Chapter 3
Horizontal-vertical anisotropy in visual space

Abstract:

We investigated the structure of visual space with a 3D exocentric pointing task. Observers had to direct a pointer towards a ball. Positions of both objects were varied. We measured the deviations from veridical pointing-directions in the horizontal and vertical plane (slant and tilt resp.). The slant increased linearly with an increasing horizontal visual angle. We also examined the effect of relative distance, i.e. the ratio of the distances between the two objects and the observer. When the pointer was further away from the observer than the ball, the observer directed the pointer in between himself and the ball, whereas when the pointer was closer to the observer, he directed the pointer too far away. Neither the horizontal visual angle nor the relative distance had an effect on the tilt. The vertical visual angle had no effect on the deviations of the slant, but had a linear effect on the tilt. These results quantify the anisotropy of visual space.

In press as:

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J., (2006). Horizontal-vertical anisotropy of visual space. *Acta Psychologica*. doi:10.1016/j.actpsy.2006.01.004

3.1 Introduction

We take it for granted that we can estimate in reasonable approximation in which direction another person is looking (Cline, 1967). But in view of the physiology of the visual system, this is not as straightforward as implicitly assumed. The three-dimensional world is first projected on a two-dimensional retina where all visual processing originates. The third dimension has to be derived from two 2D images. The brain uses heuristics to calculate the depth-values for each object that impinges on the retinae. This means that egocentric distances of objects are extracted indirectly and mistakes are easily made in judging the relative positions of two objects.

Two important questions are suggested by these observations. First of all, how do we make the translation from a 2D image to a 3D one? Secondly, are our judgments as veridical as we intuitively assume?

The first question has attracted the attention of many scientists working in the field of visual perception. Much research has been done on the cues that contribute to seeing depth, conventionally called depth cues (Gillam, 1995). Furthermore, scientists have been wondering how observers select these different cues as sources of information, and how they combine these sources to form the image of one unambiguous scene (Cutting & Vishton, 1995). This question will not be adressed in this paper.

In the research reported here, we will not try to specify the nature of specific depth cues, but we will try to answer the second question concerning the nature and veridicality of judgments of positions of objects. Over the years, researchers have tried to find out whether the size and direction of these deviations are dependent on spatial parameters. For a long time the field has been dominated by Luneburg's conjecture that visual space has a Riemannian geometry with a constant curvature (Luneburg, 1950). Luneburg assumed that visual space would have a homogeneous geometry. This means that the metric will be constant over various conditions. It also implies isotropy, i.e. that the metric is similar in all directions of the space. Many attempts have been made to find the curvature of this geometry. However, the curvature of visual space was found to vary for different observers (Battro, Di Pierro Nettro, & Rozestraten, 1976), different distances (Battro et al., 1976; Koenderink, Van Doorn, Kappers, & Lappin, 2002), different tasks (Koenderink, Van Doorn, & Lappin, 2000, Koenderink et al., 2002), and under different viewing conditions (Wagner, 1985). These experiments disprove the homogeneity assumption of Luneburg's conjecture.

The early research was done mainly in dark rooms and the heads of the observers were fixated. In this way, all information from motion and pictorial information about depth were eliminated from the visual field. Since then, vision scientists have worked with richer viewing conditions using a large variety of environments, scales and tasks. The environments that are used most are wide open fields (Kelly, Loomis, & Beall, 2004, Koenderink et al., 2000) and laboratory rooms (Cuijpers, Kappers, & Koenderink, 2000). The distances between observer and objects vary greatly, ranging from within arm's reach (Schoumans, Kappers, & Koenderink, 2002) to distances of up to 25 metres (Koenderink et al., 2000; Kelly, Loomis, & Beall, 2004). Furthermore, the tasks that are used vary widely: traditional parallel and equidistance alleys, line-bisection, horopter formation, parallelity tasks, body-pointing tasks, collinearity tasks, exocentric pointing tasks etc. However, these experiments have dealt almost exclusively with stimuli positioned in horizontal planes. Indow and Watanabe (1988) are an exception in that they examined frontoparallel planes besides the traditional horizontal planes. They used equidistance and parallel alleys to measure the curvature of visual space. They concluded that the horizontal subspace had a hyperbolic structure whereas the frontoparallel subspace was of a Euclidean nature, thus revealing the anisotropic nature of visual space. These conclusions suggest that the results of experiments done within a 2D framework cannot be generalized to a 3D space without further examination. Therefore in this paper we will report the experiments we have done in our investigation of 3D visual space. However, first of all we need to discuss a few experiments that were done on the horizontal visual subspace that involved less restricted viewing conditions than the traditional studies of visual space. These experiments are important for our understanding of the issues involved in our field of research.

