Pergamon

PII: S0042-6989(96)00213-1

Vision Res., Vol. 37, No. 6, pp. 757–764, 1997

Printed in Great Britain
0042-6989/97 $17.00 + 0.00

The Distance Used for Scaling Disparities is the Same as the One Used for Scaling Retinal Size

WIM van DAMME,*† ELI BRENNER*

Received 8 June 1995; in revised form 10 December 1995; in final form 12 April 1996

To determine the physical size and global three-dimensional (3-D) shape of an object, retinal size and retinal disparity have to be scaled in accordance with the object's distance. We examined whether the distance used for scaling retinal disparity is the same as the distance used for scaling retinal size. Subjects adjusted the 3-D shape (size and depth) of a computer-simulated ellipsoid to match a tennis ball. Analysis of the errors when only the ellipsoid was visible in an otherwise completely dark room suggests that the distance used for scaling retinal disparity is indeed the same as that used for scaling retinal size. This was confirmed by showing that the correspondence between the distance used for scaling retinal disparity and that used for scaling retinal size does not improve when more information about distance is available (room lights on), although both distances are then much closer to the simulated distance. Finally, we show that this correspondence is not due to the use of distance-invariant higher order binocular information.

Size Disparity Shape Distance Stereopsis

INTRODUCTION

An important aspect of vision is to reconstruct three-dimensional (3-D) objects from the flat images on our retinas. If the object's distance is known, the retinal image size and the retinal disparities provide information about physical size and depth, which can be combined to give a 3-D shape. As long as the estimate of distance is correct, the perceived shape of the object should not vary systematically with distance (shape constancy). However, if a wrong estimate of distance is used, the perceived 3-D shape will not be veridical.

The most evident distortions of perceived 3-D shape were found in limited-cue environments (Johnston, 1991; Tittle *et al.*, 1995). Johnston (1991) suggested that retinal disparities were being scaled by a different distance than the actual distance, the latter being specified by the accommodation and ocular vergence required to fixate the target. Glennerster *et al.* (1993, 1994) proposed that failure of shape constancy in Johnston's (1991) experiment was at least partly due to the limited-cue environment. They showed that a rich environment (or a more naturalistic viewing condition) improved shape constancy considerably. Systematic distortions of perceived distance have, however, also been reported for full-cue conditions (Wagner, 1985), so shape may not only be misperceived in laboratory situations with limited cues.

In previous studies on 3-D shape, little attention has been paid to the object's size. Retinal size provides information about physical size if the distance of the object is known. Conversely, knowing the object's physical size can help one to judge its distance (Sedgwick, 1986).

The fact that retinal size and retinal disparity have to be scaled with distance before they can be interpreted as size and depth raises the question whether size-scaling and depth-scaling use a common estimate of distance (Colett *et al.*, 1991; Rogers & Bradshaw, 1995a). This is examined in the present study.

EXPERIMENT 1

Design and procedure

We used a matching task in which subjects were instructed to set the size and 3-D shape of a computer-simulated ellipsoid to match a familiar object: a tennis ball (radius = 33 mm). They were asked to make the appearance of the simulated ball correspond to that of a tennis ball they held in their hand. The real tennis ball served as a visible and haptic example before the experimental session. During the experimental session, the room was completely darkened and the tennis ball was held in one hand below the table surface, so it only served as a haptic reference.

Subjects could manipulate the size and 3-D shape of

*Vakgroep Fysiologie, Erasmus Universiteit, Postbus 1738, 3000 DR Rotterdam, The Netherlands.
†To whom all correspondence should be addressed [*Tel* +31 10-4087569; *Fax* +31 10-4367594; *Email* damme@fys1.fgg.eur.nl].

the simulated ellipsoid by moving the computer mouse: horizontal mouse movements simultaneously changed the width and height of the simulated ellipsoid (which we will refer to as its size) and vertical mouse movements changed the disparities of the texture elements on the simulated ellipsoid's surface (which we will refer to as its depth). When subjects were satisfied with their settings, they pressed the mouse button, whereupon the settings were stored and a new trial began. Subjects did not receive any feedback on their performance.

Each subject made 50 settings, with the simulated ellipsoid at random simulated distances between 40 and 80 cm (this small distance range was chosen to avoid strong conflicts with accommodation, which was always suitable for an object at 60 cm). At the beginning of each trial, the simulated ellipsoid had a random simulated size (frontoparallel radius between 1 and 70 mm, which results in an angular extent between 0.07 and 9.9 deg, depending on the simulated distance) and a random simulated depth (radius along the line of sight between 1 and 150 mm).

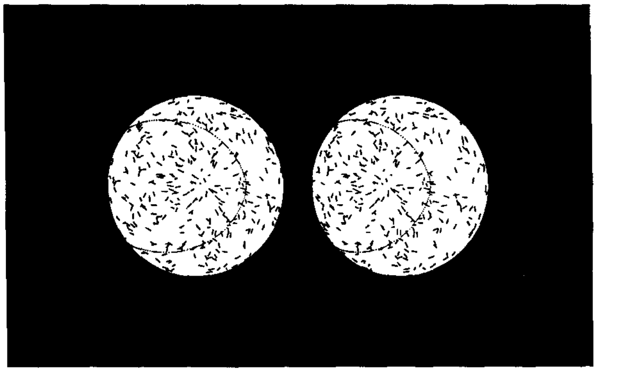

FIGURE 1. Illustration of the stimulus used in Experiments 1 and 2. The outlines of the real images on the screen were not circular, but were calculated properly, taking the solid shape of the simulated object and the viewing angle of the left and right eyes into account.

