Two eyes in action

Brenner, E.; Smeets, J.B.J.

Link to publication in VU Research Portal

RESEARCH ARTICLE

Eli Brenner · Jeroen B. J. Smeets

Two eyes in action

Abstract Do relative binocular disparities guide our movements in depth? In order to find out we asked subjects to move a 'cursor' to a target within a simulated horizontal plane at eye height. They did so by moving a computer mouse. We determined how quickly subjects responded to the target jumping in depth. We found that it took subjects about 200 ms to respond to changes in binocular disparity. Subjects responded just as quickly if the cursor was temporarily only visible to one eye near the time that the target jumped in depth, and less vigorously, though just as quickly, if the cursor jumped rather than the target, so the fastest binocular responses cannot be based directly on the *relative* retinal disparity between the target and the cursor. Subjects reacted faster to changes in the target's height in the visual field than to changes in binocular disparity, but did not react faster to changes in image size. These results suggest that binocular vision mainly improves people's everyday movements by giving them a better sense of the distances of relevant objects, rather than by relative retinal disparities being used to directly guide the movement. We propose that relative disparities only guide parts of very slow movements that require extreme precision.

Keywords Motor control · Binocular vision · Latency · Disparity · Prehension · Stereopsis · Distance · Human

Introduction

There is abundant evidence that binocular information plays an important role in human motor control (e.g. Bradshaw and Elliott 2003; Jackson et al. 1997; Mazyn et al. 2004; Mon-Williams et al. 2001b; Servos et al. 1992; Watt and Bradshaw 2003). Binocular information is used to localise the target of interest as well as to guide

E. Brenner (✉) · J. B. J. Smeets
Department of Neuroscience, ErasmusMC,
P. O. Box 1738, 3000 DR Rotterdam, The Netherlands
E-mail: e.brenner@erasmusmc.nl

the hand to the target (e.g. Bradshaw et al. 2004; Loftus et al. 2004; Servos and Goodale 1994). However, we do not really know how binocular information is used. Having two eyes, and therefore two slightly different views of the world, could aid human movements in two fundamentally different ways. The first is by helping to obtain reliable estimates of target objects' positions in space: in particular their distances. The second is by providing direct information about the moving hand's position relative to the target. For the former, the object's images in the two eyes must be combined with information about the orientations of the eyes. The latter relies on the small differences between the images in the two eyes to detect separations in depth. It seems reasonable to assume that the first mechanism is most important before the movement actually starts, whereas the second becomes more important as the hand approaches the target (Bradshaw et al. 2004). In the present study we try, and fail, to find evidence for the latter suggestion.

Relative retinal disparities

Moving in a way that gradually eliminates the relative retinal disparity between the hand and the target would appear to be a perfect way to guide the hand in depth (Morgan 1989). People are known to be extremely sensitive to the small differences between the images in the two eyes. The acuity for detecting mismatches in depth on the basis of such relative disparities is much higher than the resolution for judging distance on the basis of information about eye orientation. Moreover, relative disparities provide direct information about whether the hand is closer or further away than the target. How much nearer or further away depends on the viewing distance, because the same amount of relative disparity corresponds with a larger separation in depth at a larger viewing distance, but the separation can be estimated reasonably accurately if one knows the approximate viewing distance. Besides, information derived from the

relative disparity, such as the rate at which the mismatch in depth is decreasing, could perhaps be used to control the movement of the hand without requiring an estimate of viewing distance (Glennerster et al. 1996). So is eliminating relative disparities the strategy of choice for controlling our movements?

Relative retinal disparities are probably most suitable for guiding the last stages of goal-directed movements, because they require a close proximity between the hand and the target. In tasks that require extremely accurate localisation in depth, it is evident that relative disparities must guide the hand because there is no other way of obtaining the resolution that is required for performing the task successfully. However, the fact that such tasks are generally performed exceptionally slowly suggests that this is a special case, and that relative disparities might not guide faster movements. In the present study we examine whether relative retinal disparities directly guide the hand during moderately fast goal-directed action.

The alternatives

The binocular contribution to the on-line control of human hand movements (Bradshaw et al. 2004; Servos and Goodale 1994; Watt and Bradshaw 2002) need not be based on relative disparities. Useful information about the distance to the object of interest (and perhaps about the position of the hand) could be obtained by combining the positions of the object's (or hand's) images on the retinas (i.e. absolute retinal disparities) with information about the orientation of the eyes [for instance based on extra-retinal signals (Brenner and Smeets 2000) or vertical disparities (Brenner et al. 2001)]. Moreover, in many everyday situations people will have an abundance of alternatives to binocular information for judging the distance of the object of interest (Cutting and Vishton 1995; Sedgwick 1986). Sources such as motion parallax and height in the visual field, which provide reliable information at short distances under certain conditions, can certainly also be used to guide human reaching movements (Dijkerman and Milner 1998; Gardner and Mon-Williams 2001; Mon-Williams et al. 2001a; Watt and Bradshaw 2003). There is little reason to believe that any information that one can perceive cannot be used to guide ones' movements (except in patients with neurological disorders: Jeannerod 1986; Milner et al. 1999). However, some kinds of information may be much more suitable for the on-line control of movements than others, and may therefore normally play a much more important role than others (Rossetti 1998; Goodale and Milner 1992).

