Exp Brain Res (2010) 204:559–573
DOI 10.1007/s00221-010-2323-2

RESEARCH ARTICLE

Large perspective changes yield perception of metric shape that allows accurate feedforward reaches-to-grasp and it persists after the optic flow has stopped!

Young-Lim Lee · Geoffrey P. Bingham

Received: 30 November 2009 / Accepted: 2 June 2010 / Published online: 19 June 2010

Abstract Lee et al. (Percept Psychophys 70:1032–1046, 2008a) investigated whether visual perception of metric shape could be calibrated when used to guide feedforward reaches-to-grasp. It could not. Seated participants viewed target objects (elliptical cylinders) in normal lighting using stereo vision and free head movements that allowed small ($\approx 10°$) perspective changes. The authors concluded that poor perception of metric shape was the reason reaches-to-grasp should be visually guided online. However, Bingham and Lind (Percept Psychophys 70:524–540, 2008) showed that large perspective changes ($\geq 45°$) yield good perception of metric shape. So, now we repeated the Lee et al.'s study with the addition of information from large perspective changes. The results were accurate feedforward reaches-to-grasp reflecting accurate perception of both metric shape and metric size. Large perspective changes occur when one locomotes into a workspace in which reaches-to-grasp are subsequently performed. Does the resulting perception of metric shape persist after the large perspective changes have ceased? Experiments 2 and 3 tested reaches-to-grasp with delays (Exp. 2, 5-s delay; Exp. 3, ≈ 16-s delay) and multiple objects to be grasped after a single viewing. Perception of metric shape and metric size persisted yielding accurate reaches-to-grasp. We advocate the study of nested actions using a dynamic approach to perception/action.

Keywords Reach-to-grasp · Shape perception · Metric shape · Visual perception · Perception/action

Introduction

We investigated perception of metric shape. Geometry contains no formal notion of 'metric shape' as such. There are metric geometries (e.g. Euclidean geometry), but such geometries do not capture what we generally mean by shape. The general intuition of shape is that a large and a small circle share the same shape. They are both circles. Likewise, large and small squares are both squares. The shape is invariant although the size changes. Similarity geometry captures what we generally mean by shape. It allows isotropic changes in size that preserve aspect ratios between, for instance, the width and the height of a circle or of an ellipse of a given eccentricity. An aspect ratio can be used to measure this characteristic of shape. In this paper, this is what we mean by 'metric shape'. The metric in question is the aspect ratio.[1]

A large number of shape perception studies have shown that metric 3-D shape cannot be perceived accurately and

[1] In other recent work [for instance, Perotti et al. (1998) or Norman et al. (2006)], local measures of shape have been used, that is, measures of shape at a point on a curve or surface. Koenderink (1990, pp 319–324) developed both a measure of qualitative variations (cylindrical, ellipsoidal, saddle, etc.) in surface shape at a point, the "shape index", and a measure of quantitative variations in surface curvature at a point, the "curvedness". Such local measures are different from more global measures like the aspect ratio. A large and a small circle exhibit different 'curvedness' by definition because curvature is the inverse of the ratio of the tangent circle to a curve at a point. So, the local measure captures a different property of shape than that which corresponds to the more familiar intuition about what shape is.

Y.-L. Lee · G. P. Bingham
Indiana University, Bloomington, IN, USA

G. P. Bingham (✉)
Department of Psychology, Indiana University,
1101 East Tenth Street, Bloomington,
IN 47405-7007, USA
e-mail: gbingham@indiana.edu

that the perception of object shape is distorted [see Todd et al.'s (1995) review]. The recovery of 3-D structure has been found to be inaccurate in studies of structure-from-motion (Norman and Lappin 1992; Norman and Todd 1993; Perotti et al. 1998; Tittle et al. 1995; Todd and Bressan 1990; Todd and Norman 1991), structure from binocular or disparity (Johnston 1991; Tittle et al. 1995), structure from the combination of binocular disparity and motion (Tittle and Braunstein 1993; Tittle et al. 1995), and even structure under full cue conditions including lambertian shading, texture, and specular highlights (Norman and Todd 1996; Norman et al. 1995). The authors of these studies have argued both that metric 3-D shape cannot be perceived accurately and that the perception of the relative depth of objects is systematically distorted. For instance, Johnston (1991) investigated perceived 3-D structure from binocular disparity and showed that shapes in depth tend to be systematically compressed at larger distances and systematically expanded at shorter distances. In other words, truly circular cylinders appeared flattened at a far distance and elongated at a near distance. Similar evidence for distortions in the perception of 3-D structure has been obtained in structure-from-motion studies. Tittle et al. (1995) asked observers to adjust the eccentricity of a motion-defined cylinder until its cross-section appeared circular. The observers consistently adjusted the cylinder so that its shape was compressed in depth. Because this compressed surface appeared to have a circular cross-section, this suggested that an expansion in depth occurs during perception.

Problems exist, however, with the idea that perception simply imposes a distortion on perceived objects and space. For instance, if a circular cylinder in a structure-from-motion display was perceived to be stretched in depth (i.e., as a cylinder with an elongated elliptical cross-section), then, if this cylinder was rotated by 90° around its symmetry axis, the rotation should be perceived as non-rigid because the cylinder would still be perceived to be elongated in depth after rotation. However, such rotations are typically perceived as rigid despite the perceived constant expansion in depth and implied non-rigid shape change. In this context, Todd and Bressan (1990) suggested that only affine structure can be recovered from optical information. They found that judgments of affine structure were accurate, whereas judgments of Euclidean structure were not. The affine hypothesis was motivated by a theory of the "geometry of visual space" that had been developed by the authors and their collaborators (namely, J. Koenderink and A. van Doorn). According to the theory, the geometry of visual space entails a mapping from physical space (which is Euclidean) to a hypothesized perceptual space. The perceptual space is anticipated to represent space in a geometry that is other than Euclidean. The geometry of visual space has been hypothesized to be Elliptical (Indow et al. 1962a, b; Luneburg 1950), Hyperbolic and now Affine. In the nineteenth century, Klein categorized geometries into a hierarchy based on how different transformations affect different aspects of object structure. As one descends the hierarchy, one goes from Euclidean to similarity to affine geometry and so on. Any geometric transformation of an object changes some of its structural properties, while preserving others. Euclidean geometry allows only rigid rotation and translation, which preserves the size and shape of an object. Similarity geometry also allows the isotropic scaling of size, which preserves only the shape of the object. Affine geometry allows the scaling of size differently in different directions, that is, along frontoparallel directions and the line of sight, respectively. This preserves the shape of the object only with respect to "relief" (or relative depth structure). Accordingly, two objects related by a stretching transformation along the line of sight would not be discriminated. Two such objects are referred as being "affine equivalent along the line of sight".

The affine theory of visual space assumes that inaccuracies in shape perception reflect a general problem in perceiving space. This was challenged in a study by Bingham et al. (2004) who found that inaccuracies in location (distance and direction) perception and in shape perception, respectively, were not produced by a single continuous transformation of perceived space. That is, position perception and shape perception reflected different distortions or ambiguities. Thus, the structure of visual space cannot simply reflect a single coherent affine geometry. Instead, Bingham et al. suggested that perception is a function of the available information. Accordingly, inaccuracies in the perception of metric shape would reflect the nature of the information about shape.

The question is whether inaccuracies in metric shape perception are in fact yielded by systematic distortion. Other results suggest that perceived depths might be merely ambiguous rather than systematically distorted. Todd and Norman (2003) reported that depth judgments were highly variable and inconsistent as viewing distance, orientation or response task were changed. In a similar vein, Lind et al. (2003) tested the perception of the shape of cylindrical objects as viewing height and distance were varied. The observers adjusted the shape of an elliptical outline on a computer screen to be the same as the shape of the cross-section of the elliptical cylinders (which all exhibited different depth-to-width aspect ratios). The results differed across experiments and varied as a function of individual differences, differences in the range of shapes, and differences in the experimental tasks. If metric shape is actually ambiguous, then the systematicity of the distortions found in many of the previously reviewed studies may have been produced by contextual factors. Lee et al. (2007) investigated this hypothesis under full cue conditions, manipulating the relevant context by using two different ranges of object aspect ratios and two different tasks. First, in a 2-D task, objects were viewed so that the elliptical cross-section was frontoparallel. Then, in a subsequent 3-D task, objects were viewed so that the elliptical cross-section was in depth. Observers had to adjust the aspect ratio of an ellipse on a computer screen to be same as the cross-section of the target object. The results showed: (1) that judgments of 2-D aspect ratios were accurate, and (2) that judgments of 3-D shape changed as a function of the context, namely the range of 2-D aspect ratios that were judged before the 3-D judgment task.

