

Surface tilt (the direction of slant): A neglected psychophysical variable

KENT A. STEVENS

Massachusetts Institute of Technology, Cambridge, Massachusetts

Surface slant (the angle between the line of sight and the surface normal) is an important psychophysical variable. However, slant angle captures only one of the two degrees of freedom of surface orientation, the other being the *direction* of slant. Slant direction, measured in the image plane, coincides with the direction of the gradient of distance from viewer to surface and, equivalently, with the direction the surface normal would point if projected onto the image plane. Since slant direction may be quantified by the tilt of the projected normal (which ranges over 360 deg in the frontal plane), it is referred to here as *surface tilt*. (Note that slant angle is measured perpendicular to the image plane, whereas tilt angle is measured in the image plane.) Compared with slant angle's popularity as a psychophysical variable, the attention paid to surface tilt seems undeservedly scant. Experiments that demonstrate a technique for measuring apparent surface tilt are reported. The experimental stimuli were oblique crosses and parallelograms, which suggest oriented planes in 3-D. The apparent tilt of the plane might be probed by orienting a needle in 3-D so as to appear normal, projecting the normal onto the image plane, and measuring its direction (e.g., relative to the horizontal). It is shown to be preferable, however, to merely rotate a line segment in 2-D, superimposed on the display, until it appears normal to the perceived surface. The apparent surface tilt recorded in these experiments corresponded closely to that predicted by assuming the 3-D configurations consist of equal-length lines and perpendicular intersections.

Slant is a commonly investigated psychophysical variable in studies of surface perception. It is usually defined as the angle between the line of sight and the normal to the viewed surface. Slant angle varies over a range of 90 deg, where zero slant corresponds to a surface patch that lies perpendicular to the line of sight, that is, parallel to the image plane. Gibson (1950a, 1950b) observed correctly that surface orientation has two degrees of freedom and that slant varies both in magnitude and direction. There is a large literature, both theoretical and experimental, concerning the perception of slant magnitude (i.e., angle), but to the author's knowledge, the *direction* in which a surface slants has not been used as an independent psychophysical variable. Moreover, until recently, little theoretical study had been made of perception of slant direction. The goal of this article is two-fold: (1) to consider the directional aspect of slant as a perceptual quantity, a property of the visual world to be

determined from visual sources, and (2) to illustrate its use as a psychophysical variable.

The quantity to which Gibson (1950b) referred only as "the direction of slant" will be given a simple and indicative name: surface tilt (see Marr, 1982; Stevens, 1979). The use of tilt in this context will be shown to be consistent with its conventional usage in, for example, the tilt aftereffect (in general, tilt is an angular quantity used to describe orientation or direction measured in the frontoparallel image plane). It will be found equally appropriate to use tilt to refer to the direction of slant, which is also measured in the image plane. (In the following, tilt will be used without the qualifier "slant" when the context makes it obvious that one refers to a surface.) Since the direction of slant is seldom discussed, the term *slant* has become synonymous with slant angle or magnitude; that practice will be continued here. Hence, the two degrees of freedom of surface orientation may be described by slant and tilt, both of which are angular measures: slant is an angle measured perpendicular to the image plane, and tilt is an angle measured in the image plane.

The Two Degrees of Freedom of Surface Orientation

Only two quantities are needed to specify uniquely the orientation of a surface patch relative to a viewpoint, but many different schemes are possible. One

I gratefully acknowledge the insightful comments provided by David Marr and Shimon Ullman. This work was carried out at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology, support for which is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research Contract N00014-75-C-0643 and in part by National Science Foundation Grant MCS77-07569. I am now at the Department of Computer and Information Science, University of Oregon, Eugene, Oregon 97403.

broad approach is to describe the surface normal relative to a viewer coordinate system. Given a Cartesian coordinate system with the x-y plane coinciding with the frontal plane, the distance z to a smooth patch of surface along the line of sight is some continuous function $z = f(x,y)$. The surface normal, N_{3-D} , is given by grad f:

$$N_{3-D} = \frac{\partial f}{\partial x}i + \frac{\partial f}{\partial y}j - k.$$

The partial derivatives with respect to x and y are conventionally referred to as p and q, respectively. Note that the surface normal projects onto the x-y (image) plane as the vector

$$N_{2-D} = pi + qj$$

and that N_{2-D} "points" in the direction τ of the distance gradient, that is, the image direction in which given distance from the viewer to the surface increases most rapidly. Specifically, τ is defined by

$$\tau = \tan^{-1}(q/p). \tag{1}$$

Note also that the slant angle σ between the surface normal N_{3-D} and the line of sight (z-axis) is defined by

$$\sigma = \tan^{-1}(p^2 + q^2)^{1/2}. \tag{2}$$

The components p and q may be regarded as the Cartesian coordinates of a point in gradient space (Horn, 1975; Huffman, 1971; Mackworth, 1973). There is a one-to-one correspondence between visible surface orientations and points in gradient space: the origin corresponds to a surface parallel to the frontal plane (i.e., the surface normal coincides with the z-axis and the slant angle is zero), a slanted surface corresponds to a point in gradient space at some radial distance from the origin given by Equation 2 in the direction given by Equation 1, and a surface of 90-deg slant lies infinitely far from the origin. Any given visible surface orientation maps to some point (p,q) in gradient space, which can also be addressed by polar coordinates as $(\tau, \tan\sigma)$. That is, each pair (σ, τ) specifies a unique surface orientation, and the set of all visible surface orientations corresponds to the set $\{\sigma, \tau \mid 0 \leq \sigma \leq 90 \text{ deg}, 0 \leq \tau < 360 \text{ deg}\}$.

