Journal of Experimental Psychology:
Human Perception and Performance
1984, Vol. 10, No. 5, 734–739

Observations

Simulation of Curved Surfaces From Patterns of Optical Texture

James T. Todd
Brandeis University

Ennio Mingolla
Center for Adaptive Systems
Boston University

Previous research on the perceptual analysis of optical texture has been severely restricted by the lack of an appropriate technique for distributing texture on curved surfaces in a uniform manner. In an effort to overcome this problem, the present article presents a new algorithm for generating stochastically regular distributions of texture on any smooth surface regardless of its curvature. We also present a new technique for representing the global organization of a textured image based on the formal concept of a projected area field.

The concept of texture in perceptual psychology has two distinct meanings that are sometimes designated as *surface texture* and *optical texture.* Surface texture refers to the physical or chemical discontinuities on an object's surface that can influence the reflection of light. Optical texture, on the other hand, refers to the discontinuities of intensity or wavelength within a cone of visual solid angles. It is useful to conceive of these two types of texture as if they were composed of elementary units. The elements of surface texture can be thought of as bounded regions of one reflectance surrounded by a background of some other reflectance, whereas the elements of optical texture can be portrayed similarly as bounded regions of homogeneous luminance surrounded by a background of some other luminance (cf. Todd, 1984). Because variations in luminance within a cone of visual solid angles are directly influenced by variations in reflectance on a visible surface, the elements of surface texture and optical texture are typically in one-to-one correspondence. It is important to keep in mind, however, that their overall patterns of organization are generally quite different due to the effects of perspective.

The importance of texture for the study of human perception was first recognized by Gibson (1950) in his seminal work *The Perception of the Visual World.* Gibson assumed that most of the surfaces encountered in nature have patterns of texture that are stochastically regular (i.e., the texture elements within equal areas of a surface have comparable distributions of size, shape, and density). Whenever this assumption is satisfied, he argued, the structure of a surface in three-dimensional space can be unambiguously specified by its corresponding pattern of optical texture. In order to understand why this is so, it is useful to consider a planar cross section of the cone of visual solid angles (i.e., a picture plane). Suppose that a surface is covered with small circular dots. Although the sizes and shapes of these dots in three-dimensional space may be identical, their projected sizes and shapes on the picture plane will vary as a function of two physical variables: (a) the distance of the surface from the point of observation and (b) its orientation in depth relative to the line of sight (see Figure 1). It is this systematic variation in the sizes and shapes of the optical texture elements that provides information about an object's three-dimensional form.

The ability of human observers to make use of texture information has been studied extensively during the past 3 decades. (See Braunstein, 1976, for an excellent review). Unfortunately, almost all of this research has been confined to the perception of planar surfaces, which may not be representative of other types of objects encountered in nature. The primary determinant of textural variations in an image of a planar surface is that some parts of the surface may be farther away from the point of observation than are others. Thus, as is evident in Figure 1, the optical texture elements may vary in size but will have relatively little variation in shape. In an image of a curved surface, on the other hand, much of the variation in texture is due to the fact that some parts of the surface have different orientations than do others. The optical texture elements in this case vary in both size and shape. The differences between curved and planar

We wish to thank Myron Braunstein, James Cutting, and Joseph Lappin for their comments on an earlier draft of this manuscript.

Requests for reprints should be sent to James T. Todd, Department of Psychology, Brandeis University, Waltham, Massachusetts 02254.