This finding suggested a solution to a puzzle. The puzzle was how accurate reaches-to-grasp might be performed despite inaccuracies found in shape perception studies. The solution might be that feedback information from grasping is used to calibrate the perception of shape aspect ratios. That is, feedback information might play the role of contextual variable both to resolve and constrain the ambiguity of perceived shape and to make it accurate for perceptual control of reaches-to-grasp. Previous studies showed that accurate performance of visually guided reaching and grasping requires use of feedback information to calibrate both reaches and grasps. Bingham and Mon-Williams and colleagues found that haptic feedback corrected inaccuracy and instability of both reached distances and grasped object sizes during feedforward reach-to-grasp tasks (Bingham 2005; Bingham et al. 2000, 2007; Coats et al. 2008; Mon-Williams and Bingham 2007).

The evidence shows that both object distance and object size can be calibrated to allow accurate perceptual guidance of reaches-to-grasp. However, this would not be sufficient for accurate grasping in most cases, namely, those that involve a grip axis [or opposition axis (Iberall et al. 1986)] through the object that does not lie in a frontoparallel plane. Grasps are known to be accurately sized relative to the size of an object as the hand approaches the object. When the grasp involves contact of thumb and fingers on the front and back of the object, respectively, then the specification of the relevant extent of the object (in depth) requires combined information about metric object size in a frontoparallel plane and metric shape (characterized by the aspect ratio of object depth-to-width). Given that perception of both object size and distance can be calibrated for accurate reaches-to-grasp, the question then was whether perceived shape could likewise be calibrated. Bingham (2005) measured reaches to objects at different viewing distances to test calibration of distance, size, and shape perception. Using a virtual environment system, he measured reaches in conditions without and with terminal feedback. In the condition without feedback, observers could not see their hands, although they could see the target object both before and during the reach. In the condition with feedback, observers could see, in addition, a handheld stylus in relation to the target only at the end of each reach. The results showed that object distance and (frontoparallel) size were calibrated by feedback, but object shape (or depth) was not. Lee et al. (2008a) tested two possibilities left by this result. First, the virtual environment might have perturbed performance. Second, perhaps only grasping, not reaching, is calibrated to perceived shape. These questions were tested in a normal full cue environment by asking observers to reach-to-grasp actual cylindrical objects. These authors found that continuous online guidance (closed loop control) was required to pre-shape the hand with respect to object shape for effective grasps because haptic feedback without vision of the hand (i.e., feedforward grasping) failed to correct the inaccuracy of perceived shape. Shape perception was not calibrated by haptic feedback to allow accurate feedforward grasping.

Other studies of visually guided reaches-to-grasp have shown that feedforward grasping (visual open-loop control) is not accurate when participants use both binocular disparity and motion parallex information even though reaching is accurate (Bradshaw et al. 2000; Brenner and van Damme 1999; Hibbard and Bradshaw 2003; Melmoth and Grant 2006; Watt and Bradshaw 2003). In contrast, online guidance using binocular vision yields accurate grasping (Cuijpers et al. 2004; Melmoth and Grant 2006; Watt and Bradshaw 2000, 2003). Online binocular guidance of reaches-to-grasp is effective in the period after peak deceleration of the reach (Churchill et al. 2000) or right at the end of the reach during the grasp (Bingham et al. 2007; Melmoth and Grant 2006) because disparity matching is used in the final phases of the grasping movement to bring the fingers to the same depth plane as the object surfaces.

To summarize, shape perception is not calibrated by haptic feedback information, and poor shape perception yields inaccurate feedforward grasping. Lee et al. (2008a, b) inferred accordingly that this is the reason why we need continuous online guidance to reach-to-grasp objects effectively and accurately. Still the questions remain: Is accurate perception of metric shape ever possible under any circumstances? Relatedly, are the conditions typically studied in the laboratory environment truly representative? The situation typically being investigated in shape perception studies, whether judgment or reach-to-grasp measures are used, is current, immediate time perception by seated observers who are able to make only small head movements that generate relatively small perspective changes ($\approx$10–15° of perspective change). However, locomoting observers typically experience much larger perspective changes when looking at objects as they locomote. Thus, observers approaching a workspace in which they are to perform reach-to-grasp tasks experience large perspective changes as they enter the workspace. So, the questions are the following: Do large perspective changes potentially allow accurate perception of metric shape? If so, then for how long does this perception remain effective? In this context, another consideration is that we do not typically observe only a single isolated object in our surrounding environment as we move through it. Instead, we are exposed to cluttered space involving many objects. Someone who enters an office and sits down at a desk is confronted simultaneously with an array of objects (like a coffee cup, a pen or a phone, etc.) with which he or she may subsequently interact. So, the question is, do we perceive the shape of all the objects in the cluttered space using the large perspective changes that are available when we locomote into the workspace and does this perception remain effective when we subsequently interact with the objects in reaches-to-grasp them? Although the existing evidence indicates that we need continuous on-line guidance to reach-to-grasp effectively and accurately, is it possible that, under the more representative circumstances, we might reach-to-grasp an object accurately without continuous vision of the hand.

Bingham and Lind (2008) recently found that perspective changes of 45° or greater do indeed allow accurate perception of metric shape. Bingham and Lind used targeted reaching as a measure and found performance reflecting accurate metric shape perception when perspective changes where equal to or greater than 45°. Performance with 30° change was comparable to that found in previous studies with only 10–15° changes. Subsequently, Lee et al. (2008b) tested judgments of metric shape using computer displays combining SFM and stereo information. The amount of perspective change was systematically increased. The resulting judgments were all equally poor with changes less than 45°, but judgments became accurate as the amount of perspective change exceeded 45°. Relatedly, Brenner and van Damme (1999) had observers move a computer mouse to adjust the depth of a simulated ellipsoid to match a tennis ball held in the left hand (but not seen). Observers overestimated the depth of the ellipsoid when a static object was presented, but their judgments became accurate when the object was rotated by 60°. This study also suggested that reaches-to-grasp can be accurate without on-line guidance if the object is visually explored with greater than continuous 45° perspective changes. In this last study, shape judgments were expressed using the hand, but reach-to-grasp actions was not investigated as such.

Experiment 1

We now investigated whether large perspective changes (>45°) allow accurate feedforward reaches-to-grasp. Experiment 1 replicated the methods used by Lee et al. (2008a, b) but with the addition of large perspective changes before reaches-to-grasp. We measured the effect by using two different variables in two different types of trials. The two variables were the maximum grasp aperture (MGA), which occurs roughly halfway through a reach, and the terminal grasp aperture (TGA), which occurs at the end of a reach (reach velocity $\approx$ 0) but before the fingers contact the target object (Bingham et al. 2008; Coats et al. 2008; Mon-Williams and Bingham 2005, 2007; Mon-Williams et al. 2004). The reason we measured both the MGA and the TGA is because we found in the previous research (Lee et al. 2008a) that the TGA was constrained by the physical presence of the object on some occasions. The two trial types were feedback and probe. While participants grasped an actual wooden target in feedback trials, in probe trials, they grasped a virtual target that they only saw but could not contact. We needed feedback trials because Bingham et al. (2007) showed that the accuracy of reaches-to-grasp deteriorates without occasional feedback to calibrate performance. We added probe trials because the TGA might be biased by contact with the object in feedback trials. If both the MGA (in feedback and probe trials) and the TGA (in probe trials) aspect ratios varied reliably and accurately with target aspect ratios, then the evidence for accurate perception of metric shape and its use to guide feedforward grasping would be clear. Accurate pre-shaping of the hand for a grasp requires good perception of both metric shape and metric size. Thus, we also report data on metric size perception.

Methods

Participants

Ten adults, four men and six women, participated and they were between 18 and 27 years of age. All had normal or corrected-to-normal vision and no motor impairments. Only right-handed individuals who did not have a pacemaker or any metal in their body were recruited due to the apparatus being used. All of the participants were naïve as to the purposes of the study and were paid at $7 per hour. All procedures were approved by and conformed to the standards of the Indiana University Human Subjects Committee.