To illustrate slant and tilt, the ellipses in Figure 1 represent the projections of circles of varying spatial orientation. The ellipses in a column correspond to surface patches of constant tilt but varying slant; those in a row correspond to surface patches of constant slant but varying tilt. Note that each ellipse, regarded as a slanted circle, is ambiguous: two interpretations in depth that differ by 180 deg in tilt direc-




Figure 1. The ellipses in a column correspond to circles of constant tilt but varying slant, and those in a row, to surface patches of constant slant but varying tilt.

tion are possible. (Reversals in apparent-tilt direction will be discussed further below.) In the above discussion, the visible surface normal was defined relative to a fixed but arbitrary retinal-based coordinate system. From the expression for the projection of the normal, expressions for the direction τ and magnitude σ of slant were defined. It would be a simple matter to propose other means for describing surface orientation, for instance, by the angles the normal in 3-D (three dimensions) makes with the two image axes x and y. There are, however, computational reasons and psychophysical evidence suggesting that the slant-tilt formulation has psychological relevance (see Stevens, 1979, in press). These are reviewed briefly in the following.

Slant and Tilt as Perceptual Variables

Of the various means by which surface orientation might be encoded, the choice of slant and tilt has several advantages from a perceptual point of view. First, the formulation of surface orientation as essentially magnitude and direction allows the visual system to decompose the problem of perceiving surface orientation into two relatively independent sub-problems: the determination of "which way" and "how much." As shall be discussed, there are distinct and relatively independent sources of information about these two quantities.

Surface orientation was defined earlier relative to the distance gradient. It follows that the visual system might determine the orientation of a surface patch from the local gradient of perceived distance (provided, e.g., by stereopsis). The surface tilt corresponds to the direction of greatest variation in per-

ceived distance (and perpendicular to the orientation of least variation), and slant varies with the magnitude of the gradient according to Equation 2. Hence, perceived distance across a smooth surface patch can be converted to local surface orientation.

Surface orientation may also be determined more directly from an image by several processes. For example, as Gibson (1950b) observed, the direction of the texture gradient (quantified, say, by local texture density) corresponds to the tilt direction; the tilt is also perpendicular to the orientation in which that measure is least variable (Stevens, 1981a).¹ In addition to the gradient of texture, the direction and magnitude of image texture foreshortening carries information about surface orientation (Flock, 1964; Witkin, 1981). Other direct sources of orientation information include motion and optic flow (Hoffman, 1982; Koenderink & van Doorn, 1976; Prazdny, 1980) and certain configurations of contours lying across a surface (Stevens, 1981b). In each of these studies, the available information about surface orientation seems to be provided most straightforwardly in terms of the direction and magnitude of slant. It is parsimonious to expect that the internal encoding of surface orientation reflects the manner in which it is most directly perceivable. There are additional reasons for expecting that slant and tilt are psychologically real variables, that is, more than merely mathematical conveniences. One reason follows from the observation that slant angle is often difficult to determine accurately, more so than tilt. (More generally, it appears a more difficult problem to determine how rapidly a surface recedes from the viewer than to determine the direction in which it recedes.) This conclusion is supported by several recent theoretical studies. For example, the direction of greatest texture foreshortening reliably corresponds to the tilt direction (Witkin, 1981). The amount of slant, however, can be inferred from the magnitude of the texture foreshortening only when the foreshortening varies predictably with slant angle. Foreshortening relations are usually formulated for physical texture that lies flush with the surface (such as mottled texture on the ground cast by sunlight filtered through trees), but physical texture usually extends above and below the mean surface level (such as waves, rocks, bushes), which greatly complicates the interpretation of foreshortening (Stevens, 1981a; Witkin, 1981).² Similar geometric analysis suggests that it is easier to derive tilt than slant from the motion field, both in the case of perspective (Koenderink & van Doorn, 1976; Prazdny, 1980) and orthographic projection (Hoffman, 1982).

Since the direction and magnitude of slant may be determined independently, and the information about the two components may have differing reliability in natural situations, it is valuable to keep their per-

ceptual processing separate. The slant-tilt formulation provides this advantage, of course. There are additional advantages to encoding surface orientation by slant and tilt within the visual system. For instance, reversals in depth are easily attributed to ambiguity in surface orientation in which tilt direction is known only up to a 180-deg reversal (Stevens, 1979). Also, the polar form provides some invariance to rotations of the image plane: the tilt varies linearly with rotation (e.g., as the head is tilted) but the slant remains constant. Alternative Cartesian-based encoding schemes would suffer both components' changing with image rotations. A more complete discussion is in Stevens (in press).