Probably the most important requirement for any information to be used for controlling ongoing movements is that the delays involved in using the information are short (Brenner and Smeets 2001). A short delay means that one can quickly respond to changes in the environment. Moreover, even if other sources with longer delays provide equivalent information, the response will presumably have been initiated before such information becomes available, so that in practice the source with the shortest delay will dominate the response. Thus, for instance, it is known that distance information based on motion parallax can guide hand movements quite effectively (Dijkerman et al. 1999; Watt and Bradshaw 2003). However, it probably contributes more to correctly judging the layout than to the on-line control of the hand's movement, because motion parallax necessarily involves relatively long delays due to its dependence on motion processing (which takes quite long; Brenner et al. 1998) and on the observer moving his or her head. Information from motion parallax may even altogether be inconsequential in the presence of binocular vision (Marotta et al. 1998; Watt and Bradshaw 2003).

The importance of being the fastest source of reliable information is especially relevant for information based on the relationship with the hand, because if the hand's movement is adjusted on the basis of faster alternative information, such as information about the target's position in space, then the information about the relationship with the hand will already be incorrect by the time it becomes available. Consequently, we expect the relative disparity between the hand and the target to only be used to guide fast movements if the delay involved in using this source of information is short. We therefore decided to examine how quickly people can respond to a change in retinal disparity, and to compare this with changes to two other sources of information about distance: image size and height in the visual field.

The experiments

A common way to determine the delays that are involved in using a given kind of information is with perturbation experiments. Subjects are asked to move their hand to a target, and while they are doing so the target is displaced or changed in some other manner (e.g. Brenner and Smeets 1997; Pélisson et al. 1986; Prablanc and Martin 1992). The delay is extracted from the response to the perturbation. We have recently shown that people can even respond quickly to perturbations when using a computer mouse to guide a cursor to a target (Brenner and Smeets 2003). Surprisingly, the subjects did not rely directly on the relative positions of the cursor and the target to do so. That was for motion and displacements in the frontal plane. We here determine whether this is also true when the displacement is in depth, where the relative positions are evident from the relative disparity. We used a task involving guiding a cursor to a target because this gave us better control of the sources of information than we could hope to achieve with real hand movements. Moreover, since people cannot feel the position of the cursor in the way that they can the position of their hand, this task should be particularly suitable for revealing the use of relative disparity.

The perturbations in the present study were displacements in depth. Determining the delay to respond to displacements of targets at eye height does not necessarily tell us how quickly people can respond to changes in relative disparity. In order to determine whether the responses were based on the changes in relative disparity, rather than on changes in the judged distance to the target or in its image size, we examined whether selectively disrupting relative disparity information (by temporarily removing the cursor from one eye's image) delayed the response. We also examined whether a different perturbation that has exactly the same influence on the relative disparity (shifting the cursor rather than the target) has exactly the same effect on the response.

In addition to determining whether people based their responses on relative disparity, rather than on binocular judgements of distance, we also had to determine whether people really responded to the changes in binocular information rather than to the changes in image size that inevitably accompany changes in position in depth. To do so we simulated simultaneous changes in size and distance that maintained the size of the image on the screen and retina (i.e. the images only shifted horizontally on the screen). We did not keep the image size constant in the other conditions because we wanted to use realistic simulations of displacements in depth, avoiding unnatural conflicts that could interfere with the response. However such a condition is critical for showing that the response was really driven by binocular information.

Finally, we examined whether people can respond faster to perturbations in depth if the targets are not at eye height, so that the height in the visual field also changes when the target jumps in depth. To distinguish between responses to any change in the image and responses to changes in judged distance, we compared presentations below eye height, where upwards is further in depth, with presentations above eye height, where downwards is further away. Since height in the visual field is normally only a reliable indicator of distance for objects that are resting on a horizontal plane, we only expect it to be useful as a depth cue for targets below eye height. However, if the vertical displacement simply helps to detect the change then the position should not make any difference.

Methods

With real targets it is hard to isolate the different sources of information about distance. In the present study we therefore used virtual targets and a computer mouse (as in Brenner and Smeets 2003), but the 'cursor' did not appear to move across the frontal plane of the computer screen, but to move in a horizontal plane through the screen. The targets were also on this plane, and were sometimes suddenly displaced in depth.

Subjects

Eight subjects volunteered to take part in the experiment, including one of the authors (E.B.). They were all members of our department (all right handed; six men, two women; 25–45 years of age; inter-pupil distances of 61–72 mm). All subjects had normal (corrected) vision and none had any known neuromuscular deficits. One subject had a stereo-acuity (as tested with the RANDOT Stereo test) of only 100 s of arc (with a tendency for the two eyes to alternate in dominating the percept). The others all had a stereo-acuity of 50 s of arc or better. Since the former subject's responses in the present study were no slower or weaker than those of the other seven subjects her data were simply included in the analysis.[1]

Only the author was aware of the questions that were being studied and of the specific conditions, although obviously most subjects could guess that the study was about responses to perturbations in depth. The subjects were explicitly instructed to hit the targets with the cursor as quickly as they could. This is not the same as asking them to move as fast as they can, because they had to hit each target before the next appeared, so they could not simply ignore the accuracy.

Each subject took part in seven sessions of 500 movements. Within each session there was one kind of visual perturbation (either the target or the cursor jumped in depth). About 100 of the 500 movements were perturbed. The sessions were performed on different occasions in an arbitrary order. This research is part of an ongoing research programme that has been approved by the local ethics committee.

