Large continuous perspective transformations are necessary and sufficient for accurate perception of metric shape

GEOFFREY P. BINGHAM Indiana University, Bloomington, Indiana

AND

MATS LIND Uppsala University, Uppsala, Sweden

We investigated the ability to perceive the metric shape of elliptical cylinders. A large number of previous studies have shown that small perspective variations ($\leq 10^{\circ}$) afforded by stereovision and by head movements fail to allow accurate perception of metric shape. If space perception is affine (Koenderink & van Doorn, 1991), observers are unable to compare or relate lengths in depth to frontoparallel lengths (i.e., widths). Frontoparallel lengths can be perceived correctly, whereas lengths in depth generally are not. We measured reaches to evaluate shape perception and investigated whether larger perspective variations would allow accurate perception of shape. In Experiment 1, we replicated previous results showing poor perception with small perspective variations. In Experiment 2, we found that a 90° continuous change in perspective, which swapped depth and width, allowed accurate perception of the depth/width aspect ratio. In Experiment 3, we found that discrete views differing by 90° were insufficient to allow accurate perception of metric shape and that perception of a continuous perspective changes of 30°, 45°, 60°, and 90° and discovered that a 45° change or greater allowed accurate perception of the aspect ratio and that less than this did not. In conclusion, we found that perception of metric shape is possible with continuous perspective transformations somewhat larger than those investigated in the substantial number of previous studies.

Object shape entails both qualitative and quantitative properties. On the one hand, shape entails relief structure or relative variations in the surface conformation (e.g., ellipsoidal vs. cylindrical). As has been described by Perotti, Todd, Lappin, and Phillips (1998), this characteristic of shape can be measured locally for smoothly curved surfaces by the shape characteristic, a ratio of the principal curvatures. On the other hand, metric shape can be measured by curvedness (Koenderink, 1990) or, alternatively, the ratio of an object's width to its depth. As has been reviewed by Todd, Tittle, and Norman (1995) and shown by Perotti et al. (1998), among many others (e.g., Brenner & van Damme, 1999; Lappin & Ahlström, 1994; Scarfe & Hibbard, 2006; Tittle, Todd, Perotti, & Norman, 1995; Todd & Norman, 1991), observers do not appear to be able to perceive metric shape, or the relative depth of objects, very accurately. Perotti et al. (1998), in particular, found that observers were well able to judge the shape characteristic but that judgments of curvedness were biased and highly variable. These studies involved structure from motion or stereo computer graphic displays of objects. Lind, Bingham, and Forsell (2003) asked observers to judge the shape of textured wooden cylinders-that is, actual objects

that were about 7 cm in size. Aspect ratios (depth/width [D/W]) between 0.46 and 1.81 were tested. Observers used binocular vision with free head movement in normal lighting, and the objects sat on a tabletop within reach distance, so that the tops of the objects could be seen. The participants adjusted the shape of an elliptical outline on a computer screen to match the perceived cross-sectional shape of each object. The results were much like those in Perotti et al. (1998). Judgments were highly variable and inaccurate. Lind et al. (2003) replicated the result with variations in viewing height and distance. The participants did not reliably begin to get it right until they were essentially looking straight down at the tops of the objects, so they could match their outline to the top edge of the cylinders appearing in a frontoparallel plane. It is clear that observers cannot judge metric shape under viewing conditions that allow only relatively small perspective variations ($\leq 10^{\circ}$).

There are many studies showing dissociations between the results of judgment tasks and those of tasks involving relevant action measures. (For instance, Pagano & Bingham, 1998, found that judgment errors and reach errors were uncorrelated [lag 0, 1, or 2] when observers were asked to judge the distances of targets and then reach to them, with feedback about the reach error on each trial.) We performed a number of studies in which we used reach measures to evaluate perception of distance, size, and shape of target objects in both actual environments (Bingham, Zaal, Robin, & Shull, 2000) and virtual environments (Bingham, 2005; Bingham, Crowell, & Todd, 2004) and also comparing the two (Bingham, Bradley, Bailey, & Vinner, 2001). We found that feedback about reach error on successive trials yielded reliably accurate performance in respect to object distance and size. Distance and size perception needed to be calibrated, but once they were, performance was good (see also Bingham, 2005; Bingham, Coats, & Mon-Williams, 2007; Coats, Bingham, & Mon-Williams, in press; Mon-Williams & Bingham, 2007; Mon-Williams, Coats, & Bingham, 2004). However, feedback failed to improve performance in respect to object shape (Bingham, 2005; Lee, Crabtree, Norman, & Bingham, 2008), and position perception was found to be independent of shape perception (Bingham et al., 2004; see also Crowell, Todd, & Bingham, 2000, 2001; Loomis, Philbeck, & Zahorik, 2002).

These results were consistent with results from studies of visually guided reaches-to-grasp. Open loop reaches-tograsp performed with both stereovision and motion parallax exhibit inaccurate grasping (although reaching is accurate; Brenner & van Damme, 1999; Hibbard & Bradshaw, 2003; Melmoth & Grant, 2006; Watt & Bradshaw, 2003). In contrast, online guidance using specifically binocular vision (which affords disparity matching at the end of the reachto-grasp movement) yields accurate grasping (Bradshaw et al., 2004; Cuijpers, Smeets, & Brenner, 2004; Melmoth & Grant, 2006; Watt & Bradshaw, 2003). Bradshaw and Elliot (2003) manipulated when online binocular guidance became available during a reach and found that it was differentially effective only right at the end during the grasp (see also Hibbard & Bradshaw, 2003; Watt & Bradshaw, 2000). Binocular vision is uniquely effective in the context of reach-to-grasp actions, but not because it affords good perception of metric shape used to control feedforward portions of the actions, but because it allows the use of disparity matching to guide the fingers to the surfaces of the target object in the final phases of the grasping movement.

The question remains: Might observers ever be able to perceive metric shape correctly? We now hypothesize that correct shape perception requires sufficiently large perspective variations. The question, then, is how large? On a first pass, there appears to be a fundamental inability to relate object depths to widths. This would be consistent with the possibility that vision yields reliable apprehension only of affine properties of visual space (e.g., Koenderink & van Doorn, 1991; Todd & Bressan, 1990). If this were true, the way to establish a relation between width and depth would be to exchange them-that is, rotate an object by 90° to move the width into the depth plane and the depth into a frontoparallel plane. If we do this, can observers now finally judge metric shape correctly? Presumably, they should. However, given the limitations of metric shape perception, two discrete views of an object before and after it is rotated by 90° should be insufficient to allow reliable perception of metric shape. Instead, a continuous variation in perspective should be required. Might a perspective variation less than 90° be sufficient? Given the natural symmetries of 45° angles (namely, $\sin 45^{\circ} = \cos 45^{\circ}$), we hypothesized that a continuous perspective change of 45° should be enough to allow accurate perception of metric shape. In the following experiments, we systematically explored these possibilities. In Experiment 1, we tested two different versions of a reaching task with small changes in perspective-that is, 10°-15° of effective rotation. In both cases, we expected to replicate previous results showing inaccurate perception of metric shape. In Experiments 2 and 3, we tested perspective changes of 90°. We expected that perception of continuous rigid rotations would be essential for the accurate perception of metric shape. In Experiment 4, we investigated perspective changes less than 90° but greater than 15°, with the expectation that 45° of continuous rigid rotation would be both necessary and sufficient to yield accurate perception of metric shape.