Some vision scientists have focused on research in large open fields. Characteristic of this type of research is of course that during daylight, plenty of pictorial information is present in the visual field. Binocular information and physiological depth cues are not effective when objects are positioned more than a few metres from the observer. Often, the viewing conditions are less restricted than in laboratory environments, i.e. no chin-rests are used and observers are often simply standing at a certain position. A recent example of this kind of research is the work of Kelly, Loomis and Beall (2004). They had observers make judgments of exocentric direction using two tasks: judging a point on a distant fence collinear to a perceived line segment and orienting their body in the same orientation as a similar line segment. They concluded that the exocentric directions were misperceived during these tasks. However, this misperception of exocentric direction could not be explained by a misjudgement of egocentric distances. Levin and Haber (1993) and Foley, Ribeiro-Filho, and Da Silva (2004) have also been working with outdoor settings. They both had observers estimate egocentric and exocentric distances. Levin and Haber (1993) found that observers overestimated distances perpendicular to the line of sight (frontoparallel) but found no systematic misjudgements parallel to the line of sight. Foley et al. (2004) replicated these findings and fitted their data to the “tangle-model” (transformed angle model) that states that the visual angle undergoes a magnifying transform.

Many researchers have been focussing on polar coordinate systems to explain misjudgements of distances between objects. However, it could well be the case that people use another reference frame to code locations of objects in a visual scene. Gibson (1950), for example, proposed the idea that observers use the ground surface as a reference for representing space. Sinai, Ooi, and He (1998) confirmed this idea by showing that observers misjudge egocentric distances when a gap is present in the ground surface in between themselves and the target, or even when the texture on the ground surface is discontinuous. Thus, we should keep in mind the fact that the ground surface is an important factor in judging distances to objects. More researchers are working with alternative explanations of visual space, besides the traditional research focussed on metric properties of visual space. Yang and Purves (2003), for example, explain differences between visual and physical space by a statistical relationship between scene geometry and observer.

Koenderink and Van Doorn (1998) introduced an exocentric pointing task in a setting similar to the one used by Kelly, Loomis, and Beall (2004). In this task, an observer had to direct a pointer towards a target using a remote control. The task was done with both the pointer and the target in a horizontal plane at eye-height. In subsequent research, Koenderink, Van Doorn and Lappin (2000) found that when both objects (pointer and target) were up to 5 metres from the observer, the settings were concave for the observer. This means that the observers directed the pointer to a point that was further away from the observer than the target. When the objects were far away from the observer, 14 metres or more, the settings were convex for the observer, which means that the observers directed the pointer somewhere in between themselves and the target. These experiments were done with the pointer and target equidistant from the observer.

Cuijpers, Kappers and Koenderink (2000A) used the 2D exocentric pointing task in a horizontal plane in a laboratory environment. The walls of the room were covered with wrinkled plastic and neither floor nor ceiling was visible to the observer. In this environment, they found that the size of the deviations increased with separation angle, the angle between the lines connecting the two objects with the observer. Furthermore, the size and the sign of the deviations varied with relative distance, the ratio of the distances between the observer and the two objects. When the pointer was closer to the observer than the ball, they found an overshoot, i.e. the observer directed the pointer further away than the position of the target. When the pointer was further away from the observer than the ball, the observer undershot the target, i.e. he pointed somewhere between the ball and the observer. When the objects were equidistant, the deviations were small, but mainly positive (overshooting). This dependence on relative distance was also found by Kelly, Loomis, and Beall (2004) and Kelly, Beall, and Loomis (2004).

In our previous work (Doumen, Kappers, and Koenderink, 2005 [Chapter 2]), we expanded the experiments of Cuijpers et al. (2000), using a more elaborate context. The wrinkled plastic was removed from the walls. The observer's view of the floor and ceiling was not obscured by a cabin and the observer's head was not fixated by a chin-rest. With these changes, we tried to increase the degree of structure surrounding the observer, so that he could extract more information from the context. First of all, the dimensions of the room were visible, which could be used as an allocentric reference. Secondly, the observer was free to move his head: head-movements could provide some information. However, we found the same pattern and sizes of deviations as Cuijpers and colleagues found in their experiments. So the addition of structure that we provided did not reduce the size of the deviations nor did it change the pattern of the deviations in the pointing task.

Since very little has been written about visual space except with regard to the horizontal plane, we expanded our pointing task to a 3D task. Schoumans and Denier van der Gon (1999) have been working with a 3D exocentric pointing task in a virtual setup. They used a computer-display at 120 cm distance from the observer. The only depth information that was present was binocular disparity (stereoscopic presentation of the stimuli), linear perspective and motion cues (when the observer rotated the pointer). They concluded that visual space is isotropic. Using a real-life setup, one can overcome some problems that one faces with virtual stimuli. First of all, there is a conflict between different depth cues in virtual stimuli. Second, using a normal-sized computer screen, one can only test visual space at short distances.