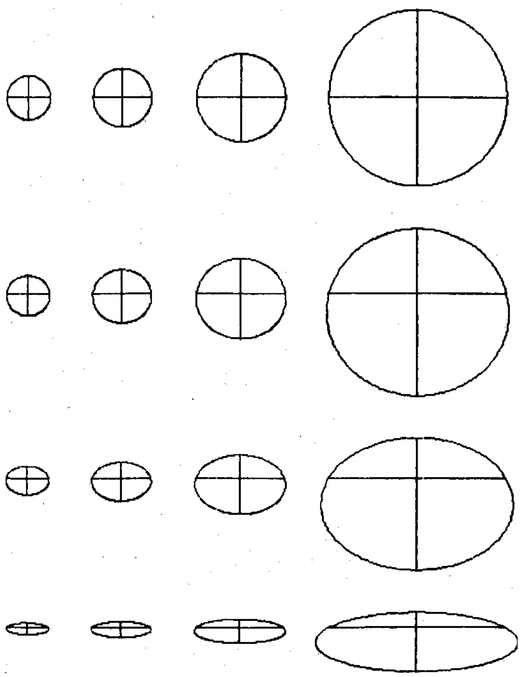

Figure 1. The effects of viewing distance and surface orientation on the polar projection of a circular texture element. (The columns from right to left depict viewing distances of 1, 2, 3, and 4 circle diameters, respectively. The rows from top to bottom depict slant angles of 0°, 25°, 50°, and 75°, respectively. The projected center of each circle is represented by the point where the horizontal and vertical lines intersect. The center point appears displaced in the projections of slanted circles at close viewing distances because of the convergence effects of linear perspective.)

surfaces are especially pronounced when viewing distance is sufficiently large to approximate a parallel projection (e.g., when a photograph is taken through a zoom lens). In that case, the proportional differences in viewing distance for different parts of a surface become negligible. The patterns of texture in an image of a curved surface still vary as a function of surface orientation (see Stevens, 1981; Witkin, 1981), but in an image of a planar surface the optical texture elements are all identical. On the basis of these observations, it is reasonable to expect that our perceptions of curved surfaces could be based on different sources of information than are our perceptions of planar surfaces. Indeed, this hypothesis has recently been confirmed by Cutting and Millard (1984).

One possible reason why perceptual psychologists have restricted themselves to the study of planar surfaces is the lack of an adequate technique for simulating texture patterns on curved surfaces.[1] In order to create a perceptually compelling visual display of a textured surface, the elements of surface texture must be distributed evenly so that all possible surface locations have an equal probability of being covered. This requirement is easily satisfied when dealing with planar surfaces. The usual procedure is to index the points on a surface with a Cartesian coordinate system and to select the coordinates of each texture element using a random number generator. When applied to a curved surface, however, this random selection of coordinate values inevitably results in an uneven distribution of texture—that is to say, some regions of the surface will have a higher concentration of texture elements than do others (cf. Braunstein, 1976).[2]

In our first attempt to solve this problem, we imagined that an object was completely covered by a long piece of string wrapped around its circumference in repeating loops. This defines a mapping relationship in which each point on an object's surface corresponds to a point somewhere along the length of the string. By selecting proportions of its total length at random, it is possible to generate an even distribution of texture. The main problem with this approach is that determining the length of a curved pieced of string requires the evaluation of a line integral. For some curves such as circles or ellipses this is a relatively simple process, but for many others the integral can be intractable. (Try it, for example, with a sinusoid.)

In an effort to avoid this difficulty, we recently devised an alternative technique that does not require the process of integration. Our procedure is analogous to tacking small pieces of construction paper at randomly selected positions on a curved surface in three-dimensional space. We assume that each piece of construction paper (i.e., a texture element) has an area S, and that its position on the surface is defined by a point at its center. It will be useful to refer to any region of the surface

---

[1] Several methods have been proposed in the computer graphics literature to create texture in smoothly shaded images by simulating small bumps and wrinkles on the depicted surface (e.g., Blinn, 1978; Blinn & Newell, 1976; Feibush, Levoy, & Cook, 1980; Haruyama & Barsky, 1984; Norton, Rockwood, & Skolmoski, 1982). There is also an algorithm recently described by Schweitzer (1983) that is appropriate for unshaded displays. Schweitzer's algorithm is similar to the one presented here but requires a much greater amount of computation time.

[2] For any curved surface whose Gaussian curvature is zero at every point (e.g., a cylinder) it is possible to adopt an alternative coordinate system for which a random selection of coordinate values produces a uniform distribution of texture (see Stevens, 1981; Todd & Mingolla, 1983). Such procedures do not exist, however, if the Gaussian curvature of any point on a surface is nonzero.