EXPERIMENT 1

We used targeted reaching to evaluate perception of the metric shape of elliptical cylinders. Three different objects with D/W aspect ratios of 0.5, 1.0, and 1.5 were viewed from two perspectives varying by 90°, yielding two additional aspect ratios of 2.0 and 0.67, respectively (see Figure 1 for an illustration). Objects were viewed in a virtual environment on a visible support surface. The participants reached to place a stylus, held vertically in the hand, tangent to the surface of the cylinder at the front, back, left, or right side of the object. Each of two different



Figure 1. (A) The way shape aspect ratios were represented by the objects. (B) Target placement for calibration and test trials. (C) Stylus placement task in the virtual environment. See the text for details in all cases.

groups of participants did this in one of two ways. The first group performed a separate reach to each given target location on each object. The second group performed a single reach to each object and then placed the stylus at each of the four locations in (random) sequence. The goals of Experiment 1 were twofold. First, we were replicating the previous results. Second, we were testing whether the two methods would yield the same results, so we could use either method, as needed, in subsequent experiments.

Method

Participants. Twenty adults 19–30 years of age participated in the experiment. Ten participated in the multiple-reach condition (3 of them male, 7 female). Ten participated in the single-reach condition (5 of them male, 5 female). The participants were paid \$7/h. All the participants had normal or corrected-to-normal eyesight (using contacts) and normal motor abilities. All were right-handed.

Apparatus. The Virtual Environment Lab consisted of an SGI Octane graphics computer, a Flock of Birds (FOB) motion measurement system with two markers, and a Virtual Research V8 stereo head-mounted display (HMD). One marker was placed on the HMD, and the other on a stylus held in the participant's hand. The stylus was a Lucite dowel 18.5 cm in length and 1 cm in diameter. Displays in the HMD portrayed a virtual target cylinder on a surface and a handheld stylus. As is shown in Figure 2, the virtual target cylinder was covered with random green phosphorescent triangular texture elements and appeared on a dark green 0.5×0.5 m horizontal virtual support surface. The front edge of the surface was positioned directly below the participant's eyes. The stylus and marker were modeled precisely and appeared as a gray virtual stylus with a blue and red marker at its bottom. The hand was not modeled, so the participants saw only the virtual stylus, but its position and motion were the same as the actual stylus. There were no shadows cast on the target by the stylus or by the target on the stylus.

Three target cylinders were presented as shown in Figure 1A. They varied in terms of the lengths of one of the principal axes of their elliptical cross section. All of the other principal axes were 7 cm in length. The major axes were 3.5, 7, and 10.5 cm, respectively. The noncircular targets were presented from two perspectives, looking along the major or the minor axis, respectively. Effectively, five targets with five different D/W aspect ratios were presented. Only circular target cylinders were used during calibration trials.

The HMD displays subtended a 60° field diagonally with complete overlap of the left and right fields. The resolution was $640 \times$



Figure 2. Illustration of the virtual objects that appeared in our displays.

480 pixels, and the frame rate was 60 Hz. The weight of the helmet was 0.82 kg. The sampling rate of the FOB was 120 Hz. As described in Bingham et al. (2001), we measured the focal distance to the virtual image, the image distortion, the phase lag, and the spatial calibration. The virtual image was at 1 m distance from the eyes. The phase lag was 80 msec. The spatial calibration yielded a resolution of about 2 mm (see Bingham et al., 2001, for additional information about the virtual environment).

Procedure. Seated participants donned the HMD and spent a few minutes moving their head and hand to explore and acclimate to the virtual environment. For Group 1, the distance of the support surface below eye height was 15 cm. For Group 2, distances of 5 and 20 cm below eye height were tested. For calibration trials, the target cylinder was placed at 50% of the participant's maximum reach distance. For subsequent test blocks, the target was placed at 70% of maximum reach distance (see Figures 1B and 1C).

The task was explained to the participants. The participants were instructed to reach to place the stylus at one of four locations relative to the surface of the target cylinder, as shown in Figure 1C. Holding the stylus vertically, they reached to place the stylus tangent to the surface of the cylinder to the front, right, left, or back. Only the virtual target object and the support surface could be seen, not the virtual stylus, except at the very end of trials on the calibration trials, at which point the virtual stylus was made visible, as will be explained below. For the first group, a separate reach from the hip was performed to each of the four locations on the target cylinder. For the second group, a single reach from the hip was performed to place the stylus successively at each of the four target locations. The participants were provided no other information about the target objects, and in particular, they did not know how many target objects they would see or what the aspect ratios of those objects would be.

At the beginning of each trial, the target appeared, and the computer announced to the participant the sequence in which locations were to be touched on the target (e.g., front, right, back, left). The participants first moved their head 10 cm side to side two or three times at preferred rates while counterrotating their head to keep the target centered in the display and looking at the targeted locus on the surface. Then, the participants reached at preferred rates. Once the participants had reached the target, they said "OK" and the 3-D coordinates of the stylus were recorded. In the test trials, this ended the trial. In the calibration trials, the virtual stylus would become visible (seen together with the target cylinder) at the same time that the 3-D coordinates of the stylus were recorded. When the stylus was made visible, the participants were allowed to move the stylus to correct its position, if necessary. For Group 2, the stylus was made invisible again before the participants moved to the next location, where, after the position was recorded, it was made visible again, and so on.

In Group 1, a block of trials consisted of reaches to each of the 4 locations on each of the five targets—that is, 20 locations visited in a random order. Three blocks of trials were performed. In Group 2, trials were first blocked by object, so that all four locations on a given target object were visited in a random order, with all objects being tested before a given object was tested again. Again, three blocks of trials were performed. Both groups performed 60 test reaches, preceded by 12 calibration trials in which the participants reached to a circular cylinder. Group 2 performed a second set of calibration and test trials at the second eye height. The order in which the 5- and 20-cm surface heights were tested was counterbalanced across participants.

Dependent measures. The method allowed us to evaluate a number of perceptual properties concurrently. Four dependent measures were computed for each set of four reaches to each object. We used Cartesian coordinates, so that depth varied along the *x*-axis and the *y*-axis lay in a frontoparallel plane. We computed the target distance as the *x* centroid of the four reaches. The difference in *y* between reaches to the left and right yielded *width* in a frontoparallel plane. *Depth* was computed as the difference in *x* between front and back. *Shape* was computed as the aspect ratio of depth to width. Width, depth, and the aspect ratio were computed for each participant and each target object, using each sequence of four successive reaches



Figure 3. (A) Data for Group 1 in Experiment 1: Mean judged width and depth (with standard error bars) plotted as a function of actual width and depth, each shown with a line fit by a least squares regression. Filled circles, width; filled squares, depth. (B) Data for Group 2 in Experiment 1: Mean judged width and depth (with standard error bars) plotted as a function of actual width and depth, each shown with a line fit by a least squares regression. Open circles, width at 5-cm eye height; filled circles, width at 20-cm eye height; open squares, depth at 5-cm eye height; filled squares, depth at 20-cm eye height. (C) Data for Group 1 in Experiment 1: Mean judged aspect ratios (with standard error bars) plotted as a function of actual aspect ratios, shown with a line fit by a least squares regression. *D/W*, depth/width. (D) Data for Group 2 in Experiment 1: Mean judged aspect ratios (with standard error bars) plotted as a function of actual aspect ratios, each shown with a line fit by a least squares regression. Filled circles, 5-cm eye height; filled squares, 20-cm eye height.

to the four locations, yielding three of each measure for each object and participant within each group and surface height.

Results and Discussion

The results were that metric shape was not perceived accurately and, in particular, object depths were poorly resolved. If metric shape was perceived accurately, the slopes and intercepts of simple regressions relating actual to judged widths and actual to judged depths should be the same. A multiple regression was performed to test whether this was the case. As is shown in Table 1 and Figure 3, the slopes were not the same. Slopes for depth, in particular, were low.