Apparatus

The apparatus used in this study was similar to that used in the previous study of Lee et al. (2008a). As shown in Fig. 1, participants sat near the corner of an L-shaped "mirror table". A wooden surface 38 cm wide and 80 cm long was placed over each arm of the L-shaped table. Each

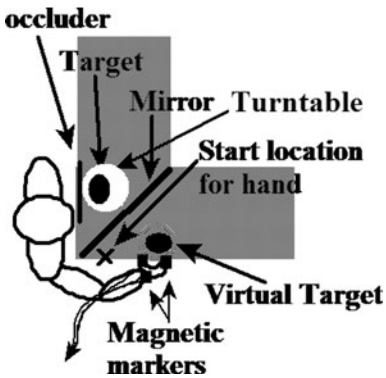

Fig. 1 Arrangement of the experimental apparatus. See the text for explanation

surface was painted matte black but was highly visible in normal room illumination. A semi-silvered mirror (which reflected 60% of the light and transmitted 40%) 33.7 cm wide and 24.3 cm high was placed between the two surfaces. The surfaces were cut diagonally to fit the bottom edge of the mirror oriented 45° to the line of sight extending across the corner of the L-shaped table. The mirror was used to create the illusion that an object was behind the mirror, although it was actually located in front of the mirror. The rear surface of the mirror had a removable panel. When this was removed, the image of the object in front of the mirror could be perfectly aligned with a physically identical object behind the mirror. When the back panel of the mirror was replaced, the visible image of the object was the same as the physical object that could be grasped without vision of the hand. A black panel was placed at the front edge of the table in front of the participant. The upper edge of this panel dropped diagonally to the right so as to allow the bottom edge of the mirror to be visible while simultaneously occluding the participant's view of the surface and target in front of the mirror.

The targets were placed on a turntable 18.8 cm in diameter and 4 cm in height located on the table surface to the left of the participant (i.e., in front of the mirror). The turntable was painted matte black and had a handle so the experimenter could move it back and forth easily. A board of the same height as the turntable was used behind the mirror to match the height of an object placed behind the mirror to the height of the object placed in front of the mirror. We used three different target objects painted matte black with green phosphorescent texture elements. The medium object had a circular shape 6.6 cm in diameter. The small object had an elliptical shape 4.4 × 6.6 cm, and the large object had an elliptical shape 6.6 × 8.2 cm. All objects were 4.5 cm in height. The objects were rotated by 90° to present each of the two principle axes in a fronto-parallel plane, so that three different object shapes yielded five different object aspect ratios; 0.667, 0.805, 1.0, 1.24,

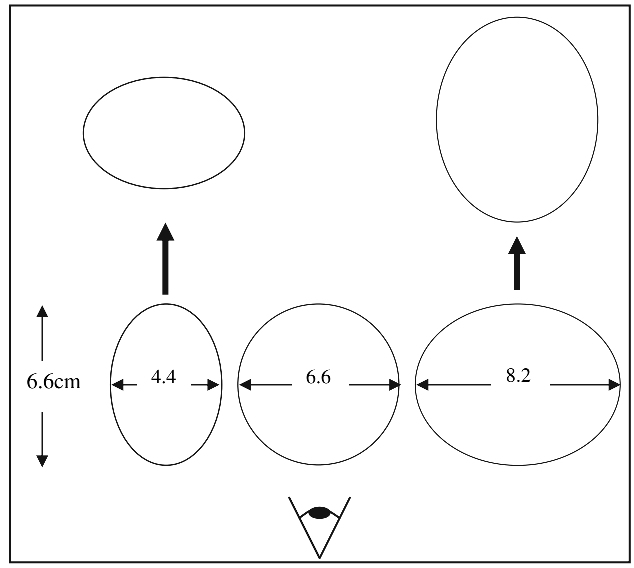

Fig. 2 A schematic representation of the objects used in this study

and 1.5 (see Fig. 2). A physically identical object was located behind the mirror at a location coincident with the virtual object seen in the mirror. The reaching distance from the starting location to the center of the object behind the mirror was 32 cm.

A three-marker Ascension miniBIRD magnetic measurement system was used to measure reach kinematics. This measurement system sampled movements at 103 Hz. Markers 1.1 × 0.8 × 0.8 cm were placed on the nail of the index finger, the thumbnail, and the wrist of the right hand using double-sided tape. The wires were gathered around the forearm with velcro. The emitter for the measurement system was located below the table and was centered in the reach space. Thick, black velveteen sound-attenuating drapes enclosed the experimental area around the L-shaped table and the participant. Black felt covered the table extending to the floor. The experimenter controlled the data collection using the computer keyboard, which was located beneath the surface. The data were stored in computer memory for subsequent analyses.

Procedure

The participants read and signed the consent forms, and then the miniBIRD markers were placed on the index finger nail, thumbnail, and wrist. By adjusting the height of the chair, the participant's eye height was adjusted to 20 cm relative to the bottom of the target objects. The task and procedure were explained to the participant. The participant sat at the table, pinched the index finger and thumb together and then lightly rested them at the starting location (located at the right bottom corner of the mirror). The participant was asked to close his/her eyes between trials, while the target objects were placed in position. Once the target objects were placed, the participant was asked to open his/her eyes. The participant was informed whether each trial would be feedback or probe. When the trial was feedback, the physically identical object was placed behind the mirror so the participant could touch the object. When the trial was probe, the object was not placed behind the mirror so the participant could not touch the object, that is, he/she grasped a virtual object (that was only visually present). When the participant was ready for a trial, the experimenter moved the turntable back and forth by 50° twice so that a 50° perspective variation the object was shown. When the experimenter finished rotating the target object in this way, the participant was informed which principle axis (i.e., either depth or width) should be grasped first. The participant was asked to move from the starting location to reach-to-grasp target objects so as to span the principle axis he/she was informed to grasp first (either the width or the depth) with the index finger and thumb. Before the participant started to reach-to-grasp, the experimenter initiated data recording. The experimenter stopped the data recording when the participant finished grasping the object. If the trial was probe, the participant was instructed to hold their thumb and index finger in place and say “done” or “ok” when he/she finished the reach-to-grasp. The participant then was told to go back to the start position and then reached-to-grasp the next axis (either the width or the depth depending on the axis he/she was informed to reach-to-grasp first). The order of axes to be grasped was randomized. These successive grasps within a trial (i.e., a single target presentation) were used to compute an aspect ratio (depth-to-width). This was the measure of interest. The participants could see the targets while reaching-to-grasp in all trials, but they could not see their hands.

The participants were allowed to practice reaching to grasp an object not used during the experiment both with haptic feedback (as in feedback trials) and without haptic feedback (as in probe trials). There were five blocks, and all aspect ratio and trial type (either feedback or probe) combinations (5 × 2) were randomly mixed in each block. We used three different protocols and the protocol was randomly selected for each participant. The participants performed 100 reaches (reaches-to-grasp object width and depth) in 50 trials.

Dependent measures

To evaluate grasping with respect to the object shapes, the TGA and the MGA were used, as described earlier. Although the TGA occurs before contact of the hand with the target object, Mon-Williams and Bingham (2005, 2010) found that it varies closely with object size. Typically, the TGA is within 0.5 cm of the object size. As described by Mon-Williams and Bingham, reaching-to-grasp is as much a collision-avoidance task as a target-acquisition task. The fingers are opened on approach to an object to avoid hitting the object as well as to prepare for the grasp. On approach, the variation in the orientation of the grasp aperture is as much a concern as is the object size itself in determining the appropriate aperture size with respect to the object. As the orientation of the grasp aperture varies away from the horizontal (shown to vary by as much as 45° from the horizontal), the size of the grasp aperture has to be increased to avoid collision with the object. Mon-Williams and Bingham found that the MGA is sized with respect to the maximum (diagonal) extent of the object or a diagonal formed by 45° slice through the object (relative to the horizontal), whichever is less. The MGA occurs at 50–70% of the total reach duration and is identified simply as the maximum aperture size between the index finger and thumb. Thus, whereas the TGA varies with object size, Mon-Williams and Bingham have shown that the MGA varies with the maximum object extent (MOE). We computed MOE as the Pythagorean of object width (or depth) and height—for instance, MOE $W = (W^2 + H^2)^5$. Thus, the D/W ratio computed using MOE was $(D^2 + H^2)^5/(W^2 + H^2)^5$. We used both the MGA and the TGA as dependent measures because, although the TGA is less variable than the MGA and varies closely with object size, we found in the previous experiment (Lee et al. 2008a) that the TGA was constrained by the physical presence of the object on some occasions. Thus, we looked for comparable results using both measures. Since the markers were placed on the index finger nail and thumbnail, the thickness of the fingers was subtracted from the TGA and MGA. After they had completed all trials of the experiment, the participants were asked to grasp the circular object (6.6 cm in diameter) width and depth with vision of the hand. To measure the thickness of the fingers, the circular object width or depth was subtracted from the final grasp aperture (FGA), which occurs when the fingers are finally in contact with the target object. Then, the thickness of the fingers was subtracted from the TGA and MGA to measure the actual distance between the two fingers.