Tilt as a Psychophysical Variable

Studies of the theoretical relationship between surface orientation and variables in the image (e.g., Purdy, Note 1, regarding texture gradients, and Attneave & Frost, 1969, regarding right angles and other geometric relations) have concentrated almost exclusively on the perception of slant (i.e., angle). Accordingly, the psychophysical experiments that supported or challenged these theoretical studies used slant as the primary psychophysical variable. Until recently, surface tilt was largely overlooked both as a perceptual variable and as a psychophysical quantity. One might conjecture that part of the reason tilt has been neglected is that the direction a surface slants is usually very obvious, perhaps more so than the particular slant angle in experimental conditions. Nonetheless, the direction of slant must also be perceived since it is implicit in the visual array.

Measuring Apparent Tilt

There are various means by which apparent tilt may be measured by adjusting the tilt of a response surface to match that of the presented stimulus surface, by rotating the stimulus image so that the apparent tilt coincides with the vertical, or by adjusting an indicator so that it points in the same direction as would the surface normal. For reasons that will be discussed, having subjects indicate apparent tilt by a pointer is likely the preferable method.

The use of a response surface to match apparent tilts, analogous to the "palm board" used to match apparent slant (Gibson, 1950b; Flock & Moscatelli, 1964; Kraft & Winnick, 1967; Purdy, Note 1), is problematic because the subject, in general, must adjust both slant and tilt when matching surface orientations. It is often important in an experiment to vary both the slant and tilt of the stimuli (see below); hence, the subject must also vary the slant of the response board. Since naive subjects are usually unaware of the fact that tilt and slant are orthogonal variables that may be adjusted independently, it frequently occurs that when the tilts have been brought

into alignment, the subject often perturbs both components alternatively in attempting to match the two surface orientations completely (even though only the slants need to be aligned). The following alternative approaches, however, allow more precision in measuring apparent tilt.

One alternative involves rotating the stimulus display so that the apparent tilt corresponds to the vertical. This technique is rather general. For example, rotate Figure 2 so that the ruled surface in 3-D appears horizontal. By instructing the subject to make careful adjustments so that the surface does not tilt or "tip" to either the left or the right, precise judgments of apparent tilt can be had. Rotation of the stimulus display is one way of dissociating apparent tilt and slant; however, it has a significant limitation: this technique does not measure the apparent tilt of the original stimulus configuration, since the subject modifies the stimulus by rotating it. Nonetheless, it provides a means for determining the apparent orientation of the normal relative to the perceived surface. For example, in the experiments reported below, apparent tilt was recorded by the angle relative to the horizontal, and the stimulus figure was displayed at some angle relative to the horizontal. Rather than adjust an indicator to point in the same direction as the imagined surface normal, the stimulus figure would be rotated so that the $\tau = 90$, that is, so the normal would point vertically. These two tasks provide comparable measures of apparent tilt relative to the perceived surface when apparent tilt varies linearly with the stimulus orientation in the image plane (see Experiment 1). Alternatively, one may set the tilt to the




Figure 2. This texture gradient can be interpreted as an oblique perspective view of a planar surface with parallel linear rulings. The tilt of the apparent surface can be judged by orienting the figure so that the surface appears horizontal (and the surface normal points vertically upward).

horizontal, so that the normal would point to the left or the right.)

A more general means for measuring tilt is to superimpose on the stimulus surface display an indicator that is adjusted by the subject so as to appear normal to the surface. In a pilot study, a 3-D apparatus consisting of a needle attached by a small universal joint to the center of a transparent plate was constructed. The plate was mounted flush against a CRT display on which stimulus surfaces were projected. In dim background illumination, the fixed end of the needle appeared to touch the depicted surface; the movable end of the needle could then be oriented so that the needle appeared to be perpendicular to the surface. To measure the tilt after the needle was positioned to the subject's satisfaction, a line was displayed on the CRT with one end fixed so that it was coincident with the fixed end of the needle, and rotated until it was precisely hidden behind the needle from the subject's point of view. (Note that parallax errors might be introduced if the rotatable needle is not close to the stimulus display or if the subject does not remain perpendicular to the plane of the display.) The apparent tilt was then measured by the angle between the displayed line and the horizontal.

Surprisingly, although the subjects had little difficulty orienting the needle so that it appeared normal, they were much more confident in merely rotating the line on the CRT, circumventing the needle apparatus entirely. (When adjusting the needle, much of the effort apparently was expended in deciding on the slant. It was common to observe a subject's moving the tip of the needle alternatively towards and away from the display while keeping the tilt substantially constant.) The implicit slant-adjustment task was avoided by removing the physical needle apparatus and instead merely rotating the displayed "needle" on the CRT. The displayed needle appeared to be oriented in 3-D relative to the depicted surface, particularly under monocular presentation.

