from disparity and texture. There are, however, three cases in which the data deviated from the predictions.

1) JMH's PSEs in the two-cue condition at 19.1 cm and base slants of -60 and $+60$ deg (Fig. 9). The weight given disparity was lower than predicted when the absolute value of the perturbed-cue slant was greater than 60 deg. In the disparity-alone $|70|$ -deg, 19.1-cm condition the disparity gradient was too high and JMH could not fuse the random dot stimulus (thus, thresholds were infinite). The same was true in the two-cue condition: slant judgments were made on diplopic images, making the task more complicated. Our model does not consider how depth judgments are made in diplopic conditions. Given this, the discrepancy between observed and two-cue data is understandable.

2) ACD's JNDs in the two-cue condition for all base slants at 19.1 cm and for the larger base slants at 57.3 cm (Fig. 11). Her two-cue thresholds were consistently lower than predicted. Moreover, ACD gave slightly more weight to disparity than predicted for base slants of ± 30 and ± 60 deg at 57.3 cm (Fig. 9). The most obvious explanation for these discrepancies is that the disparity-alone JNDs (Fig. 6) gave slight over-estimates of the variance of ACD's disparity estimator in the two-cue experiment. As described in the Methods (Fig. 2), ACD may have given some weight to the uninformative texture signal in the disparity-alone experiment for non-zero base slants. This would have caused an over-estimate of the variance of the disparity estimator whenever the disparity weight was relatively high in the two-cue experiment (which occurs when the viewing distance is 19.1 or 57.3 cm) and whenever the base slant differed significantly from zero. Thus, this discrepancy between predicted and observed data is probably caused by a slight over-estimate of the variance of ACD's disparity estimator in the single-cue experiment.

3) JMH's and ACD's disparity weights were higher than predicted at 171.9 cm when the base slant was 0 deg. We think this small discrepancy is caused by variation in binocular fusion at long distances. Both observers reported difficulty fusing the random-dot stimulus in the single-cue experiment when the viewing distance was 171.9 cm (perhaps because of the conflict between vergence and accommodation). Thus, their thresholds at 171.9 cm may have slightly over-estimated the variance of the disparity estimator at that distance. (ACD also had difficulty fusing the random-dot stimulus at 19.1 cm, which may have contributed to the apparent over-estimate of the variance of the disparity estimator as discussed under #2 above.) Both observers found it easier to fuse the Voronoi stimulus at 171.9 cm presumably because that stimulus provides contours to guide vergence eye movements. The discrepancy is most likely to show up when the base slant is 0 deg because the disparity weight is highest in that case. Thus, this discrepancy between predicted and observed behavior is probably caused by fusion difficulties in the single-cue experiment at the long distance.

The great majority of the data is consistent with the MLE predictions and strongly supports the hypothesis that observers combine the slant cues of disparity and texture in a statistically optimal fashion.

Comparison to other studies

Four studies have examined quantitatively whether cue combination is statistically optimal (Ernst & Banks, 2002; Gepshtein & Banks, 2003; Knill & Saunders, 2003; Landy & Kojima, 2001). In agreement with our results, all four found that combination of cues from different sensory modalities (haptics and vision; Ernst & Banks, 2002; Gepshtein & Banks, 2003) or different visual cues (Knill & Saunders, 2003; Landy & Kojima, 2001) was quite close to MLE predictions.

Knill and Saunders (2003) tested the MLE model for combining texture and disparity cues to surface slant. Their stimuli were slanted about a horizontal axis (tilt = 90 deg). Like us, they took advantage of the fact that the relative reliabilities of texture and disparity vary naturally with viewing geometry. They reported reasonable agreement between observed and predicted behavior. We extended their investigation by examining texture and disparity combination for surfaces slanted about a vertical axis (tilt = 0 deg) at various viewing distances. Our data are both similar and dissimilar from Knill and Saunders'. Our texture-alone data exhibited a smaller effect of base slant on JNDs (compare our Fig. 6 to their Fig. 6). Average texture-alone JNDs in our study were ~8 and ~1.5 deg at 0 and 70 deg, respectively (a ratio of 5.3). The corresponding JNDs in Knill and Saunders were ~40 and ~2 deg (ratio of 20). The fact that our JNDs were generally lower is undoubtedly due to the fact that our Voronoi patterns were more regular than Knill and Saunders'. The differing effects of base slant are most likely due to differences in how the angular subtense of the stimuli varied with slant: ours varied with base slant and theirs was constant. To hold angular size constant, Knill and Saunders had to add texture elements as slant increased, and this adds information. The added information probably explains why their slant-from-texture discrimination thresholds varied more with slant than ours did.

We also examined how viewing distance affects the weights assigned to texture and disparity and found that the weight assignment is essentially optimal. This result seems to contradict numerous reports of failures to scale veridically for distance in stereoscopic tasks. For example, Johnston (1991), Johnston, Cumming, and Parker (1993), and Bradshaw, Glennerster, and Rogers (1996) had observers judge the amount of depth in disparity-defined cylinders, spheres, and ridges when presented at difference distances. Responses were far from veridical, indicating that depth was over-estimated at near and under-estimated at far distances. How could we observe optimal weight changes as a function of distance, while previous work showed apparent failures to take distance into account? We think the answer lies in the influence of unmodeled cues. In all but one of the previous experiments (Experiment 1 in Johnston et al., 1993), the texture gradient specified a fronto-parallel plane. From our analysis, one would expect observers to report seeing less depth at long distances, not because they failed to take distance into account, but rather because they gave increasing weight to a signal specifying that the stimulus is flat. This claim is supported by the observation that making the texture gradient consistent with the disparity-specified shape generally makes judgments more veridical (Buckley & Frisby, 1993; Johnston et al., 1993). Furthermore, when the task is to adjust the shape of a surface until it appears planar and thereby consistent with the texture-specified shape, observers seem to take distance into account veridically (Rogers & Bradshaw, 1995).

Dynamic determination of cue weights

MLE cue combination has the advantage that it produces the least-variable estimate of slant given the available cues. But it requires the observer to choose weights based on the reliability of the cues. In the case of texture, the reliability clearly depends on the slant, which is what the observer is trying to estimate. Thus, the choice of weights must be made dynamically, with the possibility of varying weights from trial to trial (or from location to location within a stimulus, discussed shortly). The model suggests that on each trial the observer makes an estimate of slant from each cue, uses the value of slant for each cue along with other relevant information ("ancillary cues" such as a distance estimate, Landy, Maloney, Johnston, & Young, 1995) to determine that cue's current reliability. The relative reliabilities are then used to determine the cue weights (Eq. 2), followed by weighted cue combination (Eq. 1). In our experiments, the slant shown to the observer was selected randomly before each trial from the set of possible slants within each block. For performance to approach optimality, the weights must have been determined in a trial-by-trial dynamic fashion. In a previous study we also had clear-cut evidence of weights changing from trial to trial (Hillis et al., 2002). The reader may wonder how such dynamic computation could be accomplished in a biological system without prior knowledge of the likelihood functions associated with each slant cue. Ernst and Banks (2002) outlined a plausible neural model that could carry out the computation automatically.

Comparison of observed and expected effects of slant and distance on disparity- and texture-based JNDs

We observed three effects in the single-cue experiments—a large improvement in discrimination threshold with decreasing distance with disparity alone, a small improvement in threshold with increasing slant with disparity alone (see also Knill & Saunders, 2003), and a large improvement in discrimination threshold with increasing slant with texture alone (Knill, 1998b; Knill & Saunders, 2003). (Of course, with texture alone we did not observe a threshold change with distance because distance has no effect on the information in the stimulus.) Here, we ask whether the three observed effects are expected from the slant information in the stimulus.