Equipment

Images were presented with a Silicon Graphics Onyx Reality-Engine on a CRT monitor (120 Hz; horizontal size 39.2 cm, 815 pixels; vertical size 29.3 cm, 611 pixels; spatial resolution refined with anti-aliasing techniques). Subjects sat with their head in a chin-rest that was positioned so that their eyes were 80 cm from the screen and were aligned vertically with the screen centre. The images were viewed through liquid crystal shutter spectacles that were synchronised with the refresh rate of the monitor so that alternate images were presented to the left and right eye. Each eye received a new image every 16.7 ms (60 Hz). Every image was drawn in accordance with the way in which the objects in question would be seen from the position of the eye for which the image was intended, taking the screen distance and the viewer's inter-ocular distance into consideration. Thus both the ocular convergence when

---

[1]We also tested one stereo-blind subject (author J.S.) in the condition in which subjects were forced to rely on binocular information (the *size* condition described below). In contrast to the above-mentioned subject with a low stereo-acuity he had great difficulty performing the task and showed no response at all within 400 ms of the perturbation.

fixating the target (or cursor) and the images on the retinas were always appropriate for the simulated distance.

Stimuli

With the help of the above-mentioned shutter spectacles we were able to simulate objects in a three-dimensional space. We simulated ten spheres. Each of these spheres had a diameter of 1.5 cm and a spot with a diameter of 0.75 cm on its surface (at the side facing the subject). The cursor was black with a green spot. The target was also black but had a red spot. The other eight spheres were dark grey with a light grey spot. The grey spheres were only included to enhance the sensation of a stable three-dimensional space. They were situated 8 cm above and below the four possible target locations when the targets were at eye height. Otherwise they were 8 and 16 cm above or below the possible target locations. The rest of the screen was white. The experiment was conducted in a dimly illuminated room.

Within each session there were four possible target positions. The target was either 4.2° to the left or 4.2° to the right of the screen centre, and either 7.5 cm in front of or 7.5 cm behind the screen. Thus the target was either 6.4 cm to the left or right at a distance of 87.5 cm, or 5.3 cm to the left or right at a distance of 72.5 cm. In some cases the target was not at eye height but 8 cm below or above eye height. In those cases the four grey spheres that would normally be at that height were at eye height instead. The cursor always moved in a horizontal plane at the same height as the target. The range of possible positions was limited laterally by the width of the screen. In depth the range was limited to 45 cm (from 15 cm in front of the nearest to 15 cm behind the furthest possible target position).

The task

The subject's task was to quickly move the cursor to the target by moving the mouse. Despite simulating a three-dimensional space, the task itself was two-dimensional. The difference between this task and the normal way that we use a mouse to position a cursor on a screen is that the cursor moved in a horizontal plane, in the same direction as the mouse. Moving the mouse laterally made the cursor move laterally, while moving the mouse towards oneself or away from oneself resulted in simulated cursor motion in depth. The cursor's motion in depth was linearly related to the mouse movement. Its displacement in depth was about three times as large as that of the mouse. The cursor's lateral motion depended on its distance, because lateral mouse movements were interpreted as angular changes in cursor position. Once the cursor

touched the target the latter disappeared and a new target appeared at the other side, either at the same distance or at the other distance. Subjects were encouraged to start moving to the new target as soon as it appeared.

The perturbations

Sometimes the target jumped in depth while the cursor was moving towards it. It jumped from its current position to the other position at that side of the screen. Thus if it was originally at the near position it jumped to the far position, and vice versa. In one session the cursor jumped rather than the target. In that case, if the target was at the near position the cursor jumped to a position that was 15 cm further away (while maintaining its angular position). If the target was at a far position the cursor jumped to a position that was 15 cm nearer than its position at that moment. The jumps were triggered by the subject's action. They took place as soon as the cursor had moved one-third of the lateral distance to the target. At that time the cursor was about 5.6° to the left or right of the target, and was moving rapidly towards the target.

The seven sessions

Of the seven sessions there were four in which the target was presented at eye height, two in which it was presented lower and one in which it was presented higher. At each height there was one session in which the target simply jumped in depth, named *standard*, *low* and *high* to indicate the plane in which the targets were presented. In one of the remaining sessions in which the target was presented at eye height the cursor jumped in depth rather than the target (*cursor jumps*). In another the target jumped but the cursor was only visible to one eye when that happened (*cursor monocular*). In the last session in which the target was presented at eye height the target's simulated size changed when it jumped, so that its image did not change in size (*size*). Finally, in the second session in which the targets were presented below eye height, the target itself was only presented to one eye near the time that it jumped (*low monocular target*).

The distance information

The sessions differed slightly in order to introduce or eliminate selected sources of distance information. Whenever a target (or the cursor) jumped in depth, its retinal image shifted differently in the two eyes. The shifts correspond with a vergence angle of about 0.9° (depending on the distance between the subject's eyes). Thus the extent to which the eyes had to converge to fixate the target suddenly changed by almost 1°, as did

the relative disparity between the target and the cursor.

Beside these binocular sources of distance information, the size of the retinal image of a sphere with fixed dimensions also changes when its distance changes. Moreover, unless the sphere remains within a plane passing through the eyes, a change in distance will also be accompanied by a vertical displacement. For a horizontal plane below eye height, a sphere moving further away will rise in the visual field (and on the computer screen). For a plane above eye height a sphere moving *nearer* will rise in the visual field. In order to evaluate the importance of these sources of distance information, in comparison with binocular sources, we either kept the simulated size or the size of the retinal image constant, and either presented the targets at eye height or slightly above or below eye height.

