

Shape from stereo: A systematic approach using quadratic surfaces

SJOERD C. DE VRIES, ASTRID M. L. KAPPERS, and JAN J. KOENDERINK
Utrecht Biophysics Research Institute, Utrecht, The Netherlands

We used quadratic shapes in several psychophysical shape-from-stereo tasks. The shapes were elegantly represented in a 2-D parameter space by the scale-independent *shape index* and the scale-dependent *curvedness*. Using random-dot stereograms to depict the surfaces, we found that the shape of hyperbolic surfaces is slightly more difficult to recognize than the shape of elliptic surfaces. We found that curvedness (and indirectly, scale) has little or no influence on shape recognition.

Wheatstone (1838) was the first to observe that a dichoptic presentation of two projected images (differing slightly in the viewpoint) could induce a depth percept. This important finding reveals the ability of the human visual system to make depth-related judgments using disparity information. Generally, it is thought that stereopsis stems from the direct calculation of depth from disparities. This is a very logical assumption, because this is the inverse of the method that we normally use to generate stereograms. Recent research (Brookes & Stevens, 1989a, 1989b; Rogers & Cagenello, 1989; Stevens and Brookes, 1987, 1988) suggests that depth reconstruction of surfaces is done indirectly through surface shape descriptors such as curvatures and discontinuities of disparity fields. This indicates that shape recognition is not simply a matter of calculating some kind of depth map (Gibson, 1950). Thus, we might expect shape recognition to depend on the type of shape that is seen. Most shape-from-stereo research is done with a collection of rather arbitrary shapes (Uttal, 1987; Uttal, Davis, Welke, & Kakarala, 1988) or a very restricted set of shapes, for instance, cylinders (Johnston, 1991, Rogers & Cagenello, 1989). Here, we present a more systematic approach to the research of recognition of shape with stereo.

Shape Definitions

First of all, we need a more concise definition of "shape." Clearly, only a few objects can be identified by a descriptive name. To describe the infinitely large family of arbitrary shapes requires some kind of restriction and we will therefore consider local surface patches only.

Position and attitude of a local surface patch do not contribute to its shape, because they are dependent on the observer-object geometry only and are not an intrinsic

property of the patch. Therefore, we choose our frame of reference so as to get rid of these terms. We place the origin at the fixation point and align the z -axis parallel to the surface normal. With these coordinates, we can describe a general patch as a Taylor series expansion:

$$z = \frac{1}{2} (ax^2 + by^2) + \text{higher order terms.} \quad (1)$$

An infinite number of normal sections (the intersection of a plane containing the surface normal and the surface itself) can be drawn through the origin on the patch. Together, they totally define the patch. The curvature in a point of such a normal section is given by:

$$\kappa = \frac{z_{xx}}{(1 + z_x^2)^{3/2}}, \quad (2)$$

where the x -axis is chosen in the plane of the normal section and tangent to the surface (z_x denotes derivation of z with respect to x). Euler (see, e.g., Aleksandrov, Kolmogorov, & Lavrent'ev, 1963) showed that the curvature of all normal sections is constrained by the so-called *principal curvatures* of the surface, which are the maximum and minimum curvatures of the patch. He showed that the normal sections that have these principal curvatures (K_{\max} and K_{\min}) lie in planes that are perpendicular to each other. The principal curvatures relate to all other curvatures of normal sections in the following way:

$$\kappa(\phi) = K_{\max} \cos^2 \phi + K_{\min} \sin^2 \phi \quad (3)$$

(where ϕ is the angle between the planes containing the normal sections with curvature K_{\max} and κ). Simple combinations of the principal curvatures yield $K = K_{\max} K_{\min}$ and $H = \frac{1}{2}(K_{\max} + K_{\min})$, called the Gaussian (or total) and mean curvature, respectively. These quantities are very important descriptors of surface geometry (Aleksandrov, Kolmogorov, & Lavrent'ev, 1963). For instance, all hyperbolic surfaces have negative K , and all elliptic surfaces have positive K .

However important their geometrical properties may be, K and H are less suitable for describing shape in tangible terms. Such a shape measure should comply with

This research was supported by the SPIN project "3D Computer Vision" of the Dutch Ministry of Economic Affairs. The authors are grateful to Lon Boonen for his contributions to this research. Address correspondence to S. C. de Vries, Utrecht Biophysics Research Institute, Department of Medical and Physiological Physics, Princetonplein 5, 3584 CC Utrecht, The Netherlands.




Figure 1. The shape index scale. The shape index value of the scale corresponds to the shape index value at the center of the quadratic surfaces shown under the scale, because the shape index is a local measure and can have different values at different points on the surface.

the intuitive idea of "shape." The key observation is that spheres of different radii look the same. Although these spheres have unequal curvatures and hence unequal mean and Gaussian curvature, they "look alike." Obviously, then, perception of shape should be characterized by some kind of scale-independent quantity. Koenderink (1990) proposes such a quantity, the shape index S . It is defined by:

$$S = -\frac{2}{\pi} \arctan \frac{K_{\max} + K_{\min}}{K_{\max} - K_{\min}}. \quad (4)$$

With the shape index, shapes can be placed on a continuous one-dimensional scale. This scale can be divided into a number of meaningful sections (see Figure 1). From -1 to $-\frac{1}{2}$ the range contains concave ellipsoidal surfaces, from $-\frac{1}{2}$ to $+\frac{1}{2}$ we find the hyperbolic shapes (saddles), and from $+\frac{1}{2}$ to $+1$ the shapes are convex ellipsoidal. The boundaries at $\pm\frac{1}{2}$ are cylinders. Two objects having shape indexes of opposite sign relate much like a stamp and mold.

The amount of curvature of the patch is packed in a quantity called the curv edness, C , defined as:

$$C = \sqrt{\frac{K_{\max}^2 + K_{\min}^2}{2}}. \quad (5)$$

These definitions may seem to be rather baroque, but they are no more than polar coordinates in K_{\max} , K_{\min} space in disguise. The shape index corresponds to a direction in this space. The 0° axis is rotated over 45° in order to make the scale symmetric with respect to convexity and concavity; this rotation is apparent in the argument of the arctan. The curv edness is a measure for distance from the origin. Therefore, the distribution of shapes along the shape index scale is uniform (the arctan in Equation 4 might suggest otherwise). From the definitions of S and C , it easily follows that S is scale independent and C is scale dependent.