When the eyes are in forward gaze, as they were in these experiments, the vergence is

$$m = 2 \tan^{-1}(i/2d), \quad (10)$$

where d is viewing distance and i is the inter-ocular distance, and slant from disparity (for tilt = 0) is given to close approximation by

$$\hat{S}_d \approx -\tan^{-1}\left[\frac{1}{\mathbf{m}}\ln(HSR)\right]. \quad (11)$$

The vergence estimate (\mathbf{m}) is required to “scale” (Howard & Rogers, 1995) or “normalize” (Gårding et al., 1995) the measured disparity because the slant associated with a given HSR is roughly proportional to distance. Thus, errors in the disparity and distance estimates will both yield errors in the estimated slant. We calculated the distribution of slant estimates for different viewing conditions under the assumption that the errors in HSR and \mathbf{m} can be represented by additive, independent noises. Specifically, we conducted a Monte Carlo simulation to determine the standard deviation of slant estimates $\mathbf{s}_{\hat{S}}$ from Eq. 11. The noises were Gaussian with mean = 0. We adjusted the noise standard deviations, \mathbf{s}_m and \mathbf{s}_{HSR} , to obtain simulation JNDs similar to the observed JNDs. The simulation results are displayed in Fig. 12. The upper panel shows $\mathbf{s}_{\hat{S}}$ for various base slants and distances. The lower left panel shows $\mathbf{s}_{\hat{S}}$ as a function of distance (the curves representing different base slants) and the lower right panel shows $\mathbf{s}_{\hat{S}}$ as a function of base slant (the curves representing different distances).

The standard deviation of the slant estimate, $\mathbf{s}_{\hat{S}}$, is roughly proportional to viewing distance for all base slants (left panel). This result is expected from Eq. 11 because $d \approx i/\mathbf{m}$, so fixed additive noise in \mathbf{m} has an increasing effect with distance. We found that $\mathbf{s}_{\hat{S}}$ was proportional to distance for a wide range of \mathbf{s}_m and \mathbf{s}_{HSR} ; the key assumption is that the noise in disparity normalization is fixed and additive in vergence. The data points in the lower left panel are JNDs from observer JMH; clearly, his discrimination thresholds increased monotonically with increasing distance in much the same way as the simulation. The data from ACD were similar. Thus, the distance effect we observed in the disparity-alone experiment is expected if error in disparity normalization is additive in units of vergence.

The bottom right panel of Fig. 12 shows that $\mathbf{s}_{\hat{S}}$ is inversely related to the absolute value of slant. This relationship was observed for all values except when $\mathbf{s}_m \gg \mathbf{s}_{HSR}$. The relationship is expected from Eq. 11 because $S_d \approx k \ln(HSR)$, so fixed additive noise in HSR has progressively less effect on $\mathbf{s}_{\hat{S}}$ as base slant increases. The data points in the lower right panel are JNDs from observer JMH; ACD’s data were similar. At viewing distances of 57.3 and 171.9 cm, JMH’s discrimination thresholds decreased monotonically with slant magnitude much like the simulation’s standard deviations.

Thus, the base-slant effect we observed in the disparity-alone condition is expected if error in disparity measurement is additive in HSR . Does this assumption make sense? It does when HSR is not significantly different from 1, which was true for distances of 57.3 and 171.9 cm (see Fig. 6). However, when HSR is quite different from 1, points on the surface fall where stereo-acuity is low and

problems arise in solving the binocular correspondence problem (Burt & Julesz, 1980). HSR and the horizontal gradient of horizontal disparity are closely related,

$$DG \approx 2(1 - HSR)/(1 + HSR), \quad (12)$$

where DG is an approximation to the disparity gradient (Howard & Rogers, 2002). From Eq. 10 when d is small and S is large, HSR is quite different from 1 and thus DG will be quite different from 0. Burt and Julesz (1980) and others have shown that binocular correspondence becomes difficult when $|DG|$ deviates significantly from 0 and breaks down altogether when $|DG| = \sim 1$. Recent results indicate that this is probably a by-product of a matching process that is similar to cross-correlating the two eyes’ images to estimate the disparity in a region of the visual field (Banks et al., 2004).

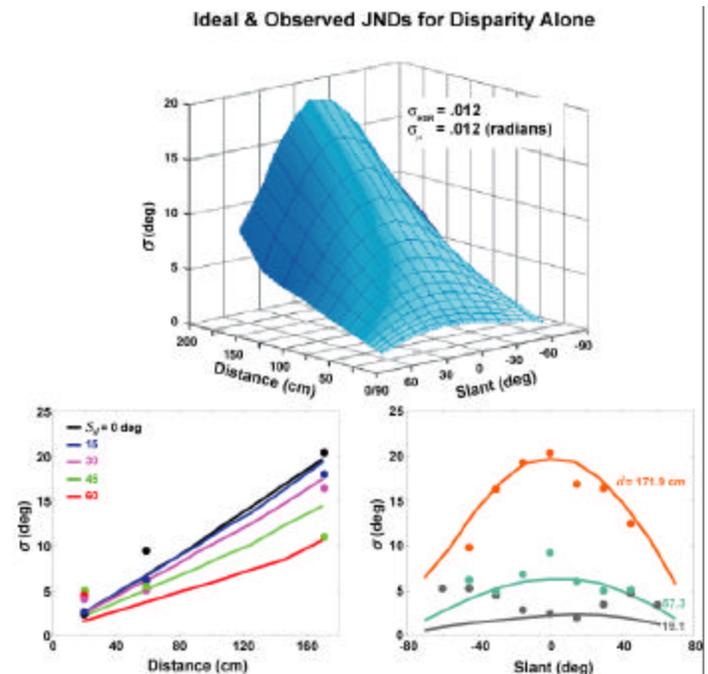


Figure 12. Results of a simulation of slant from disparity estimation. We used a Monte Carlo simulation to calculate the standard deviation of disparity-based slant estimates (Eq. 10) for different viewing conditions. We assumed that error in the slant estimates stemmed from noise in HSR and \mathbf{m} (vergence angle) and that these errors (the variances) were the same for all viewing distances. The noises were additive and Gaussian with mean = 0 and we obtained simulation results for many sets of parameters. The results for $\mathbf{s}_{HSR} = 0.012$ and $\mathbf{s}_m = 0.012$ radians, which fit the data reasonably well, are displayed in the figure. Upper panel: the standard deviations of the slant estimates for various base slants and distances. The same simulation results are plotted in the lower panels in different formats. Lower left panel: the standard deviations of slant estimates are plotted as a function of distance. Different curves represent different absolute values of base slant (i.e. different slices

through the surface displayed in the upper panel). The circles represent the observed JNDs for observer JMH at the various distances. Different colors represent different absolute values of base slant. Lower right panel: the standard deviations of slant estimates as a function of base slant. Different curves represent different viewing distances. The circles represent the observed JNDs for observer JMH. Different colors represent different distances.