We performed a multiple regression on reach widths and depths, using actual widths and depths as a continuous independent variable, width versus depth as a categorical variable (coded as ± 1), and an interaction vector. We performed this analysis separately on the Group 1 data and on the data for each eye height of Group 2. For Group 1, the result was significant [F(3,296) = 32.5, p <.001, $r^2 = .25$], and all independent variables were significant: actual width and depth (partial F = 60.2, p <.001), width versus depth (partial F = 12.2, p < .001), and the interaction (partial F = 22.1, p < .001). The results of separate simple regressions are shown in Table 1 and illustrated in Figure 3A. As indicated by the multiple regression results, the slopes for widths and depths were different. The slope for width was significantly above 1, whereas that for depth was below 1. Variations in depth were resolved poorly.

Table 1Results for Experiment 1						
Eve	Width (W)		Depth (D)		D/W	
Height (cm)	Slope	r^2	Slope	r^2	Slope	r^2
		Gr	oup 1			
15	1.63	.38	0.40	.03	0.79	.20
		Gr	oup 2			
5	1.29	.28	0.45	.06	0.48	.25
20	1.13	.17	0.40	.04	0.62	.13

The pattern of results was similar at both eye heights of Group 2. For the 5-cm eye height, the regression was significant [$F(3,296) = 57.6, p < .001, r^2 = .32$], and both the main effect of width and depth (partial F = 60.2, p < .001) and the interaction (partial F = 13.5, p < .001) were significant. For the 20-cm eye height, the regression was significant [$F(3,296) = 24.0, p < .001, r^2 = .20$], and both the main effect of width and depth (partial F = 33.8, p < .001) and the interaction (partial F = 7.6, p < .01) were significant. Simple regressions are shown in Table 1 and Figure 3B, where again the slopes for depths are low (slope ≈ 0.4), indicating that depth variations were discriminated poorly.

Next, we computed the shape D/W aspect ratios as a direct measure of perceived metric shape. We performed a multiple regression to compare the results of Group 1 and Group 2 at the 20-cm eye height. We regressed actual aspect ratios on aspect ratios derived from reach data, using group as a categorical independent variable (coded as ± 1), together with an interaction vector. The result was significant $[F(3,296) = 21.8, p < .001, r^2 = .18]$, but only actual aspect ratio was significant (partial F = 61.2, p <.001). There was no difference in results between the two methods. As is shown in Table 1 and Figures 3C and 3D, both yielded low slopes of about 0.70. However, we performed a similar regression comparing the two eye heights in Group 2, and after removing the nonsignificant categorical factor (using a procedure described in Pedhazur, 1982), we found that the smaller eye height yielded a significantly lower slope, meaning that the aspect ratios were even more poorly discriminated. The overall regression was significant $[F(2,297) = 36.4, p < .001, r^2 = .20]$, and both the actual aspect ratio (partial F = 56.2, p < .001) and the interaction (partial F = 16.5, p < .001) were significant.

EXPERIMENT 2

Koenderink and van Doorn (1991) suggested that vision may detect only information allowing the affine structure of the 3-D surroundings to be apprehended (see also Todd et al., 1995). This would mean that observers simply cannot relate widths to depths. If this is true, a perspective change that exchanged width and depth may allow depth and width to be accurately apprehended and compared. We used the method tested in Group 1 in Experiment 1 to allow us to manipulate the availability of the information in such a perspective change. Observers in Experiment 2 performed a separate reach in random-ordered trials to touch each location (front, left, right, or back) on each of five target shapes. We randomly varied whether the 90° perspective change was allowed on each trial. The result was a variation in the number of 90° perspective changes from zero to four for a given block of four reaches to a given target shape. We examined the accuracy of the aspect ratios as a function of this variation.

Method

Participants. Ten adults 18–28 years of age participated in the experiment. Four were male and 6 were female. The participants were paid \$7/h. All the participants had normal or corrected-to-normal eyesight (using contacts) and normal motor abilities. All were right-handed.

Procedure. Both the apparatus and the procedure were the same as those for Group 1 in Experiment 1, with the following changes. The participants sat in a desk chair with wheels. An actual 1-m square wooden table was placed underneath the location of the virtual support surface. This was to help to orient both the experimenters and the participants, since the participants were actively moved around the virtual target objects to change their perspective on the objects by 90° (see Figure 4). On 50% of the trials, the participants began the trial with one perspective on the object and then were moved around the corner of the table to the neighboring side to afford a change in perspective by 90°. The experimenters simply pushed the chair. The participants viewed the virtual target object continuously during this transition. Such trials were chosen randomly, given the constraint that half the trials for each target be perspective change trials and half not. Each of the 20 target locations (4 locations on each of five objects) was touched once in a random order in each block for 6 blocks of reaches. Each participant performed 120 reaches plus calibration trials, yielding 30 blocks of data for each.



Figure 4. Illustration of the change in viewing positions used in Experiment 2.

Results and Discussion

The result was that performance improved with an increase in the number of 90° perspective changes that occurred, from zero to all four reaches to a given object. We computed the widths and depths and the aspect ratios for each sequence of four reaches to each object by each participant as before. Given the design, the targets were touched in a given sequence of four reaches with varying numbers of perspective switches from zero to four, with the proportions of switch trials being normally distributed across the five possibilities (i.e., zero to four), as is shown in Figure 5. We divided the D/W aspect ratios by the actual target D/W aspect ratios. Means and standard deviations of these normalized ratios were computed and are plotted in Figure 5 as a function of the number of switches. Means dropped from values of about 1.4 (i.e., about 40% overestimation) for zero or one switch only to values near 1 for two or more switches. Also, the standard deviations dropped by about half their value in parallel with the change in the means.

As in Experiment 1, we performed multiple regressions regressing actual widths and depths on reach widths and depths, with width versus depth as a categorical independent variable (coded as ± 1) and an interaction vector. Each width or depth could occur with zero, one, or two switches. We performed the analysis for each case and found that the perception of metric shape increased in accuracy with the increase in the number of switches.

The result for zero switches was significant [$F(3,169) = 13.1, p < .001, r^2 = .19$], and all three independent vari-





Figure 5. Results for judgments of aspect ratios in Experiment 2, plotted as a function of the number of 90° perspective switches for a given target object. Filled circles, mean normalized aspect ratios (i.e., judged aspect ratios divided by actual aspect ratios); open circles, standard deviation of normalized aspect ratios; filled triangles, proportion of objects judged with each number of switches. The line at 1 signifies accurate normalized aspect ratios. The proportion of trials with a given number of switches was normally distributed around a mean at two switches.

Figure 6. Mean judgments of width and depth in Experiment 2 plotted as a function of the actual size of the width or depth, with lines fit by a least squares regression. The top, middle, and bottom panels show results for judgments made with zero, one, or two perspective switches, respectively. Filled circles, width; filled squares, depth.



Figure 7. Mean judged aspect ratios in Experiment 2 plotted as a function of actual aspect ratios, with lines fit by a least square regression. The upper left panel shows means for objects judged with zero or one perspective switch. The upper right panel shows means for objects judged with two perspective switches. The lower panel shows means for objects judged with three or four perspective switches. D/W, depth/width.

ables were significant: actual width and depth (partial F = 27.1, p < .001), width versus depth (partial F = 10.3, p < .01), and the interaction (partial F = 10.0, p < .01). The slopes for widths and depths were different, as is shown in Figure 6 and Table 2, and the slope for depth was low (≈ 0.40), meaning that depths were poorly discriminated. The difference in slope was 1.16.

The result for one switch was significant $[F(3,322) = 31.6, p < .001, r^2 = .23]$, and again all three independent variables were significant: actual width and depth (partial F = 82.7, p < .001), width versus depth (partial F = 11.5, p < .001), and the interaction (partial F = 12.8, p < .001). The slopes for widths and depths were different, as is shown in Table 2 and Figure 6, and again the slope for depth was low (≈ 0.55). The difference in slope, however, was 0.72—that is, less than what it was for zero switches.