Aim of Study

We have now obtained the tools we need for systematic shape psychophysics. We have restricted ourselves to second-order form description, but this restriction still permits a broad area of research. In the research presented here, we have centered our attention on shape recognition tasks and are interested in the relation of curv edness and shape recognition. In the definition of shape, S can

be chosen independently of C , but is the recognition of shape independent of curv edness? In other words, is shape recognition scale independent? And, for future discrimination experiments, we would like to know whether we can use distance along the S scale as a measure of the discriminability of two shapes.

We point out here that although the present research may look superficially like that of Uttal (1987) and Uttal et al. (1988), by no means are the studies similar. Although all are concerned with shape recognition from stereo-defined objects (some of them identical), there are very important differences. For instance, Uttal et al. used a set of discrete shapes, whereas we use a continuous family of shapes. When we calculate the shape index of the surfaces in the origin, we see that in their shapes three values are present: $S = 1$, $S = \frac{1}{2}$, and $S = 0$ (the plane has an indeterminate shape index value). Also, setup, tasks, parameters, and goals are quite distinct.

GENERAL METHOD

Apparatus

To depict the stereo images, we used an ATARI 1040 ST computer connected to a pair of LCD shutter spectacles. The frame rate of the monitor (P4 white phosphor) was 70 Hz. Screen dimensions were 20.75×13.0 cm ($18.7^\circ \times 12.2^\circ$), with 640×400 pixels. The LCD spectacles were toggled in synchronization with the vertical retrace of the monitor in such a way that after every even vertical retrace, the monitor showed the left-eye image and after every odd retrace, the right-eye image. We used a chinrest to immobilize the head. Viewing distance was 60 cm. The experiments were done in a dark room; average luminance of the screen was 45 cd/m².

Stimuli

The images we presented were random-dot stereograms (Julesz, 1960, 1971) so most monocular cues were avoided. The small density cue present was diminished by selecting half of the number of points from a uniform random distribution in the right-eye image, projecting them on the surface of the object to be depicted, and then projecting them back to the left-eye image and vice versa. In this way, the density cue became virtually invisible and tests showed that, monocularly, no shape information whatsoever could be extracted. Mean pixel density was 15%.

Depicted in the random-dot stereograms were quadratic shapes given by:

$$z = \frac{1}{2} (K_1 x^2 + K_2 y^2). \quad (6)$$

This form was chosen so that the principal curvatures are given

by K_1 and K_2 , as can be checked easily with Equations 2 and 3. K_1 and K_2 can be calculated from Equations 4 and 5. In this way, we are able to generate shapes with a specific shape index and curvedness. We stress here that shape index and curvedness are purely local measures. Every point on a surface has its own S and C value. Normally, the shape index and curvedness change all over the surface. An example of an exception is the sphere, which has the same S and C value everywhere. A paraboloid of rotation with $C = 0.5 \text{ cm}^{-1}$ has in its top at the origin a shape index value of 1, while at a distance of about 3 cm from this point, the shape index value drops to 0.65, getting closer to a cylindrical shape. (When referring to the shape index of an object, like the one defined in Equation 6, we always mean the value at the origin.) From the generating equation, it follows that the tangent plane to the surface in this point is frontoparallel.

The quadratic shapes subtended a circular domain of 3.5 cm (3.3°) radius. The simulated quadratic surfaces were translated 7.5 cm in front of the screen to ensure that no part of the shape would ever lie behind the screen, as this could yield an unwanted cue. Starting at 1.5 cm (1.4°) from the center, the shape was blurred by background pixels. The ratio of background to object pixels increased as a Gaussian function of distance from the unblurred region in such a way that, at the boundary, about 90% of the pixels were background. Again, this was done to remove an unwanted shape cue, in this case, the telltale boundary contour.

Due to the discrete character of the screen, disparities on the screen could only change stepwise, as is shown in Figure 2a. In order to smooth the surface, the disparity of a point on the screen could be increased by 1 pixel, with a probability proportional to the rounding error that was made by truncating the calculated disparity of that point. In this way, the average height of a small neighborhood of a point was about the same as it would be in the continuous case. There is some evidence that this averaging does indeed take place (Parker & Yang, 1989; Westheimer & Levi, 1987). A disadvantage of this method is that the depth of the surface appears slightly blurred. A typical result is shown in Figure 2b.

Subjects

Four subjects (L.B., male, age 24; S.D.V., male, age 27; C.D.J., female, age 28; and A.K., female, age 32) took part in some or all of the experiments. The subjects had normal or corrected-to-normal vision. They were checked for normal stereo vision by means of partially decorrelated random-dot test images (Julesz, 1971). Furthermore, we did some preliminary testing of the response time in the actual setup and we made sure that all the subjects were able to perceive the quadratic shapes within 3 sec (after a few trials). The stereo images had to be calculated for two interocular distances only (6.5 cm for Subjects L.B. and S.D.V. and 5.8 cm for Subjects A.K. and C.D.J.). The subjects were knowledgeable about the purposes of the experiments.

EXPERIMENT 1

Categorization Task With Shapes of Constant Curvedness

Procedure

For this experiment, we generated 501 stereo images distributed evenly along the shape index scale. The distance between two succeeding images was 0.004 (on the S scale). The images were selected randomly from this pool and the subjects assigned them to one of eight possible categories—Category I: $[-1.0, -0.75]$, II: $(-0.75, -0.5]$, III: $(-0.5, -0.25]$, IV: $(-0.25, 0]$, V: $(0, 0.25]$, VI: $(0.25, 0.5]$, VII: $(0.5, 0.75]$, and VIII: $(0.75, 1.0]$ —where parentheses indicate that the boundary is excluded, and brackets indicate that it is included. Before the experiments, we showed the subjects wire-frame pictures of quadratic surfaces of several shape index values to familiarize them with the scale. Furthermore, we