Fig. 13 plots the disparity gradient as a function of slant for the three distances we used. $|DG|$ increases rapidly as a function of slant at the short distance, so we expect performance to be worse at that distance for large slants. JMH's data exhibited this effect. His discrimination thresholds at 19.1 cm increased with slant, which is inconsistent with the assumption that the sole source of error in disparity measurement is additive in *HSR* (Fig. 12, lower-right panel, gray curve). They were higher than predicted for $|\text{slant}| \geq 30$ deg which corresponds to a higher disparity gradient ($|DG| \geq 0.19$, $HSR = 1.2$) than occurs at 57.3 and 171.9 cm. Thus, the base-slant effect in the disparity-alone experiment is expected if error in disparity measurement is additive in *HSR* except when *HSR* deviates significantly from 1 where problems arise in solving correspondence.

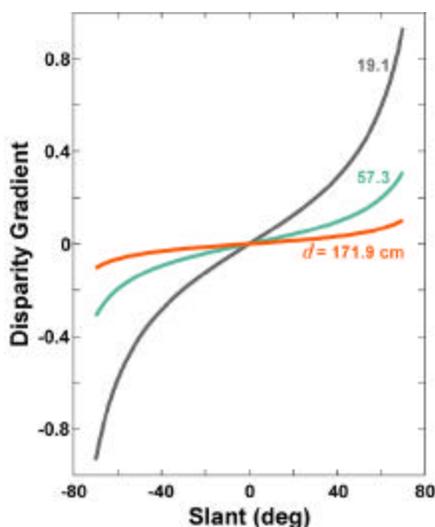


Figure 13. Disparity gradient as a function of slant at different viewing distances. The disparity gradient was calculated from Eq. 12.

We also compared our observed texture-alone thresholds with those expected from the information in the various slant cues associated with the texture gradient. Knill (1998a,b) described ideal observers for slant from texture when presented Voronoi stimuli like the stimuli in our experiments. The stimulus parameters in our experiment differed from those in his modeling and experiments in two ways.

First, the Voronoi patterns in our stimuli were more regular than in his. From this one would expect the texture-gradient cue to be more reliable in our experiment than in Knill's.

Second, the angular subtense of our stimuli varied with slant (even though there was a random element to the angular width so as to make the width an unreliable cue to slant), so the average number of texture elements was constant across slant. In order to keep the angular subtense constant, Knill added texture elements as slant increased, which adds information. This added information explains why his observed and predicted discrimination thresholds varied more with slant than ours did.

Despite the differences in stimulus parameters, it is informative to compare ideal thresholds with our observers' thresholds. The curve in Fig 13 shows the standard deviation of slant estimates from Knill's foreshortening ideal observer (Fig. 5A in Knill, 1998a). There is a striking effect of base slant. The data points are JMH's thresholds in the texture-alone experiment; ACD's data were similar. The data exhibit a base-slant effect like the ideal observer's, but the effect is smaller in our data for reasons described above. Therefore, the variation we observed in texture-based slant thresholds is by and large expected from the information content of the stimulus.

Ideal & Observed JNDs for Texture Alone

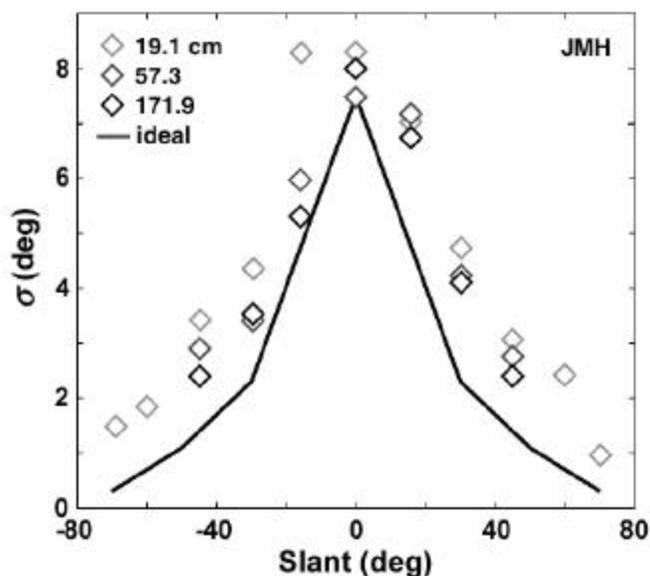


Figure 14. Ideal and measured JNDs for slant from texture as a function of base slant. The solid line represents the standard deviations of the slant estimates of the foreshortening ideal observer for Voronoi stimuli (Fig. 5A in Knill, 1998a). The diamonds represent discrimination thresholds in the texture-alone condition for observer JMH. Light gray diamonds are thresholds at 19.1 cm, medium gray at 57.3 cm, and dark at 171.9 cm.

We conclude that the effects of distance and slant on JNDs can be expected from the information present in the stimuli. These effects are summarized in Movie 1.

What other variables might affect cue weights?

Presumably, the visual system takes the disparity and texture variances into account across many viewing situations. To do so, however, is complex because many viewing properties will affect the likelihood functions associated with disparity and texture cues. Here we list the most obvious properties and suggest how the relative weights assigned to disparity and texture ought to be affected.

1. Regularity of texture. The slant information contained in the texture gradient can be divided into three cues: 1) scaling, the change in the projected sizes of texture elements, 2) foreshortening, the change in projected shapes of texture elements, and 3) density, the change in the number of elements per unit area in the projection (Blake et al., 1993; Cutting & Millard, 1984; Knill, 1998a). The reliability of scaling as a slant signal depends on the variation in the sizes of the texture elements on the surface. With greater size variation, the cue's reliability decreases (Knill, 1998b). The reliability of foreshortening depends on the variation in the shapes of the elements on the surface. For regular shapes, like circles, reliability is greater than for irregular shapes, such as ellipses with variable aspect ratios (Knill, 1998b; Young et al., 1993). The reliability of the density cue depends on the number of elements and the regularity of their positioning on the surface. Presumably, many elements placed regularly (i.e., in a grid) yield more reliable estimates than few elements placed randomly. All three cues are affected by the field of view, particularly in the tilt direction, so slant discrimination from texture is more precise with large than with small stimuli (Blake et al., 1993; Knill, 1998b). If the visual system takes the varying reliability of the texture gradient cue into account, all of these stimulus properties will affect the relative weights assigned to disparity- and texture-based signals.

2. Surface tilt The direction of slant affects the amount of perceived slant in stereograms (Howard & Rogers, 2002), so it is worthwhile considering how the signals specifying slant vary with tilt. The disparity signal for surfaces slanted about a vertical axis (tilt = 0 deg) is the horizontal gradient of horizontal disparities. We have quantified this signal as the horizontal-size ratio (*HSR*). The disparity signal for surfaces slanted about a horizontal axis (tilt = 90 deg) is the vertical gradient of horizontal disparities. This disparity pattern is often referred to as horizontal-shear disparity (Ogle & Ellerbrock, 1946; Banks et al., 2001). Random-element stereograms simulating a slanted plane with tilt = 0 deg generally produce less perceived slant than planes with tilt = 90 deg (Gillam & Ryan, 1992). Likewise, the amount of depth seen in curved disparity-defined surfaces varies with tilt: more depth is

seen when the disparities specify curvature about a horizontal axis (Buckley & Frisby, 1993). These tilt-dependent variations in the amount of perceived depth are called *slant anisotropy*. The phenomenon is most striking when the texture gradient specifies a frontoparallel plane, as is usually the case with random-element stereograms. The phenomenon is not observed when disparity and texture signal the same depth variation, as occurs with real surfaces (Bradshaw, Hibbard, van der Willigen, Watt, & Simpson, 2002; Buckley & Frisby, 1993). These observations strongly suggest that slant anisotropy is caused by conflicting disparity and texture signals in conventional random-element stereograms. They also suggest that texture is generally given more weight for tilt = 0 deg (as in our experiments) and less weight for tilt = 90 deg (as in Knill and Saunders, 2003). Indeed, we observed higher texture weights than Knill and Saunders at 57.3 cm, which is similar to their distance of 60 cm. By the argument presented here, this may be due to reduced disparity reliability for tilt 0 than for tilt 90 because there is no obvious reason for the reliability of the monocular texture cue to depend on tilt. There may, however, be differences in the steps required to combine the texture and disparity signals for different tilts. The issues involved in transforming texture-gradient signals into the same coordinates for combination with disparity signals are taken up in Appendix A.