In our real-life task, the pointer could be placed on various pillars, so that its height could be adjusted. The balls could also be hung at different heights from the ceiling. The observers had, just like in our previous experiments, an unobstructed view on the experimental room. They could rotate their heads and upper-bodies freely to approach viewing conditions as in everyday vision. We were wondering whether, just as in the 2D task, the relative distance and the horizontal separation angle have an effect on the pointing direction. A third variable was included in this experiment: the vertical separation angle. We were interested not only in the orientation of the settings in the horizontal plane (the slant), but also in the orientation of the settings in the vertical plane (the tilt). We expected that the horizontal separation angle and the relative distance would show the same effect on the slant as in the 2D experiments. However, since very little research has been done in a 3D visual scene, no prediction based on the literature could be made for either the dependence of the slant on the vertical separation angle or the dependence of the tilt on any spatial measure.

An initial assumption might be that visual space is isotropic. If this is the case, one would expect the tilt to depend on the vertical separation angle in the same way as the slant depends on the horizontal separation angle. Furthermore, the relative distance would influence the slant and the tilt in a similar way.

The present paper reports four experiments: in experiment A we varied the horizontal separation angle, in experiment B the relative distance, and in experiment C the vertical separation angle. Experiment D is an extra experiment concerning the vertical separation angle.

3.2 General Methods

Observers

Six undergraduates (three males and three females) participated in the first three experiments. Four different observers (two males and two females) participated in the last experiment. They all had little or no experience as observers in psychophysical experiments, were naive as to the purpose of the experiments and were paid for their efforts. They had normal or corrected-to-normal sight and were tested for stereo-acuity (all had an acuity of more than 60 arcseconds).

Experimental set-up

The experimental room measured 6 m by 6 m by 3.5 m. See Figure 3.1 for top- and side-views of the experimental room. The square represents the walls of the room in the graphs in the top-row of Figure 3.1. The wall on the left side of the observer contained radiators below four blinded windows. The wall in front of the observer was white with visible structure on it and some electrical sockets near the floor. The wall on the right side of the observer contained two light grey doors.

From the ceiling a horizontal iron grid was suspended below oblong fluorescent lights. The grid was 3 m above the ground. Green balls that were used as targets were hung from this grid. The balls had a diameter of 6 cm and could be hung at different heights above the ground.

Metal strips were taped on the floor to position the pointer. These strips were visible to the observer. The pointer consisted of a 45 cm long orange rod that was connected to a motor that the observer used for rotating the pointer in the vertical plane. The motor was

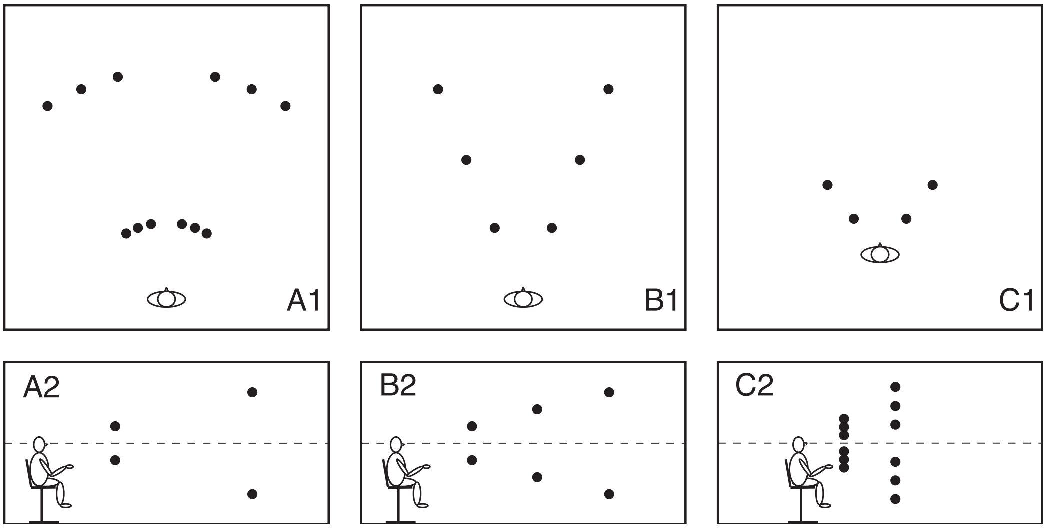

Figure 3.1
A1, B1 and C1 are top-views of the experimental room for experiments A, B and C, respectively. The squares represent the walls of the room, the black dots the positions of the objects in the azimuthal plane, and the small figure the position of the observer.
A2, B2 and C2 are side-views of the experimental room for experiments A, B and C respectively. The seated figure gives the position of the observer, the dashed line the eye-height of the observer and the black dots the positions of the objects.

covered by a white sphere with a diameter of 14 cm. The pointer was attached to a vertical iron rod. This rod was positioned on a circularly shaped foot. This foot contained a second motor that generated the rotations in the horizontal plane. The second motor was concealed by a cylindrical screen. The height of the centre of the pointer was 60 cm. Pillars were used to increase the height of the pointer. The observer could rotate the pointer using two small remote controls; one was for rotating in the horizontal plane, and one was for rotating in the vertical plane.

The observer was seated on a revolving chair that could be adjusted in height so that each observer could have an eye-height of 150 cm. This was exactly halfway between the floor and the grid (see Figure 3.1, bottom row of graphs).