Stimuli and apparatus

The stimulus was a red computer-simulated ellipsoid, speckled with approximately 1000 small black random lines (Fig. 1). The lines were randomly distributed on the surface of the ellipsoid with random orientations. As the lines were distributed uniformly over the surface of the ellipsoid, the texture and density gradients changed according to the set depth of the ellipsoid. Removal of these monocular cues would create a conflict between depth cues, which we preferred to avoid. Beside the random lines, the characteristic tennis ball curve was drawn along the surface of the ellipsoid, with the "ball" in a random orientation.

The images were generated on a Silicon Graphics GTX-210 Computer and displayed on a HL69SG monitor. The size of the screen was 34.0 × 27.0 cm with 1280 × 492 pixels (width × height). Hardware anti-aliasing techniques increased the effective resolution (intermediate colours were computed for the eight neighbouring pixels of each pixel in a line).

The images for the left and right eyes were presented in perspective projection and were displayed in alternation at a rate of 120 Hz (thus, each pair of images was displayed at 60 Hz). The images were viewed through LCD shutter spectacles which were synchronised with the monitor to ensure that each eye received the appropriate images. Red stimuli were used because the LCD shutter spectacles work best at long wavelengths (about 33% transmission when "open" and 0.3% when "shut"). Subjects sat with their head in a chin-rest at 60 cm from the screen.

Subjects

Initially, nine subjects, all members of the department, took part in the experiments. They all had normal or corrected to normal (monocular and binocular) vision. Two of them are the authors (WD and EB). The others were naïve as to the purpose of the experiment.

The results of three of the subjects were not included in further analysis because they did not appear to be using the disparities to make their settings. These subjects often set the depth of the simulated ellipsoid to the smallest possible value in our computer simulation (1 mm), while reporting that the simulated ellipsoid was spherical. Although these subjects had normal binocular vision, they were possibly confused by other cues in the stimulus. None of the remaining six subjects ever set the depth to this smallest possible value. None of the subjects ever reported that they were unable to find a setting that corresponded to that of a tennis ball.

Analysis

Once the subjects had made their settings, we determined three measures of distance: *vergence-distance* d_v (the distance that corresponds to the angle between the two lines of sight. It can also be called the "actual" distance or "simulated" distance); *size-distance* d_s (the distance at which the set retinal size would be appropriate for a tennis ball); and *disparity-distance* d_d (the distance at which the set retinal disparity would be appropriate for a tennis ball). Figure 2 illustrates these distances.

We discuss the results in terms of these three distances d_v, d_s and d_d. For a tennis ball with radius R, we can derive expressions for d_s and d_d in terms of the set size R_{xy} the set depth R_z of the simulated ellipsoid and the simulated distance d_v. Note that R_{xy}, R_z and d_v are the simulated rather than the perceived size, depth and distance.

The size-distance d_s is the distance at which the tennis ball would have to be for its retinal image to match the size set by the observer:

$$d_s = \frac{R}{R_{xy}} d_v \tag{1}$$

Similarly, the disparity-distance d_d is the distance at which the set disparity α would be that of a tennis ball. In