Results and discussion

Previous experiments showed that large perspective change (>45°) allows accurate judgment of metric shape. We now in Experiment 1 investigated whether this information would be used to guide accurate feedforward reaches-to-grasp. Results showed that TGA aspect ratios and MGA aspect ratios varied reliably and accurately with the target aspect ratios and the target MOE aspect ratios, respectively in both feedback and probe trials. Furthermore, both TGA and MGA also varied reliably and accurately with target widths and depths, respectively. These results suggest that large perspective changes can yield accurate feedforward grasping. In the previous experiment (Lee et al. 2008a), the results showed that the TGA aspect ratio varied reliably and accurately with the target aspect ratio in feedback trials, but not in probe trials. In a regression of target aspect ratio on TGA aspect ratios, the slope was near 1 for feedback trials, but significantly less than 1 ($\approx$0.6) for probe trials. For the MGA, the slope was significantly less than 1 ($\approx$0.6) for both feedback and probe trials. Because the TGA aspect ratios for feedback trials were contaminated by the presence of the actual objects, Lee et al. (2008a) concluded that feedforward grasping does not reflect accurate perception of metric shape when only small perspective changes were available. Compared to these results, the current experiment showed that large perspective changes yielded accurate feedforward grasping, reflecting accurate perception of metric shape (see Fig. 3 and Table 1).

Because Lee et al. (2008a) found that the TGA aspect ratios were biased by the contact with the object in feedback trials, we excluded the TGA aspect ratios in the feedback trials from our analyses. As can be seen in Fig. 3, there was a difference in slope between the previous experiment with small perspective changes and the current experiment with large perspective changes for both the TGA aspect ratios (in probe trials) and the MGA aspect ratios (in both probe and feedback trials). To confirm these results, we performed multiple regressions to test differences in slopes between the two experiments separately for the TGA and the MGA data.[2] For the TGA data, we analyzed the probe trials only, and thus there were in this case three independent variables: target aspect ratios (a continuous variable), Experiment (that is, the previous experiment with small perspective changes vs. the current experiment with large perspective changes) (coded as $\pm$1), and an interaction vector (computed as the product of the first two vectors) (Pedhazur 1982). There was a main effect of experiment [$t(427) = 3.6$, $P < 0.001$], as well as a significant interaction [$t(427) = 3.6$, $P < 0.001$], showing that there was a difference in TGA slope between the two experiments. For the MGA data, the feedback trials were also considered and thus we added two more independent variables: trial type (i.e., feedback vs. probe trials) (coded as $\pm$1) and an interaction vector (computed as the product of the trial type categorical variable and the continuous variable). We found neither a main effect of trial type nor a significant interaction such that there was no difference between the feedback and the probe trials. However, there was a main effect of the experiment [$t(953) = 5.7$, $P < 0.001$], as well as a significant interaction [$t(953) = 7.3$, $P < 0.001$], showing that there was a difference in MGA slope between the two experiments (see Fig. 4).

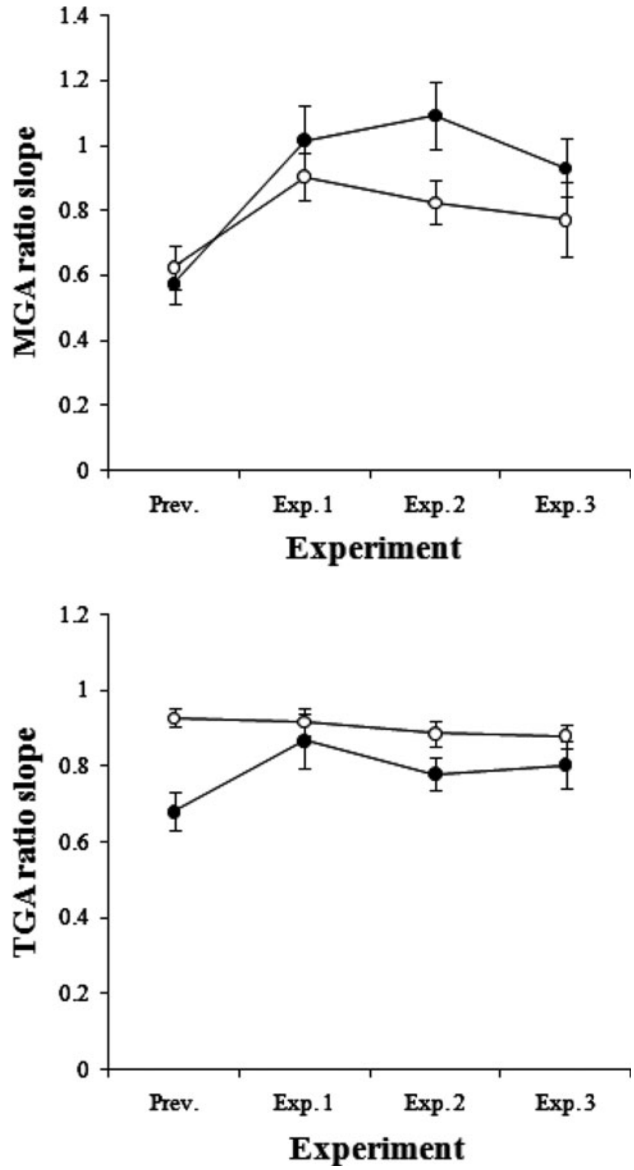

Fig. 3 Comparison of all experiments including all three experiments in the current study and the previous study (Lee et al. 2008a, b) in terms of slopes. The top panel shows the slopes of mean MGA aspect ratios and the bottom panel shows the slopes of mean TGA aspect ratios with standard error bars representing between-subjects variability. The open circles represent the feedback trials and the filled circles represent the probe trials

Whereas small perspective changes yielded slopes $\approx$ 0.6, large perspective changes (>45°) produced slopes $\approx$ 0.9 or better meaning that object shape aspect ratios

[2] We also analyzed these data by performing linear regressions separately for each participant regressing actual aspect ratios on grasp aspect ratios and recording the slope in each case. We then performed Anovas on the collected slope values. There was a single between subject factor, Experiment. In addition, there was a repeated measures factor with three levels: MGA Feedback MGA Probe and TGA Probe. The results of these analyses were the same as those in the multiple regressions reported for all experiments.

Table 1 The slope and r^2 of the relation between target aspect ratios and TGA aspect ratios computed for feedback and probe trials, respectively

		Feedback trials		Probe trials	
		Slope	r^2	Slope	r^2
Previous Exp.	TGA	0.93	0.84	0.69	0.72
	MGA	0.62	0.55	0.6	0.53
Exp. 1	TGA	0.92	0.88	0.87	0.71
	MGA	0.9	0.58	1.01	0.59
Exp. 2	TGA	0.88	0.85	0.78	0.74
	MGA	0.81	0.54	1.08	0.6
Exp. 3	TGA	0.87	0.82	0.8	0.63
	MGA	0.79	0.33	0.93	0.52

Also, the slope and r^2 of the relation between target MOE aspect ratios and MGA aspect ratios computed for feedback and probe trials, respectively. A slope was computed for the means of each participant. The previous Experiment refers to Experiment 3 of Lee et al. (2008a, b), which was conducted under conditions of small perspective changes

were accordingly reproduced in reaches-to-grasp the objects. The previous research had shown that feedforward grasping is poorly guided with respect to object shape (Lee et al. 2008a). As we mentioned in the previous paper (Lee et al. 2008a), shape perception is not well calibrated by haptic feedback without vision of the hand, although size and distance perception can be calibrated by haptic feedback (Bingham 2005; Bingham et al. 2000; Coats et al. 2008; Mon-Williams and Bingham 2007). However, our results in this experiment showed large perspective changes can yield accurate metric shape perception and thus accurate feedforward grasping.

Next, we tested the results for metric size, that is, width and depth. We subtracted actual depth or width from the respective MGA or TGA aperture values. (MOE values were used for MGA.) With poor perception of actual sizes, these difference values would be expected to vary strongly as a function of actual sizes because grasp apertures would not be adjusted well as actual size varied. If perception of size is accurate, then variation of difference scores with actual size should be weak because the apertures are being adjusted as a function of the size of the object. The means are shown in Fig. 5 where it can be seen that perception of metric sizes was poor in the previous experiment with small perspective changes. In contrast, perception of metric size was accurate in the current experiment with large perspective changes. In each case (that is, separately for depth and width for each measure), we performed a multiple regression of actual size (width or depth) on the respective difference scores with Experiment as a categorical independent variable (coded as +/– 1) and an interaction vector to test for slope differences. In each and every case, all three factors were significant ($P < 0.01$ or better). The significant interaction vector meant that the slopes were significantly different in the two experiments. The slopes were quite low (≈ -0.10) in the current experiment with large perspective changes and large in the previous experiment with small perspective changes (≈ -0.30). Apertures with large perspective changes were on average within about 1 cm to 1.5 cm of object size, whereas with small perspective changes, they varied on average from as much as 3 cm to as little as -0.5 cm with respect to object size. In short, perception of metric size was good with large perspective changes but not without them.