Two refinements were later incorporated. It was found that if the orientation of the 2-D needle is within the range that is geometrically consistent with being the projection of the surface normal (see Stevens, 1981b, for geometry), it will tend to appear perpendicular to the surface. It is therefore difficult to choose the subjectively "best" judgment of apparent perpendicularity within this range when all choices seem more or less satisfactory. This tendency is particularly strong in viewing simple line drawings; it becomes difficult to choose the subjectively "best" judgment of apparent perpendicularity when many orientations seem satisfactory. Moreover, suppose the fixed end of the 2-D needle is placed on the screen at the point of intersection of two lines that appear perpendicular in 3-D. When the needle is correctly oriented relative to the figure, the display corresponds to a right trihe-

dron in 3-D, that is, three mutually perpendicular lines. But if the configuration is not consistent with the projection of a line perpendicular to two perpendicular lines, there is ambiguity as to which angle is not a right angle: either the lines are perpendicular and the needle is not normal or the needle is perpendicular to the surface and the lines do not intersect at a right angle. The ambiguity is largely avoided by having the needle flash while the subject continuously views the stimulus surface: the subject visualizes the surface normal protruding above the surface and then flashes the needle to compare its orientation with that visualized. The needle acts as a probe without significantly perturbing the 3-D impression.

The second refinement concerned the length of the displayed needle. A 3-D needle would be expected to change its projected length as it rotates in 3-D (reflecting changes in foreshortening relative to the viewer), but the displayed needle originally remained constant in length. This was disconcerting and distracting; to some subjects, the needle appeared to change length as it rotated. This distraction was avoided satisfactorily simply by having the needle extend beyond the field of view.

Controlling for Depth Reversals

Apparent-depth reversals are a common phenomenon, especially apparent with figures in orthographic projection, such as in the Necker cube, the Schroeder staircase, and the Mach illusion. An ellipse interpreted as a slanted circle is a particularly simple example: the bottom edge of the ellipse in Figure 3a can be interpreted as being either nearer or farther than the upper edge. The two interpretations are illustrated in Figures 3b and 3c. (Note again that the interpretations correspond to reversals in the apparent *direction* of the tilt, that is, that they differ by 180 deg.)




Figure 3. (A) The lower edge of the disk could either be nearer or farther than the upper edge. (B) The normal corresponds to the depth interpretation where the lower edge is nearer. (C) The reversal. Note the 180 deg reversal in tilt direction.

Experiments on apparent tilt have to cope with spontaneous reversals in depth that radically change the apparent three-dimensionality of the figure. Depth reversals may occur even when the normal is displayed (examine Figure 3c, for example). Spontaneous depth reversals often distract the subject and should therefore be considered in the experimental design.

One approach is to orient the figure to bias one depth interpretation over another. It is well known that, given two vertically separated points in an image, the lower point usually corresponds to a closer point in 3-D, because of our conventional perspective relative to objects on the ground. Hence, the ellipse in Figure 3a is usually interpreted as in Figure 3b (and Figure 3c would correspond to viewing a ceiling fixture). More generally, if the tilt is roughly vertical, the upward-pointing normal is usually preferred over the downward. (If the display is oriented so that the tilt is roughly horizontal, the surface is interpreted with the normal pointing to the left or the right with roughly equal preference—rotate Figure 3a to demonstrate this.) Thus, to minimize the depth reversals, one may rotate the stimulus figure so that the tilt is always within, say, ± 45 deg of vertical. (When viewing a surface with approximately vertical tilt the subject might favor a tilt judgment that corresponds to the exact vertical, but this bias has not been observed experimentally.) A less restrictive approach is to familiarize the subjects with the reversals by having them observe how apparent-tilt reverses direction along with reversals in depth. The subjects are then instructed to maintain a depth interpretation that places the normal in, for example, the second quadrant. If the depth reverses spontaneously, the subjects would regain the original depth interpretation before continuing with the experiment.

Surface-Tilt Judgments From Simple Geometric Figures

Two experiments that demonstrate the use of tilt as a psychophysical variable are described below. The experiments involved viewing oblique crosses (Figure 4a) and parallelograms (Figure 4b), which appear in 3-D as consisting of right-angle intersections. The parallelogram may be interpreted as the orthographic projection of a rectangle (in perspective, it would generally project as an irregular quadrilateral, and as a trapezoid from certain viewpoints). Many studies have been made of the apparent slant of a rectangle in 3-D, usually for the special case of a rectangle projecting as a trapezoid (see, e.g., Braunstein & Payne, 1969; Flock, 1965; Freeman, 1966; Olson, 1974). But to the author's knowledge, the apparent tilt of a slanted rectangle has not been studied experimentally. The following experiments examine the apparent tilt of two related figures—a rectangle projected ortho-




Figure 4. The oblique crosses (A) and the corresponding parallelograms (B) are interpreted as having similar orientations in 3-D. Both can be parameterized in terms of three variables: the obtuse angle of intersection, the ratio of the two line lengths, and the figure's orientation in the image plane. The 3-D orientation can be determined if each is regarded as the projection of two equal-length lines meeting at a right angle in space.

graphically as a parallelogram, and a cross (both of which can be parameterized by two lengths and an included obtuse angle). Hering and Wundt recognized that a cross defined by two obliquely intersecting lines appears strikingly three-dimensional (see Luckiesh, 1965; Robinson, 1972), but this very simple figure has received little experimental attention.