3. Reliability of estimated distance and azimuth.

To estimate slant from the measured disparities, the visual system must "normalize" the disparities with a distance estimate and "correct" the disparities with an azimuth estimate (Gårding et al., 1995). Relaxing the assumption of forward gaze in Eq. 11, slant about a vertical axis (tilt = 0) is

$$S \approx -\arctan\left(\frac{1}{m} \ln HSR - \tan g\right), \quad (13)$$

where m is vergence, and g is the azimuth (the angle between the head's median plane and the Cyclopean line of sight) (Backus et al., 1999). m is estimated both from extra-retinal signals concerning the eyes' vergence and from the horizontal gradient of vertical disparity (Rogers & Bradshaw, 1995). When vertical disparities are large, as occurs with large stimuli at close range, they are the predominant means for estimating distance. However, when vertical disparities are small because the stimulus is small (Rogers & Bradshaw, 1995), or unmeasurable because the texture contains no horizontal contours (Helmholtz, 1910), the eyes' vergence becomes the predominant means of estimating distance and the accuracy of disparity normalization drops (Rogers & Bradshaw, 1995).

The azimuth angle g is used to correct disparities; it is estimated from extra-retinal, eye-position signals and from the magnitude of vertical disparities (Backus et al., 1999). When vertical disparities are large, as occurs with near stimuli subtending a large angle, they are the predominant

means of estimating azimuth. When the stimulus is short or when vertical disparities are unmeasurable, eye position becomes the predominant means and the accuracy of disparity correction suffers (Backus et al., 1999).

Similar arguments apply for slant estimation with tilt = 90 deg. In this case, slant around a surface point is

$$S \approx -\arctan\left(\frac{1}{m} \ln HSh - t\right), \quad (14)$$

where m is again the vergence angle, HSh is horizontal shear disparity (a quantification of the vertical gradient of horizontal disparity) (Kaneko, 1994).

For our present purposes, when the viewing situation reduces the reliability of the estimates of the normalizing and/or correcting signals, the disparity estimate will become more variable. This will occur, for example, when the stimulus subtends a small angle, when the surface markings make the measurement of vertical disparity unreliable, and when the stimulus is distant. If the visual system takes such changes into account, the weight given to disparity should decrease in those circumstances.

In the two-cue experiments, disparity cues had to be normalized and corrected to be converted into slant estimates, which is required before those estimates can be combined with information from the texture gradient (see *Control experiments and procedures to validate single-cue measurements*, Fig. 4). Thus, the variance of the disparity-based estimate of slant included components due to measuring the horizontal disparities as well as variability in the estimates of distance and azimuth. In the single-cue disparity experiment, azimuth and distance were not varied in any block of trials, and hence slant discrimination could have been carried out using the raw disparity gradients without invoking disparity normalization and correction. If this occurred, we could have overestimated disparity reliability. Given the accuracy of our two-cue predictions, this suggests either that the noise from the estimation of distance and azimuth is swamped by the variability of estimating the disparity gradient, or that observers applied disparity normalization and correction in both the single- and two-cue experiments.

4. Duration. Van Ee and Erkelens (1998) showed that the slant perceived from disparity-defined planes increases with stimulus duration. Their random-element stereograms contained the texture gradient associated with a frontoparallel plane, so their results are consistent with a model in which the weight given to disparity relative to texture increases over time. Presumably, disparity and texture estimates both become more precise with increases in stimulus duration, but the increase may be slower for disparity. Thus, stimulus duration may also affect the relative weights given to disparity- and texture-based slant estimates.

Are cue weights computed locally?

It is interesting to consider whether the visual system determines one set of weights for each surface or whether the weights are calculated locally. That is, can the weights vary from one surface patch to another? If they are calculated locally, there are situations in which a cue-conflict stimulus specifying a plane should appear curved. Here we explain why this should happen and report that the predicted curvature is in fact observed.

The left panel of Fig. 15 shows how surface slant and distance vary with azimuth. The center of the stimulus plane is straight ahead at distance d . The angle between the surface normal and the Cyclopean line of sight is the slant S . When the line of sight is rotated through angle g (azimuth), the slant of the portion of the stimulus plane intersected by that line is

$$S_g = S - g, \quad (15)$$

and the distance to the intersection of the line and plane is

$$d_g = \frac{d}{\cos g + \sin g \tan S}. \quad (16)$$

The left and middle panels of Fig. 16 show how local slant S_g and distance d_g vary as a function of azimuth for different base slants and a viewing distance of 19.1 cm. Because the local slant and distance vary with azimuth, the statistically optimal weights for the texture and disparity cues should vary with azimuth.

Now consider the cue-conflict stimulus in the middle panel of Fig. 15. For rightward gaze ($g < 0$), slants S_{lg} and S_{rg} approach zero and distance d_g decreases. Our data (Figs. 9 and 10) show that texture weight is relatively low when the absolute value of slant is ~ 0 and distance is short. Thus, if the weights used in combining slant estimates (Eq. 2) are determined locally, one would expect the texture weight in this situation to be lower on the right than straight ahead. (The changes in local slant and distance with changes in azimuth are unaffected by the direction in which the eyes are looking; they are determined only by the positions of surface points relative to the head. Thus, when we say "on the right" or "straight ahead", we refer to the azimuth of a line of sight from the Cyclopean eye and not necessarily the azimuth of fixation.) For leftward azimuth ($g > 0$), the slants become increasingly negative and the distance increases. Thus, the texture weight in this situation should be higher on the left than straight ahead. If the relative disparity and texture weights change across the surface, then the slant estimate should change when the disparity- and texture-specified slants differ. This is depicted in the right panel of Figure 14. The thin red line between the disparity- and texture-specified surfaces represents the slant estimate if the disparity and texture weights were equal throughout. However, if the disparity weight increased with increasingly rightward azimuth, the slant estimate should approach the disparity-specified slant toward the right (and the texture-

specified slant toward the left). This is depicted by the clockwise rotation of the thick red line segments relative to the thin line moving from left to right. As a consequence, this particular cue-conflict stimulus should appear concave.