When the size of the target's retinal image was kept constant our simulation corresponded to a target that changed size whenever it jumped. We never kept the cursor's image size constant because we were worried that doing so would disrupt judgements of its direction of motion. When the targets and the cursor were not at eye height, they were 8 cm below or above eye height. In that case a jump in depth was associated with a vertical displacement of 1.1°. In terms of angular displacements relative to each eye, this vertical shift is only about twice as large as the horizontal shift that is responsible for the binocular information.

We wanted to determine whether subjects respond to changes in the judged distance of the target (and cursor), or whether they use the relative disparity between cursor and target directly to guide the former to the latter. In order to do so we included sessions in which it was impossible to use relative disparity directly to adjust ones movements to the perturbations. This was achieved in a rather artificial manner. In one session with targets at eye height, we did not present images of the *cursor* to one of the eyes near the moment that the target jumped. In a second session with targets below eye height, we did not present images of the *target* to one of the eyes near the moment that the target jumped (so that initially only the vertical displacement and change in image size indicated that the target's distance had changed). In these two cases the cursor or target was no longer rendered for one eye from the moment that the cursor had moved one quarter of the lateral distance to the target, which turned out to be about 60 ms before the target jumped. The target remained invisible to one eye until 100 ms after the perturbation. If either the cursor or the target is only visible to one eye there is obviously no relative disparity between the two. The eye that did not receive images depended on the moment at which the information was removed, which depended on the precise moment at which the cursor had moved one quarter of the lateral distance to the target.

Data analysis

Mouse co-ordinates were obtained at the frame rate of 120 Hz. We only analysed the mouse movements towards and away from the subject: the movements that normally give rise to vertical displacements of the cursor on the screen, but which in our experiments gave rise to displacements of the cursor in depth. We will refer to such sagittal mouse movements as movements in depth. The resolution in this direction was about 0.7 mm of cursor displacement or a movement of the mouse of about 0.2 mm. The mouse co-ordinates were transformed into velocity signals by taking the differences between consecutive values and dividing them by the 8 ms interval between the frames. The jumps could occur when the cursor was moving to the left or to the right. However, since we were only interested in the depth component of the movement, we ignored this distinction.

Ignoring whether the cursor was moving to the left or to the right still left us with four possible trajectories because the cursor could start at a near or a far position (depending on where the previous target had been) and the target could initially either be near or far. For each of these four combinations we had about 100 trials with no perturbation and 25 with a perturbation. We aligned all these trials with respect to the moment that the cursor moved one-third of the distance to the target: the moment that the perturbation would or did occur. By then calculating the average velocity (in depth) on trials with and without perturbations, and subtracting one from the other, we were able to isolate the response to the perturbation. We assigned a positive sign to motion of the mouse that moved the cursor in the direction of the perturbation, irrespective of the actual direction of perturbation, so that we could average the values for the four combinations to get a single response for each subject and each session. In the figures we refer to this response as a *difference in velocity*. It is given in centimetre per second of movement of the hand (or mouse).

In order to be able to evaluate the data statistically we determined the onset of the response for each subject and session, and compared these values between the sessions with paired *t*-tests. In order to limit the number of statistical tests we compared each session with the *standard* session in which the targets were presented at eye height. The onset of the response was determined by finding the first moment at which the difference in velocity was higher than a threshold of 0.25 cm/s and remained above that threshold until the peak in the response.

Results

On average it took subjects 878 ms to reach the target when there was no perturbation, and 950 ms when there was a perturbation. This is the total time from the mo-
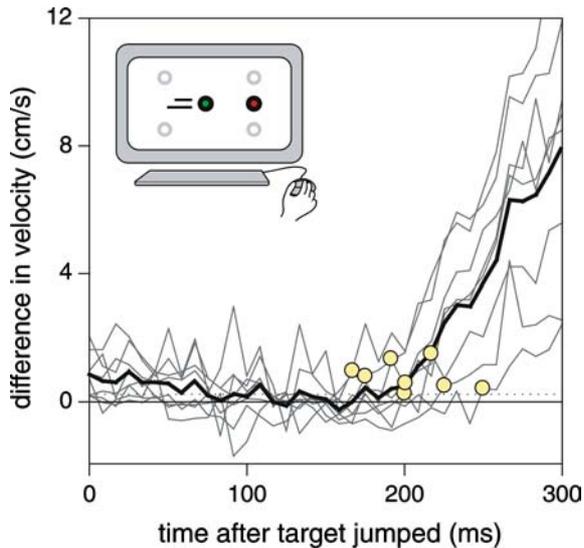
Fig. 1 Response to the perturbation in the standard sessions: targets jumping in depth at eye height. The *thick line* shows the overall average difference between the sagittal velocity of the hand on trials with and without a perturbation, as a function of the time after the perturbation. The *thin lines* show the eight individual subjects' responses with the estimated latencies indicated by *circles*. The latter were the points at which the difference in velocity rose and remained above a threshold of 0.25 cm/s (*dotted line*)

Fig. 2 Response to the perturbation in the four sessions in which the targets were at eye height. Overall average differences between the sagittal velocity of the hand on trials with and without a perturbation, as a function of the time after the perturbation. The response was almost identical to that in the standard session (*thick black line*) if the target's image size did not depend on its distance (*thin grey line*) or if the cursor was only visible to one eye near the time of the perturbation (*thin black line*). The response was weaker but the latency was the same if the cursor jumped in depth rather than the target (*thick grey line*). The *small figures* along the time axis show the means and standard errors of the estimated latencies of the eight subjects' responses

ment that a target appeared to the moment that the cursor touched it.[2] Figure 1 shows the average response to the perturbation during the *standard* session for each of the eight subjects. The figure shows the mean difference between the velocity in depth on trials with and without a perturbation, as a function of the time after the perturbation. The circles indicate the onsets of the responses that were determined for each subject. The overall average response is shown as a thick line. It took subjects about 200 ms to respond to the displacement (mean onset 203 ms; standard deviation 27 ms).