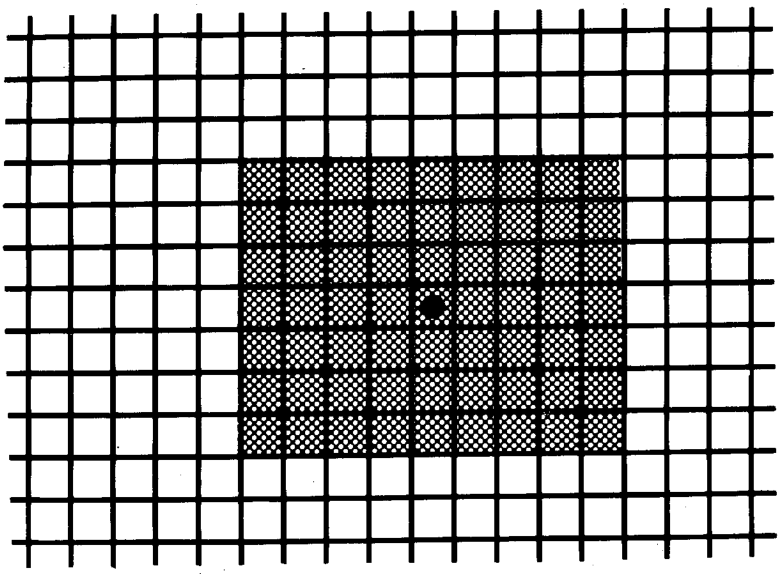

Figure 2. A network of bounded regions as might be produced by projecting a grid of picture elements onto a curved surface in three-dimensional space. (The shaded regions represent those that are covered by a texture element, the center of which is represented by a small dot. Our analysis assumes that the curvature of a surface within a local neighborhood is sufficiently small so that all of the bounded regions within that neighborhood have approximately the same area D. Because a texture element has an area S, it covers approximately S/D regions.)

that is covered by a texture element as *foreground* and any region that is not covered by a texture element as *background*. Our goal is to ensure that the probability q of being part of the background is the same for every unit area over the entire surface.

The first step in the procedure is to partition the surface into a network of bounded regions. This is accomplished by filling the display screen with a lattice grid of small rectangles called *picture elements* or *pixels* and projecting that grid onto the simulated surface. Note that the projective relationship between the picture plane and the surface is bidirectional—we can think of the surface as projecting onto the picture plane, or, alternatively, we can think of the picture plane as projecting onto the surface.[3] Because of the digital nature of computer graphics displays, an individual picture element cannot be partially filled; it is either on or off. To maintain consistency, the same constraint must also be applied to the simulated surface—that is to say, each of its bounded regions must be either completely covered by a texture element (with a probability $1 - q$) or completely uncovered (with a probability q).

Let us now compute the probability p that the center of a texture element is contained within a region R, which has an area D. We shall assume that the local curvature of the surface is not too large in the vicinity of R so that all of its neighboring regions will have approximately the same area. A possible configuration that could arise under these conditions is shown in Figure 2. The shaded area in this figure represents a texture element, the center of which is represented by a small dot. Because a texture element has an area S, any element that is positioned near R will cover approximately S/D regions. Similarly, there are approximately S/D regions where an element cannot be positioned for R to remain uncovered. Each of these S/D regions has a probability $1 - p$ of not containing the center of a texture element. Thus, because q is the probability that R is part of the background, it follows from the above that

$$q = (1 - p)^{S/D}$$

By rearranging terms we get

$$p = 1 - q^{D/S}$$

[3] It is interesting to note that the ratio of foreground to background remains invariant under this projective transformation. If every unit area of an object's surface has a probability $1 - q$ of being covered by surface texture, then it is also the case that every unit area of the picture plane has a probability of $1 - q$ of being covered by optical texture.

Note that all of the terms on the right side of this equation are given in the analysis. The values of q and S are defined as constants for an entire surface, and the value of D for the projected region of any given picture element is easily computed from the surface equation (see Appendix for details).[4] On the basis of this analysis it is possible to create a perceptually compelling visual display of a curved surface using the following procedure: For the surface projection of each picture element on the display screen we compute the value of p, select a random number between zero and one, and draw the appropriate projection of a texture element whenever the random number is less than p.