Given these good results for both metric size and metric shape, it is important to understand that these are not redundant measures. First, each measure has associated variability. It is possible for the depth and the width to each vary, but for the corresponding aspect ratio to remain invariant. This is exactly what is allowed in Similarity geometry. Here, of course, it is possible that metric depth and width are primary, and the aspect ratio is derived. In this case, the latter measure would be redundant. However,

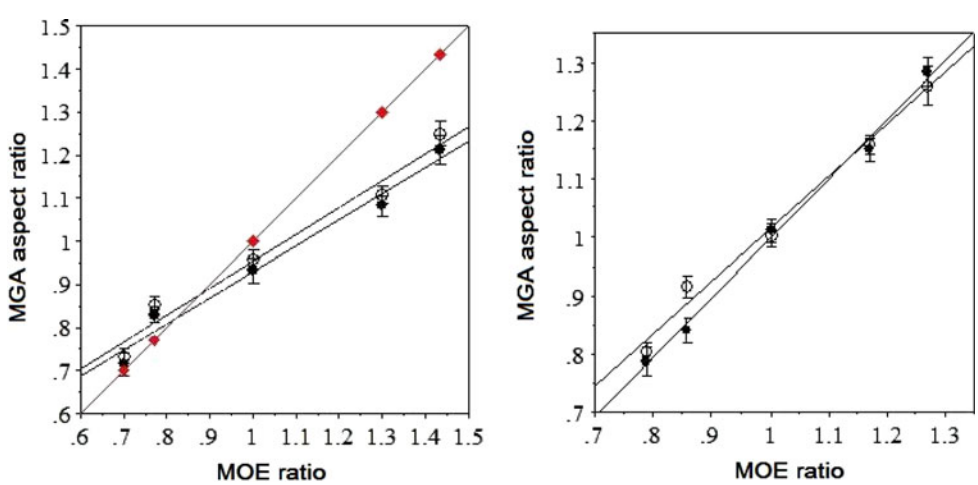

Fig. 4 Mean MGA aspect ratios with standard error bars representing between-subjects variability. The *left panel* shows the mean MGA aspect ratios in the previous study (Lee et al. 2008a, b), and the *right panel* shows the mean MGA aspect ratios in Experiment 1 of the current study. The *open circles* represent the feedback trials and the *filled circles* represent the probe trials. The *line* connecting *red diamonds* represents the correct target aspect ratio

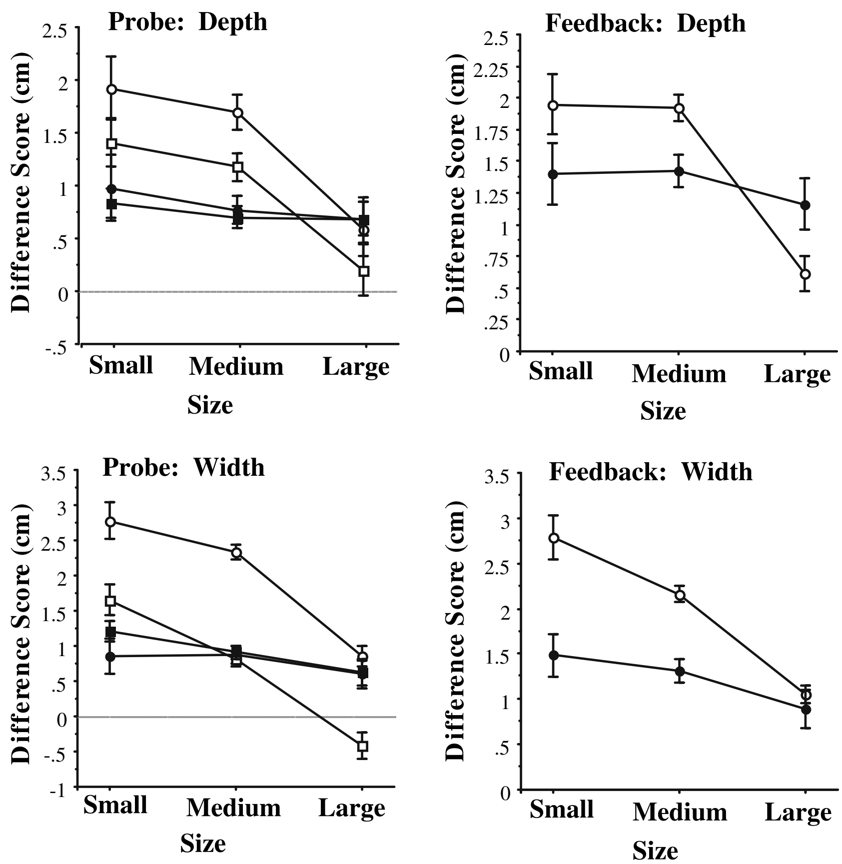

Fig. 5 Mean difference scores (with standard error bars) from the previous experiment with small perspective changes (*open symbols*) and Experiment 1 with large perspective changes (*filled symbols*). MGA means (*circles*); TGA mean (*squares*). Upper panels are depths, and lower panels are widths. *Left panels* are probe trials and include means for both MGA and TGA measures. *Right panels* are for feedback trials and include means only for MGA. See the text for additional details

if depth and width were independent variables and the aspect ratio was derived from them, then the variability of the aspect ratio (measured by its coefficient of variation or COV) would be greater than that of the corresponding depth or width. We simulated this a number of times. We added zero mean Gaussian random values to 30 sets of three depth values (4, 6, and 8 akin to those in our study) to create a data set. The added random values were scaled to produce a COV for each depth value of about 15% (as in our data). We did the same for a set of widths with random values that were uncorrelated with those added to the depths. Using these depths and widths, we then computed 30 sets of five aspect ratios like those in our study and then computed the COV for each of the five aspect ratios. The COV for the aspect ratios was reliably greater than that for either the depths or the widths by about 6% (that is, it was about 21%).

The question, then, was whether the COV for aspect ratios in our study was greater than that for the depths and widths. We found that it was not, meaning that the aspect ratio was primary, not derived from width and depth. For each participant in Experiment 1, we computed the COV for each of the three depths and the three widths as well as for each of the five aspect ratios. For each participant, we then computed a mean COV for size (that is, width and depth) and a mean COV for aspect ratio. Separately for MGA and TGA, we compared COV's for size and aspect ratio using a two-tailed paired t-test. The difference was significant in both cases, MGA: $t(9) = -2.3$, $P < 0.05$; TGA: $t(9) = -4.1$, $P < 0.01$. For MGA, the size COV was 13.8%, and the aspect ratio COV was 12.4%. For TGA, the size COV was 14.6% and the aspect ratio COV was 11.5%. The variation for the aspect ratio was LESS than that for the depth and width, not greater. This is exactly what we expected. When a participant viewed a 3D object, the metric shape had to be perceived together with the metric width (this is what is usually meant by 'size') to derive the metric depth. The perception of metric shape was primary.

In the following two experiments, these results were replicated, that is, the COV for the aspect ratio was less than that for either width or depth both for MGA and TGA. The values were comparable to those reported for Experiment 1, that is, about 15% for size and about 11% for the aspect ratio.

Experiment 2

In Experiment 1, we concluded that if information from large perspective changes is available, accurate perception of metric shape allows accurate feedforward grasping. We now asked how long such information from large perspective changes remains effective. Hu et al. (1999) found that feedforward grasping deteriorated after a 5-s delay between viewing the objects and initiating the reach. In their study, however, no visual information was provided after initial viewing of the target object. Observers were unable to see either their hand or the target object during the delay interval before the reach or during the reach itself. In contrast, we provided visual information (static image structure) during the delay interval and during the reach in this experiment. In each trial, participants first viewed large perspective change optic flow, then optic flow was stopped leaving only static image structure during the delay and the subsequent reach (because they also could not see their moving hand). Thus, in Experiment 2, we investigated whether dynamic perspective (optic flow) information remains effective once optic flow has ceased and only image structure remains available. To do this, we replicated the methods used in Experiment 1, but with the insertion of a 5-s delay between cessation of perspective changes and execution of the reach.