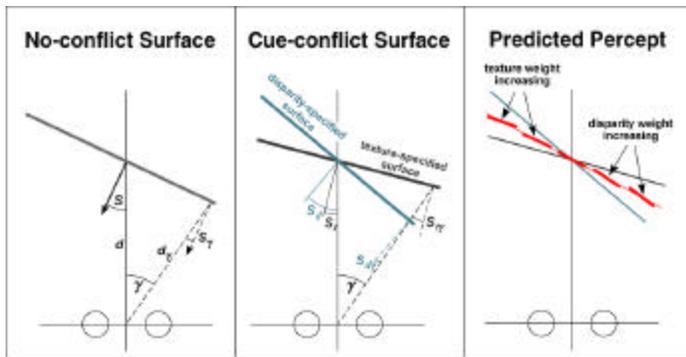


Figure 15. Change in slant with azimuth. Left panel: Definitions of surface slant and distance and local slant and distance. A plane is positioned straight ahead of the observer's Cyclopean eye. The surface normal at the intersection of the plane and the head's median plane is shown. Slant S is the angle between the normal and the Cyclopean line of sight to that point and d is the distance to that point. The azimuth to another point on the plane is g . The slant with respect to that point is S_g and the distance is d_g . Middle panel: Slants and local slants for a cue-conflict stimulus. The disparity-specified slant of the stimulus is $S_d = -25$ deg and the texture-specified slant is $S_t = -10$ deg. The local slants at azimuth g are S_{dg} and S_{tg} . Right panel: Predicted perceived slants across the cue-conflict stimulus if cue weights are determined locally. The thin red line represents the predicted slant if cue weights are fixed across the surface. The thick red line segments represent the predicted change in slant if the weights are determined locally. The surface should appear concave in this particular case.

We calculated how local slant S_g should change for observer JMH. The cue-conflict stimulus in the calculation had a disparity-specified slant S_d of -25 deg and texture-specified slant S_t of -10 deg. Distance d to the midpoint was 19.1 cm. If weights are determined locally for each azimuth g then the observer must associate local disparity-defined slant S_{dg} with its corresponding variance s_{dg}^2 and the texture-defined slant S_{tg} with its variance s_{tg}^2 . From Eq. 2, the weights as a function of azimuth are

$$w_{dg} = \frac{r_{dg}}{r_{dg} + r_{tg}} \tag{17}$$

$$w_{tg} = \frac{r_{tg}}{r_{dg} + r_{tg}},$$

where r_i are the reciprocals of the variances. The right panel of Fig. 16 shows the expected change in slant as a

function of azimuth. With leftward azimuth, the slant estimate approaches the texture-specified slant, and with rightward azimuth, it approaches the disparity-specified slant. The result would be an apparently concave surface as schematized in the right panel of Fig. 15. If the disparity- and texture-specified surfaces were swapped ($S_d = -10$; $S_t = -25$ deg), the result would be an apparently convex surface. If the disparity and texture specified the same slant, as they would with most real surfaces, the result should be an apparently planar surface.

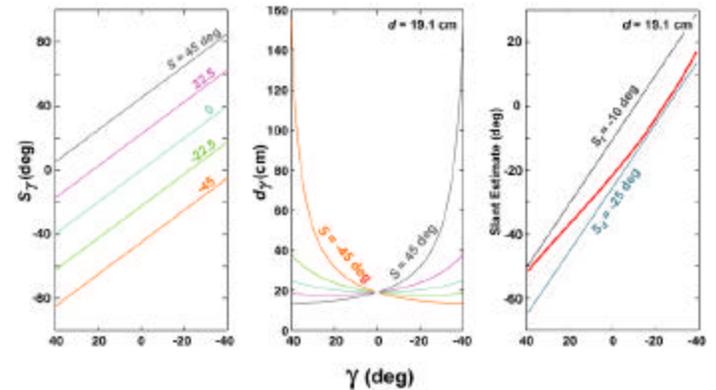


Figure 16. Change in local slant, local distance, and estimated slant with azimuth. Left panel: Local slant S_g (see Fig. 15 for definition) as a function of azimuth g . The lines from top to bottom show S_g for different slants S . Middle panel: Local distance d_g as a function of azimuth g . The gray, magenta, blue, green, and orange curves show d_g for slants of 45 , 22.5 , 0 , -22.5 , and -45 deg, respectively. Right panel: Predicted local slant estimates from the weighted sum (Eq. 1) if the cue weights are determined locally. Disparity-specified slant (S_d , blue) is -25 deg and texture-specified slant (S_t , gray) is -10 deg. Distance (d) is 19.1 cm. Using the cue weights as a function of base slant at a distance of 19.1 cm for observer JMH, we calculated the estimated slant at each azimuth. Those local slants (S_g) are represented by the red curve.

In doing these calculations, we assumed that the cue weights were determined by the local disparity and texture slants only and not by the distance (because the texture-specified distance is undefined). The predicted curvature would have been somewhat greater if we had included the changes in disparity-defined distance in the calculation.

Do people see the predicted curvature? To answer this, we set up a large-field display. A Voronoi pattern was back-projected onto a screen; the texture gradient specified different slants relative to the projection screen. The disparity-defined slant was identical to that of the projection screen (and hence was consistent with other cues such as blur and accommodation). Observers viewed the display binocularly. Azimuth was manipulated by having them stand at an oblique position relative to the screen center. The situation in Fig. 15 and in the right panel of Fig. 16 was recreated as follows. Observers stood 25 deg to

the right of center at a distance of 19.1 cm and viewed a Voronoi pattern whose texture gradient specified a slant of +15 deg relative to the projection screen. Thus, $g = -25$, $S_d = -25$, and $S_t = -10$ deg. At this azimuth, the field of view was ~70 deg wide. The Voronoi pattern was clipped by an elliptical window so that the outline shape was an unreliable cue to slant. The room was completely dark except for the display so the screen's frame could not be seen. By the argument above, this cue-conflict stimulus should appear concave. We also created a viewing situation in which the surface should appear convex— $g = -25$, $S_d = -25$, and $S_t = -40$ deg—and another in which the surface should appear planar— $g = -25$, $S_d = -25$, and $S_t = -25$ deg. Seven observers (five naïve) viewed the displays and reported whether they appeared concave, planar, or convex. Five of the seven (three naïve) reported that the stimuli predicted to look concave and convex actually looked that way; the other two said that the stimuli all appeared concave or planar but that the one predicted to look concave appeared the most concave. We asked them to order the three stimuli according to the amount of perceived concavity and all seven ordered them in the predicted order.

We conclude that the weights assigned to disparity and texture are estimated locally. As a consequence, extended cue-conflict stimuli in which the two cues specify planes can appear curved. Again, one would not predict curvature for stimuli with consistent slant cues, such as occurs with most real surfaces. In the main experiments, observers did not notice any distortions in the surfaces in the cue-conflict stimuli. This is not surprising because the conflict size was relatively small and the stimuli were far smaller than the large-field displays just described.

Ideal observer models

We have based our modeling on an ideal Bayesian or statistical observer, assuming uncorrelated (or conditionally independent) cues and a negligible effect of the prior distribution. Here, we consider the impact and validity of these assumptions.

Oruç et al. (2003) developed an ideal observer that allowed for cue correlation. The resulting cue-combination rule remains a linear one (a weighted average), but the optimal weights do not satisfy Eq. 2. Rather, they must be “corrected” for the degree of cue correlation. Oruç et al. studied slant estimation using two cues that should be highly correlated (texture gradient and linear perspective, both pictorial cues). Subjects displayed a range of behaviors including optimal cue combination consistent with independent cues, optimal cue combination consistent with correlated cues, and sub-optimal linear cue combination. In the present study, the agreement between the predicted and observed two-cue data is quite good, which indicates that the correlation between the noises of the texture and disparity estimators is probably small. One might expect a strong correlation between the noises

associated with disparity- and texture-based slant estimators because the two estimators must share some processing (e.g., noise in the stimulus itself, eye movements, retinal-image formation, and retinal processing). The fact that we observe (as did Knill & Saunders, 2003) the improvement in JNDs expected from combining two conditionally independent estimates implies that the dominant noises are independent. Thus, those noises probably arise in separate processes such as comparing the two eyes' images, normalizing for distance, and correcting for azimuth.