One important problem with this particular variation of our procedure is that it is computationally expensive. It is not uncommon for a graphics display to have millions of individual picture elements; to calculate the value of p for each projected region would take many hours on a typical laboratory computer. Because of the probabilistic nature of our algorithm, however, it is also possible to create appropriately textured displays by sampling only a small subset of these regions. The value of p for each region in the sample is then determined by

$$p = 1 - q^{D/Sm},$$

where m is the proportion of regions for which the calculations are performed. Figure 3 shows a simulated ellipsoid and a simulated hyperboloid of one sheet that were generated with this modified procedure. The following parameter values were used in the simulation: The display screen had a resolution of 640×480 pixels; the center of each object was located in the picture plane at a distance of 700 pixel units from the point of observation; for the ellipsoid, the x-, y-, and z-semiaxes had lengths of 260, 200, and 230 pixel units, respectively, while for the hyperboloid, they had lengths of 100, −100, and 100, respectively; the individual texture elements had dimensions of 15×15 pixel units and were positioned on the surface at randomly selected orientations; in each case the value of q was 0.8; and the value of m was 0.0625. The entire generation process took approximately 10 min using an LSI-11/23 processor.

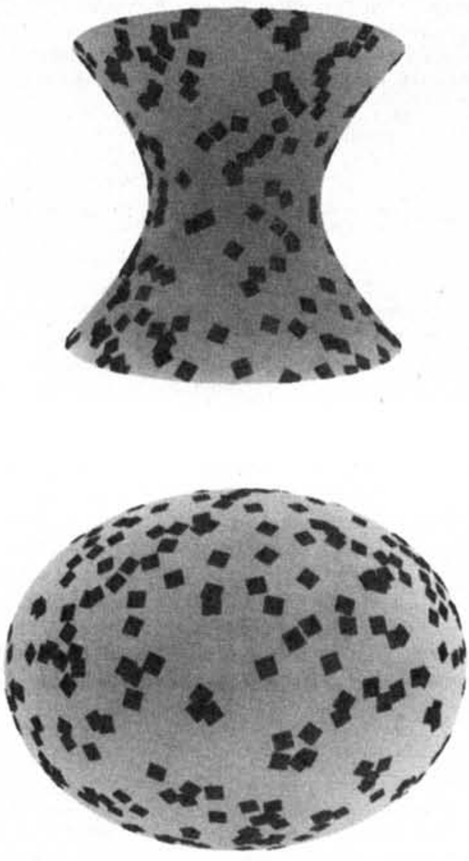

Figure 3. A computer simulation of a textured ellipsoid and a hyperboloid of one sheet.

To better understand the conceptual underpinnings of our proposed procedure, it is important to keep in mind that the value of p for any given region is uniquely determined by three physical variables: its area (D), the proportion of the surface that is not covered by texture elements (q), and the sizes of those elements (S). Note that the latter two of these variables reflect global properties of a surface that do not vary as a function of screen position. Thus, all variations in p are due entirely to the fact that some pixels project to larger regions of the observed surface than do others. It is also interesting to note in this regard that the values of D for the projected regions of individual picture elements form a two-dimensional scalar field. In other words, for every picture element on the display screen, there is a corresponding scalar quantity that represents the area of an observed surface to which it projects. Consider, for example, the two images depicted in Figure 3. The area field for each of these images is represented in Figure 4 by an interconnecting network of solid and dashed lines. The solid lines are called isoarea

[4] A simpler, though somewhat less accurate method of computing D has been described by Schweitzer (1983). It is also possible to determine the value of S/D directly by calculating the projected size of a texture element in each region of the display screen. This latter method is practical, however, only if all of the texture elements have a regular shape such as a square or a circle.

contours because they connect sets of pixels whose projected surface areas are all identical. The innermost isoarea contour of each figure represents an area that is 1.25 times larger than the smallest value of D. Each successive contour moving outward represents an additional multiplicative increase of 1.25. The isoarea contours are bunched closely together near the boundaries of the figures because the value of D is changing very rapidly in those regions. The directions of maximum change in D, called *area gradients,* are represented in the figures by dashed lines, which are everywhere locally orthogonal to the isoarea contours.