Methods

Participants

Ten adults, four men and six women, participated and they aged between 19 and 29 years. All had normal or corrected-to-normal vision and no motor impairments. Only right-handed individuals who did not have a pacemaker or any metal in their body were recruited due to the apparatus being used. All of the participants were naïve as to the purposes of the study and were paid at $7 per hour. All procedures were approved by and conform to the standards of the Indiana University Human Subjects Committee.

Apparatus and procedure

The apparatus was the same as in Experiment 1. The procedure was the same as in Experiment 1, with one exception. In each trial, the participants had to wait before reaching-to-grasp for 5 s after the experimenter finished rotating the target object back and forth by 50° twice. During the time between cessation of perspective changes and execution of the reach, however, the stationary target object was still viewed by the participants. Five blocks of trials were tested and all aspect ratio and trial type (either feedback or probe) combinations (5×2) were randomly mixed in each block. We used three different protocols, and the protocol was randomly selected for each participant. The participants performed 100 reaches (reaches-to-grasp object width and depth) in 50 trials.

Results and discussion

We concluded in Experiment 1 that large perspective changes could yield accurate metric shape perception and thus allow accurate feedforward grasping. In Experiment 2, we investigated whether such information from large perspective changes would continue to yield accurate metric shape perception and accurate grasping after a delay during which static image structure projected from a target object remained available (although perspective changes or optic flow had ceased). Results showed that TGA aspect ratios and MGA aspect ratios varied reliably and accurately with the target aspect ratios and the target MOE aspect ratios, respectively, in both feedback and probe trials. Furthermore, there was no significant difference between results of Experiments 1 and 2. Like results of Experiment 1, results of Experiment 2 were significantly different from those of the previous study with small perspective changes (See Fig. 3). Together, these results suggested that the information from large perspective changes remained effective for accurate feedforward grasping after a period of delay (5 s in this case).

We compared Experiment 2 with Experiment 1 and the previous small perspective change experiment, respectively, performing multiple regressions using the same design as in Experiment 1. For the TGA data in probe trials, we found no difference between Experiment 1 and Experiment 2 and therefore no deterioration in feedforward grasping occurred after a delay. Furthermore, we found that the results of Experiment 2 were different from those of the previous small perspective change experiment. The regressions yielded a main effect of experiment [$t(425) = 2.4$, $P < 0.02$], as well as a significant interaction [$t(425) = 2.0$, $P < 0.05$], showing that there was a difference in TGA slope between the two experiments. The results of the MGA were consistent with those of the TGA. For the MGA data, we also found that there was no difference between Experiment 1 and Experiment 2. When we compared Experiment 2 with the previous small perspective change experiment, the regressions yielded a main effect of experiment [$t(953) = 5.6$, $P < 0.001$], as well as a significant interaction [$t(953) = 6.9$, $P < 0.001$], showing that there was a difference in MGA slope between the two experiments. Thus, information from large perspective changes was effective for feedforward grasping after a delay. In the regressions on the MGA, we found that there was a difference between feedback and probe trials when Experiment 2 was compared with Experiment 1, but not when compared with the previous small perspective change experiment. In Experiment 2, MGA slopes for probe trials were slightly greater than 1 while slopes for feedback trials were somewhat less than 1. These variations were enough to yield a significant difference between the two conditions. We do not know what might have produced this difference, but both were close to 1, especially when compared to the results with small perspective changes.

As in Experiment 1, we also analyzed results for metric size comparing performance in Experiment 2 with the previous small perspective change experiment. The pattern of results was the same and the slopes in the comparison of difference scores with actual sizes in Experiment 2 were on average ≈ -0.13, that is, essentially the same as found in Experiment 1. Direct comparison of Experiments 1 and 2 using the same multiple regression design performed separately for each measure for depth and for width yielded no statistically significant differences between the experiments.

Taken together, these results showed that dynamic perspective (optic flow) information remains effective once optic flow has ceased and only image structure remains available. Thus, we concluded that large perspective changes could continue to yield accurate perception of metric shape and allow accurate feedforward grasping even after a 5-s delay.

Experiment 3

In Experiment 2, we concluded that, following large perspective change optic flow, a delay before reaching during which only static image structure was available failed to prevent accurate grasping. In Experiment 3, we simulated initial locomotion into a cluttered reach workspace during which large perspective changes would occur. Would subsequent reaches-to-grasp by a seated actor be accurate? This experiment introduced viewing of multiple target objects during perspective change followed by sequential reaches-to-grasp the targets yielding an effective delay of $\approx$16 s before the reaches to the last object.

Methods

Participants

Ten adults, six men and four women participated and they were between 18 and 26 years of age. All had normal or corrected-to-normal vision and no motor impairments. Only right-handed individuals who did not have a pacemaker or any metal in their body were recruited due to the apparatus being used. All of the participants were naïve as to the purposes of the study and were paid at $7 per hour. All procedures were approved by and conform to the standards of the Indiana University Human Subjects Committee.

Apparatus and procedure

The apparatus was the same as in Experiment 1 and 2 with one exception. In this experiment, we used a bigger and lighter turntable (25.5 cm in diameter and 2.5 cm in height) to allow room for three objects at the same time. While the turntable used in previous Experiments had a handle so the experimenter could move it back and forth by 50° easily, the turntable used in this Experiment had no handle because the objects were rotated by 360°.

The procedure was the same as in Experiments 1 and 2 with the following exceptions. The experimenter placed three target objects on the turntable and one object behind the mirror at the right location while the participants had their eyes closed. There were eight different templates for object placement (see Fig. 6). A green sticker was placed on the top of either the left or right object among the three target objects (the yellow object in Fig. 6). The sticker indicated the order of objects to be reached-to-grasp. When the sticker was placed on the left object, the order was left-middle-right, and when the sticker was placed on the right object, the order was right-middle-left. The sticker also indicated whether a trial was feedback or probe. Only the object on which the sticker was placed would be a feedback trial and the other two objects without stickers would be probe trials. In other words, the object to reach-to-grasp first was always a feedback trial and other two objects were probe trials. After setting up for each trial, the participants were asked to open their eyes and then, the experimenter rotated the turntable by 360° twice. Right after the rotation was completed, the participants were asked to reach-to-grasp the first object on which the sticker was placed. The participants were informed before reaching which axis should be grasped first, depth or width of the object. When the participants finished grasping depth (or width) of the object, they went back to the starting place and then, they reached-to-grasp the other axis of the object. After grasping both depth and width of the first object, the participants did the same for the other two objects. The participants always started reaching-to-grasp from the starting location, that is,

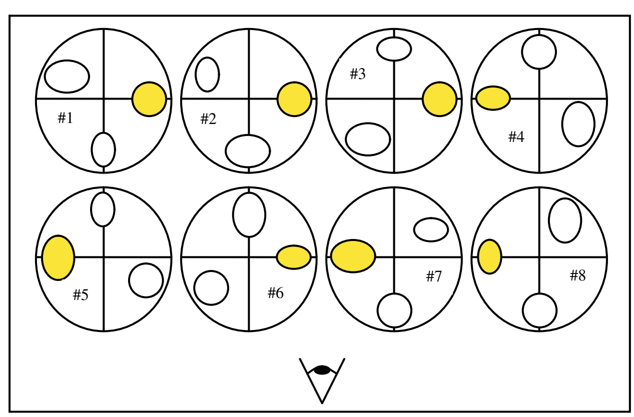

Fig. 6 A schematic representation of 8 patterns of object layout used in Experiment 3

they always return to the starting location after grasping each axis of the object. There was no delay between cessation of perspective changes and execution of the reach but sequential reaches-to-grasp the targets yielded an effective delay of ≈16 s before a reach to the last object. It took ≈4 s on average for each reach-to-grasp. Thus, a time delay of ≈8 s resulted before a reach to the second object and ≈16 s before a reach to the last object. There were two blocks, and eight templates were randomized in each block. The participants performed 96 reaches (reaches-to-grasp width and depth) in 48 trials (3 objects × 8 paradigms × 2 repetitions).

Results and discussion

We concluded in Experiments 1 and 2 that large perspective changes can yield accurate perception of metric shape and that such information could remain effective even after a 5-s delay. In Experiment 3, we investigated whether the information from large perspective changes remains effective and yields accurate feedforward grasping in a sequence of reaches-to-grasp multiple objects. Results showed that the TGA aspect ratios and MGA aspect ratios varied reliably and accurately with the target aspect ratios and the target MOE aspect ratios, respectively in both feedback and probe trials. Similarly, MGA and TGA apertures were scaled accurately to the relevant metric object extents. Furthermore, results showed that there was no difference between Experiment 3 and Experiment 1 and that the results of Experiment 3 were significantly different from those of the previous experiment with small perspective changes (see Fig. 3). Thus, these results suggested that large perspective changes yielded accurate perception of metric shape of multiple objects in a scene, that remained effective with continued viewing from a single perspective, and that allowed accurate control of subsequent feedforward reaches-to-grasp.