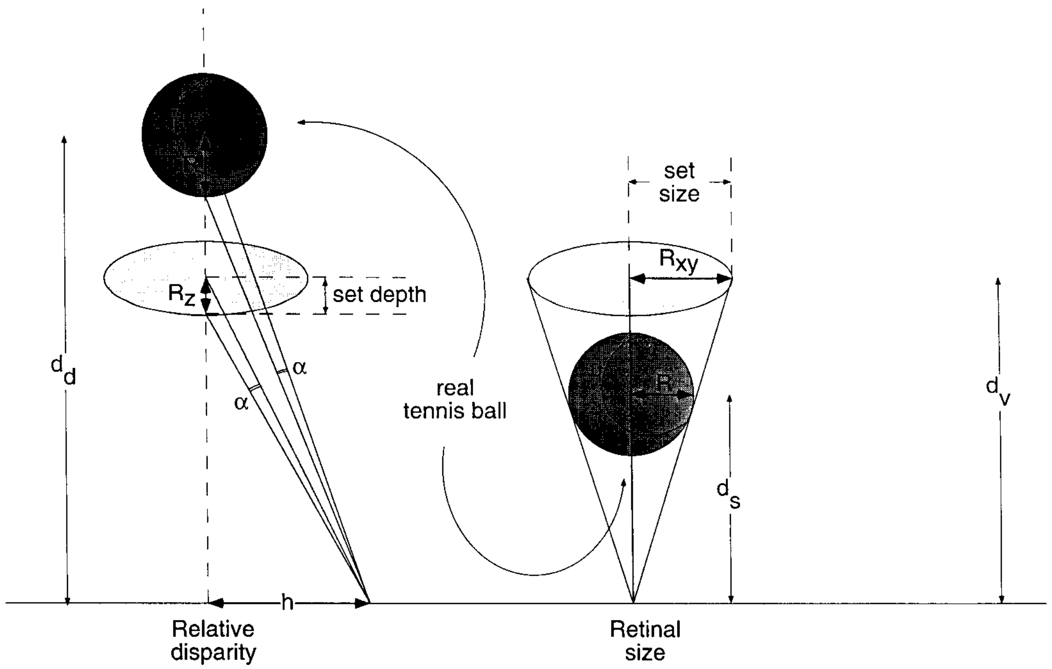

FIGURE 2. Size-distance d_s, disparity-distance d_d and simulated distance d_v and their relation to the set size R_{xy}, the set depth R_z, and the actual size of a tennis ball R. The disparity-distance d_d is the distance at which a real tennis ball would give the same disparity α as the ellipsoid set by the observer at distance d_v. The size-distance d_s is the distance at which a real tennis ball would give the same retinal image size as the ellipsoid set by the observer at distance d_v.

that case, we have:

$$\begin{aligned}\alpha &= \arctan\left(\frac{d_d}{h}\right) - \arctan\left(\frac{d_d - R}{h}\right) \\ &= \arctan\left(\frac{d_v}{h}\right) - \arctan\left(\frac{d_v - R_z}{h}\right)\end{aligned} \tag{2}$$

where $2h$ is the inter-ocular distance. From this expression d_d can be derived:

$$d_d = \frac{1}{2}R + \sqrt{\left(\frac{1}{4}R^2 - h^2 + \frac{R}{R_z}(d_v^2 - h^2) - d_v R\right)} \tag{3}$$

Results

Figure 3(a,b,c) shows the results for one subject (MZ) in terms of the three possible combinations of the three distances. The solid line with slope 1 represents perfect correspondence: if a data point is on this line, the two distances are the same. Figure 3(d,e,f) shows the results for a second subject (ST).

Figure 3(a) shows d_s as a function of d_v. It is clear that d_s and d_v do not correspond very well. Instead, d_s only increases slightly when d_v is increased. In other words, the size set by this subject does not reflect a correct scaling of retinal size with simulated distance.

Figure 3(b) shows d_d as a function of d_v. There is more scatter than in the previous figure, but the same systematic deviation from perfect correspondence can be seen. Evidently, disparity is not scaled by a veridical measure of distance either, despite suggestions in the literature of extra-ocular information about the orientation of the two eyes—i.e., the vergence-distance d_v—influencing the neural analysis of retinal disparities (Trotter *et al.*, 1992).

Finally, Fig. 3(c) shows d_d as a function of d_s. There is still a considerable amount of variation, but the deviation from perfect correspondence appears to be smaller and less systematic than for the other comparisons.

Qualitatively similar results were obtained for the other subjects, although there were large quantitative differences [see Fig. 3(d,e,f)]. Clearly, the "deviation from perfect correspondence" has to be quantified before any conclusion is justified. For each point, the absolute value of the separation from the ideal line was taken as a measure of how well the two distances correspond [see Fig. 3(b) for an illustration]. The distribution of separations was not symmetrical, because d_s and d_d do not depend linearly on the set size and set depth. In such cases, the median of the separations is a suitable measure of the level of discrepancy between the two distances being compared. If the median is zero, we have an ideal setting and the lowest level of discrepancy (highest level of correspondence) between the two distances. The larger the median, the higher the level of discrepancy.

The levels of discrepancy for each subject and for each combination of distances (d_s vs d_v, d_d vs d_v and d_d vs d_s) are shown in Fig. 4.

As can be seen from Fig. 4, for all subjects the discrepancy level was much smaller for the combination d_d, d_s than it was for the other combinations (d_d, d_v and d_s, d_v). This was confirmed with paired t-tests which showed that the level of discrepancy did not differ significantly between the combinations d_s, d_v and d_d, d_v ($P = 0.94$),