These two figures demonstrate an apparent tendency to interpret an obtuse intersection of two line segments as a foreshortened right-angle intersection in 3-D. Moreover, there is a tendency to interpret the line segments as equal length in 3-D; that is, the parallelogram is interpreted as a slanted square (Attneave, 1972; Attneave & Frost, 1969). Thus far, the major support for these two assumptions (equal length and perpendicularity) have come from slant experiments; in these, apparent-tilt judgments have provided additional support (see also Perkins, 1972).

EXPERIMENT 1

The first experiment examined the apparent tilt as a function of the image orientation of an oblique cross and a parallelogram. If tilt were determined by the geometry of the figure independently of its orientation, a linear association between tilt judgments and image orientation (with unity slope) would be expected.

Method

Stimuli. Line-drawn crosses and parallelograms were projected on a Knight raster scan CRT display in a darkened room; the lines were luminous. Both the cross and the parallelogram were parameterized by the ratio (R) of the two line lengths, the obtuse angle of intersection (β), and the orientation (α) of the figure on the screen (Figure 5b). In Experiment 1, the values were $R = .27$ and $\beta = 110$ deg; the experimental variable was α , and the independent

variable was tilt (τ). (Arbitrarily define $\tau = 0$ deg, to correspond to a surface patch whose normal points to the right.)

Since spontaneous reversals in depth interpretation were expected if the total rotation exceeded 90 deg, the rotation of the figures was restricted to a range of 70 deg, specifically, $\alpha = 10, 20, 40, 60,$ and 80 deg, placing the tilt into the second and fourth quadrants (depending on the depth interpretation—see Figure 5a). The figures subtended roughly 7 deg of visual angle and were viewed monocularly through a 25-mm-diam circular aperture positioned 50 cm from the display.

A computer program controlled the stimulus presentation, the data collection, and the flashing and rotation of the tilt needle. Rotation was stepped clockwise and counterclockwise in 1- and 5-deg increments by individual keystrokes. By a separate keystroke, the needle would appear on the screen for roughly 250 msec.

Procedure. A stimulus figure was displayed continuously on the screen, and the subject was instructed to interpret the figure as a 3-D surface and to visualize the normal to the surface. The subject was introduced to the depth reversals, and for uniformity of the data, was instructed to maintain a depth interpretation that placed the normal in the fourth quadrant (otherwise, tilt judgments would form two distributions separated by 180 deg). The experimental task was to superimpose the flashed needle over the visualized normal by successive adjustments. The angle of the needle on the display was then recorded, and the next stimulus figure was presented. The experiment involved randomized presentations of the two types of figures at five orientations. Each of the 10 presentations was given once with unlimited viewing time. Six unpaid, volunteer graduate students (five male, one female) were the subjects.

Results

The data across subjects were tabulated separately for the parallelogram and oblique-cross figures (Table 1). In both cases, the linear association between τ and α was significant: $r = .98$ [$t(30) = 27.736, p < .05$] for the cross, and $r = .94$ [$t(30) = 14.473, p < .05$] for the parallelogram. The computed slopes of simple linear regression lines were $.96$ ($SE = .035$) for the




Figure 5. The apparent tilt (τ) of the surface suggested by the obtuse intersection (A) was measured for varying values of image orientation (α), angle of intersection (β), and relative line length (R), as shown in (B).

Table 1
Values of Mean Apparent Tilt τ for Cross and Parallelogram
as a Function of Image Orientation α

α	Predicted τ	τ for Cross		τ for Parallelogram	
		Mean	SD	Mean	SD
10	281.45	285.33	1.51	287.50	11.31
20	291.45	294.50	2.59	290.80	9.44
40	311.45	316.67	8.57	309.00	8.32
60	331.45	333.33	5.61	329.67	4.03
80	351.45	352.50	3.39	352.67	10.71

Note—Both the cross and the parallelogram had fixed values of $R = .27$ and $\beta = 110$ deg. The predicted values of τ derive from the geometric constraint of perpendicularity and equal-length lines in 3-D. Note that the subjects interpreted the normal as pointing into the fourth quadrant.

cross, and .95 (SE = .066) for the parallelogram. Neither slope was significantly different from 1.0 [$t(30) = .785$, $p > .2$, and $t(30) = 1.126$, $p > .2$, respectively].

The data for both types of figure for each subject were then analyzed individually, and the correlation coefficients were all significant: the least significant finding was $r = .94$ [$t(3) = 4.007$, $p < .05$]. For the cross, the slopes of the linear regression lines for each subject ranged from .88 to 1.05. In comparing these slopes with 1.0, none of the differences reached significance ($p > .2$). For the parallelogram only, the slopes for two subjects were significantly different from 1.0.