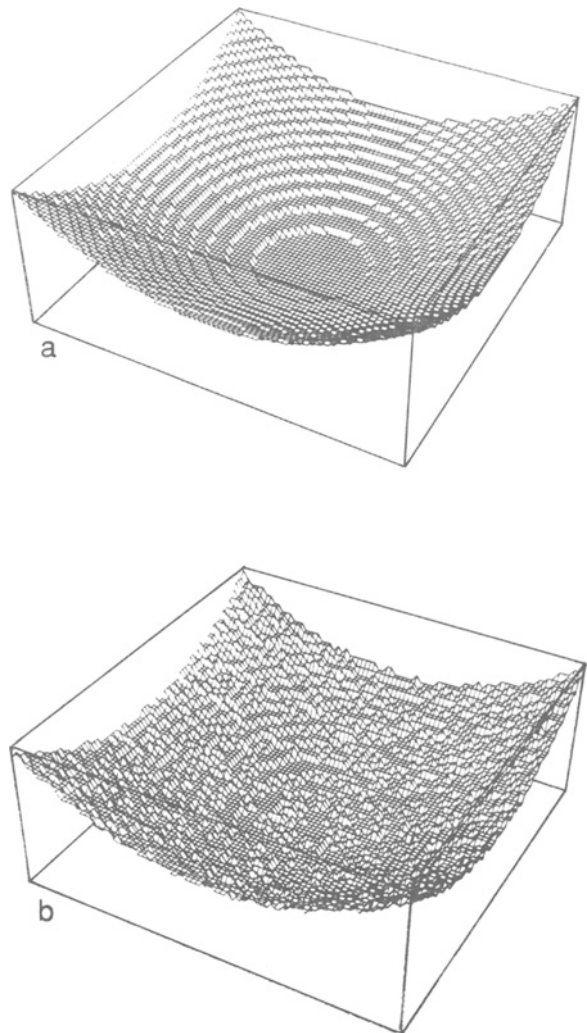


Figure 2. Disparities present on the screen in the case of a shape with $S = -1$ and $C = 0.5$. The undistorted figure is shown in (a). Each step corresponds to a 1-pixel disparity. Size and shape of the equidisparsity areas depend on S and C . The smoothed figure is shown in (b). An error of 1 pixel is added with a probability proportional to the distance to the next lower step. This not only smooths the shape, but also conceals the steps, which could be used as a shape cue otherwise.

presented the actual stereo images in a few sessions, with feedback to accustom them to the scale.

Stimulus display time was 3 sec. In some of the later measurements, the subjects had the opportunity to interrupt the presentation of the stimulus, at their request. The subjects typically responded within 2 sec; longer display times did not improve their performance, provided that they interrupted the stimulus only when they were sure that their response would not change anymore. The response times were recorded, although it was not our primary goal to measure them; they served as additional information.

Each of the 501 shapes was rotated around the z -axis over a random angle, so that the direction of the principal curvatures did not yield a cue in the categorization task (convex cylinders generated with Equations 4, 5, and 6 have an orientation that is perpendicular to concave cylinders). This also provided a means to get rid of a directional anisotropy (Mustillo, Francis, Oross, Fox, & Or-




Figure 3. Cross section through a family of symmetric ellipsoids ($S = -1$). The curvednesses shown here cover the range used in Experiments 1-4. The values are (in upward direction): 0.3, 0.5, 0.75, 1.0, 1.25, and 1.5 cm^{-1} .

ban, 1988; Rogers & Cagenello, 1989; Rogers & Graham, 1983). Because of the large number of available shapes, each category had a wide range of different orientations and any existing orientational effect would be averaged out.

In each session, the subjects categorized 200 shapes; this took about 30 min. Each complete experiment consisted of three such sessions. The experiments were done with four values of the curvedness: 0.3, 0.5, 0.8, and 1.25 cm^{-1} (see Figure 3).

Results

The results for Subjects L.B. and S.D.V. (with $C = 0.5 \text{ cm}^{-1}$) are shown in Figures 4a and 4b, respectively. The horizontal scale on the right marks the stimulus categories, the horizontal scale on the left marks the response categories, and the vertical data show the responses as a percentage of the total number of stimuli in a particular category. The results of the other subjects and with other curvednesses were qualitatively similar. Clearly, answers are clustered around the stimulus category. The distribution is unimodal. Percentages of correct answers or scores are, on average, lower for the hyperbolic surfaces than for convex or concave. From the response time data of Figure 5, it follows that hyperbolic shapes take longer to categorize. Small peaks are visible at the borders between Categories 1 and 2 and between 7 and 8, indicating that responding to stimuli from the vicinity of these borders also takes more time. A closer look at the time data reveals that it takes longer (500-1000 msec) to give an incorrect answer than a correct one.

In Figure 6a, scores averaged over all shape categories are plotted against curvedness. In most of the cases in this graph, scores are higher when curvedness is higher. If we consider the shape categories individually, this trend is not so clear. For certain shape categories, curvedness $C = 0.3 \text{ cm}^{-1}$ yields the highest score (see, e.g., Figure 6b). This behavior was strongly subject dependent. An analysis of variance (ANOVA) shows a significant overall trend in the averaged data for all the subjects, except for Subject A.K. ($C.D.J.$, $p = .0003$; L.B., $p = .0003$; S.D.V., $p = .0013$; and A.K., $p = .19$). Tests on the data ignoring $C = 0.3 \text{ cm}^{-1}$ show no significant differences for Subjects L.B. ($p = .92$), S.D.V. ($p = .05$), and A.K. ($p = .23$). Subject C.D.J. still has significant differences in her scores ($p = .023$).

Discussion

We can make a few observations on the basis of the data. First, the subjects were able to use the shape index reasonably well. The results indicate that distance on this scale can be used as a discrimination measure for shape, because confusion of shapes decreases monotonically with distance on the scale.

Second, hyperbolic surfaces are somewhat harder to recognize than elliptic ones, although this effect was not equally strong among the subjects. The dip in the score in the hyperbolic region is not symmetric. Category 4 scores are higher than Category 5 scores except for 1 subject (L.B.). This suggests that the difference between Categories 4 and 5 especially is hard to perceive and that the subjects' responses are biased in the relatively large region around the border where the difference between Categories 4 and 5 is not evident. L.B.'s bias is probably the reverse of that of the others.