The result for two switches was significant [$F(3,157) = 25.8, p < .001, r^2 = .33$], and only the main effect for actual width and depth was significant (partial F = 74.5, p < .001). The slopes for widths and depths were not dif-

ferent, as is shown in Table 2 and Figure 6, and the slope for depth, in particular, was closer to 1.

We computed simple regressions of actual D/W on reach D/W for blocks with zero or one switch, two switches, and two or more switches. The results are shown in Figure 7 and Table 2. With fewer than two switches, performance

Table 2

	Width (W)		Depth (D)		
Rotation	Slope	r^2	Slope	r^2	
0	1.55	.45	0.38	.02	
1	1.28	.42	0.55	.07	
2	1.06	.41	1.37	.29	
	D/V	V			
	Slope	r^2			
0-1	1.53	.25			
2	1.06	.27			
≥ 2	1.05	.27			

remained inaccurate. With two switches or more, performance became accurate on average. This was the first time in any of our many experiments that we obtained performance reflecting accurate perception of metric shape.

The results were consistent with the affine hypothesis, which is that the problem in accurate perception of metric shape lies in relating the scales of distance in the frontoparallel and depth planes. Continuous perspective changes that exchange object dimensions, switching depth into width and vice versa, enable observers to relate the scales so as to be able to perceive the shape correctly.

EXPERIMENT 3

The results of Experiment 2 indicated that a perspective change of 90° enabled observers to perceive metric shape correctly. However, it remains unclear whether this simply entailed two discrete views or, instead, required a continuous perspective change. Ostensibly, simply viewing the major and minor axes of an elliptical object, each in a frontoparallel plane, would allow judgment of the aspect ratio. The obvious problem, however, is that an observer could be fooled and never know the difference. For instance, if, in the second discrete view, the original object was replaced by a different object with a different depth, the observer would never know, because perception of the depth viewed in depth is ambiguous. More to the point, an observer would not be able to detect a failure to rotate an object between two discrete views. With a continuously viewed transformation, a failure to rotate would be obvious, and a change in the shape of the object would result in a nonrigid transformation that an observer presumably could detect. To test the possibility that observers might be able to use discrete views to perform accurate judgments of metric shape, we used a nonrigid transformation to produce a failure to rotate. A nonrigid sliding of the surface texture (as if on a sleeve) over and around a cylindrical object yields a transformation very similar to rotation, without the object itself actually being moved. In fact, for a circular cylinder, this transformation is identical to rigid rotation.

The question was whether deception of this sort could be perceptually detected. The sliding sleeve transformation would yield a constant flow display-that is, a structure-from-motion display in which the optic flow at any given point in the image would be constant. Perotti, Todd, and Norman (1996) used constant flow displays to investigate whether human observers use information available over more than two frames-that is, greater than first-order flow. They found that observers did not and, thus, that observers could not distinguish nonrigid constant flow from rigid rotation. However, Perotti et al.'s (1996) displays were generated using orthographic projection (i.e., parallel perspective). Subsequently, Blair, Wickelgren, and Bingham (2001) investigated whether orthographic projection is a good model of perspective for other than small angle vision (>4° of visual angle; see also Börjesson & Lind, 1996; Eagle & Hogervorst, 1999; Hogervorst & Eagle, 2000). They found that observers

could reliably distinguish rigid rotation from constant flow displays subtending 8° or more with polar perspective. Therefore, for objects similar to a coffee cup within reach distance, observers should be able to detect an attempt to deceive them with a nonrigid nonrotation event.

Finally, perspective changes can occur either because the observer moves or because objects move. In Experiment 2, we tested the former. In the case of observer movement, the rigidity of objects in the surroundings might be expected, because object transformations are not geared to observer motions (except, perhaps, in the paranoid novels of Philip K. Dick!). In contrast, object movements of many kinds occur in the surroundings, including animate motions that are nonrigid. Hence, the ability to detect the relative rigidity of motion is likely to be an important part of metric shape perception in this case.

In Experiment 3, we tested (1) whether, indeed, a 90° perspective transformation (rigid rotation) would allow accurate perception of metric shape, (2) whether observers could detect an attempt to deceive them by substituting a nonrigid constant flow transformation for rigid rotation, (3) whether performance would be comparable for both passive object rotation and active rotation of the observer around the object, and (4) whether performance would fail with a discrete change between two views separated by either a rigid 90° rotation or a constant flow transformation equivalent to 90°. Note that in this last discrete views condition, when the objects are rigidly rotated, participants observe the short and long axes of the elliptical cylinders successively in a frontoparallel plane. If successive discrete views of the two axes are sufficient to allow accurate perception of metric shape, judgments in this condition should be accurate. If they are not, discrete views are not enough.

Method

Participants. Thirty adults 18–29 years of age participated in the experiment. A separate group of 10 participated in each of three conditions: passive continuous rotation (5 of them male, 5 female), active continuous rotation (4 of them male, 6 female), and discrete rotation (3 of them male, 7 female). The participants were paid \$7/h. All the participants had normal or corrected-to-normal eyesight (using contacts) and normal motor abilities. All were right-handed.

Procedure. Both the apparatus and the procedure were the same as those for Group 2 in Experiment 1, with the following changes. The table was used. Before entering the virtual environment, the participants were shown demonstrations illustrating the difference between rigid rotation and the sliding sleeve events. The participants were shown a wooden elliptical cylinder in a paper sleeve similar to the virtual cylinders viewed in the experiment. The paper sleeve was black with a random texture of white triangular patches. The cylinder was held up before the participants and rotated to illustrate rigid rotation. To illustrate nonrigid constant flow, the cylinder was held up and the paper sleeve was slid around the unmoving object. A circular cylinder was used to illustrate again these two transformations and to show that only rigid rotation occurred in that case. This was important. It was possible that the observers might try to distinguish nonrigid events by labeling events perceived to be rigid rotation of a circular cylinder as nonrigid. For the observers truly to be able to distinguish the nonrigid constant flow displays, they had to distinguish them from rigid rotations, including those of circular cylinders

On each trial, the target cylinder appeared on the virtual support surface 15 cm below eye height. The observer viewed the transformation and then performed the reach. After the stylus had been recorded in each of the four locations around the object, the participant verbally judged whether the object had rotated rigidly or exhibited nonrigid change. The experimenter recorded the judgment.

In the passive continuous rotation group, the participants viewed the object in the first perspective for 5 sec; then they watched while the object either rotated by 90° or exhibited the nonrigid change. Both transformations occurred over 5 sec, yielding a rotation rate for the rigid rotation of 18°/sec. In the active continuous rotation group, the participants sat in the wheeled chair and grabbed the table to wheel themselves around to the other edge of the table, with the assistance of the experimenter. The participants viewed the target object continuously while moving in the chair. For nonrigid constant flow displays, the object rotated with the observers so as to maintain the same orientation to the observers' eyes. The surface texture (the sleeve) stayed in place, and the object rotated underneath the surface texture. Thus, if the observers began by looking down the long axis of an elliptical cylinder, that axis simply rotated to track the observers as they moved around the table while the texture stayed in place. The procedure was otherwise the same as that for the passive rotation group. The procedure for the discrete rotation group was the same as that for the passive continuous rotation group, except that the object became invisible during the transformations. The object simply disappeared for 5 sec and then reappeared.

The participants reached to each of the five target objects six times each, three times with rigid rotation and three times with nonrigid constant flow, for a total of 30 trials following calibration trials. Objects and transformations were randomly ordered. On each trial, four locations were visited, for a total of 120 recorded placements of the handheld stylus, relative to the target objects.

Results and Discussion

We first will describe the results for the rigidity judgments and then will address the reaching results. The proportion of *rigid* judgments is shown in Figure 8 for each of the target objects and each type of display (rigid rotation vs. constant flow). Each of the three conditions is shown in a separate graph. We also computed d' for each target object (i.e., each aspect ratio) in each condition. The d's are shown in Table 3. We used the rule of thumb that a $d' \ge 1$ signified an ability to distinguish rigid rotation from constant flow correctly, whereas d' < 1 reflected an inability to do this. If the task was being done correctly, circular cylinders (aspect ratio = 1) should have been judged as rigid rotation in all cases.