We again compared Experiment 3 with Experiment 1 and the previous small perspective change experiment, respectively, performing multiple regressions using the same design as before. For the TGA data in probe trials, we found that there was no difference between Experiment 1 and Experiment 3 and therefore there was no deterioration in accuracy of grasping multiple objects. Furthermore, we found that the results of Experiment 3 were different from those of the previous small perspective change experiment. There was a main effect of experiment [$t(489) = 2.6$, $P < 0.01$], as well as a significant interaction [$t(489) = 2.3$, $P < 0.02$], showing that there was a difference in TGA slope between the two experiments. The results of the MGA were consistent with the results of the TGA. For the MGA data, we found that there was no difference between Experiment 1 and Experiment 3. In contrast, when we compared Experiment 3 with the previous small perspective change experiment, the regressions yielded a main effect of experiment [$t(933) = 3.9$, $P < 0.001$], as well as a significant interaction [$t(933) = 4.9$, $P < 0.001$], showing that there was a difference in MGA slope between the two experiments.

As in Experiments 1 and 2, we also analyzed results for metric size comparing performance in Experiment 3 with the previous small perspective change experiment. The pattern of results was the same and the slopes in the comparison of difference scores with actual sizes in Experiment 3 were on average ≈–0.11, that is, again essentially the same as found in both Experiments 1 (≈–0.10) and 2 (≈–0.13) and of course, different from the previous small perspective change experiment (≈–0.30). Direct comparison of Experiments 1 and 3 using the same multiple regression design performed separately for each measure for depth and for width yielded no statistically significant differences between the experiments.

Thus, we concluded in Experiment 3 that large changes in perspective on a set of objects followed by a delay during which only static visual image information remained available still yielded accurate feedforward grasping even when the final delay was as long as ≈16 s, and participants were unable to attend and concentrate throughout on only a single object to prepare a reach-to-grasp to that object.

Given the results in comparison of pairs of experiments, we found that the results of all three experiments in the current study where large perspective changes were available were significantly different from those of the previous experiment where only small perspective changes were available (Experiment 3 in Lee et al. 2008a). We also found that the results of the three experiments in this study, however, were not different from each other. Thus, we found that information from large perspective changes yielded accurate perception of metric shape that, coupled with accurate perception of metric size, yielded accurate feedforward grasping. Furthermore, such information remained effective for accurate feedforward grasping even after a 5-s delay and also in a sequence of reaches-to-grasp multiple objects. To confirm these results, we examined slope errors by subtracting slope of each condition for each observer from slope of 1 and then taking absolute values. We performed a mixed design ANOVA with grasp aperture type (MGA in feedback trials, MGA in probe trials, and TGA in probe trials) as a within subject factor and experiment (the previous small perspective change experiment versus all three experiments in the current study) as a between subject factor. We found a main effect of experiment [$F(1, 38) = 13.9$, $P < 0.001$] but neither a main effect of grasp aperture type nor an interaction. In short, accurate feedforward grasping can be yielded by large perspective changes but not by small perspective changes.

General discussion

Previous research had shown that metric shape cannot be perceived accurately from binocular disparity (Johnston 1991; Tittle et al. 1995) or structure-from-motion (Norman and Lappin 1992; Norman and Todd 1993; Perotti et al. 1998; Tittle et al. 1995; Todd and Bressan 1990; Todd and Norman 1991) or combinations thereof (Norman and Todd 1996; Norman et al. 1995; Tittle and Braunstein 1993; Tittle et al. 1995), and that the perception of object shape is frequently distorted [see Todd et al.'s (1995) review]. More recent research, however, showed that the inaccurate perception of 3-D object shape might occur not because of systematic distortion, but because of ambiguity (Lind et al. 2003; Todd and Norman 2003), and that the systematicity of perceptual distortions is produced by contextual variables (Lee et al. 2007).

The poor shape perception as found in these studies is perplexing in the context of visually guided reaching and grasping because we normally perform reaches-to-grasp well and accurate shape perception would be required for accurate reaches-to-grasp in many cases, namely, those performed using feedforward control. Object size and shape information must be combined for accurate grasping when the grasp involves contact of thumb and fingers on the front and back of an object, as it so often does. Thus, Lee et al. (2008a, b) investigated how we might perform reaches-to-grasp well, despite poor shape perception. They first hypothesized that haptic feedback information might constrain the ambiguity of perceived shape and calibrate information about shape so as to allow accurate reaches-to-grasp. However, the results were consistent with the results of Bingham (2005) in which shape perception was not calibrated by haptic feedback, while distance and size perception were. Lee et al. (2008a, b) concluded that feedforward grasping is poorly guided with respect to metric object shape and thus online guidance (closed-loop control) is needed to pre-shape the hand with respect to object shape for effective and accurate grasps.

However, more recent research has shown that accurate metric shape perception is possible using information in continuous large perspective changes (>45°) (Bingham and Lind 2008; Lee et al. 2008a, b). Thus, we now investigated whether such information would allow accurate feedforward reaches-to-grasp and if so, the duration over which such dynamic perspective information remains effective once optic flow has ceased and only image structure remains available. Experiment 1 replicated the methods used by Lee et al. (2008a, b) but with the addition of large perspective changes before reaches-to-grasp. The results showed that the TGA and the MGA aspect ratios varied reliably and accurately with actual object aspect ratios in both feedback and probe trials. Thus, we concluded that large perspective changes can yield accurate perception of the metric shape and thus, given concurrent accurate perception of metric size, allow accurate feedforward grasping.

Next, we made the observation that large perspective changes, akin to those studied in Experiment 1, usually occur only when a person locomotes into a workspace in which reaches-to-grasp are performed subsequently (for instance, walking to the kitchen counter or to sit at a desk). We next investigated the stability of metric shape perception after the large perspective changes have ceased. Does accurate perception of the metric shape remain effective in supporting accurate feedforward grasping despite cessation of optic flow and continued availability of only image structure? Previous research had suggested that it might not. Hu et al. (1999) had examined the effect of a 5-s delay on grasping. They found that the maximum grip aperture was significantly larger when observers waited 5 s to initiate reaching after viewing the target object. Because observers were not able to see their hand or the target object during the delay interval before the reaches-to-grasp and during the reaches themselves, participants had to rely on stored visual information to guide their grasping after delay. However, the reason for the deterioration in performance in that study might have been because no information was provided after the initial viewing of the object. In contrast, in Experiment 2, we provided visual image structure both during the delay interval and during execution of the reach (after optic flow information had ceased). The methods replicated those of Experiment 1, but with the insertion of a 5-s delay between cessation of perspective changes and execution of the reach. The results were not different from those of Experiment 1 in which there was no delay. Accurate perception of metric shape yielded by large perspective changes remains effective for accurate feedforward grasping at least over a 5-s delay.

The problem was then that the representative situation would involve the simultaneous viewing of multiple objects as one walked into the workspace and subsequently reached-to-grasp multiple objects. For instance, a person might walk to a kitchen counter, then reach-to-grasp a glass to move it out of the way, then reach-to-grasp a mug for coffee followed by the sugar bowl. In Experiment 3, we simulated this natural situation in which longer delays would occur and the actor cannot simply attend and concentrate on a single object to be grasped. Participants viewed three objects undergoing a single common set of large perspective changes (that is, rigid motion) and then reached-to-grasp each of the objects in sequence. We found that the results of Experiment 3 were not different from those of Experiments 1 and 2 with respect to the accuracy of perceived metric shape and size and grasping. Large changes in perspective on a set of objects followed by a delay during which only static visual image structure remained available continued to yield accurate feedforward grasping even when the final delay before a reach-to-grasp was of 16 s or more.

These are remarkable results with important theoretical and methodological implications. The results of Hu et al. (1999) have been interpreted in support of the two visual systems hypothesis of Milner and Goodale (1996). They hypothesized that one of two visual systems, namely, the 'dorsal' channel serves perceptually guided actions like reaches-to-grasp and that the channel is distinguished, in part, from another 'ventral' channel (that performs object recognition) by a lack of visual memory. Hu et al. had shown that the accuracy of reaches-to-grasp deteriorates rapidly once the availability of visual information is removed. However, we have now found that the accuracy of reaches-to-grasp persists beyond removal of the relevant optical flow information, if static visual information is allowed to remain. The static image structure appears to provide a form of embodied memory for the optical flow information. The result shows how optical flow vision and image-based vision might functionally interact to yield great adaptive advantage. Optical flow intrinsically provides good information about the depth structure of surrounding surfaces (motion parallax, structure-from-motion, progressive occlusion) and the observer's relation to those surfaces (e.g. Cutting and Vishton 1995). Image structure is weaker in this regard. Figure-segregation is a difficult problem in static image vision (e.g. Gillam 1995; Peterson 2003). On the other hand, optical flow is fleeting while static image structure persists. Here, we see that optical flow can effectively calibrate the information in image structure about 3D surface layout allowing the information to persist and actions to remain stably accurate.