Third, there seems to be an influence of curvedness on shape recognition, but this is mostly due to the lower scores at $C = 0.3 \text{ cm}^{-1}$. A lower curvedness implies a smaller height range, so it could be assumed that the lower scores might be due to simple height detection thresholds. This, however, is not the case.

Among his subjects, Howard (1919) found the best results of stereoacuity to be about $2''$ of arc and the worst




Figure 4. Results of a categorization experiment with constant curvedness, in this case $C = 0.5 \text{ cm}^{-1}$. Upper panel (a) shows results for Subject L.B., and lower panel (b), for Subject S.D.V. Right scale is category number of the presented stimuli, left scale is response category, and vertical data show scores in percent.

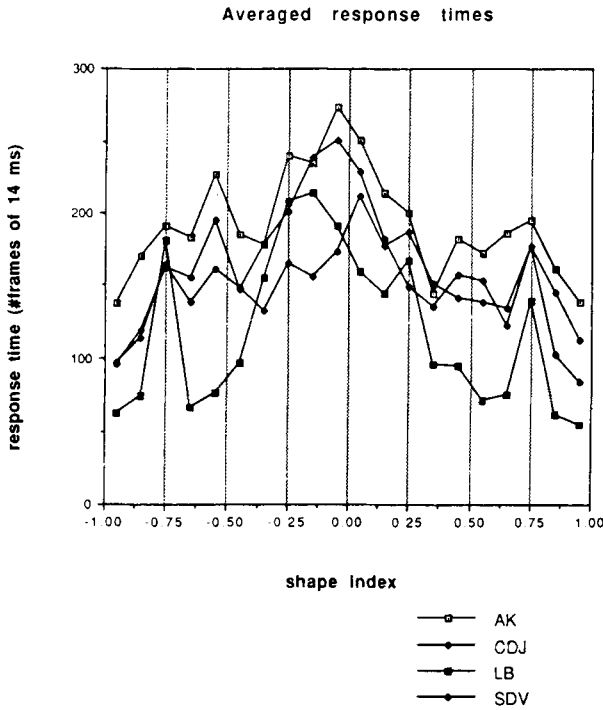


Figure 5. Response times averaged over all the experiments, when measured. Each point is the average over a range of 0.1 on the shape index scale.

to be 10"–136" of arc. Our stepsize is about 140" of arc, so this is certainly suprathreshold. The number of height levels does not constrain the performance either. To define a curvature, we need at least three points (see Equation 2). We chose the value of 0.3 cm^{-1} as a minimum curvedness value, because in its unblurred region of 1.5 cm it had just enough steps to define a curvature. With blurred and unblurred regions combined, quadratic shapes with this curvedness had a minimum of six to eight steps (depending on the shape index) in each direction, which should be enough for an ideal observer.

A third-order two-dimensional polynomial fit to a set of points like that of Figure 2b generated a surface whose principal curvatures, calculated by Equation 2, deviated from the real values by only 1%. Even a very sparse sampling of 5×5 points was enough for a correct calculation. So it seems that neither the stimulus nor early stereo processing is responsible for the lower scores at $C = 0.3 \text{ cm}^{-1}$. Hence, we could assume that it is subsequent 3-D shape processing that determines the results.

EXPERIMENT 2 Categorization Task With Shapes of Random Curvedness

Procedure

To investigate further whether or not curvedness influences shape recognition, we used a slightly modified version of Experiment 1. In this experiment, the shapes had random curvedness values ranging from 0.3 to 0.8 cm^{-1} . This forced the subjects to use relative curvature measures instead of absolute curvature measures, which they

might have done previously. The ratio of K_{\max} and K_{\min} is dependent on S only, and has a unique value for each S . Hence, in principle, it should be possible to determine the shape index independent of curvedness. The program producing the above-mentioned fit to the data points did not require knowledge about the curvedness of the object in order to reconstruct its shape index.

Results

The response histograms were qualitatively the same as those of Figure 4. The curvedness range was split into

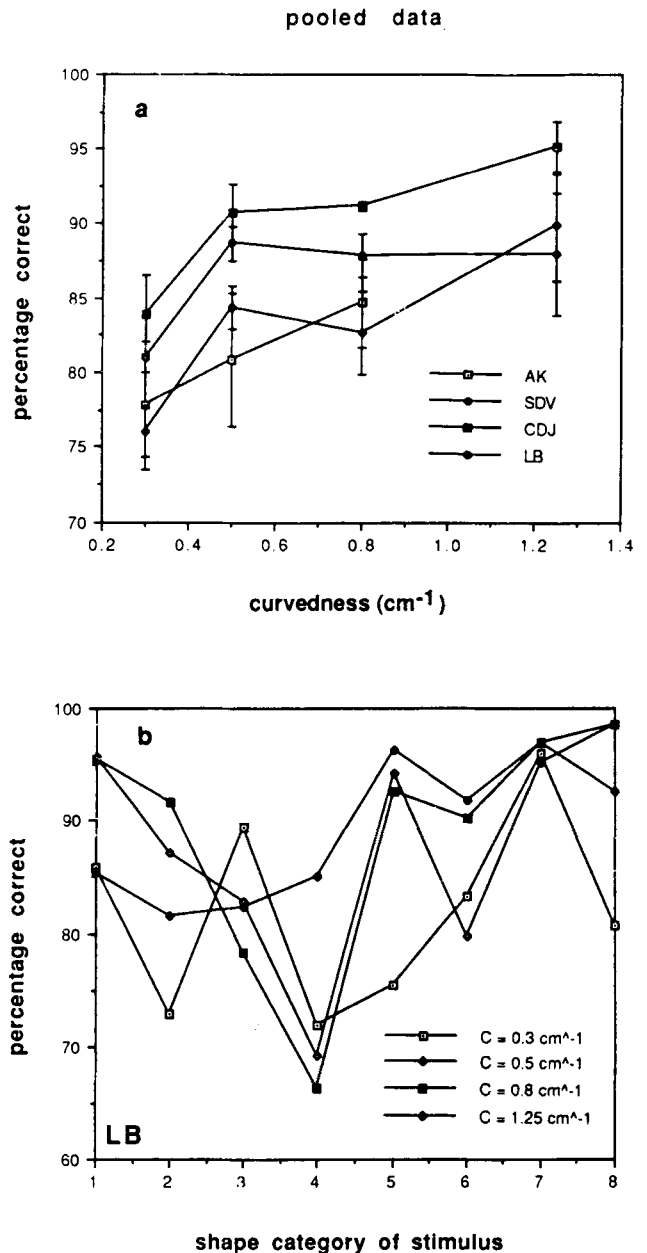


Figure 6. Results of the categorization experiments with constant curvedness. Averaged scores for 4 subjects as a function of stimulus curvedness are shown in (a). Scores are the percentage of correct answers (usually the values of the central peaks, like those of Figure 4). Scores of 1 subject (L.B.) as a function of stimulus shape category parametrized by curvedness are shown in (b).