Procedure

This paper reports the results of four experiments. The first three experiments (A, B and C) were measured in one block of sessions. For two of these experiments (A and B) the trials were measured together, whereas the third experiment (C) was measured separately. This was done because during experiment C the observer sat at a different position. Three observers started with C and did A and B after completing experiment C. The other observers did A and B first and then C. Experiment D was measured separately from the other experiments with different observers.

Each observer took about seven hours to complete the first three experiments; the experiments were performed in sessions of an hour each. Experiment D took about three hours. The observers did one hour per day, or they did two hours with a break between the two sessions.

The first session started with a description of the task. The observers were simply instructed to rotate the pointer towards the target. After instruction they could try the remote controls and had a practice trial. After this trial, the experiment started. After each trial the observer was told to close his eyes while the experimenter was busy reading the protractors of the pointer and rearranging the objects. The observer also had to rotate the pointer slightly in the horizontal and the vertical plane so that he had no information about the previous trial when starting the next one. The observers did not get any feedback during the experiment. After the last trial of each session, the observer closed his eyes again, so that the experimenter could read the values on the protractors and remove the ball from the raster. After the ball had been removed from the grid, the observer could open his eyes again and leave the room. In this way, the observer never saw the pointer and the ball together from any other position than the chair.

Analysis

For all experiments we looked at two dependent variables: the deviations from veridical settings of both the slant and the tilt. The slant is the orientation of the pointer in the horizontal plane, the tilt is the orientation in the vertical plane. We used these two variables because the slant and the tilt were also measured separately: one of the remote controls was for the orientation in the horizontal plane, the other for the orientation in the vertical plane.

3.3 Experiment A: Horizontal separation angle

In experiments with a 2D exocentric pointing task (Cuijpers et al., 2000; Doumen et al., 2005 [Chapter 2]), it was found that the deviations depended linearly on the horizontal separation angle. Thus, a logical step was to see whether this dependence was still present in a 3D task. In particular, we expected the settings for the slant to be comparable. The settings for the tilt, on the other hand, do not have a counterpart in our previous work. Thus, we were interested in the effect of the horizontal separation angle on the deviations of the slant and the tilt. With regard to the slant we can hypothesise that the deviations will increase with increasing separation angle. There are no indications in the literature with regard to the tilt. The most logical prediction, also in line with Luneburg's conjecture, is that with an increasing separation angle, and thus with increasing distance between the objects, the deviations of the tilt will also increase.

Methods

The objects were positioned either 150 or 450 cm from the observer in the azimuthal plane (See Figure 3.1 A1). For each position two different heights were used, one below and one above eye-height symmetrical around eye-height. The objects were positioned such that the vertical visual angle was the same for both the 150 cm and the 450 cm distance positions (See Figure 3.1 A2). The vertical separation angle was 23°, which corresponds to heights of 120 and 180 cm at the position 150 cm away from the observer, and heights of 60 and 240 cm above the ground for the positions that were 450 cm away from the observer. For each distance from the observer, the objects could be positioned at three different horizontal angles from the midsagittal plane of the observer so that we could examine the effects of varying the horizontal separation angle (HSA, see Figure 3.2 A). To do this, we used separation angles of 20°, 40° and 60° all symmetric around the midsagittal plane of the observer. The observer always had to direct the pointer from the right to the left side or vice versa, from a distance of 150 cm to 450 cm or vice versa and from a low position to a high position or vice versa.

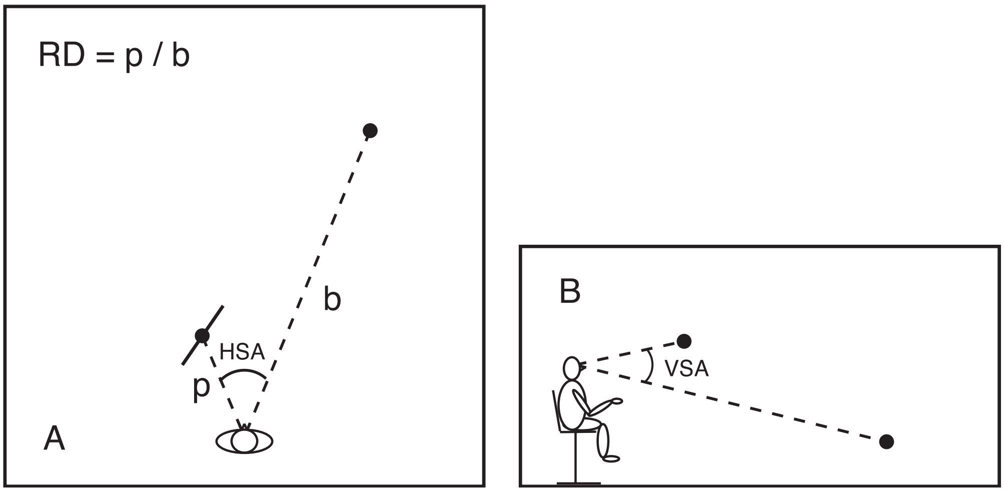

Figure 3.2
A graphical illustration of the parameters that were used. Figure A shows the horizontal separation angle (HSA) and the relative distance (RD). The relative distance is defined by the ratio of the distance from the pointer (p) and the distance from the ball (b). Figure B shows the vertical separation angle (VSA).