The methodological (and theoretical) implication of this work is that perception/action cannot be safely segregated into discrete and distinct actions for study, like walking *versus* reaching-to-grasp, actions partially distinguished by the fact that they entail use of lower versus upper limbs, respectively. For one, such actions are often nested. See Wann et al. (1993) and Anderson and Bingham (2010) for studies of walking-to-reach. Alternatively, such actions often occur in sequence and as we have found here, they can interact, even in a sequence where a number of reaches-to-grasp are performed after a person has locomoted into a workspace. The implication for future study of space perception and guided actions is that it requires consideration of nested actions at different spatial–temporal scales.

References

Anderson J, Bingham GP (2010) Proportional rate control of walking to reach. (submitted)

Bingham GP (2005) Calibration of distance and size does not calibrate shape information: comparison of dynamic monocular and static and dynamic binocular vision. Ecol Psychol 17:55–74

Bingham GP, Lind M (2008) Large continuous perspective transformations are necessary and sufficient for accurate perception of metric shape. Percept Psychophys 70:524–540

Bingham GP, Zaal FT, Robin D, Shull JA (2000) Distortions in definite distance and shape perception as measured by reaching without and with haptic feedback. J Exp Psychol Hum Percept Perform 26:1436–1460

Bingham GP, Crowell JA, Todd JT (2004) Distortions of distance and shape are not produced by a single continuous transformation of reach space. Percept Psychophys 66:152–169

Bingham GP, Coats RO, Mon-Williams M (2007) Natural prehension in the absence of haptic feedback when calibration is allowed. Neuropsychologia 45:288–294

Bingham GP, Hughes K, Mon-Williams M (2008) The coordination patterns observed when two hands reach-to-grasp separate objects. Exp Brain Res 184:283–293

Bradshaw MF, Parton AD, Glennerster A (2000) The task-dependent use of binocular disparity and motion parallex information. Vision Res 40:3725–3734

Brenner E, van Damme WJM (1999) Perceived distance, shape and size. Vision Res 39:975–986

Churchill A, Hopkins B, Rönnqvist L (2000) Vision of the hand and environmental context in human prehension. Exp Brain Res 134:81–89

Coats R, Bingham GP, Mon-Williams M (2008) Calibrating grasp size and reach distance: interactions reveal integral organization of reaching-to-grasp movements. Exp Brain Res 189:211–220

Cuijpers RH, Smeets JB, Brenner E (2004) On the relation between object shape and grasping kinematics. J Neurophysiol 91:2598–2606

Cutting JE, Vishton PM (1995) Perceiving layout and knowing distances: the integration, relative potency, and contextual use of different information about depth. In: Epstein W, Rogers S (eds) Perception of space and motion. Academic Press, San Diego

Gillam B (1995) The perception of spatial layout from static optical information. In: Epstein W, Rogers S (eds) Perception of space and motion. Academic Press, San Diego

Hibbard PB, Bradshaw MF (2003) Reaching for virtual objects: binocular disparity and the control of prehension. Exp Brain Res 148:196–201

Hu Y, Eagleson R, Goodale MA (1999) The effects of delay on the kinematics of grasping. Exp Brain Res 126:109–116

Iberall T, Bingham GP, Arbib MA (1986) Opposition space as a structuring concept for the analysis of skilled hand movements. Experimental brain research series 15. Springer, Heidelberg

Indow T, Inoue E, Matsushima K (1962a) An experimental study of the Luneburg theory of binocular space perception: 1. The 3- and 4-point experiments. Jpn Psychol Res 4:6–16

Indow T, Inoue E, Matsushima K (1962b) An experimental study of the Luneburg theory of binocular space perception: 2. The alley experiments. Jpn Psychol Res 4:17–24

Johnston EB (1991) Systematic distortions of shape from stereopsis. Vis Res 31:1351–1360
Koenderink JJ (1990) Solid shape. MIT, Cambrige
Lee Y, Lind M, Bingham GP (2007) Shape perception is merely ambiguous, not systematically distorted. J Vis 7(9):842a
Lee Y, Crabtree CE, Norman JF, Bingham GP (2008a) Poor shape perception is the reason reaches-to-grasp are visually guided online. Percept Psychophys 70:1032–1046
Lee Y, Lind M, Bingham GP (2008b) Metric shape perception requires a 45° continuous perspective change. J Vis 8(6):759a
Lind M, Bingham GP, Forsell C (2003) Metric 3D structure in visualizations. Inf Vis 2:51–57
Luneburg RK (1950) The metric of binocular visual space. J Opt Soc Am 40:627–642
Melmoth DR, Grant S (2006) Advantages of binocular vision for the control of reaching and grasping. Exp Brain Res 171:371–388
Milner AD, Goodale MA (1996) The visual brain in action. Oxford University Press, Oxford
Mon-Williams M, Bingham GP (2005) Task constraints alter prehension movements qualitatively and quantitatively. J Vis 5(8):124a
Mon-Williams M, Bingham GP (2007) Calibrating reach distance to visual targets. J Exp Psychol Hum Percept Perform 33:645–656
Mon-Williams M, Bingham GP (2010) Stop, fly-through and slam: affordances and the spatial structure of reaches-to-grasp. (submitted)
Mon-Williams M, Coats R, Bingham GP (2004) Reaching with feeling. J Vis 4(8):411a
Norman JF, Lappin JS (1992) The detection of surface curvatures defined by optical motion. Percept Psychophys 51:386–396
Norman JF, Todd JT (1993) The perceptual analysis of structure from motion for rotating objects undergoing affine stretching transformations. Percept Psychophys 53:279–291
Norman JF, Todd JT (1996) The discriminability of local surface structure. Perception 25:381–398
Norman JF, Todd JT, Phillips F (1995) The visual perception of surface orientation from multiple sources of optical information. Percept Psychophys 57:629–636
Norman JF, Todd JT, Norman HF, Clayton AM, McBride TR (2006) Visual discrimination of local surface structure: slant, tilt, and curvedness. Vis Res 46:1057–1069
Pedhazur EJ (1982) Multiple regression in behavioral research: explanation and prediction, 2nd edn. Holt, Rinehart & Winston, New York
Perotti VJ, Todd JT, Lappin JS, Phillips F (1998) The perception of surface curvature from optical motion. Percept Psychophys 60:377–388
Peterson MA (2003) On figures, grounds, and varieties of surface completion. Perceptual organization in vision: Behavioral and neural perspectives. In: Kimchi R, Behrmann M, Olson CR (eds) Perceptual organization in vision: behavioral and neural perspectives. Lawrence Erlbaum Associates, Mahwah, pp 87–116
Tittle JS, Braunstein ML (1993) Recovery of 3-D shape from binocular disparity and structure from motion. Percept Psychophys 54:157–169
Tittle JS, Todd JT, Perotti VJ, Norman JF (1995) Systematic distortion of perceived three-dimensional structure from motion and binocular stereopsis. J Exp Psychol Hum Percept Perform 21:663–678
Todd JT, Bressan P (1990) The perception of 3-dimensional affine structure from minimal apparent motion sequences. Percept Psychophys 48:419–430
Todd JT, Norman JF (1991) The visual perception of smoothly curved surfaces from minimal apparent motion sequences. Percept Psychophys 50:509–523
Todd JT, Norman JF (2003) The visual perception of 3-D shape from multiple cues: are observers capable of perceiving metric structure? Percept Psychophys 65:31–47
Todd JT, Tittle JS, Norman JF (1995) Distortions of three-dimensional space in the perceptual analysis of motion and stereo. Perception 24:75–86
Wann JP, Edgar P, Blair D (1993) Time-to-contact judgment in the locomotion of adults and preschool children. J Exp Psychol Hum Percept Perform 19:1053–1065
Watt SJ, Bradshaw MF (2000) Binocular cues are important in controlling the grasp but not the reach in natural prehension movements. Neuropsychologia 38:1473–1481
Watt SJ, Bradshaw MF (2003) The visual control of reaching and grasping: binocular disparity and motion parallax. J Exp Psychol Hum Percept Perform 29:404–415