In the modeling we assumed that the prior distribution ($P(\hat{S})$ in Eq. 3) has a negligible effect. This assumption is usually justified (e.g., Ernst & Banks, 2002) by assuming that the variance of the prior is much greater than the variances of the likelihoods. Is this the case in the estimation of surface slant? It is reasonable to assume that the distribution of surface slants in the world is uniform, particularly for tilt = 0. But if that distribution is uniform, the probability of observing slant S at the retina will be proportional to $\cos(S)$ because steeply slanted surfaces project to smaller retinal images. Eqs. 7-9 show how the observers' judgments will be affected by the stimulus values and weights. If we add the prior into those equations, Eq. 9 becomes

$$w_d = \frac{(1 - w_p)d}{\Delta}. \quad (18)$$

The value of w_p depends on its inverse variance, or reliability, relative to the estimator reliabilities (Eq. 5). As long as the prior's variance is large relative to the estimators' variances, w_p will be small and will have no discernible effect on the data. The prior distribution is proportional to $\cos(S)$, which is defined from -90 to 90 deg. Such a half cosine has a standard deviation of ~40 deg. The standard deviations of the disparity and texture estimators (Fig. 6) ranged from 1-20 deg. They were the highest when distance = 171.9 cm and slant = 0 deg. Then the standard deviations of the disparity and texture estimators were 20.4 and 8.0 deg for JMH and 16.1 and 5.3 deg for ACD. We can use Eq. 5 to calculate the expected w_p for those conditions: $w_p = 0.034$ and 0.016 for JMH and ACD, respectively. Those represent the largest possible influence of the prior distribution on the results. Even in that case w_p is quite small, so we conclude that the prior had no discernible influence on our results.

Decision noise

An important underlying assumption in our analysis is that single-cue thresholds accurately reflect the observer's uncertainty about slant from the cue in question. In reality, discrimination thresholds are affected by other sources of uncertainty, such as high-level decision noise. Here we examine the consequences of decision noise on the predicted and observed results. We assume that the decision noise is additive, and has a mean of zero and

standard deviation of \mathbf{s}_n . We also assume that \mathbf{s}_n has the same value in the single-cue and two-cue experiments.

First consider the PSE data (Figs. 9 and 10). Let $\tilde{\mathbf{s}}_d^2$ and $\tilde{\mathbf{s}}_t^2$ represent the variances we measured in the disparity- and texture-alone experiments:

$$\tilde{\mathbf{s}}_d^2 = \mathbf{s}_d^2 + \mathbf{s}_n^2 \quad (19)$$

$$\tilde{\mathbf{s}}_t^2 = \mathbf{s}_t^2 + \mathbf{s}_n^2.$$

We used those measured variances to generate predictions for the weights given disparity and texture. Thus, Eq. 2 becomes

$$w_{dP} = \frac{\tilde{r}_d}{\tilde{r}_d + \tilde{r}_t} \quad (20)$$

$$w_{tP} = \frac{\tilde{r}_t}{\tilde{r}_d + \tilde{r}_t},$$

where \tilde{r}_d and \tilde{r}_t represent the measured reliabilities from the disparity-alone and texture-alone experiments (the measured reliabilities include the effects of decision noise). In the two-cue experiment, we measured the weight observers actually assigned to disparity and texture and those weights were presumably affected only by the visual system's estimates of the uncertainties of the disparity and texture estimators. In other words, Eq. 2 rather than Eq. 20 describes what the observed weights should be. Decision noise should, therefore, affect the predicted weights (Eq. 20) and not the observed weights in the two-cue experiment. To determine the consequences of decision noise, we calculated the predicted and observed weights for a variety of situations. We set the sum of estimator variances to one ($\mathbf{s}_d^2 + \mathbf{s}_t^2 = 1$) and varied \mathbf{s}_t^2 from ~ 0 to ~ 1 . \mathbf{s}_n^2 was set to 0, 0.1, 0.32, or 1. The left panel of Fig. 17 shows the results. The predicted texture weight (w_{tP}) is plotted as a function of the actual texture weight (w_t). Naturally, the prediction (diagonal dashed line) is perfect when $\mathbf{s}_n^2 = 0$ because Eqs. 2 and 20 are then identical. For $\mathbf{s}_n^2 > 0$, the predicted weights deviate from the observed. When w_t is greater than 0.5 (and hence greater than w_d), w_{tP} is less than w_t . When w_t is less than 0.5, the opposite occurs. When $w_t \approx w_d$ ($w_t = \sim 0.5$), the effect of decision noise is negligible. Thus, if decision noise were sufficiently large in our experiments, it should cause error in the PSE data when w_t is either much larger or much smaller than w_d . This circumstance occurred when the viewing distance was 19.1 cm and the base slant was 0 deg and when the viewing distance was 171.9 cm (Figs. 9 and 10). With the exception of distance = 171.9, base slant = 0, the agreement between predicted and observed PSEs is excellent in these cases. This implies that uncertainty due to decision noise (and other additive noises) was small relative to the uncertainty of the underlying slant estimators.

Knill and Saunders (2003) also considered the effects of decision noise on observed and predicted weights. In their simulation, they assumed that \mathbf{s}_n^2 was equal to the variance of the combined estimate (Eq. 6), so \mathbf{s}_n^2 varied with w_t ; in particular it had lower values when $w_t = \sim 0$ or ~ 1 . Thus, their simulations showed very little effect of decision noise on predicted weights. Knill and Saunders also modeled constant-variance noise like we did, but they did not show the results of that analysis.

Now consider the JND data (Fig. 11). Again $\tilde{\mathbf{s}}_d^2$ and $\tilde{\mathbf{s}}_t^2$ represent the variances we measured in the disparity-alone and texture-alone experiments, respectively (Eq. 19). Using those measured values and Eq. 6, we generate a prediction for the two-cue JND (ignoring division by $\sqrt{2}$ to convert from 2-IFC thresholds to \mathbf{s} 's):

$$\mathbf{s}_{cP} = \sqrt{\frac{\tilde{\mathbf{s}}_d^2 \tilde{\mathbf{s}}_t^2}{\tilde{\mathbf{s}}_d^2 + \tilde{\mathbf{s}}_t^2}}. \quad (21)$$

Now we make the two-cue measurements in order to compare the observed and predicted JNDs. In the two-cue experiment, the visual system would weight the cues as in Eqs. 1 and 2, and the decision noise would again affect the threshold measurement. Thus, the JND we measure in the two-cue experiment is

$$\tilde{\mathbf{s}}_c = \sqrt{\frac{\mathbf{s}_d^2 \mathbf{s}_t^2}{\mathbf{s}_d^2 + \mathbf{s}_t^2} + \mathbf{s}_n^2}. \quad (22)$$

Eqs. 21 and 22 are equal to one another when $\mathbf{s}_n = 0$, but when $\mathbf{s}_n > 0$, the decision noise has different effects on the predicted and observed two-cue JNDs. To determine how additive decision noise could affect the interpretation of the JNDs, we calculated the ratio of observed JND (Eq. 22) divided by the predicted JND (Eq. 21). In doing the calculations we again set the sum of estimator variances to one and varied \mathbf{s}_t^2 from ~ 0 to ~ 1 . \mathbf{s}_n^2 was again set to 0, 0.1, 0.32, or 1. The right panel of Fig. 17 shows the results. The JND ratio is plotted as a function of the texture weight (from Eq. 2). The dashed horizontal line represents the ratio when $\mathbf{s}_n^2 = 0$; the ratio is 1 because observed and predicted JNDs are the same. As \mathbf{s}_n^2 increases from 0 to 1, the observed JND becomes larger than the predicted. The ratio is largest at ~ 1.3 when $w_t = 0.5$ and $\mathbf{s}_n^2 = 1$. The JNDs we observed in the two-cue experiment were generally quite close to the predicted values (Fig. 11), so this analysis suggests that we can rule out the presence of decision noise whose variance is greater than approximately half the sum of the estimator variances.