Figure 7. Response scores as a function of curviness, in the random curviness experiment (see text for details).

15 subranges. For each subrange, the percentage of correct answers was calculated; Figure 7 shows the results. In this figure, a trend with curviness is not visible and a calculation of product-moment correlations shows that all correlation coefficients are insignificant (S.D.V., $r = .482$, $p > .05$; C.D.J., $r = .332$, $p > .10$; L.B., $r = .220$, $p > .10$; A.K., $r = .294$, $p > .10$). Mean scores of Experiments 1 and 2 were S.D.V., 81.0% and 78.2%; C.D.J., 88.6% and 89.9%; L.B., 86.1% and 79.2%; AK, 81.1% and 81.7%, respectively.

Discussion

The results do not reveal an effect of curviness on shape recognition. However, correlations are useful for testing whether a linear model underlies the data, which, in this case (finite y range), clearly cannot be the case.

The chaotic behavior of the scores versus curviness (Figure 7) might indicate that the slight trend we noticed in Experiment 1 is just a coincidence. So, in the region of $0.3 \leq C \leq 1.25 \text{ cm}^{-1}$, shape recognition is not strongly dependent on the curviness.

A comparison of the mean scores for the constant curviness experiment with those for the random curviness experiment shows that scores are lower for 2 subjects (2.8% and 6.9%, respectively) and higher for the 2 others (by only 1.3% and 0.6%). This could mean that Subject L.B. was indeed using absolute curvature measures.

EXPERIMENT 3 Shape Index Estimation

Procedure

This experiment was the same as Experiment 1, except that the subjects estimated the shape index value instead of assigning it to

a category. They could respond as precisely as they liked, entering their response at the numerical keypad of the ATARI. In practice, this meant entering one or two decimals, though occasionally three decimals were used. Two subjects (S.D.V. and L.B.) took part in this experiment. Measurements were done at $C = 0.5, 0.8$ (S.D.V. only), and 1.0 cm^{-1} .

Results

Figure 8 shows the results for Subjects L.B. (top) and S.D.V. (bottom). Clearly, Subject L.B. deviates less from the perfect score (straight line) than does Subject S.D.V. In Table 1, mean absolute deviations from correct response are shown. An ANOVA reveals that means do not differ significantly for Subject L.B. ($p = .3$ and $p = .09$ for $C = 0.5$ and 1.0 cm^{-1} , respectively). For Subject S.D.V., the difference in means is significant ($p = .0006$, $p = 0$, and $p = 0$ for $C = 0.5, 0.8$, and 1.0 cm^{-1} , respectively). This is mainly due to the positive hyperbolic surfaces. Without these, means do not differ significantly in two of the three cases ($p = .01$, $p = .32$, and $p = .76$, respectively).

Discussion

The results for Subject L.B. reveal no differences in estimation errors between surface types. This is rather different from what we might expect after the results of the first experiment with this subject, when there was always at least an 8% score difference between the highest and the lowest scoring surface type. This might be due to a learning effect. The results of Subject S.D.V. are more like those in Experiment 1. The high deviation at the positive hyperbolic surfaces is reflected in the dip in the scores for this area.

A striking feature of the graphs of Figure 8 is their stair-like character. It appears that there are gauge shapes toward which perceived shapes are shifted. The assumed existence of these gauge shapes suggests that estimation of shape is done by means of some "natural categories," like the cylinder.

EXPERIMENT 4 Matching of Shapes

Procedure

In the previous experiments, the subjects translated their notion of shape into some abstract number. In order to circumvent this, a new setup was devised.

Instead of responding with a category number or a shape index value, the subjects adjusted a test shape in such a way that it most resembled a reference shape. They were shown a reference and a test object (in this order), each for 3 seconds. Again, the display time could be shortened by the subjects if desired. Adjustments to the test shape could then be made by increasing or decreasing the shape index of the test shape by three different step sizes (0.2, 0.02, and 0.004). After each alteration, the test object was shown again. The reference image could be presented again by pressing a button. The subjects were allowed to see the reference and test shapes as often as they liked and could make as many adjustments as they considered necessary. The value of the shape index of the test shape was not known to them. The curviness of the reference shape was not necessarily the same as that of the test shape. In this way, independence of curviness and shape recognition could be tested.