Is the response really based on binocular information?

Figure 2 shows the overall average responses for all the sessions in which the targets were presented at eye height. Comparing the response in the session in which the change in image size that normally accompanies a change in distance was removed (*size*), with that in the session in which it was not (*standard;* reproduced from Fig. 1), shows that the change in size makes very little difference. This means that subjects can respond to binocular information within about 200 ms. It does not necessarily mean that subjects refrain from interpreting

[2]An analysis of variance on the eight subjects' average values showed that it took longer to reach the target when there was a perturbation ($P<0.01$) and longer to reach the target in some conditions than in others ($P<0.05$; performance was faster when the targets were below eye height), but that the influence of the perturbation was similar for all conditions (no significant interaction).
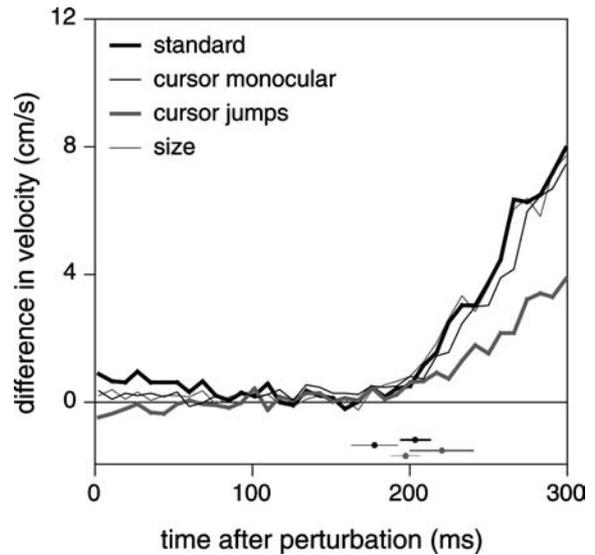
a change in image size as a change in distance, but it shows that they do not respond to such information faster than to binocular information.

Is the response based on relative disparity?

Removing relative disparities at the critical moment, by only presenting the cursor to one eye from about 60 ms before until 100 ms after the target jumped, made very little difference (compare *standard* with *cursor monocular* in Fig. 2). If the cursor jumped rather than the target, the response latency was similar, but the response itself was weaker (compare *standard* with *cursor jumps*). Thus when relative disparities were removed without changing the binocular information about the target's distance (*cursor monocular*), the response did not change. Conversely, for a given change in relative disparity, the response depended on whether the binocular information about the target's distance had changed (which it did in the *standard* session but not when the *cursor jumps*). We can therefore conclude that the response was not based on relative disparity but on judged distances. The weaker response to changes in the distance of the cursor (*cursor jumps*) than to changes in the distance of the target (*standard; size; cursor monocular*) could have to do with the fact that the cursor was moving quite fast at the moment of the perturbation, or to subjects (presumably) fixating the target most of the time.

Can responses be faster?

Figure 3 shows the overall average responses in the three sessions in which the targets were not at eye height. The response in the *standard* session is reproduced again to help compare the sessions. Responses were considerably faster when the plane containing the cursor and targets was placed 8 cm below eye level so that the height in the visual field changed when the target's distance changed (compare *low* with *standard*). When the targets were below eye level the responses were even quite fast when binocular information about the target jump was temporarily removed (*low monocular target*), showing that the response was really to the height in the field, rather than the vertical displacements somehow helping subjects to determine the distance from binocular information. Placing the plane slightly above eye level did not speed up the response (*high*). It may even have slowed it down. The fact that the faster responses depended on a specific relationship between the vertical and depth directions of the displacement confirms that height in the visual field is effective as a source of information about (relative) distance, rather than the perceived vertical displacement simply indicating to subjects that they should switch their response to the other target location.
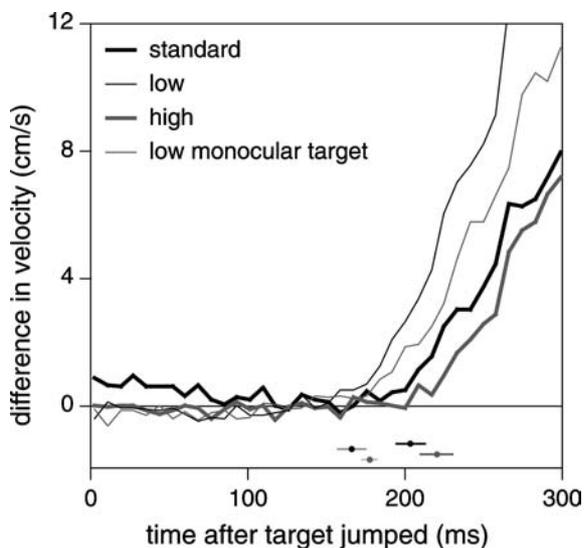