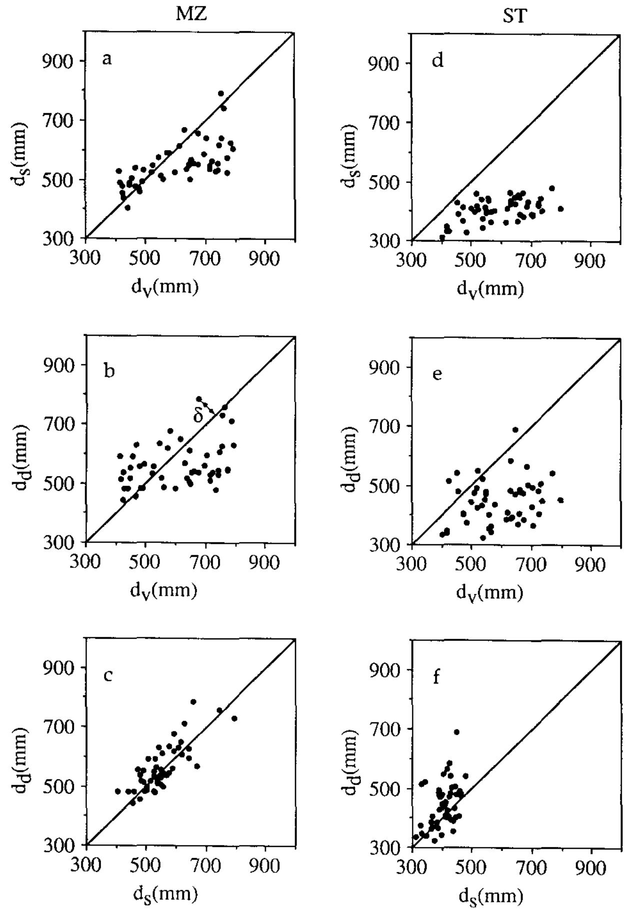

FIGURE 3. Settings of size and depth for two subjects (MZ and ST), expressed as size-distance d_s, disparity-distance d_d and vergence or simulated distance d_v. The solid line represents perfect correspondence between the two measures. In Fig. 3(b), δ denotes the separation of one data point from the solid line. We used the median of the absolute values of such separations as an overall measure of the level of discrepancy of the two distances that were being compared.

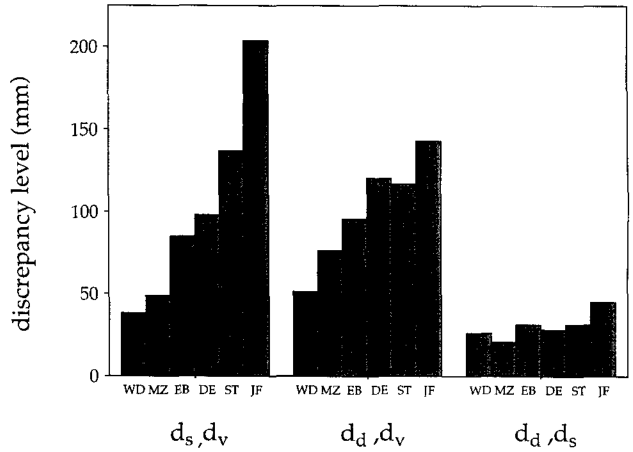

FIGURE 4. Results of Experiment 1. This figure shows the level of discrepancy for each subject for the three possible combinations of vergence-distance, size-distance and disparity-distance. Room lights are off.

but did differ significantly between d_s, d_v and d_d, d_s ($P = 0.02$), and between d_d, d_v and d_d, d_s ($P < 0.01$).

A large level of discrepancy for d_s, d_v and d_d, d_v could indicate that the simulated distance d_v hardly influences the perceived size and depth. However, it could also arise from systematic under or overestimation of the distance. We therefore determined the linear fit to the d_s, d_v and the d_d, d_v settings. Figure 9 shows the slopes of these fits for all the experiments. The regression analysis of the data gave an average slope (over subjects) of 0.32 for both d_s, d_v and d_d, d_v. The average value of d_s was 47 cm and that of d_d 48 cm. The average simulated distance (d_v) was 60 cm. Thus, variations in the simulated distance were underestimated and the ellipsoid was seen closer than it was.

The main finding is that the level of correspondence between size-distance and disparity-distance was significantly higher than for the other distance combinations. In other words, it is likely that the distance used to scale retinal image size is the same as the distance used to scale disparity.

Uncertainty concerning the size of the real tennis ball will always lead to some variation. For example, a 10% error in assumed ball size leads to errors of about 60 mm in d_s and 30 mm in d_d. Considering this source of error, and limitations of the accuracy with which the subjects can set the desired size and disparity, the variability in, for example, Fig. 3(c) is not too surprising. However, this variability could also be due to different measures of distance being used, but the measures being similar under such limited-cue conditions.

EXPERIMENT 2

Glennerster *et al.* (1994) showed that the use of a limited-cue environment (such as was used by Johnston, 1991 and in Experiment 1) leads to poor shape constancy. Limited information about the target's distance could account for the variations in the simulated distance being underestimated in Experiment 1. The better correspondence between size-distance and disparity-distance suggests that the same measure of distance was used (although it was often incorrect) to scale retinal size and retinal disparity.

In the second experiment we increased the information about distance by turning the room lights on. With the room lights on, additional objects such as the monitor become visible. Such objects' distances can be determined from a variety of sources (including familiar size). The distance of the simulated ellipsoid could be estimated in relation to such objects on the basis of relative disparity. Thus, turning on the room lights should improve the correspondence between the simulated and the perceived size and depth. We expect the size-distance and disparity-distance to be closer to the simulated distance, and thus the discrepancy between d_s and d_v and between d_d and d_v to become smaller. Does this change the correspondence between the size-distance and the disparity-distance? If a single estimate of distance is used to scale disparity and size, the correspondence between d_s and d_d should not improve (discrepancy should not decrease). If separate estimates of distance are used,