An area field diagram, such as the ones depicted in Figure 4, can provide considerable insight into the patterns of optical texture within a visual display. Along each isoarea contour, for example, the optical texture elements are evenly distributed, and there are no systematic variations in their sizes and shapes. (Random variations may still be observed, however, as is evident in Figure 3.) The foreshortening of an optical texture element is always aligned with the area gradient in the direction that D is changing most rapidly. Moreover, as we move along these gradient contours, the optical texture elements vary continuously in size, shape, and density. This global pattern of optical structure may not be noticeable in an actual textured image due to the random variations that are inherent in our generation procedure. When an image is represented with an area field diagram, in contrast, its global structure becomes immediately apparent.

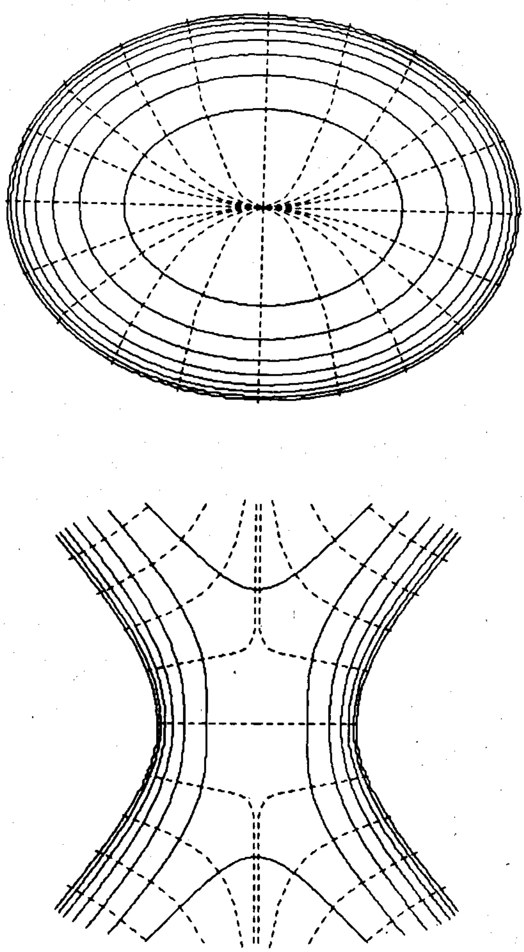

Figure 4. The projected area fields for the two images depicted in Figure 3.

Summary

It has been over 30 years since Gibson (1950) first suggested that patterns of optical texture provide useful information for the perception of three-dimensional form, yet we still know very little about how this information is utilized by human observers (cf. Stevens, 1981). By developing techniques for generating texture patterns on arbitrary curved surfaces and for representing the global organization of a textured image based on the formal concept of an area field, we have attempted to eliminate some of the stumbling blocks that have restricted previous investigations. We believe that these new tools provide a solid conceptual foundation for the mathematical and psychophysical analysis of optical texture, and we hope that they will stimulate future research.

References

Blinn, J. F. (1978). Simulation of wrinkled surfaces. *Computer Graphics, 12,* 286–292.

Blinn, J. F., & Newell, M. E. (1976). Texture and reflection in computer generated images. *Communications of the ACM, 19,* 542–547.

Braunstein, M. L. (1976). *Depth perception through motion.* New York: Academic Press.

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

Feibush, E. A., Levoy, M., & Cook, R. L. (1980). Synthetic texturing using digital filters. *Computer Graphics, 14,* 294–301.

Gibson, J. J. (1950). *The perception of the visual world.* Boston: Houghton Mifflin.

Haruyama, S., & Barsky, B. A. (1984). Using stochastic modeling for texture generation. *IEEE Computer Graphics and Applications, 4,* 7–19.

Mingolla, E., & Todd, J. T. (1984). Computational techniques for the graphic simulation of quadric surfaces. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 740–745.

Norton, A., Rockwood, A. P., & Skolmoski, P. T. (1982). Clamping: A method of antialiasing textured surfaces by bandwidth limiting in object space. *Computer Graphics, 16,* 1–8.

Schweitzer, D. (1983). Artificial texturing: An aid to surface visualization. *Computer Graphics, 17,* 23–29.

Stevens, K. A. (1981). The information content of texture gradients. *Biological Cybernetics, 42,* 95–105.