Fig. 3 Response to the perturbation when the targets were presented at different heights. Overall average differences between the sagittal velocity of the hand on trials with and without a perturbation, as a function of the time after the perturbation. The response was fastest when the cursor and targets were below eye height (*thin black line*) and slowest if they were presented above eye height (*thick grey line*). The response was faster for targets below eye level that were only seen with one eye near the time of the perturbation (*thin grey line*) than for ones that were seen at eye height by both eyes (standard session; *thick black line*). The *small figures* along the time axis show the means and standard errors of the estimated latencies of the eight subjects' responses

Statistics

The dots and error bars below the curves in Figs. 2 and 3 show the means and standard errors of the eight subjects' response onsets. The planned paired *t*-tests established that presenting the targets below eye height decreased the response latency (by 38 ms; $t_7 = 2.54$; $P = 0.039$). The reason for this must be that subjects responded to the vertical displacement, because a decreased response latency (26 ms) was even found if the target was only visible to one eye when it jumped ($t_7 = 3.00$; $P = 0.020$). An additional *t*-test showed that the response latencies in the two conditions in which targets were presented below eye level did not differ significantly ($t_7 = 0.94$; $P = 0.38$). A planned *t*-test also revealed that presenting the targets above eye level increased the response latency (by 17 ms; $t_7 = 2.83$; $P = 0.026$). The other three planned comparisons showed that changing the target's size when it jumps ($t_7 = 0.78$; $P = 0.46$), seeing the cursor with only one eye when the target jumps ($t_7 = 1.31$; $P = 0.23$) and letting the cursor jump rather than the target ($t_7 = 0.81$; $P = 0.45$) all failed to influence the response latency in a systematic manner across subjects.

Discussion

We can draw two conclusions from our data. The first is that the binocular contribution to the on-line control of movements in depth will seldom be based on relative disparities. Relative disparities probably only contribute to the on-line control of the last stages of very slow movements that require exceptional accuracy. The second is that the on-line control of movements in depth will not always be based on binocular information, even when binocular information is available (Watt and Bradshaw 2000, 2002). In particular, when movements are directed towards objects lying on horizontal surfaces, which they often are, height in the visual field may provide a useful alternative for guiding the final part of the movement.

Fast responses based on binocular information

We estimated that it takes about 200 ms to respond to a change in disparity. This is evidently the latency to respond to a change in a target's binocularly specified distance. Direct responses to the relative retinal disparity between the target and the cursor would have been delayed if the cursor was invisible to one eye when the target jumped, and would have been as vigorous if the cursor jumped as if the target jumped (because the change in relative disparity is the same). Thus the responses that we found are not direct responses to a change in relative disparity. This does not necessarily mean that responses cannot be guided by relative disparity. What it means is that if relative

disparities guide human movements, the latency for doing so is more than 200 ms. Thus, at least for fast movements with a rather abstract tool, the computer mouse, people respond to changes in the binocularly specified distance before they can respond to changes in relative disparity. Before speculating about whether this is also likely to be the case in other, more natural tasks, we will discuss how this conclusion fits with previous findings.

The idea that relative disparities are not responsible for the on-line control of movements is not inconsistent with earlier studies, because previous studies did not distinguish between using binocular disparity to judge the target's position and using relative retinal disparities to guide the hand to the target. The fact that performance improves if binocular information is introduced as little as 240 ms before the end of the movement (Bradshaw and Elliott 2003) is consistent with our estimated latency of 200 ms. Misaligning feedback about the hand's position as it moves towards a target also leads to substantial on-line corrections (Sarlegna et al. 2004; Saunders and Knill 2003) just as the cursor displacements did in the present study (and in Brenner and Smeets 2003). These findings show that vision of the moving cursor or hand contribute to the control of the movement, but this does not necessarily mean that relative retinal disparities are involved. The findings are just as consistent with the notion that separate judgements of the positions of the target and cursor are used, as with the idea of using direct judgements of one relative to the other.

The idea that binocular judgements of distance are responsible for the on-line control of movements is consistent with another line of research. When forced to rely on binocular information, people make systematic perceptual errors that are known to arise from misjudging the distance (Brenner and van Damme 1999). When forced to rely on binocular information to reach for objects, people make grasping errors that are consistent with these perceptual biases (Bradshaw and Hibbard 2003; Hibbard and Bradshaw 2003). This suggests that binocular vision contributes to human actions by helping to judge positions in depth. However, since the subjects in the latter studies were unable to see their hand while reaching for the objects, those studies could not exclude the possibility that relative disparities are normally used for guiding the hand during the last stages of the movement, which is where they would be expected to be most useful (Bradshaw et al. 2004).

Faster responses for targets presented below eye height

The responses in our study were fastest when the targets were below eye height. This was not just because people detected the change in binocularly specified distance faster when accompanied by a vertical displacement, because if so the response would not have been fast when the binocular information was absent at the critical moment (*low monocular target*) and would have been just as fast when the targets were above eye height (*high*) as when they were below eye height. The responses were possibly slightly faster when both height in the visual field and binocular information were present (*low*) than when only the former was present (*low monocular target*), but this difference (which was not statistically significant) could easily be explained by the fact that in the latter case the target was only visible on half of the frames. Thus the responses were presumably faster when objects that were further away were also higher in the visual field because height in the visual field was used as a source of information about distance (as in Gardner and Mon-Williams 2001; Mon-Williams et al. 2001a).