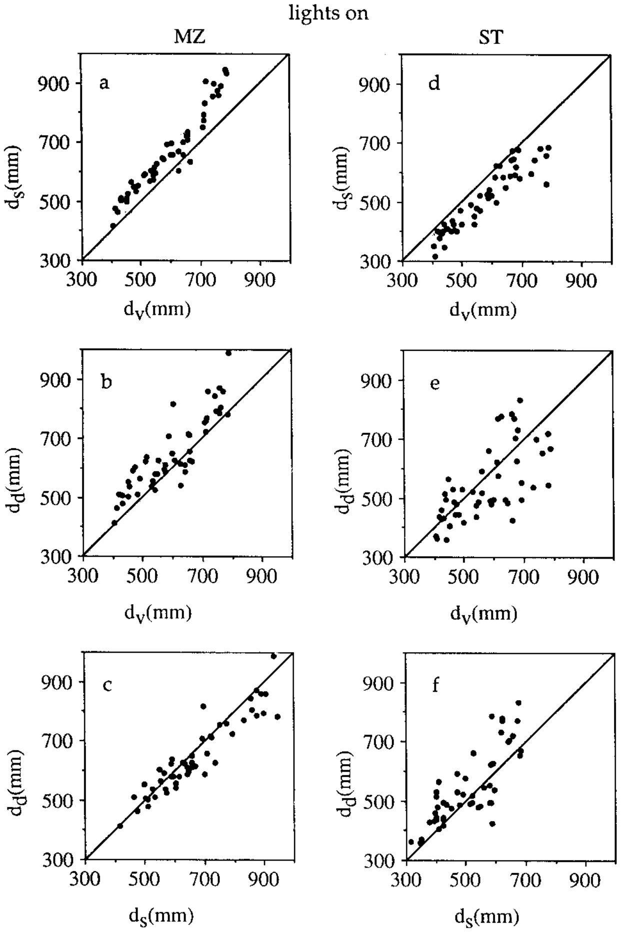

FIGURE 5. Settings of size and depth for subjects MZ and ST in Experiment 2. Format as in Fig. 3. Room lights are on.

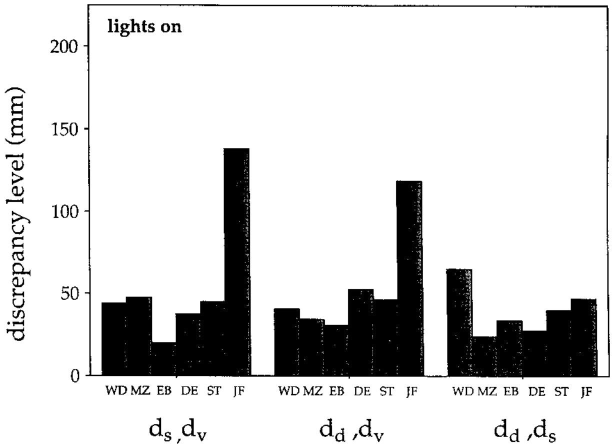

FIGURE 6. Results of Experiment 2. This figure shows the level of discrepancy for each subject for the three possible combinations of vergence-distance, size-distance and disparity-distance. Room lights are on.

bringing each estimate closer to the simulated distance should lead to a higher correspondence (discrepancy should decrease) when the room lights are on. To examine whether the correspondence between d_s and d_d changes we repeated Experiment 1 (with the same procedure, task, subjects and stimuli), but now with the room lights on. Although the room was no longer completely dark, the real tennis ball could not be seen during the experiment because it was held under the table.

Results

Figure 5 shows the results for subjects MZ and ST in the same format as in Fig. 3. Figure 6 shows the level of discrepancy in the same format as in Fig. 4. The only difference between the experiments was that the room lights were on.

As expected, the levels of discrepancy for the combinations d_s, d_v and d_d, d_v are much smaller than they were when the experiment was conducted in the dark (on average, 62 and 53% of the previous values; $P = 0.03$ and $P < 0.01$, respectively; see Fig. 4). The distances used for scaling retinal size and retinal disparity are closer to the simulated distance because there is more information about distance available when the room lights are on. The average slope (over subjects) of d_s vs d_v increased from 0.32 to 0.78 and that of d_d vs d_v increased from 0.32 to 0.77 when the room lights were turned on (see Fig. 9). The average value of d_s increased from 47 to 54 cm and that of d_d increased from 48 to 57 cm, which is in both cases closer to the average simulated distance of 60 cm.

The level of discrepancy of the d_d, d_s combination in the "lights on" condition was not smaller than that in the "darkness" condition. In fact, we found an average increase of 33%, but the increase was not statistically significant ($P = 0.19$).

The level of correspondence between size-distance and disparity-distance did not improve, although each on its own got closer to the simulated distance. This is what one expects if the same distance is used for scaling retinal size and retinal disparity.

EXPERIMENT 3

In Experiments 1 and 2, the subjects set the size and depth of the simulated ellipsoid so that its 3-D shape matched a tennis ball. We assumed that the subjects performed this task by scaling disparities and retinal size by some measure of distance. The results of Experiments 1 and 2 showed that the same distance is used for scaling retinal disparity and retinal size.