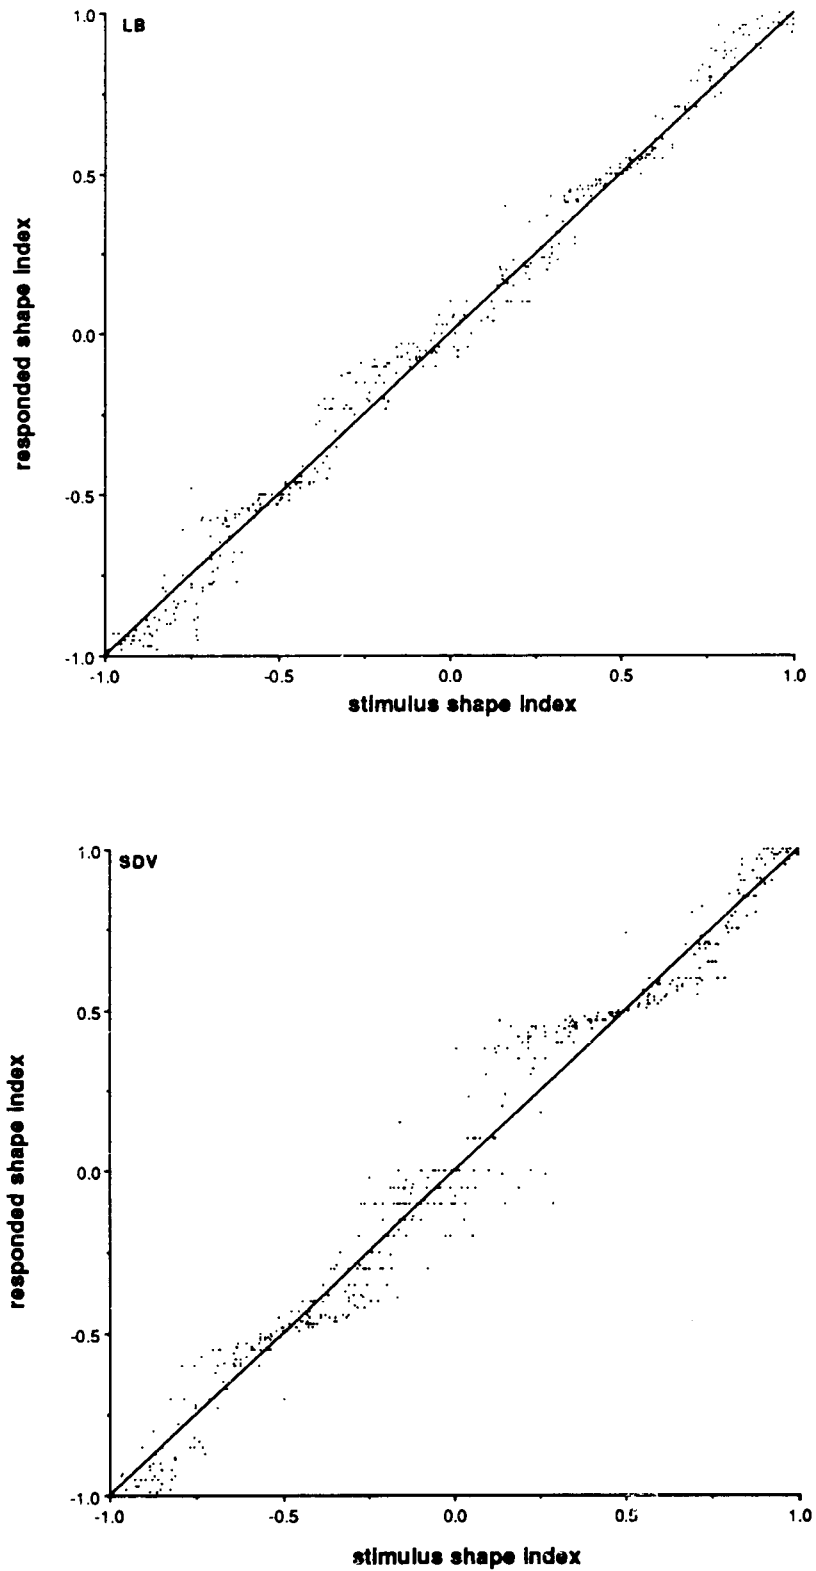


Figure 8. Responded shape index values against stimulus shape index for Subjects L.B. (top) and S.D.V. (bottom). Due to the number of stimuli (S.D.V. has 550, L.B., 400), all curvedness values are represented here with the smallest possible plot symbol. Nevertheless, in the regions around the straight line (perfect score), overlap is considerable.

Table 1
Absolute Deviations (*D*) From the Perfect Response
and the Standard Deviation (*SD*) Therein

Surface Type	Subject S.D.V.						Subject L.B.			
	<i>C</i> = 0.5 cm ⁻¹		<i>C</i> = 0.8 cm ⁻¹		<i>C</i> = 1.0 cm ⁻¹		<i>C</i> = 0.5 cm ⁻¹		<i>C</i> = 1.0 cm ⁻¹	
	<i>D</i>	<i>SD</i>	<i>D</i>	<i>SD</i>	<i>D</i>	<i>SD</i>	<i>D</i>	<i>SD</i>	<i>D</i>	<i>SD</i>
Negative Elliptic	0.06	0.04	0.08	0.05	0.05	0.05	0.05	0.04	0.06	0.06
Negative Hyperbolic	0.08	0.06	0.07	0.05	0.06	0.06	0.05	0.03	0.07	0.05
Positive Hyperbolic	0.12	0.08	0.15	0.12	0.11	0.07	0.05	0.04	0.05	0.03
Positive Elliptic	0.05	0.03	0.06	0.04	0.05	0.04	0.04	0.04	0.05	0.04

Note—Surface type *Negative Elliptic* corresponds to $S = [-1, -\frac{1}{2}]$, that is, the concave surfaces; *Negative Hyperbolic* to $S = (-\frac{1}{2}, 0]$; *Positive Hyperbolic* to $S = (0, \frac{1}{2}]$, and *Positive Elliptic* to $S = (\frac{1}{2}, 1]$, that is, the convex surfaces.

Reference shape index values were $-0.9, -0.7, -0.5, -0.3, -0.1, 0.1, 0.3, 0.5, 0.7, \text{ and } 0.9$. The initial value of the test shape was randomized, but was within a distance of 0.5 of the reference shape to prevent superfluous buttonpresses. Because we had a fixed set of reference shapes, it was important that the subjects not know the exact shape index value of the test object that they were manipulating, otherwise they could have set it to any exact value they liked just by counting buttonpresses.

Of course, the end of the scale offered the opportunity to gauge the shape index value of the test object. Pressing one of the "increase shape index" buttons when the object's shape was at the positive end of the scale would not result in a different shape and the subjects would then know the exact value of the test object. In order to prevent this, we allowed the variable representing the current test shape index value to exceed the boundary values, instead of setting it to ± 1 when the subjects reached the scale ends. So, all that the subjects would know at that moment (when the object was not changing anymore) is that they had reached the end of the scale; they had not set the variable to a specific value. The test shape started changing when the shape index value in the variable was adjusted to fall between -1 and $+1$ again. It was not known to the subjects precisely where on the scale this was, because it depended on the initial "overshoot," which was also unknown. So, gauging was only possible when the subjects chose to use the smallest adjustment steps around the bounds. Responses with shape index values out of bounds were registered as ± 1 .