Relying on objects that are further away being higher in the visual field is quite reasonable, because objects that people manipulate are often supported by a surface below eye height. However, one should obviously not always respond to elevation in this manner. For instance, when targets are in a horizontal plane above eye height they are lower in the visual field when they are further away. Our subjects took considerably longer to respond to the change in height in this condition, but they did not respond in the wrong direction (see Fig. 3). Thus, although the subjects did not consider the whole geometry of the scene, they did not automatically respond to an increase in height in the visual field as an increase in distance either. Presumably they either automatically ignored height in the visual field as a source of information about distance when the objects were above eye height, or else they quickly learnt not to rely on objects being further away when they were higher in the visual field in the session in which the cursor moved upwards whenever it moved closer (for evidence of people learning to select reliable sources of information for a task see Marotta and Goodale 1998). In either case it is evident that people only rely on this source of information about distance under certain conditions, as indeed they should.

A completely different interpretation of the finding that the responses were faster when the targets were below eye height is that this is just because people are used to pushing the mouse away to move the cursor upwards, and pulling the mouse to move it downwards. We cannot dismiss this possibility, but point out that the converse could also be true. Our willingness to use height in the visual field as a source of information about distance could explain why we find the use of a computer mouse so intuitive. At present we cannot distinguish between these two possibilities, but in either case it is evident that the fastest response that a person can make to a displacement in depth is not necessarily driven by binocular disparities.

Generality of our findings

Of course, it is always difficult to extrapolate findings to different conditions, in particular if one's study involves

a limited set of conditions and is conducted under rather unusual circumstances. In such cases subjects may make use of information that is not generally valid or useful. We conducted our experiment in the manner that we did because it made the data easy to analyse and interpret. In this section we will discuss the problems in generalisation that this may have introduced. In particular, we will discuss the consequences of the fact that there were only two target distances in our experiments, that the relevant movement was that of the cursor rather than of the hand, and that we only examined fast movements.

In our study, the direction of the jump (if there was a jump) was always predictable, because there were only two possible distances. Nevertheless we believe that our subjects responded to the change in apparent distance, rather than simply switching distances in response to any detectable change, because if subjects had simply switched distances they would have reacted just as fast when the targets were above as when they were below eye height. Moreover they would probably have responded particularly slowly when the image size did not change at the moment that the binocular depth changed (*size*), because the change in image size (by about 20%) should be very easy to detect. Wann et al. (2001) have already shown that a change in size is not automatically interpreted as a change in distance when grasping virtual balls. Thus, the fact that our subjects did not appear to respond to the change in image size supports the notion that they were responding to information about distance, rather than just to any change. Our subjects' not responding to changing size also confirms that they were not responding to motion in depth but to the changed distance, because perceived motion in depth is largely determined by the change in image size (Brenner et al. 1996).

We used a rather unusual cursor manipulation task to examine how binocular information helps guide our actions, because this gave us very good control of the visual information that the subjects had at their disposal. Of course, an important difference between a cursor task and moving one's hand to an object is that one must use visual information to localise the cursor, whereas one can feel the position of one's hand. However, people have been shown to respond to visual feedback about their own hand being displaced during 'real' pointing movements (Saunders and Knill 2003; Sarlegna et al. 2004). Moreover, having to rely on visual information about the cursor makes it more likely that relative retinal disparities are used, rather than less likely, so we consider generalising our findings to conclude that relative disparities normally cannot guide fast movements to be justified.

Another aspect of our task that may contribute to the kind of information that is used is that we emphasised performance speed, rather than accuracy. Doing so is clearly likely to place more emphasis on the fastest sources of information. In fact, subjects had to hit the targets, so the movements were not extremely fast. On average, it took almost a second from the moment that a target appeared until the cursor had moved the 20 cm to the target (for which the hand had to move even less far). This is certainly not slow, but it is also not much faster than many of our everyday actions. Moreover, the cursor was visible to both eyes during all this time, whereas the hand may often only come into both eyes' view after some time in normal grasping. However, we are aware that other movements with other requirements could rely on completely different information. Tasks that require extreme precision will probably rely on relative retinal disparities despite the long latencies for on-line control that relying on relative disparities introduces. That is probably one of the reasons why the final stages of movements that require extreme precision, such as threading a needle, are conducted extremely slowly. Thus this does not discourage us from concluding that relative retinal disparities do not contribute to the on-line control of the hand in most daily tasks.

Our study shows that when on-line control is essential, because the errors are large, the initial response is unlikely to be based on relative disparities, and may not even be based on binocular information. However, this does not mean that binocular judgements of distance or even relative disparities do not kick in a little later. Cancelling relative disparities may be the only way to get rid of the final small errors. However, considering that the delay involved in doing so must be more than 200 ms, and that the relative disparity changes constantly as the hand moves, we consider it to be unlikely that cancelling relative disparities contributes to many human everyday movements. Thus we conclude from this study that in everyday life the depth component of the on-line control of our actions may often not even primarily be guided by binocular vision, and that when binocular vision does guide our actions such guidance is usually based on independently judged positions of target and hand, and not on the changing relative disparity between the hand and the target.