In the passive continuous rotation condition, the participants were able to identify rigid rotation and constant flow correctly in all cases except for the circular cylinders (d' = 0.47), which were judged 69% of the time as exhibiting rigid rotation. Thus, the observers were well able to do this task correctly.

In the active continuous rotation condition, performance was comparable to that in the passive continuous rotation condition. Again, rigid rotation and constant flow were distinguished for all aspect ratios except that for the circular cylinders (d' = 0.75), which were judged 72% of the time as exhibiting rigid rotation.

In the passive discrete rotation condition, the observers were unable to identify constant flow and distinguish it from rigid rotation. All d's were close to 0. All aspect ratios were judged as rigid rotation—69% of the time, on average. Thus, discrete views were not enough to allow the observers to discriminate whether objects were actually rotated or not. (Note that discrete views were not static. The observers were allowed to move their heads from side to side by about 10 cm to generate optic flow while observing with stereovision.) Continuous rotations, both passive and active, did enable the observers to recognize actual rigid rotation and distinguish it from the nonrigid change.

Next, we will report the reach measure results for each condition. For the passive continuous rotation condition, we found that when target objects were judged as having rigidly rotated (especially when judged correctly), reach widths and depths were produced accurately, with the result that the D/W aspect ratios were also correct and accurate. When objects were judged to have moved nonrigidly, neither widths nor depths nor aspect ratios were produced accurately. First, we compared reach widths and depths for target objects judged to have rotated rigidly. As is shown in Figure 9 and Table 4, the slopes in this case were not different and were equal to about 0.80. We performed a multiple regression, regressing actual widths and depths on reach widths and depths, with width versus depth as a categorical independent variable (coded as ± 1 to test intercept difference) and an interaction vector (to test slope difference). The result was significant [F(3,324) = 28.8], $p < .001, r^2 = .21$], and only the actual width and depth factor was significant (partial F = 85.7, p < .001). Thus, the respective slopes and intercepts shown in the upper left panel of Figure 9 and in Table 4 were not different. As is shown in Table 4, the results were the same when we isolated objects judged correctly as exhibiting rigid rotation. The D/W aspect ratio yielded slopes near 1 in both cases also. When objects were judged to have moved nonrigidly, slopes for reach widths and depths were low (≈ 0.50), as was the slope for the D/W aspect ratio (= 0.69). We performed a multiple regression to compare slopes and intercepts for D/W aspect ratios for objects judged rigid versus nonrigid. On D/W ratios derived from reaches, we regressed actual D/W ratios, a categorical independent variable coding rigid versus nonrigid judgments as ± 1 , and an interaction vector. The result was significant $[F(3,296) = 52.1, p < .001, r^2 = .35]$, and all three independent variables were significant: actual D/W (partial F = 148.5, p < .001), the rigid versus nonrigid variable (partial F = 6.0, p < .02), and the interaction (partial F =6.7, p < .01). The slopes shown in Table 4, 1.07 versus 0.69, were significantly different. So, metric shape was judged accurately when rigid rotation was also accurately recognized and not otherwise.

We performed the same analyses for the active continuous rotation condition, and the results were essentially the same. When objects rotated rigidly, they were correctly judged to have done so, and as a result, widths, depths, and D/W aspect ratios were correctly produced by reaching, as is shown in Figure 10. When objects were judged to move nonrigidly, shape was not produced accurately. We performed a multiple regression, regressing actual widths and depths on reach widths and depths, with width versus depth as a categorical independent variable (coded as ± 1) and an interaction vector. The result was significant



Figure 8. Proportions of objects judged to be rigidly rotating in Experiment 3 by judgment condition and object aspect ratio, plotted separately for objects moving rigidly and nonrigidly. The top, middle, and bottom panels show results for the passive continuous movement condition, the active continuous movement condition, and the passive discrete movement condition, respectively. Black bars, judgments of rigid motions; white or gray bars, judgments of non-rigid motions; *D/W*, depth/width.

Table 3 Rigidity Judgment Results of Experiment 3					
		Г	arget Obje	et	
Measure	1	2	3	4	5
	Pa	ssive Conti	nuous		
Aspect ratio	0.56	0.70	1.00	1.44	1.78
d'	2.78	3.63	0.47	2.39	2.94
	А	ctive Conti	nuous		
Aspect ratio	0.56	0.70	1.00	1.44	1.78
d'	2.82	1.84	0.75	1.97	2.82
	Ι	Passive Dis	crete		
Aspect ratio	0.56	0.70	1.00	1.44	1.78
d'	-0.09	0.00	0.00	-0.09	0.09



Figure 9. Data from the passive continuous movement condition in Experiment 3. Lefthand panels: Mean judged widths and depths plotted as a function of the actual size of the widths and depths. The top panel shows the means for objects judged to be moving rigidly. The bottom panel shows the means for objects judged to be moving nonrigidly. Right-hand panel: Mean judged aspect ratios plotted as a function of actual aspect ratios, with lines fitted by a least squares regression. Filled circles, objects judged as moving rigidly; filled squares, objects judged as moving nonrigidly; D/W, depth/width.

Table 4 Reach Measure Results of Experiment 3						
	Width (W)		Depth (D)		D/W	
	Slope	r^2	Slope	r^2	Slope	r^2
Passi	ve Contir	uous I	Rotation			
Judged rigid	0.83	.30	0.73	.15	1.07	.44
Correctly judged rigid	0.83	.33	0.76	.16	1.09	.49
Judged nonrigid	0.60	.21	0.42	.05	0.69	.22
Correctly judged nonrigid	0.61	.23	0.32	.03	0.66	.23
Activ	e Contin	uous R	lotation			
Judged rigid	0.79	.15	0.87	.16	0.87	.41
Correctly judged rigid	0.95	.24	0.93	.19	1.05	.56
Judged nonrigid	0.47	.07	0.27	.02	0.37	.05
Correctly judged nonrigid	0.36	.04	0.27	.02	0.36	.05
Pas	sive Disc	rete Ro	otation			
Judged rigid	0.95	.31	0.51	.08	0.64	.28
Correctly judged rigid	0.93	.28	0.49	.06	0.71	.31
Judged nonrigid	0.65	.22	0.59	.14	0.87	.35
Correctly judged nonrigid	0.42	.09	0.50	.08	0.71	.38



Figure 10. Data from the active continuous movement condition in Experiment 3. Lefthand panels: Mean judged widths and depths plotted as a function of the actual size of the widths and depths. The top panel shows the means for objects judged to be moving rigidly. The bottom panel shows the means for objects judged to be moving nonrigidly. Right-hand panel: Mean judged aspect ratios plotted as a function of actual aspect ratios, with lines fitted by a least squares regression. Filled circles, objects judged as moving rigidly; filled squares, objects judged as moving nonrigidly; *D/W*, depth/width.

 $[F(3,360) = 22.4, p < .001, r^2 = .16]$, and only the actual width and depth factor was significant (partial F = 66.2, p < .001). Slopes and intercepts were not different. Especially when object motions were judged correctly, slopes of reach widths (= 0.95), depths (= 0.93), and the D/Waspect ratios (= 1.05) were near 1. When objects were judged to have moved nonrigidly, the corresponding slopes were low: widths, 0.36; depths, 0.27; and aspect ratios, 0.36. On D/W ratios derived from reaches, we regressed actual D/W, a categorical independent variable coding rigid versus nonrigid judgments as ± 1 , and an interaction vector. The result was significant [F(3,294) = 27.9, p <.001, $r^2 = .22$], and all three independent variables were significant: actual D/W (partial F = 64.8, p < .001), the rigid versus nonrigid variable (partial F = 9.9, p < .002), and the interaction (partial F = 9.6, p < .005). The slopes for objects judged rigid versus nonrigid were different. Thus, performance was accurate and correct when target objects rotated rigidly, but not when they merely exhibited nonrigid constant flow.