Results

First we will discuss the data obtained for the slant, the orientation of the pointer in the horizontal plane. The deviations of the slant were positive when the observer directed the pointer towards a position that was further away from himself than the actual position of the ball. Such an overshoot is shown in Figure 3.3A and B. The deviations were negative when the pointer was oriented in between the position of the observer and the ball. We will call this an undershoot (see Figure 3.3A and B). Figure 3.4 gives the deviations from veridical settings of the slant in degrees, plotted against the separation angle. A negative separation angle means that the pointer is positioned on the left side of the observer and the ball on the right, whereas a positive separation angle means that the pointer is positioned on the right side and the ball on the left side. The lines represent four different conditions (pointing upwards and away from the observer, downwards and away from the observer, upwards and towards the observer, and downwards and towards the observer). Each point gives the mean of the values of the six observers. The error-bars give the inter-observer standard deviations. When the pointer was closer to the observer than the ball (RD < 1), the deviations from veridical settings are mainly positive, which means overshooting the position of the ball. Furthermore, the size of the deviations increases with the absolute horizontal separation angle. When the pointer was further away from the observer than the ball (RD >1), there is a tendency towards negative deviations, which means that the observers were directing the

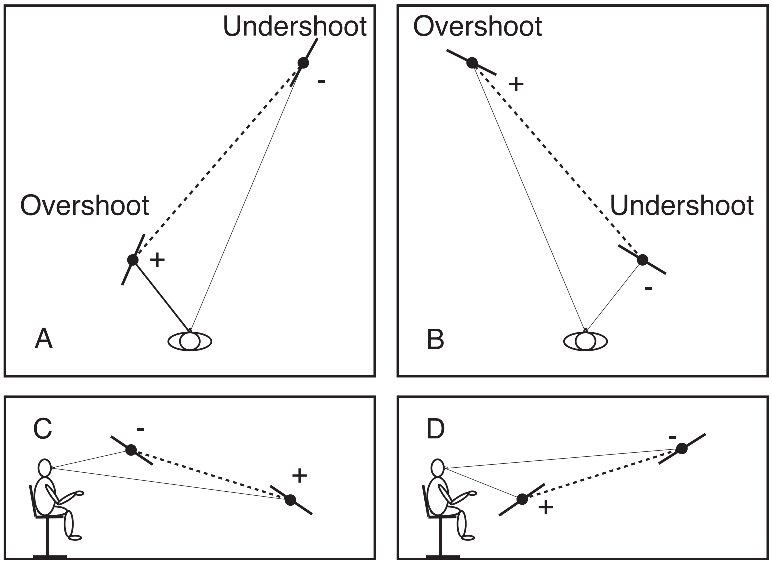

Figure 3.3
Examples of settings. The black dots represent the positions of the objects, the dashed lines the veridical pointing direction, and the small thick lines the pointing-orientation. The plus and minus signs represent positive and negative deviations. Figure A is a top-view of the room, with typical settings of the slant. Figure B is a top-view of the room, with a-typical settings of the slant. Figure C and D are side-views of the room with settings of the tilt.

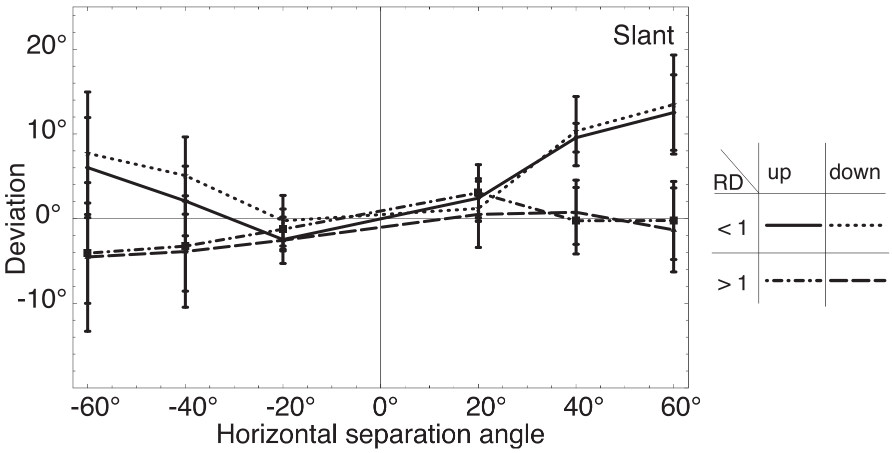

Figure 3.4
The data of the slant for experiment A. The deviation of the slant in degrees is plotted against the horizontal separation angle in degrees. The four lines represent four different conditions: pointing upwards with a relative distance smaller and larger than one and pointing downwards with a relative distance smaller and larger than one. Each data-point represents the mean of the values for the six observers. The error-bars represent the standard deviations.

pointer somewhere in between themselves and the ball (undershooting). This is depicted in Figure 3.3A.