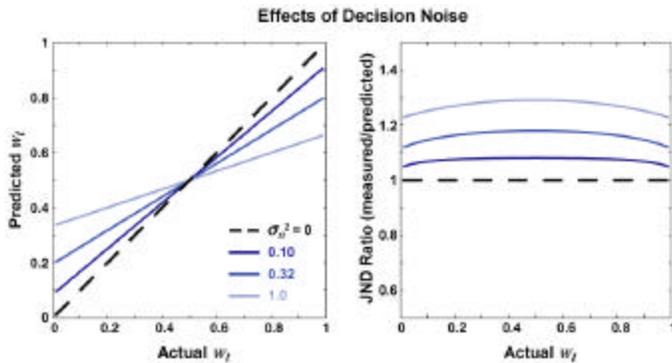


Figure 17. Results of a simulation of the effects of decision noise on PSEs and JNDs. Both graphs plot a measure of performance against the actual value of the texture weight w_i . We added decision noise to the single-cue estimates (Eq. 16). The decision noise was Gaussian with mean = 0 and variances $\sigma_d^2 = 0, 0.1, 0.32,$ and 1.0 where $\sigma_d^2 + \sigma_i^2 = 1$. Left panel: The value of w_i we would predict from the single-cue measurements (corrupted by decision noise) is plotted as a function of the actual value of w_i . The different lines represent different amounts of decision noise as indicated by the legend. Right panel: The JND ratio—the JNDs we measure in the two-cue experiment divided by the JNDs we predict from the single-cue experiments—plotted as a function of the actual value of w_i . The different curves represent the results with different amounts of decision noise as indicated by the legend in the left panel. See text for further explanation.

We conclude that uncertainty due to additive, high-level decision noise was small in comparison to the uncertainties of the underlying slants from the two cues.

Conclusion

We performed two quantitative tests of a maximum-likelihood estimator model for combining the slant cues of texture and disparity. Our results indicate that the visual system combines texture and disparity information in a statistically optimal fashion, thereby reducing the variance of slant estimates that could be achieved otherwise. To do this, the relative reliabilities of each cue have to be determined dynamically, on a trial-by-trial basis, suggesting the presence of well-developed circuits for combining depth-cue information. The success of the MLE model in this and other studies indicates that perceptual systems have gone to some trouble to incorporate all available information into the estimation process, and that they have done so in a manner that maximizes the precision of perceptual estimates.

Acknowledgements

This research was supported by NIH Research Grant EY12851 (MSB), AFOSR Research Grant F49620 (MSB),

and NIH Research Grant EY08266 (MSL). We thank Alison Dilworth for spending many hours as an observer. Part of this work was presented at the European Conference on Visual Perception in 2002.

Appendix

To express surface orientation, we need a coordinate system. The convention in the stereo literature is to place the origin at the Cyclopean eye, the point on the Vieth-Müller Circle, half way between the eyes. The X-axis is parallel to the interocular axis. The Z-axis lies in the plane of fixation and is perpendicular to the interocular axis. The Y-axis is perpendicular to the other two axes.

Consider a homogeneously textured fronto-parallel plane in front of the observer. The eyes are converged on the plane at the midline and the vergence angle is m . The plane's slant is $S_C = 0^\circ$ and its tilt t_C is undefined. From the viewpoint of the left eye (equivalent to placing the origin there),

$$S_L = \frac{m}{2} \text{ and } t_L = 0^\circ.$$

From the right eye's viewpoint,

$$S_R = -\frac{m}{2} \text{ and } t_R = 0^\circ$$

or, equivalently,

$$S_R = \frac{m}{2} \text{ and } t_R = 180^\circ.$$

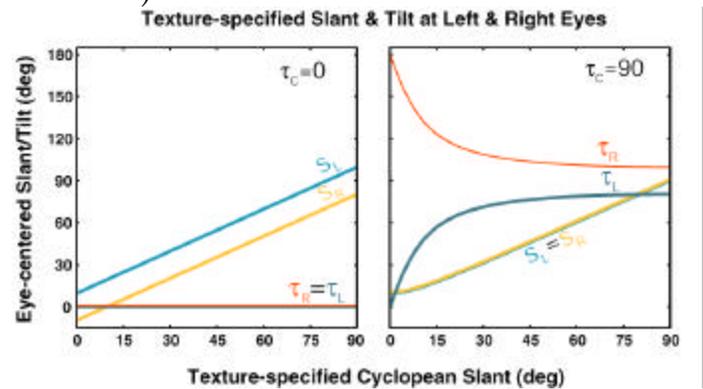


Figure 1A. Monocular texture-specified slants and tilts as a function of Cyclopean slant for a 19 cm viewing distance and 6.5 cm inter-pupillary distance. The left panel shows left and right eye slants and tilts when the Cyclopean tilt is 0° ($t_C = 0^\circ$). The right panel shows left and right eye slants and tilts when the Cyclopean tilt is 90° ($t_C = 90^\circ$).

Hence, the texture gradient specifies different surface orientations at the two eyes. The disparity-specified slant is, of course, 0 deg with the origin at the Cyclopean eye. The texture-gradient signals must be converted into the same Cyclopean coordinates as the disparity signal. The

transformation required to convert the texture-gradient signals from left- and right-eye coordinates into Cyclopean coordinates is a potential source of error. That error is presumably not measured in the single-cue, texture-alone condition because monocular stimuli were used in that condition. Here we examine the properties of the required transformation and show that they vary with tilt.

When a frontoparallel plane is rotated about the vertical axis ($S_C = \Delta$ and $t_C = 0^\circ$), the texture-specified slants and tilts in the left and right eyes vary, but the tilts do not. Technically, for rotations about a vertical axis, the tilt parameter for eye-centered viewpoints flips between 0 and 180° . This corresponds to a sign change in slant. For simplicity, we use a sign change in slant, rather than changing the tilt parameter, for vertical-axis rotations. The eye-centered slants specified by the texture gradients are

$$S_L = S_C + \frac{m}{2} \text{ and } S_R = S_C - \frac{m}{2},$$

so they differ by m . The eye-centered, texture-specified tilts are $t_L = t_R = 0$. The left panel of Fig. 1A shows eye-centered slants and tilts as a function of Cyclopean slant for $t_C = 0^\circ$.

Now consider a fronto-parallel plane rotated about the horizontal axis ($S_C = \Delta$ and $t_C = 90^\circ$). For simplicity, we now use positive slants and let the tilt flip from 0 to 180° when necessary. The texture-specified slants are equal in the two eyes:

$$S_L = S_R = \cos^{-1}\left[\cos\left(\frac{m}{2}\right)\cos(S_C)\right].$$

But the tilts differ:

$$t_L = \cos^{-1}\left[\frac{\sin\left(\frac{m}{2}\right)}{\sin(S_L)}\right] \text{ and } t_R = \cos^{-1}\left[\frac{\sin\left(-\frac{m}{2}\right)}{\sin(S_R)}\right].$$

The right panel of Fig. 1A shows these eye-centered slants and tilts.