Finally, we analyzed the data from the passive discrete rotation condition. In this case, the observers judged most of the target objects to have rotated rigidly. Even when they happen to have judged rigid rotations correctly, they were unable to produce accurate depths or D/W aspect ratios, as is shown in Figure 11. The results indicate that discrete 3-D views of widths and depths (each appearing in a frontoparallel plane) do not allow accurate apprehension of metric shape. We performed a multiple regression on trials judged as rigid rotation, regressing actual widths and depths on reach widths and depths, with width versus depth as a categorical independent variable (coded as ± 1) and an interaction vector. The result was significant [F(3,406) = 31.0, p < .001, $r^2 = .19$], and all three independent variables were significant: actual widths and depths (partial F = 84.5, p <.001), width versus depth (partial F = 7.3, p < .01), and the interaction (partial F = 7.4, p < .01). Thus, the slope for depths (= 0.51) was different from that for widths (= 0.95). Widths were produced accurately,



Figure 11. Data from the passive discrete movement condition in Experiment 3. Left-hand panels: Mean judged widths and depths plotted as function of actual size of the widths and depths. The top panel shows the means for objects judged correctly to be moving rigidly. The bottom panel shows the means for objects judged correctly to be moving nonrigidly. Right-hand panel: Mean judged aspect ratios plotted as a function of actual aspect ratios, with lines fitted by a least squares regression. Filled circles, objects judged correctly as moving rigidly; filled squares, objects judged correctly as moving nonrigidly; *D/W*, depth/width.

but not depths, even when rigid rotation was judged correctly. When rigid rotation and nonrigid constant flow were each judged correctly, the results for D/W aspect ratios were the same. The slopes were the same and low (= 0.71 in both cases). Discrete views failed to allow accurate perception of metric shape.

EXPERIMENT 4

The results of Experiment 3 showed clearly the ineffectiveness of discrete views for accurate shape perception. The participants produced shape aspect ratios inaccurately even when they correctly judged objects to have been rigidly rotated by 90°. In contrast, when the observers viewed either passively or actively generated continuous 90° rigid rotations, they produced accurate shape aspect ratios. The potential effectiveness of a full 90° rotation that takes the depth axis into a frontoparallel plane was suggested by an affine analysis of shape perception. However, the question remained whether, in fact, a full 90° rotation is required to yield sufficient continuous perspective change to allow accurate perception of metric shape. In Experiment 4, we parametrically varied the amount of rotation to investigate how much of a perspective change is required to yield accurate shape perception. The free head movements allowed in the previous experiments effectively tested perspective changes of about 12°, showing that these were inadequate to yield accurate shape. We now tested rotations by 30°, 45°, and 60°, as well as 90°.

Method

Participants. Thirty adults 18–28 years of age participated in the experiment. A separate group of 10 participated in each of three conditions: 30° rotation (4 of them male, 6 female), 45° rotation (3 of them male, 7 female), and 60° rotation (5 of them male, 5 female). The participants were paid \$7/h. All the participants had normal or corrected-to-normal eyesight (using contacts) and normal motor abilities. All were right-handed.

Procedure. Both the apparatus and the procedure were the same as those for Group 2 in Experiment 1, with the following changes. Only rigid rotations of the five target objects were tested

at a 20-cm eye height. In each case, the object appeared in canonical orientation (for instance, for aspect ratios greater than 1, the long axis of the object appeared parallel to the line of sight) and then, after the object had been viewed by the observer for 5 sec, it was rotated rigidly to the left through the given number of degrees, stopped, and then rotated back to the canonical orientation. Then the participants would perform a reach to touch the four locations in the assigned random order. All three groups were tested so as to replicate the previous results with 90° rotations: Group 1 did 30°, then 90°; Group 2 did 45°, then 90°; and Group 3 did 60°, then 90°. Each aspect ratio was tested three times in a random order for each rotation amount. Each participant reached to perform 120 recorded touches.

Results and Discussion

The results for all three conditions are shown in Table 5 and Figure 12, where the following can be seen: (1) the correct aspect ratios were consistently replicated in all three groups with 90° rotations, (2) 30° rotations did not yield production of accurate aspect ratios, but (3) rotations by 45° and 60° did yield good aspect ratios. When

Table 5Results of Experiment 4					
	Group				
30°	45°	60°			
y = 0.67x + 0.30 $r^2 = .33$	y = 1.10x - 0.11 $r^2 = .50$	y = 1.07x - 0.17 $r^2 = .51$			
Replication With 90° Rotation					
y = 1.07x - 0.02 $r^2 = .41$	y = 1.14x - 0.23 $r^2 = .65$	y = 1.00x - 0.07 $r^2 = .50$			

mean aspect ratios derived from reaches were compared with actual aspect ratios in simple regressions, slopes near 1.00 resulted for rotations of 45°, 60°, and 90°, but not for 30°, which yielded a slope of 0.67. For the data from each group, we performed a multiple regression on D/W aspect ratios derived from reach data, regressing on them the actual D/Ws, a categorical coding (±1) of the amount of rotation (i.e., 30°, 45°, or 60° vs. 90°), and an interaction vector.



Figure 12. Data from Experiment 4: Mean aspect ratios plotted as a function of actual aspect ratios, with standard error bars and lines fitted by a least squares regression. The top left panel shows results for the 30° rotation group. The top right panel shows results for the 45° rotation group. The bottom panel shows results for the 60° rotation group. Filled squares, 90° rotation in all cases; filled circles, 30° , 45° , or 60° rotation; D/W, depth/width.

For the 30° rotation condition, the multiple regression was significant [$F(3,296) = 61.4, p < .001, r^2 = .38$], and two of the independent variables were significant: actual D/W (partial F = 170.9, p < .001) and the interaction (partial F = 8.6, p < .005). The amount of rotation was marginal (p = .055). As is shown in Table 5, the slope for 90° rotation was 1.07, whereas for 30° rotation it was 0.67, and these were significantly different.

For the 45° rotation condition, the analysis was significant [F(3,296) = 128.3, p < .001, $r^2 = .57$], but only the actual D/W variable was significant (partial F = 382.4, p < .001). As is shown in Table 5, the slope for 90° rotation was 1.14, and for 45° rotation, it was 1.10. These were not significantly different.

For the 60° rotation condition, the analysis was significant [F(3,296) = 96.7, p < .001, $r^2 = .50$], and only the actual D/W variable was significant (partial F = 289.6, p < .001). As is shown in Table 5, the slope for 90° rotation was 1.00, and for 60° rotation, it was 1.07. These were not significantly different.

Finally, we also analyzed the precision in judging widths and depths. For each participant in each condition, we computed the coefficient of variation (CV) for depth and for width (i.e., the standard deviation divided by the mean of the three repetition trials). We found that mean CVs for 45°, 60°, and 90° rotations were about 22%, whereas those for 15° and 30° rotations were about 29%. The mean size of the objects was 5 cm, so a CV of 22% is about 1 cm. This means that 70% of the reaches to a given locus (left, right, front, or back) varied only within about 0.5 cm on average, and 96% were within 1 cm. Ultimately, the problem was not precision but accuracy.