We conducted a linear regression analysis, with weighted least squares because the data violated the homoscedasticity assumption. We did separate analyses for each condition

Table 3.1 *The F, p and R^2 for the regression analyses of the data of the slant of experiment A*

Relative distance	Height	Pointer position	$F_{(1,6)}$	p	R^2
< 1	Up	Left	18.41	<.001*	.535
< 1	Up	Right	38.25	<.001*	.705
< 1	Down	Left	10.03	.006*	.385
< 1	Down	Right	41.98	<.001*	.724
> 1	Up	Left	1.69	.212	.096
> 1	Up	Right	2.60	.126	.140
> 1	Down	Left	0.45	.510	.028
> 1	Down	Right	0.42	.528	.025

** The slope deviated significantly from 0*

because we were merely interested in the effect of the horizontal separation angle. For each line in Figure 3.4 we did two analyses, one for pointing from left to right (HSA < 0), and one for pointing from right to left (HSA > 0). We found a significant linear increase with increasing absolute horizontal separation angle when the pointer was closer to the observer than the ball (see Table 3.1 for the F and p values). The R^2 values were in between .35 and .71. These values are quite low, but this is due to individual differences, which are quite common in these tasks. For the condition in which the observer points towards himself, no significant effects were found (see Table 3.1 for the exact values).

The results of the tilt are plotted in Figure 3.5. The deviations of the tilt are plotted against the horizontal separation angle (in degrees), similar to Figure 3.4. A positive deviation means that the pointing-direction was higher than the physically veridical direction, whereas a negative deviation indicates a lower pointing direction than the veridical one. The different lines represent the same four conditions as in Figure 3.4 (pointing upwards and away from the observer, downwards and away from the observer, upwards and towards the observer, and downwards and towards the observer). Each point represents the mean of the values of the six observers. The error-bars give, just as in Figure 3.4, the inter-observer standard deviations. The deviations from veridical settings are smaller than the deviations found for the slant. A clear result in this figure is that the deviations are negative when pointing downwards and positive when pointing upwards (see Figure 3.3 C and D for a graphical view). Most points representing the data for pointing upwards were significantly different from zero or showed a trend to a difference from zero (one exception), whereas for pointing downwards half of the points showed (a trend to) a difference from zero. Thus, we can speak of a trend for pointing too high when pointing upwards and pointing too low when pointing downwards. However, no effect of horizontal separation angle is to be found in Figure 3.5.

Besides plotting the data separately for the slant and the tilt, we looked at the total deviations. However, since the deviations of the tilt are small (and not dependent on the horizontal separation angle) the total deviations resembled the data of the slant so closely that showing these plots would be redundant.

Discussion

When the pointer was positioned closer to the observer than the ball, the deviations of the slant were mainly positive and dependent on the horizontal separation angle. This dependence was not present when the ball was closer to the observer than the pointer. In the latter case, the deviations were mainly negative (see Figure 3.3 A). This part of our results is comparable to the results of our previous experiments in which we used a 2D exocentric pointing task, i.e. pointing in the horizontal plane at eye-height (Doumen et al., 2005). However, there was a difference between the two experiments when the relative distance was larger than 1: in the previous experiment, a linear dependence on horizontal separation angle was present when the pointer is further away from the observer than the ball, whereas no linear dependence was found in the present experiment for this condition.

There was no effect of horizontal separation angle on the settings of the tilt. A point worth mentioning, however, is the fact that when pointing upwards, the observers pointed too high, whereas when pointing downwards, they pointed too low (as depicted in Figure 3.3 C and D). These two observations complement each other and can be explained by an overestimation of the vertical separation angle. An overestimation of the horizontal separation angle has been reported earlier in the literature (Foley et al., 2004).

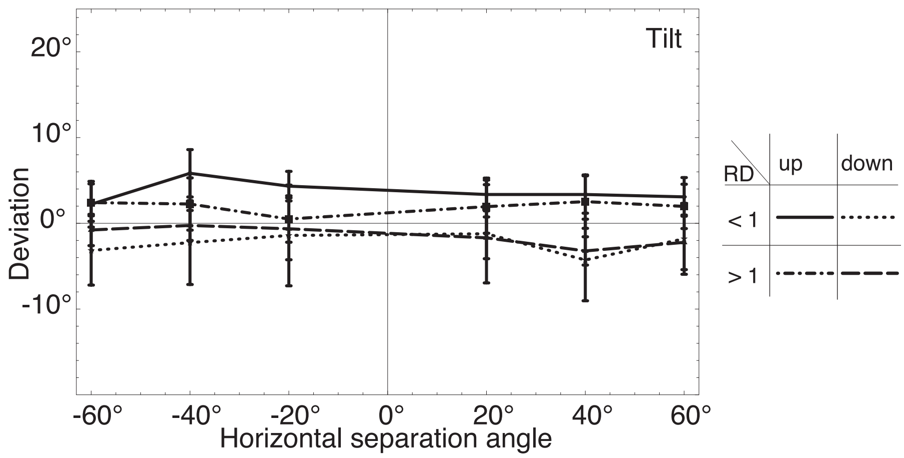

Figure 3.5
The data of the tilt for experiment A. The deviation of the tilt is plotted against the horizontal separation angle. The four lines represent the same conditions as in Figure 4. Each data-point represents the mean of the values for the six observers. The error-bars represent the standard deviations.