Rogers and Cagenello (1989) suggested that local surface curvature could be determined by a measure that does not depend on the viewing distance. This measure, the second order spatial derivative of the disparity field (which they called disparity curvature) could be used to estimate local curvature without the risk of scaling disparities with a wrong estimate of distance. Brookes and Stevens (1989) also suggested that second order spatial derivatives may serve as surface curvature measures, which could be used in the reconstruction of continuous 3-D surfaces.

Using distance-independent measures for reconstructing 3-D shape would bypass distortions in perceived distance, and could result in a higher degree of shape constancy under some conditions (see below). However, shape must then be recovered from local curvature rather than from global depth and width. This is not simple: even for a ball, local measures of curvature may depend on the position on the surface (equal disparity curvature does not imply equal intrinsic curvature) and must depend on the size of the ball (a larger ball is less curved). Howard and Rogers (1995) point out that some measure of distance is required to determine global shape from local surface curvature measures, no matter what order of disparity is used for determining these local curvature measures. In our task, however, observers could have used *a priori* knowledge of the disparity curvature of a certain part of the real tennis ball (for instance the apex of the visible surface). The correspondence between d_s and d_d in Experiments 1 and 2 could be the result of matching the disparity curvature of the chosen part of the simulation to the known disparity curvature of the corresponding part of a tennis ball. For our example of using the centre of the visible surface, setting too large an ellipsoid ($d_s < d_v$) has to be compensated for by making the ellipsoid more elongated ($d_d < d_v$) if the local curvature is to be maintained. This could lead to the covariation of d_s and d_d that was obtained in Experiments 1 and 2. We therefore repeated the previous experiments with a stimulus that eliminates the possibility of using higher order disparity information for the estimation of local curvature.

Stimulus

In this experiment a different stimulus was used. The new stimulus was a simulation of a smooth red solid ellipsoid with a single small black spot on its surface (Fig. 7).

Such a stimulus contains one measure of relative disparity (the disparity of the spot relative to the outline of the ellipsoid). Derivatives of this disparity field are, therefore, ill defined, but the relative depth of the spot is well defined. Within each trial, the simulated size of the spot was kept constant, so that the retinal size varied with the depth of the spot relative to the disk. Before each trial, the simulated radius of the spot (the spot-size) was set at random from values between 0.66 and 1.33 mm. This prevents the spot's size relative to that of the disk from acting as a monocular cue for the shape of the simulated object (but could account for some additional variability if a constant dot size is assumed; Sedgwick, 1986). Note that the shape of the simulated object was not uniquely defined. However, because the width and the depth of the simulated object were well defined, subjects were, in principle, able to perform the task. None of the subjects reported any difficulties in perceiving a 3-D shape in the simulation. Also note that this stimulus avoids the monocular (texture and density) cues that were present in the former experiments.

Procedure

Subjects were instructed to adjust the size and the 3-D shape of the simulated object to match the tennis ball.

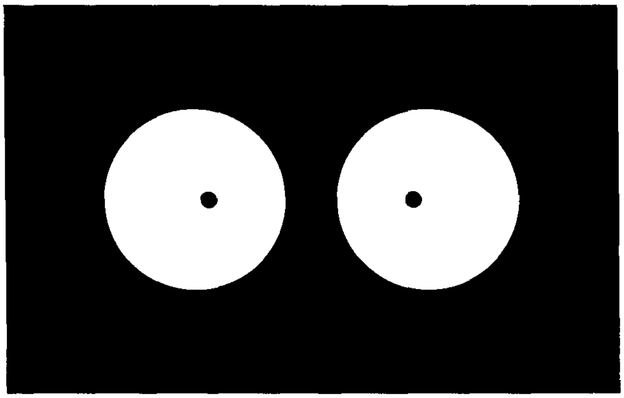

FIGURE 7. Illustration of the stimulus used in Experiment 3.

Some subjects reported that they imagined that they were placing the black spot, that was floating in depth, onto the surface of a sphere instead of deforming an object with a black spot attached to it. Both strategies should lead to the same performance. If subjects used higher order derivatives of the disparity field to set the correct shape, then they should perform worse with the new stimulus than in the previous experiments. Conversely, if subjects can perform the task as well with the new stimulus, we could conclude that they do not use higher order derivatives of the disparity field to determine 3-D shape, but indeed scale disparity with some measure of distance.

Except for the above-mentioned change in stimulus, this experiment was identical to the previous two experiments. Each subject made 100 settings of size and depth, 50 in complete darkness and 50 with the room lights on. The same subjects that took part in Experiment 1 and 2 also acted as subjects in this experiment.

Results

The results of Experiment 3 are shown in Figs 8(a) (darkness) and 8(b) (lights on). Each part shows the level of discrepancy for the three distance combinations.

As can be seen from a comparison of Figs 8(a) (the "darkness" condition) and 4, the level of discrepancy is similar for a "single spot" stimulus and for the "random lines" stimulus of Experiment 1. There is a tendency for the level of discrepancy to be smaller in the "single spot" stimulus, but only one of the differences was significant (d_d, d_v: $P = 0.03$).