When the Cyclopean tilt is between 0 and 90 deg, the eye-centered slants and tilts will both differ.

In summary, the transformation required to convert texture signals into Cyclopean coordinates depends on tilt. This fact raises the interesting possibility that different errors occur in that transformation for one tilt as opposed to another. Perhaps an explanation for slant anisotropy could be derived from an understanding of these transformations.

References

- Backus, B. T., & Banks, M. S. (1999). Estimator reliability and distance scaling in stereoscopic slant perception. *Perception, 28*, 217-242.
- Backus, B. T., Banks, M. S., van Ee R., & Crowell, J. A. (1999). Horizontal and vertical disparity, eye position, and stereoscopic slant perception. *Vision Research, 39*, 1143-1170.
- Banks, M. S., Gepshtein, S., & Landy, M. S. (2004). Why is spatial stereoresolution so low? *Journal of Neuroscience*, in press.
- Banks, M. S., Hooge, I. T. C., & Backus, B. T. (2001). Perceiving slant about a horizontal axis from stereopsis. *Journal of Vision, 1*, 55-78.
- Blake, A., Bülthoff, H. H., & Sheinberg, D. (1993). Shape from texture: Ideal observers and human psychophysics. *Vision Research, 33*, 1723-1737.
- Bradshaw, M. F., Glennerster, A., & Rogers, B. J. (1996). The effect of display size on disparity scaling from differential perspective and vergence cues. *Vision Research, 36*, 1255-1264.
- Bradshaw, M. F., Hibbard, P. B., van der Willigen, R., Watt, S. J., & Simpson, P. J. (2002). The stereoscopic anisotropy affects manual pointing. *Spatial Vision, 15*, 443-458.
- Buckley, D., & Frisby, J. P. (1993). Interaction of stereo, texture and outline cues in the shape perception of three-dimensional ridges. *Vision Research, 33*, 919-933.
- Burt, P., & Julesz, B. (1980). A disparity gradient limit for binocular fusion. *Science, 208*, 615-617.
- Cochran, W. G. (1937). Problems arising in the analysis of a series of similar experiments. *Journal of the Royal Statistical Society (Supplement), 4*, 102-118.
- Cutting, J. E., & Millard, R. T. (1984). 3 gradients and the perception of flat and curved surfaces. *Journal of Experimental Psychology: General, 113*, 198-216.
- de Berg, M., van Kreveld, M., Overmars, M., & O. Schwarzkopf (2000). *Computational Geometry: Algorithms and Applications* (2nd Ed.). New York: Springer-Verlag.
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*, 429-433.
- Frisby, J. P., Buckley, D., & Horsman, J. M. (1995). Integration of stereo, texture, and outline cues during pinhole viewing of real ridge-shaped objects and stereograms of ridges. *Perception, 24*, 181-198.

- Gårding, J., Porrill, J., Mayhew, J. E. W., & Frisby, J. P. (1995). Stereopsis, vertical disparity and relief transformations. *Vision Research*, *35*, 703-722.
- Gepshtein, S., & Banks, M. S. (2003). Viewing geometry determines how vision and haptics combine in size perception. *Current Biology*, *13*, 483-488.
- Gillam, B., & Ryan, C. (1992). Perspective, orientation disparity, and anisotropy in stereoscopic slant perception. *Perception*, *21*, 427-439.
- Green, D. M., & Swets, J. A. (1974). *Signal Detection Theory and Psychophysics*. New York: Robert E. Krieger.
- Heath, G. G. (1956). The influence of visual acuity on accommodative responses of the human eye. *American Journal of Optometry*, *33*, 513-523.
- Helmholtz, H. (1910). *Physiological Optics*. New York: Dover.
- Hillis, J. M., & Banks, M. S. (2001). Are corresponding points fixed? *Vision Research*, *41*, 2457-2473.
- Hillis, J. M., Ernst, M. O., Banks, M. S., & Landy, M. S. (2002). Combining sensory information: Mandatory fusion within, but not between, senses. *Science*, *298*, 1627-1630.
- Howard, I. P., & Kaneko, H. (1994). Relative shear disparities and the perception of surface inclination. *Vision Research*, *34*, 2505-2517.
- Howard, I. P., & Rogers, B. J. (2002). *Seeing in Depth. Volume 2: Depth Perception*. Toronto: I Porteous.
- Johnston, E. B. (1991). Systematic distortions of shape from stereopsis. *Vision Research*, *31*, 1351-1360.
- Johnston, E. B., Cumming, B. G., & Parker, A. J. (1993). Integrations of depth modules: Stereopsis and texture. *Vision Research*, *33*, 813-826.
- Knill, D. C. (1998a). Surface orientation from texture: Ideal observers, generic observers and the information content of texture cues. *Vision Research*, *38*, 1655-1682.
- Knill, D. C. (1998b). Discrimination of planar surface slant from texture: Human and ideal observers compared. *Vision Research*, *38*, 1683-1711.
- Knill, D. C., & Saunders, J. A. (2003). Do human optimally integrate stereo and texture information for judgments of surface slant? *Vision Research*, *43*, 2539-2558.
- Landy, M. S., & Kojima, H. (2001). Ideal cue combination for localizing texture-defined edges. *Journal of the Optical Society of America A*, *18*, 2307-2320.
- Landy, M. S., Maloney, L. T., Johnston, E. B., & Young, M. J. (1995). Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research*, *35*, 389-412.
- Legge, G. E. (1984). Binocular contrast summation I. Detection and discrimination. *Vision Research*, *24*, 373-383.
- Mamassian, P., Landy, M. S., & Maloney, L. T. (2002). Bayesian modeling of visual perception. In R. P. N. Rao, B. A. Olshausen, & M. S. Lewicki (Eds.), *Probabilistic Models of the Brain: Perception and Neural Function* (pp. 13-36). Cambridge, MA: MIT Press.
- Mather, G., & Smith, D. R. (2002). Blur discrimination and its relation to blur-mediated depth perception. *Perception*, *31*, 1211-1219.
- Ogle, K. N. (1950). *Researches in Binocular Vision*. Philadelphia, PA: W. B. Saunders.
- Ogle, K. N., & Ellerbrock, V. J. (1946). Cyclofusional movements. *AMA Archives of Ophthalmology*, *36*, 700-735.
- Oruç, I., Maloney, L. T., & Landy, M. S. (2003). Weighted linear cue combination with possibly correlated error. *Vision Research*, *43*, 2451-2468.
- Rogers, B. J., & Bradshaw, M. F. (1995). Disparity scaling and the perception of frontoparallel surfaces. *Perception*, *24*, 155-179.
- van Beers, R. J., Sittig, A. C., & Denier van der Gon, J. J. (1998). The precision of proprioceptive position sense. *Experimental Brain Research*, *122*, 367-377.
- van Beers, R. J., Wolpert, D. M., & Haggard, P. (2002). When feeling is more important than seeing in sensorimotor adaptation. *Current Biology*, *12*, 834-837.
- van Ee, R., & Erkelens, C. J. (1998). Temporal aspects of stereoscopic slant estimation: An evaluation and extension of Howard and Kaneko's theory. *Vision Research*, *38*, 3871-3882.
- Young, M. J., Landy, M. S., & Maloney, L. T. (1993). A perturbation analysis of depth perception from combinations of texture and motion cues. *Vision Research*, *33*, 2685-2696.