Discussion

The linear association between τ and α suggests that (for the simple cross and parallelogram) the perception of surface tilt is essentially fixed relative to the stimulus configuration and is only linearly dependent on its orientation. In general, by holding constant all stimulus variables except for the orientation (α) of the display, one can examine to what extent perceived surface orientation is determined by the geometry of the stimulus independently of rotational attitude on the image plane. There are important effects on apparent tilt caused by rotation of the stimulus display, however. One such effect concerns the tendency to have spontaneous depth reversals, as discussed above, and to choose the depth interception that causes the normal to point upward rather than downward. There is also significant hysteresis in this regard, such that one can continuously rotate a figure and hold the same depth interpretation (i.e., the apparent tilt rotates with the figure).

When the depth reverses, in theory the apparent tilt should change by 180 deg, but pilot studies reveal slight deviations from a strict reversal in tilt direction. Hence, one cannot simply adjust the data by 180 deg, as needed, to place all tilt judgments in the same quadrant so that they may be analyzed together.

Incidentally, the tilt data for the cross and parallelogram were similar in this experiment, and in turn close to the tilt predicted by the equal-length and perpendicularity constraints. (An expression for the predicted τ as a function of β and R is derived in the Appendix; the data are compared with the prediction later, in the discussion of Experiment 2.) Despite the similarity between the data for the cross and the parallelogram, one should not conclude that crosses and parallelograms having identical obtuse angles and line-length ratios are seen to have identical tilts. Subsequent experiments have found differences in the apparent tilt of parallelograms and crosses having identical β , R , and α . Probably one cause is that the parallelogram, despite being correct for orthographic projection, displays a disconcerting lack of perspective, which tends to reduce the apparent slant and therefore to influence the tilt (since the two quantities are coupled for a given image configuration; see Appendix). There are clearly other factors governing our perception of surface orientation from simple figures than the geometric parameters β and R , but they are beyond the scope of this article.

EXPERIMENT 2

The second experiment examined the functional dependence of apparent tilt on the two parameters of the oblique cross—the ratio (R) of line lengths and the included obtuse angle of intersection (β). The experimental design was the same as that in Experiment 1.

Method

Stimuli. The cross figures were examined with three values of angle of intersection ($\beta = 110, 130$, and 170 deg) and three length ratios ($R = .272, .455$, and $.727$). So that the presentations would appear varied, two image orientations ($\alpha = 20$ and 60 deg) were used.

Procedure. The total of 18 presentations was performed with successive presentations alternating between $\alpha = 20$ and 60 deg. The sequence was randomized in terms of β and R . The subjects made tilt judgments in the second quadrant. Each presentation was given once; however, the data from the two image orientations would effectively provide two data points for each combination of β and R , after adjusting for α . Five unpaid, volunteer graduate students (four male, one female) were the subjects. One subject had participated in Experiment 1.

Table 2
Analysis of Variance: Mean Tilt (Combined Data From $\alpha = 20$ and 60 deg) Examined According to Effects of
Obtuse Angle β and Length Ratio R

Source	SS	df	MS (SS/df)	MSR
Between β	1340.188	2	670.094	23.805
Between R	1351.438	2	675.719	24.005
β - R Interaction	404.390	4	101.098	3.591
Residual	2280.047	81	28.149	

Note—All MSR's reach .05 significance.

Table 3
Values of Mean Tilt τ for Two Image Orientations, $\alpha = 20$ and 60 deg, Over Nine Combinations of Obtuse Angle β and Length Ratio R

β	R	Predicted τ	τ for $\alpha = 20$ deg		τ for $\alpha = 60$ deg		Comparison
			Mean	SD	Mean	SD	
170	.27	110.68	110.73	1.53	111.13	1.76	$p > .2$
170	.45	111.69	110.33	3.06	111.13	3.69	$p > .2$
170	.73	113.45	112.73	2.82	113.13	4.59	$p > .2$
130	.27	112.12	112.93	2.00	113.33	6.86	$p < .05\ddagger$
130	.45	115.96	116.33	4.60	119.90	4.09**	$p > .2$
130	.73	124.91	124.93	6.92	127.13	6.53	$p > .2$
110	.27	111.45	111.53	5.60	117.13	7.31*	$p > .2$
110	.45	114.48	117.73	3.34*	120.13	10.86	$p < .05\ddagger\ddagger$
110	.73	124.88	123.70	5.66	131.10	4.27†	$p < .05$

Note—The means in the column for $\alpha = 60$ deg were computed after reducing the data by 40 deg, so that the data could be compared at a single image orientation. The last column shows the results of comparison of the means at the two values of α . In comparing the two means, if the variances were not significant, then a *t* test was performed. Each mean was also compared with the corresponding theoretic value, and except where noted, the differences did not reach significance ($p > .02$).

*.2 < p < .1. ** .05 < p < .1. † p < .05. †† Variances significantly different by *F* test.

Results

In light of the observation that the tilt judgments were essentially fixed relative to the stimulus figure, independently of its orientation (α) in the image plane (Experiment 1), the τ data collected at $\alpha = 60$ deg were reduced by 40.0 deg in order to compare all data at $\alpha = 20$ deg. The values of τ for each image orientation were then tabulated for each of the nine combinations of β and R. The results of a two-way analysis of variance with equal replications are given in Table 2.