3.4 Experiment B: Relative distance

Not only the linear effect of the separation angle, but also the ratio of the distances between the two objects and the observer affect the deviations in a 2D pointing task (Cuijpers et al., 2000; Doumen et al., 2005 [Chapter 2]; Kelly, Loomis, & Beall, 2004). If the pointer is closer to the observer than the ball, the observers overshoot the position of the ball; and if the pointer is further away from the observer than the ball, the observers undershoot. In experiment B we tried to replicate these findings with the slant-data. If visual space is isotropic we should find similar effects for the slant and the tilt.

Methods

The trials for experiment B were measured together with the trials for experiment A. For this experiment we used only one horizontal and one vertical separation angle, 40° and 23° respectively (see Figure 3.1 B1 and B2 for a graphical view). As in experiment A, the observer always had to direct the pointer from left to right or from right to left, and from a position above eye-height to a position below eye-height and vice versa. The objects were 150, 300 and 450 cm away from the observer in the azimuthal plane. From each point on the right of the observer, the observer had to direct the pointer to each point on the left of the observer (and vice versa), which resulted in relative distances of 0.3, 0.5, 0.7,1, 1.5, 2 and 3.

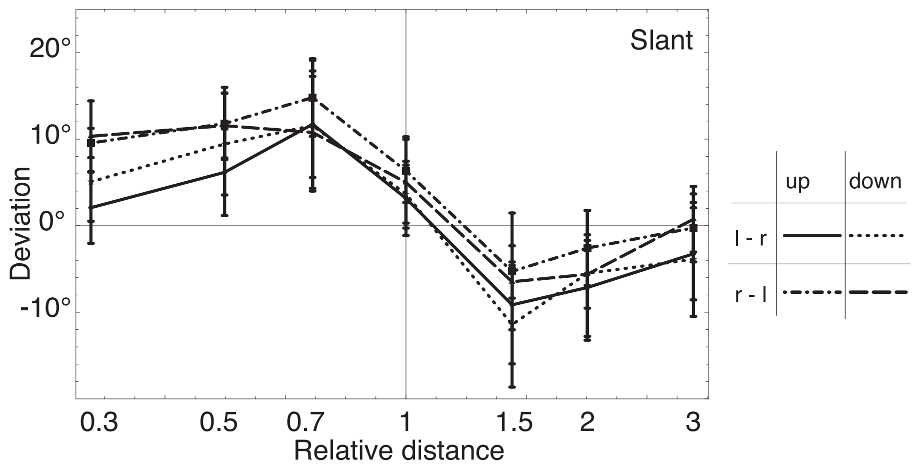

Figure 3.6
The data of the slant for experiment B. The deviation of the slant is plotted against the relative distance in logarithmic scale. The four lines represent four different conditions: pointing upwards and downwards from left to right and pointing upwards and downwards from right to left. Each data-point represents the mean of the values for the six observers. The error-bars represent the standard deviations.

Results

In Figure 3.6 the deviations of the slant are plotted against the relative distance in logarithmic scale. The relative distance is defined as the ratio between the distance between pointer and observer and the distance between ball and observer (see Figure 3.2 A). The different lines represent four conditions that were measured: pointing upwards and from left to right, pointing downwards and from left to right, pointing upwards from right to left, and pointing downwards from right to left. Each point represents the mean of the values of all observers, each error-bar the inter-observer standard deviation. When the relative distance is 1 the deviations are small and mainly positive. When the relative distance is smaller than 1, i.e. the pointer is closer to the observer than the ball, the deviations are positive. When the relative distance is larger than 1, the deviations are mainly negative. In Figure 3.3 A examples are given for settings when the relative distance was smaller and larger than 1. A striking feature of Figure 3.6 is that the deviations first increase and then decrease when the relative distance moves away from 1 in either direction. The four lines all follow the same pattern and are not qualitatively different.

For the tilt we plotted the deviations against the relative distance (in logarithmic scale) in Figure 3.7. The four lines represent the same conditions as for the slant-data. The deviations were quite small: most deviations were smaller than 5°. No effect of relative distance was found for the tilt. We find mainly positive deviations for upward pointing and mainly negative deviations for downward pointing (Figure 3.3 C and D). The majority of these deviations was either significantly different from zero or showed a trend towards a difference.