Todd, J. T. (1984). Formal theories of visual information. In W. H. Warren & R. E. Shaw (Eds.), *Persistence and change: Proceedings from the first international conference on event perception* (pp. 87–102). Hillsdale, NJ: Erlbaum.

Todd, J. T., & Mingolla, E. (1983). Perception of surface curvature and direction of illumination from patterns of shading. *Journal of Experimental Psychology: Human Perception and Performance, 9,* 583–595.

Witkin, A. P. (1981). Recovering surface shape and orientation from texture. *Artificial Intelligence, 17,* 17–45.

Appendix

The analysis presented below provides a specific method for computing the projected area of a picture element. The analysis begins with two points: the point of observation O, and a point P' on the picture plane. From the coordinates of these points we determine the point P on the simulated surface that is colinear with O and P' (see Mingolla & Todd, 1984, for details). We then calculate the distance z between O and P and its angle of intersection γ with the simulated surface. We also calculate the distance z' between O and P' and its angle of intersection γ' with the picture plane.

Following these preliminary calculations, we construct an arbitrarily small circular cone that has its vertex at O and is centered around the line segment OP at a visual angle α. We then calculate the area of intersection A_P between the cone and the simulated surface and the area of intersection $A_{P'}$ between the cone and the picture plane. The projected area D of a picture element at P' is defined as the ratio $A_P/A_{P'}$.

The value of D is most easily computed when the simulated surface is smooth. In that case it is possible to make the angle α small enough so that the region of intersection between the cone and the simulated surface is approximately planar. When a cone intersects a planar surface in this context, their region of intersection forms the shape of an ellipse. Let M be the center of this ellipse and let a and b be the minor and major axes, respectively. These are depicted in Figure A1 together with three other variables (c, d, and e) employed in the computation. Using the law of sines, we can show from the figures that

$$e = z \sin \alpha / \sin(\alpha + \gamma)$$

$$b + d = z \sin \alpha / \sin(\alpha + \pi - \gamma)$$

$$d = (b + d - e)/2$$

$$c = z \tan \alpha.$$

By placing the origin at M, we can define the region of intersection by the familiar equation for an ellipse (see Figure A1).

$$x^2/a^2 + y^2/b^2 = 1.$$

Because b is known, we can solve for a by letting $x = c$ and $y = d$ to yield

$$a = c\sqrt{1 - (d/b)^2}$$

The area of the ellipse is then given by

$$A_P = \pi ab.$$

To determine the projected area D of a picture element, we must perform these calculations twice: once for the point P on the simulated surface to obtain A_P and once for the point P' on the picture plane to obtain $A_{P'}$. By forming a ratio of these two areas, we obtain the value of D.

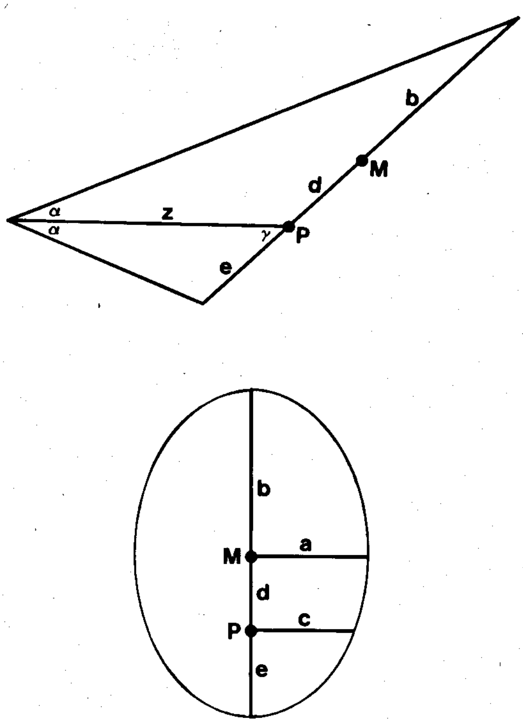

Figure A1. The intersection of a planar surface with a cone of visual solid angles as viewed parallel to the surface (top) and perpendicular to the surface (bottom).

Received February 6, 1984
Revision received May 26, 1984 ■