The data from $\alpha = 20$ deg were compared with the adjusted data from $\alpha = 60$ deg for a further test of whether there is a functional dependence of τ on the

image orientation. The results are shown in Table 3. The differences between the two sample means reached significance in only three instances ($\beta = 130$ deg, R = .27; $\beta = 110$ deg, R = .73); however, the actual differences are .4, 2.4, and 7.4 deg respectively. The mean tilt judgments are shown in Figure 6.

Discussion

A strong functional dependence of τ on both β and R was found. The values of τ were compared with the corresponding values of tilt that would be predicted if the lines in 3-D were perpendicular and of equal length. These values are given in Table 3. The judgment means did not differ significantly from those predictions, except where indicated.

The close correspondence between hypothetical and apparent tilt in the oblique cross supports the hypothesis that the 3-D interpretation incorporates constraints of perpendicularity and equal length. As noted above, the apparent tilt of parallelograms is not predicted as well. The parallel edges of the parallelogram may be partly responsible, since parallel lines on a slanted plane usually project in perspective to converging lines. The parallel edges in the image, therefore, may affect the apparent 3-D orientation of the plane (see also Attneave & Frost 1969).

GENERAL DISCUSSION

The use of a needle rotated in the image plane seems to be a useful probe for measuring the apparent tilt of a stimulus surface. Subjects find it natural and easy to perceive the needle as perpendicular to the surface in 3-D, and rapidly forget that the needle is actually in 2-D. The pilot experiments found that having the needle visible only when requested by the subject, and then only for a flash, re-




Figure 6. These figures show the mean judgments of surface tilt as a function of relative line length (R) and angle of intersection (β).

sults in more reliable data, probably because there is less opportunity for its presence to perturb the apparent surface orientation. Particularly when the subject visualizes how a surface normal would appear, when the needle is flashed momentarily one can make subtle judgments of the disparity between the imagined normal and the probe so that the two can be brought into close alignment.

The scheme is not without limitations, however. One potential problem is posed by the instructional set. Since spontaneous depth reversals may cause the apparent tilt to reverse direction in some stimuli, and since the reversal is not necessarily a precise reversal of 180 deg, one must either treat the data from the two depth interpretations separately or ignore one set of data, which is inefficient, or instruct the subjects to attend to only one of the two interpretations. Since usually it can be arranged so that the data will fall in one quadrant or its reflection (the first and third, or second and fourth), depending on the depth interpretation, one approach is to familiarize the subject with the reversals and their effect on apparent tilt, and to provide specific instructions to keep the interpretation in the desired quadrant. Similarly, the stimuli may be oriented so that the subject need only restrict the tilt responses to the first or second quadrants (recall that upward-pointing normals are preferred over downward). One might also criticize the practice of explicitly instructing the subject to restrict his range of responses, but the apparent tilt direction has been found usually to be compelling, so that restricting the subject to even a single quadrant is unlikely to influence the tilt judgments.

In closing, the surface tilt probe (and the technique of rotating the stimulus figure demonstrated by Figure 2) should have useful application for the psychophysical study of tilt, the neglected direction of slant.

REFERENCE NOTE

1. Purdy, W. C. *The hypothesis of psychological correspondence in space perception* (No. R60ELC56). Schenectady, N.Y.: General Electric Technical Information Series, 1960.

REFERENCES

- ATTNEAVE, F. Representation of physical space. In A. W. Melton & E. Martin (Eds.), *Coding processes in human memory*. New York: Wiley, 1972.
- ATTNEAVE, F., & FROST, R. The determination of perceived tridimensional orientation by minimum criteria. *Perception & Psychophysics*, 1969, 6, 391-396.
- BRAUNSTEIN, M., & PAYNE, H. W. Perspective and form ratio as determinants of relative slant judgments. *Journal of Experimental Psychology*, 1969, 81, 584-590.
- FLOCK, H. R. A possible optical basis for monocular slant perception. *Psychological Review*, 1964, 71, 380-391.
- FLOCK, H. R. Optical texture and linear perspective as stimuli for slant perception. *Psychological Review*, 1965, 72, 505-514.
- FLOCK, H. R., & MOSCATELLI, A. Variables of surface texture and accuracy of space perceptions. *Perceptual and Motor Skills*, 1964, 19, 327-334.
- FREEMAN, R. B. The effect of size on visual slant. *Journal of Experimental Psychology*, 1966, 71, 96-103.
- GIBSON, J. J. *The perception of the visible world*. Boston: Houghton Mifflin, 1950. (a)
- GIBSON, J. J. The perception of visible surfaces. *American Journal of Psychology*, 1950, 63, 367-384. (b)
- HOFFMAN, D. D. Inferring local surface orientation from motion fields. *Journal of the Optical Society of America*, 1982, 72, 888-892.
- HORN, B. K. P. Obtaining shape from shading information. In P. H. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill, 1975.
- HUFFMAN, D. A. Impossible objects as nonsense sentences. *Machine Intelligence*, 1971, 6, 295-323.
- KOENDERINK, J. J., & VAN DOORN, A. J. Local structure of movement parallax of the plane. *Journal of the Optical Society of America*, 1976, 66, 717-723.
- KRAFT, A. L., & WINNICK, W. A. The effect of pattern and texture gradient on slant and shape judgments. *Perception & Psychophysics*, 1967, 2, 141-147.
- LUCKIESH, M. *Visual illusions: Their causes, characteristics and applications*. New York: Dover, 1965.
- MACKWORTH, A. K. Interpreting pictures of polyhedral scenes. *Artificial Intelligence*, 1973, 4, 121-137.
- MARR, D. C. *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman, 1982.
- OLSON, R. K. Slant judgments from static and rotating trapezoids correspond to rules of perspective geometry. *Perception & Psychophysics*, 1974, 15, 509-516.
- PERKINS, D. N. Visual discrimination between rectangular and nonrectangular parallelepipeds. *Perception & Psychophysics*, 1972, 12, 396-400.
- PRAZDNY, K. Egomotion and relative depth map from optical flow. *Biological Cybernetics*, 1980, 36, 87-102.
- ROBINSON, J. O. *The psychology of visual illusion*. London: Hutchinson University Library, 1972.
- STEVENS, K. A. Representing and analyzing surface orientation. In P. H. Winston & R. H. Brown (Eds.), *Artificial Intelligence: An MIT perspective*. Cambridge: MIT Press, 1979.
- STEVENS, K. A. The information content of texture gradients. *Biological Cybernetics*, 1981, 42, 95-105. (a)
- STEVENS, K. A. The visual interpretation of surface contours. *Artificial Intelligence*, 1981, 17, 47-73. (b)
- STEVENS, K. A. Slant-tilt: The visual encoding of surface orientation. *Biological Cybernetics*, 1982, in press.
- WITKIN, A. P. Estimating shape from texture. *Artificial Intelligence*, 1981, 17, 17-45.