For each group, we performed a repeated measures ANOVA with rotation amount (30°, 45°, or 60° vs. 90°), width versus depth, and size (small, medium, or large) as factors. For the 60° group, rotation was not significant (p > .4), and the mean CVs were 24% and 22% for 60° and 90° rotations, respectively. Only width versus depth was significant [F(1,18) = 7.8, p < .02], and the means were width = 20% and depth = 26%. For the 45° group, rotation was not significant (p > .6), and the mean CVs were 27% and 26% for 45° and 90° rotations, respectively. (Two of the 10 participants were exceptionally variable in their reaching. When the analysis was performed on the remaining 8 participants, the results were the same, but mean CVs were 23% and 23%, respectively.) Only size was significant [F(2,36) = 5.4, p < .01], and the means were 30%, 29%, and 20% for small, medium, and large, respectively. For the 30° group, rotation was significant [F(1,18) = 4.5, p < .05], and the mean CVs were 29% and 21% for 30° and 90° rotations, respectively. No other factors were significant. Finally, we compared CVs for the 30° rotation with the CVs from Experiment 1 (i.e., for about 15° rotation). The results were the same for the two eye height conditions in Experiment 1. Rotation was not significant (p > .6). The mean CVs were 28% and 29% for 15° and 30° rotations, respectively.

The obvious conclusion from these results was that a 90° rotation is not required for accurate and precise perception of metric shape. Rather, a rotation by at least 45° is both sufficient and necessary.

GENERAL DISCUSSION

As was reviewed in the introduction to this article, a very large number of studies have shown that the modest perspective variations afforded by stereovision and the motion parallax or structure-from-motion information generated by small head movements do not allow human observers to perceive metric shape accurately. This has been shown in various judgment studies and in studies using relevant action measures, including targeted reaching and reach-to-grasp movements. The latter have shown that grasping is often accurate, not because people are able to perceive metric shape properties, but because disparity matching is afforded by stereovision and is used to guide the hand and fingers to targeted object surfaces during the final phases of the reach-to-grasp movement.

In the present set of experiments, we started, in Experiment 1, by replicating poor metric shape perception results in the context of a targeted reaching task, the same task as that used in the subsequent Experiments 2–4. We tested two versions of the task and found that the inaccuracy and imprecision were not produced by motor variability. The results were the same when the participants performed a separate reach (i.e., a large movement) to each targeted locus on a given object (front, back, left, and right sides) or a single reach to touch the four locations in sequence.

The question that remained was whether it is ever possible for human observers to perceive metric shape. If observers are ever going to be able to do this, they should be able when they can compare object depth with width by rotating an object by 90° to bring the axis first observed in depth into a frontoparallel plane. The problem made evident in previous studies and shown in Experiment 1 is that perception of the extent of an object in depth is poor. So, bringing that extent up into a frontoparallel plane should afford good apprehension of its extent in relation to the object extent along a perpendicular axis (i.e., the width) also observed in a frontoparallel plane. We tested this in Experiment 2, where we also varied the frequency of such large perspective change information as participants reached to touch each of the targeted locations on a given object encountered on different occasions. The results showed that the 90° perspective change enabled participants to perceived metric shape correctly, and the greater the frequency in the given task, the better their performance.

Another question was raised by these results, however. Would discrete views from two perspectives differing by 90° be sufficient to yield perception of metric shape? The problem incurred by the inability to perceive metric shape (or depth) in a single view (or with small perspective variations) is that observers would not be able to recognize reliably whether the perspective had, in fact, changed in a second view or, relatedly, whether an object of different shape was substituted in a second view. In either of these instances, observers certainly would not perceive metric shape correctly. The point is that there is ambiguity and uncertainty entailed necessarily in any attempt to perceive metric shape via two discrete views. We designed Experiment 3 to preserve that uncertainty and to remove any artifactual certainty that might be imported by a contract,

implicit or explicit, between the experimenter and the participants-namely, that the experimenter would not try to fool the participants. The participants experienced two types of trial-namely, trials on which the target object rigidly rotated by 90° and trials on which it did not. In the latter case, a nonrigid sliding of the surface texture over the rigid object occurred. With continuous viewing, the participants had the ability to discriminate the two types of events. With discrete views, the observers remained uncertain about the change in perspective. Nevertheless, on trials on which the object rigidly rotated by 90°, the participants were able to see the target object from two different perspectives differing by 90°. The participants were required to judge whether a rigid rotation had occurred. They did this reliably in continuous-viewing conditions, but not in discrete-viewing conditions. Nevertheless, they did get it right sometimes in the latter case. Even when they judged correctly that a rigid rotation by 90° had occurred, the observers were unable to perceive metric shape correctly from two discrete views differing by 90°.

This is an important result. It means that continuous perspective changes are required for accurate perception of metric shape. Discrete views of objects are used in many studies of object perception and, in particular, in studies of object recognition. This result makes it clear that metric shape cannot be perceived reliably under such conditions even when views differing by 90° are shown but that there *are* conditions under which metric shape *can* be perceived. Those conditions entail the viewing of continuous perspective changes that enable one to perceive that the motion is rigid. This, in turn, requires relatively wide angle (>8°) polar perspective, in addition to large changes in perspective.

Finally, having investigated the most obvious means by which people might successfully perceive metric shape, we next investigated whether continuous perspective variations less than 90° but larger than the 15° would allow per-



Figure 13. An elliptical cylinder is oriented with the long axis of the elliptical cross section viewed in depth. If the cylinder is then rotated by 45° around a vertical axis through its center, the lengths of the two principal axes in the object can be directly compared, because $\sin(45^{\circ}) = \cos(45^{\circ})$.

ception of metric shape. We hypothesized that a 45° change would be both sufficient and necessary. A special symmetry, as illustrated in Figure 13, is yielded by 45°-namely, $sin(45^\circ) = cos(45^\circ)$. Given this, when the major and minor axes of an elliptical cylinder are rotated by 45° from canonical orientation, each projects an equivalent component into the optics, and therefore, the two can be directly related and compared. When any two perpendicular axes through an object are viewed at orientations of 45° to either side of the gaze axis and then rotated by 45° so that one becomes the depth axis (parallel to the gaze axis) and the other becomes the width axis (lying in a frontoparallel plane), the aspect ratio of those two axes should be specified to the observer. In Experiment 4, we tested this, in particular, for the major and minor axes of elliptical cylinders, which were viewed first in canonical orientation (i.e., one of the axes was aligned with the gaze axis); then the object was rotated by 45° out and then 45° back to canonical orientation. We tested rotations of 30°, 45°, and 60°, each with a different group of participants, and then, to be sure that the groups were otherwise comparable, we also tested rotations of 90° with all three groups. The result was that rotation by 30° yielded performance that was as poor as that found with small perspective changes in Experiment 1. Performance levels with rotations of 45°, 60°, and 90° were equally good and correct, as well as precise. The conclusion was that our hypothesis was supported. Rotation by at least 45° (or more) is both sufficient and necessary for accurate and reliable perception of metric shape.

Looking back at the literature, a related result can be found in Brenner and van Damme (1999), although the authors apparently did not quite recognize it for what it was. Brenner and van Damme tested metric shape perception in three viewing conditions, all of which included stereovision: no rotation, rotation by 30°, and rotation by 60° . Judgments were inaccurate in the first two conditions but accurate in the last one, 60° rotation (see Brenner & van Damme, 1999, p. 980, Figure 5A).

It remains possible that gradual improvements in judgments of metric shape might be found with increasing amounts of rotation between 30° and 45° . Although we would not expect this, we are currently investigating the possibility. The main conclusion of the present work is that perception of metric shape is possible with perspective variations that are somewhat larger than those generally studied heretofore—namely, 45° .

Of what practical value might this information be? Perspective transformations of 45° would commonly occur as one enters a work space—that is, coming into the kitchen or the office in the morning and sitting down to work at a table or a desk, respectively. The question is, How stable would the perceptions enabled by these transformations be? Would this initial experience act to calibrate subsequent perception and action within the given work space? It remains to be seen. In any case, visual space cannot be entirely affine.