References

Bradshaw MF, Elliott KM (2003) The role of binocular information in the 'on-line' control of prehension. Spat Vis 16:295–309

Bradshaw MF, Hibbard PB (2003) Reaching for virtual objects: binocular disparity, retinal motion and the control of prehension. Arq Bras Oftalmol 66:53–61

Bradshaw MF, Elliott KM, Watt SJ, Hibbard PB, Davies IR, Simpson PJ (2004) Binocular cues and the control of prehension. Spat Vis 17:95–110

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

Brenner E, Smeets JBJ (1997) Fast responses of the human hand to changes in target position. J Motor Behav 29:297–310

Brenner E, Smeets JBJ (2000) Comparing extra-retinal information about distance and direction. Vis Res 40:1649–1651

Brenner E, Smeets JBJ (2001) We are better off without perfect perception. Behav Brain Sci 24:215

Brenner E, Smeets JBJ (2003) Fast corrections of movements with a computer mouse. Spat Vis 16:365–376

Brenner E, van den Berg AV, van Damme WJ (1996) Perceived motion in depth. Vis Res 36:699–706

Brenner E, Smeets JBJ, de Lussanet MHE (1998) Hitting moving targets: continuous control of the acceleration of the hand on the basis of the target's velocity. Exp Brain Res 122:467–474

Brenner E, Smeets JBJ, Landy MS (2001) How vertical disparities assist judgements of distance. Vis Res 41:3455–3465

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, New York, pp 69–117

Dijkerman HC, Milner AD (1998) The perception and prehension of objects oriented in the depth plane II Dissociated orientation functions in normal subjects. Exp Brain Res 118:408–414

Dijkerman HC, Milner AD, Carey DP (1999) Motion parallax enables depth processing for action in a visual form agnosic when binocular vision is unavailable. Neuropsychologia 37:1505–1510

Gardner PL, Mon-Williams M (2001) Vertical gaze angle: absolute height-in-scene information for the programming of prehension. Exp Brain Res 136:379–385

Glennerster A, Rogers BJ, Bradshaw MF (1996) Stereoscopic depth constancy depends on the subject's task. Vis Res 36:3441–3456

Goodale MA, Milner AD (1992) Separate visual pathways for perception and action. Trends Neurosci 15:20–25

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

Jackson SR, Jones CA, Newport R, Pritchard C (1997) A kinematic analysis of goal-directed prehension movements executed under binocular, monocular, and memory-guided viewing conditions. Vis Cogn 4:113–142

Jeannerod M (1986) The formation of finger grip during prehension A cortically mediated visuomotor pattern. Behav Brain Res 19:99–116

Loftus A, Servos P, Goodale MA, Mendarozqueta N, Mon-Williams M (2004) When two eyes are better than one in prehension: monocular viewing and end-point variance. Exp Brain Res 158:317–327

Marotta JJ, Goodale MA (1998) The role of learned pictorial cues in the programming and control of grasping. Exp Brain Res 121:465–470

Marotta JJ, Kruyer A, Goodale MA (1998) The role of head movements in the control of manual prehension. Exp Brain Res 120:134–138

Mazyn LI, Lenoir M, Montagne G, Savelsbergh GJ (2004) The contribution of stereo vision to one-handed catching. Exp Brain Res 157:383–390

Milner AD, Paulignan Y, Dijkerman HC, Michel F, Jeannerod M (1999) A paradoxical improvement of misreaching in optic ataxia: new evidence for two separate neural systems for visual localization. Proc R Soc Biol Sci 266:2225–2229

Mon-Williams M, McIntosh RD, Milner AD (2001a) Vertical gaze angle as a distance cue for programming reaching: insights from visual form agnosia II (of III). Exp Brain Res 139:137–142

Mon-Williams M, Tresilian JR, McIntosh RD, Milner AD (2001b) Monocular and binocular distance cues: insights from visual form agnosia I (of III). Exp Brain Res 139:127–136

Morgan MJ (1989) Vision of solid objects. Nature 339:101–103

Prablanc C, Martin O (1992) Automatic control during hand reaching at undetected two-dimensional target displacements. J Neurophysiol 67:455–469

Pélisson D, Prablanc C, Goodale MA, Jeannerod M (1986) Visual control of reaching movements without vision of the limb II Evidence of fast unconscious processes correcting the trajectory of the hand to the final position of a double-step stimulus. Exp Brain Res 62:303–311

Rossetti Y (1998) Implicit short-lived motor representations of space in brain damaged and healthy subjects. Conscious Cogn 7:520–558

Sarlegna F, Blouin J, Vercher JL, Bresciani JP, Bourdin C, Gauthier GM (2004) Online control of the direction of rapid reaching movements. Exp Brain Res 157:468–471

Saunders JA, Knill DC (2003) Humans use continuous visual feedback from the hand to control fast reaching movements. Exp Brain Res 152:341–352

Sedgwick HA (1986) Space perception. In: Boff KR, Kaufman L, Thomas JP (eds) Handbook of perception and human performance, vol 1. Sensory processes and perception. Wiley, New York, pp 21.1–21.57

Servos P, Goodale MA (1994) Binocular vision and the on-line control of human prehension. Exp Brain Res 98:119–127

Servos P, Goodale MA, Jakobson LS (1992) The role of binocular vision in prehension: a kinematic analysis. Vis Res 32:1513–1521

Wann JP, Mon-Williams M, McIntosh RD, Smyth M, Milner AD (2001) The role of size and binocular information in guiding reaching: insights from virtual reality and visual form agnosia III (of III). Exp Brain Res 139:143–150

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 (2002) Binocular information in the control of prehensile movements in multiple-object scenes. Spat Vis 15:141–155

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