NOTES

1. Slant is harder to interpret from a texture gradient, however, for several reasons (see Stevens, 1981a). For instance, the magnitude of the texture density gradient is not related to slant in the simple manner that has been suggested (Gibson, 1950b; Purdy, Note 1), because it confounds variations in distance and surface curvature.
2. A proposal by Flock (1964) involves counting texture features in the direction of and perpendicular to the texture gradient, and inferring slant from their ratio. This alternative suffers the same problem of interpretation unless the image features are in reliable correspondence to discrete physical surface features, such as individual wave crests or rocks.

APPENDIX

An expression for tilt (τ) will be derived as a function of the obtuse angle of intersection (β) and the ratio of line

lengths (R) for the cross or parallelogram figure (as in Figure 5b). Note that orthographic (parallel) projection is assumed in this derivation. Sufficient constraint is provided when the corresponding lines in 3-D are perpendicular and of equal length. The approach is to regard the two lines in 3-D as vectors defining a plane whose cross product is the normal to the plane. Without loss of generality, rotate the image coordinate system so that the positive x -axis aligns with the longer line in the image. Tilt is then computed by projecting the normal onto the image plane and measuring the angle counterclockwise from the positive x -axis to the projected normal.

Call the 3-D vector corresponding to the longer line U , and that to the shorter, V . For convenience, consider the projection of U to be a unit vector:

$$U = \{1, 0, a\} \text{ and } V = \{R \cos\beta, R \sin\beta, b\}.$$

The quantities a and b are the unknown components of U and V along the z -axis (perpendicular to the image plane). (Assume that a is nonzero, i.e., that the vector U is foreshortened; if $a = 0$, then immediately $\tau = 90$ deg.) The surface normal N is the cross product

$$N = \{-aR \sin\beta, aR \cos\beta - b, R \sin\beta\}, \tag{1}$$

so the tilt τ is given by

$$\tau = \tan^{-1} \frac{N_y}{N_x}, \tag{2}$$

where N_x and N_y are the two components of the normal vector projected onto the image plane.

The expression for the normal in Equation 1 carries two unknowns, reflecting the two degrees of freedom of surface orientation when no restrictions are imposed. Now, since

the intersection is perpendicular, the dot product of U and V is zero, from which it follows:

$$b = \frac{-R \cos\beta}{a}. \tag{3}$$

Since U and V are equal length,

$$a^2 = R^2 + b^2 - 1. \tag{4}$$

Solving Equation 4 for a^2 , with substitution of Equation 3 to eliminate b , gives

$$a^2 = \frac{(R^2 - 1) + [(R^2 - 1)^2 + 4R^2 \cos^2\beta]^{1/2}}{2}.$$

Substituting Equation 3 into Equation 1 gives

$$N = \{-aR \sin\beta, \frac{a^2 + 1}{a} R \cos\beta, R \sin\beta\},$$

so that the tilt, from Equation 2, is

$$\tau = \tan^{-1} \left[-\cot\beta \frac{a^2 + 1}{a^2} \right].$$

The above derivation results in a rather complicated expression relating tilt to the stimulus variables β and R . Of course, it should not be construed from this that the perceptual processes in human vision need explicitly to employ such computations (see Attneave and Frost, 1969, for similar discussion and a similar mathematical treatment).

(Manuscript received August 20, 1982;
revision accepted for publication October 4, 1982.)