With the room lights on [Fig. 8(b)], one subject (ST) had difficulty interpreting the stimulus. His level of discrepancy for two of the three comparisons is very high. For all other subjects, the levels of discrepancy for the "single spot" stimulus were similar to those for the "random lines" stimulus (none of the differences were significant).

Figure 8(a) shows that, in the dark, the level of discrepancy is lower for the combination d_d, d_s than for the two other combinations (d_s, d_v vs d_d, d_v: $P = 0.31$; d_s, d_v vs d_d, d_s: $P = 0.03$ and d_d, d_v vs d_d, d_s: $P = 0.08$). When the room lights were on, the levels of discrepancy decreased, as was to be expected, for all but the

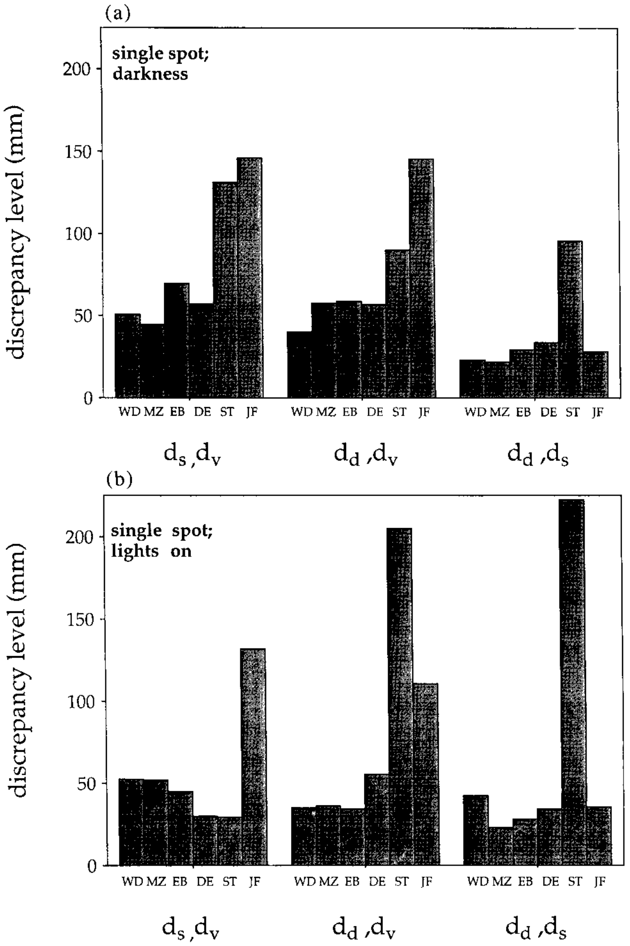

FIGURE 8. Results of Experiment 3. This figure shows the level of discrepancy for each subject for the three possible combinations of vergence-distance, size-distance and disparity-distance. Procedure as in Experiments 1 and 2, but now with the "single spot" stimulus. (a) Room lights off; (b) Room lights on.

combination d_d, d_s. For the latter combination, the level of discrepancy stays the same (except for that of subject ST, for whom it increased).

Figure 9 shows that the influence of the simulated distance on d_s and d_d again gets larger when the lights are turned on: the slopes get closer to 1. With the room lights off, the variation in simulated distance d_v was underestimated. The underestimation was comparable to that found in Experiment 1. The average value of d_s was 48 cm and that of d_d 51 cm. Thus, again the ellipsoid was seen closer than it was, very much as in Experiment 1. When the lights were turned on, the average values increased to 57 and 66 cm, respectively, which is, in both cases, closer to the average simulated distance.

In summary, when we compare the results from Experiments 1 and 2 with those of Experiment 3, we find that the levels of discrepancy were more or less equal for a stimulus with a single spot (Fig. 8) and for one with random lines (Figs 4 and 6). Moreover, in both cases, the size-distance and the disparity-distance came closer to the simulated distance when the room lights were turned on, but the level of discrepancy for the combination size-distance, disparity-distance remained the same.

GENERAL DISCUSSION

The results of the present experiments confirm earlier reports describing a systematic distortion in the perception of three-dimensional shape from binocular stereopsis under reduced cue conditions (Johnston, 1991). The results of Experiment 1 show that, if we assume that retinal size and retinal disparity are each scaled by some distance, they are probably scaled by the *same* distance. This was confirmed by the results of Experiment 2: when the settings were performed with the room lights on, the size-distance and the disparity-distance were both closer to the simulated distance, but the level of discrepancy between size-distance and disparity-distance did not change, and was still always lower than the level of discrepancy between any other distance combination. This is what one would expect if both size and depth settings were based on the same estimate of distance. These results are in line with preliminary results reported by Rogers and Bradshaw (1995b) who also conclude that retinal size and disparity are scaled by the same distance. They found similar magnitudes of distance scaling for