We used two curvedness values for the reference shape, 0.5 and 1.25 cm⁻¹. For each curvedness setting of the reference shape, we used the following curvedness values for the test shapes: 0.3, 0.5, 0.75, 1.0, 1.25, or 1.5 cm⁻¹. In one session, the curvedness of the reference and test shapes remained constant. Each of the above-mentioned reference shapes was shown five times, in random order; therefore, the subjects had to adjust 50 test shapes for a total of 600 test shapes. Each adjustment took 1–2 min; hence, the entire experiment took about 15 h for each subject. A mean and a standard deviation were calculated from the five settings.

Results

A regression analysis performed on the final settings of the test shape index versus reference shape index reveals that there is an almost perfect linear relationship between these values. The constant term is low (it ranges from -0.04 to 0.01) and the linear coefficient is close to 1 (it ranges from 0.93 to 1.08). Coefficients of determination range from .977 to .995. Because of this, plots of the results tend to be very cluttered; so, to obtain better visualization of the data, we parametrize plots here

by reference shape index instead of test curvedness. Figure 9 shows the results for 1 subject (S.D.V.) and the two reference curvedness values used (0.5 and 1.25 cm⁻¹ for 9a and 9b, respectively). The vertical axis depicts response, and the horizontal axis denotes the curvedness value of the test shape. For example, when the reference shape has $S = 0.5$ and shape recognition is independent of curvedness, we expect a straight line at $S = 0.5$.

Clearly, the influence of curvedness is small again. Only 13 out of 40 responses correlate significantly with curvedness. The responses of the 2 subjects do not correlate (only 2 out of 20 correlations surpass the level of significance). The pattern of errors is different for the subjects. The correlations of error and curvedness are predominantly negative (31 out of 40), indicating that errors decrease with increasing curvedness. However, only 9 out of 40 are significant.

We calculated the mean value of the absolute deviation (response shape index minus reference shape index) for the four surface types (negative elliptic, etc.) for each combination of reference and test curvedness. We found almost no statistical difference in deviations between them. An ANOVA (with four categories and 10×5 responses for each combination) shows that the means of only 4 cases out of 24 (number of curvedness values of the test objects \times number of subjects \times number of curvedness values of the reference object) differ significantly (S.D.V., $C_{\text{ref}} = 0.5 \text{ cm}^{-1}$, $C_{\text{test}} = 1.5 \text{ cm}^{-1}$, $p = .0099$; S.D.V., $C_{\text{ref}} = 1.25 \text{ cm}^{-1}$, $C_{\text{test}} = 0.75 \text{ cm}^{-1}$, $p = .0084$; L.B., $C_{\text{ref}} = 0.5 \text{ cm}^{-1}$, $C_{\text{test}} = 0.3 \text{ cm}^{-1}$, $p = .0266$; L.B., $C_{\text{ref}} = 1.25 \text{ cm}^{-1}$, $C_{\text{test}} = 1.5 \text{ cm}^{-1}$, $p = .0099$). In half of these cases we find that elliptic surfaces score best, and in the other half hyperbolic surfaces perform best. But in only 5 out of 24 cases do we find the lowest deviation value in the hyperbolic region; the highest value is found there in 12 cases.

Discussion

The influence of curvedness on shape matching seems to be small. There is some evidence that accuracy improves with increasing curvedness, but the dichotomy resulting from Experiments 1 and 2 concerning the depen-

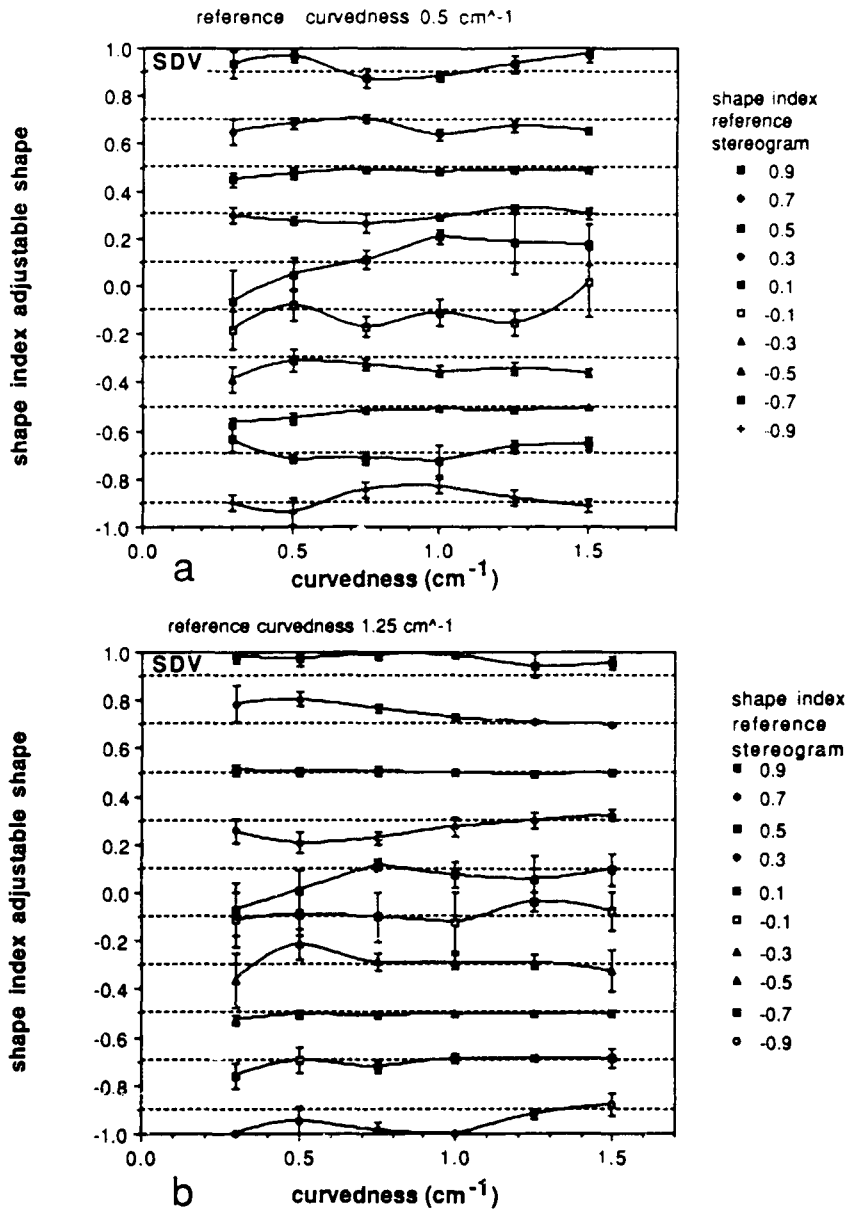


Figure 9. Value of the shape index of the adjustable test shape that was matched to a reference shape, as a function of the test shapes' curvedness, parametrized by the shape index of the reference shape. Reference $C = 0.5 \text{ cm}^{-1}$ (a), and Reference $C = 1.25 \text{ cm}^{-1}$ (b). Both panels show the results of Subject S.D.V.