AUTHOR NOTE

This work was supported in part by Grant R01 EY11741-01A3 from the National Eye Institute. Please send all correspondence to G. P. Bingham, Department of Psychological and Brain Sciences, Indiana University, 1101 East Tenth Street, Bloomington, IN 47405-7007 (e-mail: gbingham@indiana.edu).

REFERENCES

- BINGHAM, G. P. (2005). Calibration of distance and size does not calibrate shape information: Comparison of dynamic monocular and static and dynamic binocular vision. *Ecological Psychology*, **17**, 55-74.
- BINGHAM, G. P., BRADLEY, A., BAILEY, M., & VINNER, R. (2001). Accommodation, occlusion and disparity matching are used to guide reaching: A comparison of actual versus virtual environments. *Journal of Experimental Psychology: Human Perception & Performance*, 27, 1314-1344.
- BINGHAM, G. P., COATS, R., & MON-WILLIAMS, M. (2007). Unnatural prehension to virtual objects is not inevitable if calibration is allowed. *Neuropsychologia*, 45, 288-294.
- BINGHAM, G. P., CROWELL, J. A., & TODD, J. T. (2004). Distortions of distance and shape are not produced by a single continuous transformation of reach space. *Perception & Psychophysics*, 66, 152-169.
- BINGHAM, G. P., ZAAL, F., ROBIN, D., & SHULL, J. A. (2000). Distortions in definite distance and shape perception as measured by reaching without and with haptic feedback. *Journal of Experimental Psychol*ogy: Human Perception & Performance, 26, 1436-1460.
- BLAIR, N., WICKELGREN, E., & BINGHAM, G. P. (2001, June). How representative is parallel projection of structure-from-motion? Paper presented at the 11th International Conference on Perception and Action, Storrs, CT.
- BÖRJESSON, E., & LIND, M. (1996). The effect of polar projection on the perception of euclidean structure from motion. *Perception & Psychophysics*, 58, 871-882.
- BRADSHAW, M. F., & ELLIOT, K. M. (2003). The role of binocular information in the 'on-line' control of prehension. *Spatial Vision*, 16, 295-309.
- BRADSHAW, M. F., ELLIOT, K. M., WATT, S. J., HIBBARD, P. B., DA-VIES, I. R. L., & SIMPSON, P. J. (2004). Binocular cues and the control of prehension. *Spatial Vision*, **17**, 95-110.
- BRENNER, E., & VAN DAMME, W. J. M. (1999). Perceived distance, shape and size. *Vision Research*, **39**, 975-986.
- COATS, R., BINGHAM, G. P., & MON-WILLIAMS, M. (in press). Calibrating grasp size and reach distance: Interactions reveal integral organization in reaching-to-grasp movements. *Experimental Brain Research*.
- CROWELL, J. A., TODD, J. T., & BINGHAM, G. P. (2000, November). Distinct visuomotor transformations for visually guided reaching. Paper presented at the 41st Annual Meeting of the Psychonomic Society, New Orleans.
- CROWELL, J. A., TODD, J. T., & BINGHAM, G. P. (2001, May). Distinct perceptual representations for visually guided reaches. Paper presented at the annual meeting of the Vision Sciences Society, Sarasota, FL.
- CUIJPERS, R. H., SMEETS, J. B. J., & BRENNER, E. (2004). On the relation between object shape and grasping kinematics. *Journal of Neu*rophysiology, 91, 2598-2606.
- EAGLE, R. A., & HOGERVORST, M. A. (1999). The role of perspective information in the recovery of 3D structure-from-motion. *Vision Re*search, **39**, 1713-1722.
- HIBBARD, P. B., & BRADSHAW, M. F. (2003). Reaching for virtual objects: Binocular disparity and the control of prehension. *Experimental Brain Research*, 148, 196-201.
- HOGERVORST, M. A., & EAGLE, R. A. (2000). The role of perspective effects and accelerations in perceived three-dimensional structure-frommotion. *Journal of Experimental Psychology: Human Perception & Performance*, 26, 934-955.

KOENDERINK, J. J. (1990). Solid shape. Cambridge, MA: MIT Press.

- KOENDERINK, J. J., & VAN DOORN, A. J. (1991). Affine structure from motion. *Journal of the Optical Society of America A*, 9, 377-385.
- LAPPIN, J. S., & AHLSTRÖM, U. B. (1994). On the scaling of visual space from motion—in response to Pizlo and Salach-Golyska. *Perception & Psychophysics*, 55, 235-242.
- LEE, Y.-L., CRABTREE, C. E., NORMAN, J. F., & BINGHAM, G. P. (2008). Poor shape perception is the reason reaches-to-grasp are visually guided online. Manuscript submitted for publication.
- LIND, M., BINGHAM, G. P., & FORSELL, C. (2003). Metric 3D structure in visualizations. *Information Visualization*, 2, 51-57.
- LOOMIS, J. M., PHILBECK, J. W., & ZAHORIK, P. (2002). Dissociation between location and shape in visual space. *Journal of Experimental Psychology: Human Perception & Performance*, 28, 1202-1212.
- MELMOTH, D. R., & GRANT, S. (2006). Advantages of binocular vision for the control of reaching and grasping. *Experimental Brain Research*, **171**, 371-388.
- MON-WILLIAMS, M., & BINGHAM, G. P. (2007). Calibrating reach distance to visual targets. *Journal of Experimental Psychology: Human Perception & Performance*, 33, 645-656.
- MON-WILLIAMS, M., COATS, R., & BINGHAM, G. P. (2004). Reaching with feeling [Abstract]. *Journal of Vision*, 4, 411a.
- PAGANO, C. C., & BINGHAM, G. P. (1998). Comparing measures of monocular distance perception: Verbal reaching errors are not correlated. *Journal of Experimental Psychology: Human Perception & Performance*, 24, 1037-1051.
- PEDHAZUR, E. J. (1982). Multiple regression in behavioral research: Explanation and prediction (2nd ed.). New York: Holt, Rinehart & Winston.
- PEROTTI, V. J., TODD, J. T., LAPPIN, J. S., & PHILLIPS, F. (1998). The perception of surface curvature from optical motion. *Perception & Psychophysics*, **60**, 377-388.
- PEROTTI, V. J., TODD, J. T., & NORMAN, J. F. (1996). The visual perception of rigid motion from constant flow fields. *Perception & Psychophysics*, 58, 666-679.
- SCARFE, P., & HIBBARD, P. B. (2006). Disparity-defined objects moving in depth do not elicit three-dimensional shape constancy. *Vision Research*, 46, 1599-1610.
- TITTLE, J. S., TODD, J. T., PEROTTI, V. J., & NORMAN, J. F. (1995). Systematic distortion of perceived three-dimensional structure from motion and binocular stereopsis. *Journal of Experimental Psychology: Human Perception & Performance*, 21, 663-678.
- TODD, J. T., & BRESSAN, P. (1990). The perception of 3-dimensional affine structure from minimal apparent motion sequences. *Perception* & *Psychophysics*, 48, 419-430.
- TODD, J. T., & NORMAN, J. F. (1991). The visual perception of smoothly curved surfaces from minimal apparent motion sequences. *Perception* & *Psychophysics*, **50**, 509-523.
- TODD, J. T., TITTLE, J. S., & NORMAN, J. F. (1995). Distortions of threedimensional space in the perceptual analysis of motion and stereo. *Perception*, 24, 75-86.
- WATT, S. J., & BRADSHAW, M. F. (2000). Binocular cues are important in controlling the grasp but not the reach in natural prehension movements. *Neuropsychologia*, 38, 1473-1481.
- WATT, S. J., & BRADSHAW, M. F. (2003). The visual control of reaching and grasping: Binocular disparity and motion parallax. *Journal of Experimental Psychology: Human Perception & Performance*, 29, 404-415.

(Manuscript received December 7, 2006; revision accepted for publication November 7, 2007.)