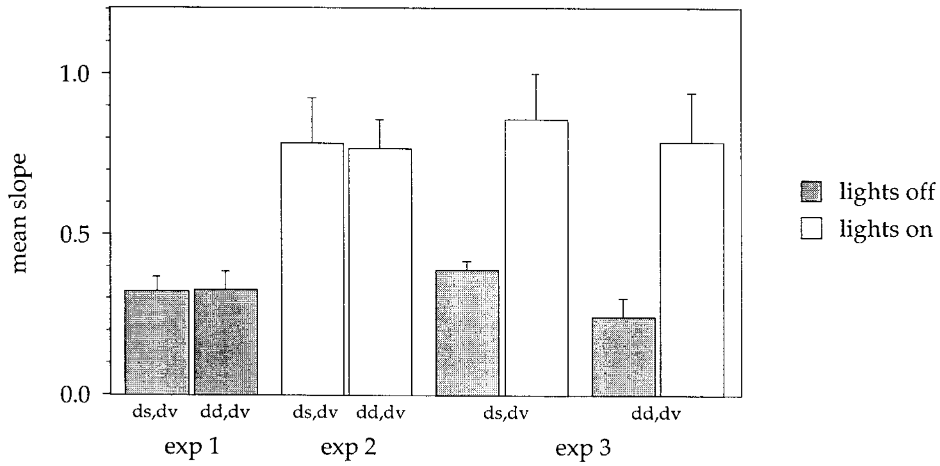

FIGURE 9. Average slopes (over six subjects; with standard errors) from linear regression analyses of the distance combinations d_s, d_v and d_d, d_v. The average intercepts were 28 cm (Experiment 1); 10 cm (Experiment 2); 27 cm (Experiment 3, lights off) and 13 cm (Experiment 3, lights on).

size, depth and shape in experiments with a much larger range of simulated distances.

Experiment 3 showed that matching 3-D shape using a single spot results in comparable performance to matching 3-D shape with a stimulus containing a rich disparity field (the textured ellipsoid). A rich disparity field provides information about local 3-D curvature that could be used to improve the perception of global 3-D shape. In our experiments, the additional presence of such information does not lead to a more veridical perception of global shape. Performance is almost the same for both types of stimuli, which suggests that the same strategy is followed in both cases; i.e., that higher order derivatives are not used to improve the perception of 3-D shape.

For a stimulus with a single spot, the only meaningful strategy is to estimate depth by scaling the single relative disparity with distance. It is, therefore, reasonable to conclude that the perception of 3-D shape in the presence of a well defined disparity field is accomplished by scaling retinal disparity and retinal size by the same estimate of distance, and that using disparity curvature or other distance-independent measures for local surface curvature does not improve the reconstruction of 3-D shape.

REFERENCES

Brookes, A. & Stevens, K. A. (1989). The analogy between stereo and brightness. *Perception, 18*, 601–614.

Colett, T. S., Schwarz, U. & Sobel, E. C. (1991). The interaction of oculomotor cues and stimulus size in stereoscopic depth constancy. *Perception, 20*, 733–754.

Glennerster, A., Rogers, B. J. & Bradshaw, M. F. (1993). The constancy of depth and surface shape for stereoscopic surfaces under more naturalistic viewing conditions. *Perception, 22* (Suppl.), 118.

Glennerster, A., Rogers, B. J. & Bradshaw, M. F. (1994). The effects of (i) different cues and (ii) the observer's task in stereoscopic depth constancy. *Investigative Ophthalmology and Visual Science, 35*, 2112.

Howard, I. P. & Rogers, B. J. (1995). *Binocular vision and stereopsis. Oxford Psychology Series No. 29.* Oxford: Oxford University Press.

Johnston, E. B. (1991). Systematic distortions of shape from stereopsis. *Vision Research, 31*, 1351–1360.

Rogers, B. J. & Bradshaw, M. F. (1992). Differential perspective effects in binocular stereopsis and binocular disparity. *Investigative Ophthalmology and Visual Science, 34*, 1333.

Rogers, B. J. & Bradshaw, M. F. (1995a) Disparity scaling and the perception of frontoparallel surfaces. *Perception, 24*, 155–179.

Rogers, B. J. & Bradshaw, M. F. (1995b) Binocular judgements of depth, size, shape and absolute distance: is the same "d" used for all judgements? *Investigative Ophthalmology and Visual Science, 36*, 230.

Rogers, B. J. & Cagenello, R. (1989). Disparity curvature and the perception of three-dimensional surfaces. *Nature, 339*, 135–137.

Sedgwick, H. A. (1986). Space perception. In Boff, K. R., Kaufman, L. & Thomas, J. P. (Eds), *Handbook of perception and human performance: Vol. 1. Sensory processes and perception* (pp. 21.1–21.57). New York: Wiley.

Tittle, J. S., Todd, J. T., Perotti, V. J. & Norman, J. F. (1995). The systematic distortion of perceived 3-D structure from motion and binocular stereopsis. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 663–678.

Trotter, Y., Celebrini, S., Stricanne, B., Thorpe, S. & Imbert, M. (1992). Modulation of neural stereoscopic processing in primate area V1 by the viewing distance. *Science, 257*, 1279–1281.

Wagner, M. (1985). The metric of visual space. *Perception and Psychophysics, 38*, 483–495.

Acknowledgements—This work was supported by the Dutch Foundation for Scientific Research (NWO). We wish to thank Brian Rogers and the anonymous second reviewer for their helpful comments on a previous version of this article.