dence of shape recognition on curvedness could not be resolved. Also, the question of the influence of surface type on shape recognition remains undecided.

Overall, performance was good and was somewhat better than in the case of the numerical responses, as could be expected (mean absolute deviations are 0.05 and 0.065, respectively). Although there appears to be an influence of surface type, this was shown to be not significant.

GENERAL DISCUSSION

Our results show that shape index and curvedness are convenient measures. Our subjects learned to use them quickly, so they seem to be valuable tools for shape recognition tasks. We should stress here that we used second-order shapes only. Local and global shapes are self-similar in these kinds of surfaces and the subjects might have as-

sessed the local shape measures by means of the global shape only. As a matter of fact, local shape does not exist in these kinds of discrete images (see General Method). Only when some sort of surface reconstruction, as proposed by Stevens and Brookes (1987), is done by the visual system will there be local shape, built up from interpolation or fitting. We intend to address the problem of local shape recognition by considering recognition of local shape in higher order (third-order or higher) surfaces.

Another problem connected with using second-order shapes is that they are symmetric. The symmetry axes lie in the same normal sections as the principal curvatures, which makes it easier to find them and to use them to calculate local shape measures. Again, the use of higher order surfaces will prevent this.

Other experiments that use the *S/C* paradigm on quadratic surfaces in combination with categorization tasks to investigate shape from shading (Erens, Kappers, & Koenderink, 1991) show that—in this very restricted case—the visual system is better at processing shape from stereo than shape from shading (using diffuse shading). Of course, one must be cautious about drawing such a conclusion from these rather dissimilar experiments.

The confusion of the parabolic arch with the paraboloid of rotation and the saddle, as reported in Uttal (1987) and Uttal et al. (1988), was not found. All the experiments show that confusion appears only with shapes close to each other on the shape index scale and diminishes with distance. This difference in results is probably due to a difference in dot density. Uttal et al. used 2–64 dots per stereogram, whereas we used approximately 4,000 dots.

REFERENCES

- ALEKSANDROV, A. D., KOLMOGOROV, A. D., & LAVRENT'EV, M. A. (Eds.). (1963). *Mathematics, its content, methods and meaning*. Providence, RI: American Mathematical Society.
- BROOKES, A., & STEVENS, R. A. (1989a). The analogy between stereo-depth and brightness. *Perception*, **18**, 601-614.
- BROOKES, A., & STEVENS, R. A. (1989b). Binocular depth from surfaces versus volumes. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 479-484.
- ERENS, R., KAPPERS, A. M. L., & KOENDERINK, J. J. (1991). Limits on the perception of local shape from shading. In P. J. Beek, R. J. Bootsma, & P. C. W. van Wieringen (Eds.), *Studies in perception and action* (pp. 72-75). Amsterdam: Rodopi.
- GIBSON, J. J. (1950). *The perception of the visual world*. Cambridge, MA: Riverside Press.
- HOWARD, H. J. (1919). A test for the judgment of distance. *American Journal of Ophthalmology*, **2**, 656-675.
- JOHNSTON, E. B. (1991). Systematic distortions of shape from stereopsis. *Vision Research*, **31**, 1351-1360.
- JULESZ, B. (1960). Binocular depth perception of computer-generated patterns. *Bell System Technical Journal*, **39**, 1125-1162.
- JULESZ, B. (1971). *Foundations of cyclopean perception*. London: University of Chicago Press.
- KOENDERINK, J. J. (1990). *Solid shape*. Cambridge, MA: MIT Press.
- MUSTILLO, P., FRANCIS, E., GROSS, S., FOX, R., & ORBAN, G. A. (1988). Anisotropies in global stereoscopic orientation discrimination. *Vision Research*, **28**, 1315-1321.
- PARKER, A. J., & YANG, Y. (1989). Spatial properties of disparity pooling in human stereo vision. *Vision Research*, **29**, 1525-1538.
- ROGERS, B. J., & CAGENELLO, R. (1989). Disparity curvature and the perception of three-dimensional surfaces. *Nature*, **339**, 135-137.
- ROGERS, B. J., & GRAHAM, M. E. (1983). Anisotropies in the perception of three-dimensional surfaces. *Science*, **221**, 1409-1411.
- STEVENS, K. A., & BROOKES, A. (1987, June). Depth reconstruction in stereopsis. *Proceedings of the first IEEE International Conference on Computer Vision* (pp. 682-686). Washington, DC: IEEE Computer Society.
- STEVENS, K. A., & BROOKES, A. (1988). Integrating stereopsis with monocular interpretations of planar surfaces. *Vision Research*, **28**, 371-386.
- UTTAL, W. R. (1987). *The perception of dotted forms*. Hillsdale, NJ: Erlbaum.
- UTTAL, W. R., DAVIS, S. N., WELKE, C., & KAKARALA, R. (1988). The reconstruction of static visual forms from sparse dotted samples. *Perception & Psychophysics*, **43**, 223-240.
- WESTHEIMER, G., & LEVI, D. M. (1987). Depth attraction and repulsion of disparate foveal stimuli. *Vision Research*, **27**, 1361-1368.
- WHEATSTONE, C. (1838). On some remarkable, and hitherto unobserved, phenomena of binocular vision. *Philosophical Transactions of the Royal Society of London*, **128**, 371-394.

(Manuscript received January 27, 1992;
revision accepted for publication June 29, 1992.)