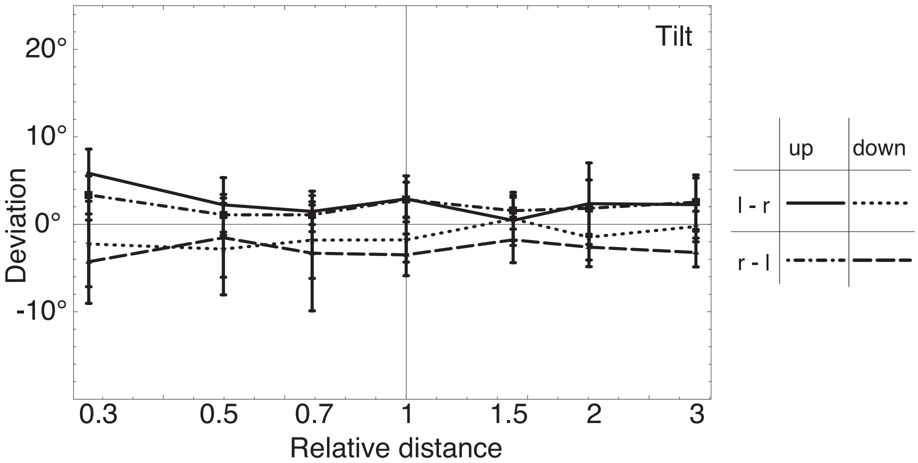

Figure 3.7
The data of the tilt for experiment B. The deviation of the tilt is plotted against the relative distance in logarithmic scale. The four lines represent four different conditions: pointing upwards and downwards from left to right and pointing upwards and downwards from right to left. Each data-point represents the mean of the values for the six observers. The error-bars represent the standard deviations.

Discussion

The part of our data concerning the deviations in the horizontal plane confirms the results found in previous research in the horizontal plane (Cuijpers et al., 2000; Doumen et al., 2005 [Chapter 2]; Kelly, Loomis, & Beall, 2004). When the pointer is closer to the observer than the ball, there is an overshoot (the observer pointed further away than the ball actually was). On the other hand, when the ball is closer to the observer than the pointer, the observer pointed in between the observer and the position of the ball.

The deviations from veridical settings in the vertical direction (the tilt) are not dependent on the relative distance. According to the isotropy-hypothesis, however, one would expect to find them to be dependent on the relative distance. We can therefore conclude that visual space is not isotropic. A point that we would like to stress, however, is the change in sign of the deviations when pointing upwards and downwards. These observations can both be explained by an overestimation of the vertical separation angle as described in the discussion of experiment A.

3.5 Experiment C: Vertical separation angle 1

Since a 3D task introduces an extra dimension to the setup, there is an extra parameter that can be varied. In this case it is the height of the objects and with this the separation angle in the vertical plane. So in experiment C we kept the horizontal separation angle and the relative distances constant, but we varied the vertical separation angle. If visual space is isotropic, the tilt would show the same dependence on the vertical separation angle as the slant does on the horizontal separation angle. Since in our previous work (Doumen et al. [Chapter 2], 2005) we found a linear increase of the slant with increasing horizontal separation angle, we hypothesised that the tilt would increase linearly with increasing vertical separation angle.

Methods

Experiment C was measured separately from experiments A and B, but with the same observers. In this experiment the distances from the observer were kept smaller than in the first two experiments, so that we could use larger vertical separation angles. The azimuthal distances used were 80 and 160 cm, the horizontal separation angle 60°. For each horizontal distance from the observer, three distances from the object to the horizontal plane at eye-height were chosen so that we had three different heights above and below eye-height. This resulted in vertical separation angles of 20°, 40° and 60°. The vertical separation angle is negative with downward pointing, and positive with upward pointing. See Figure 3.1 C1 and C2 for a graphical view of the setup of experiment C.

Results

In Figure 3.8 we plotted the deviations of the slant against the vertical separation angle. The various lines represent the following conditions: the pointer closer to the observer than the ball with a pointing-direction from left to right and from right to left, and the pointer further away from the observer than the ball with a pointing-direction from left to right and from right to left. Each point gives the mean of the values of all observers and each error-bar gives the inter-observer standard deviation. Neither of the lines shows an effect of the vertical separation angle. However, the size of the deviations is observer-dependent, which causes the large error-bars. All points deviated significantly from zero except for the points for which the pointing direction was upwards and from right to left with the pointer further away from the observer than the ball. As in experiment A and B, we see negative deviations when the relative distance is larger than 1 (pointer is further away than the ball) and positive deviations when the relative distance is smaller than 1 (pointer is closer than the ball).

The data for the tilt are given in Figure 3.9. The deviations from veridical settings for the tilt are plotted against the vertical separation angle for the four different conditions. As in Figure 3.8, the vertical separation angle is negative when pointing downwards, and positive with upward pointing. Although the deviations from veridical settings were rather small, there is a slight dependency on the vertical separation angle. Once again, two linear regression analyses (one for pointing upwards and one for pointing downwards) were done for the four different conditions. For the downward pointing conditions, all analyses revealed a significant linear effect of vertical separation angle (see Table 3.2 for the F, p and R^2

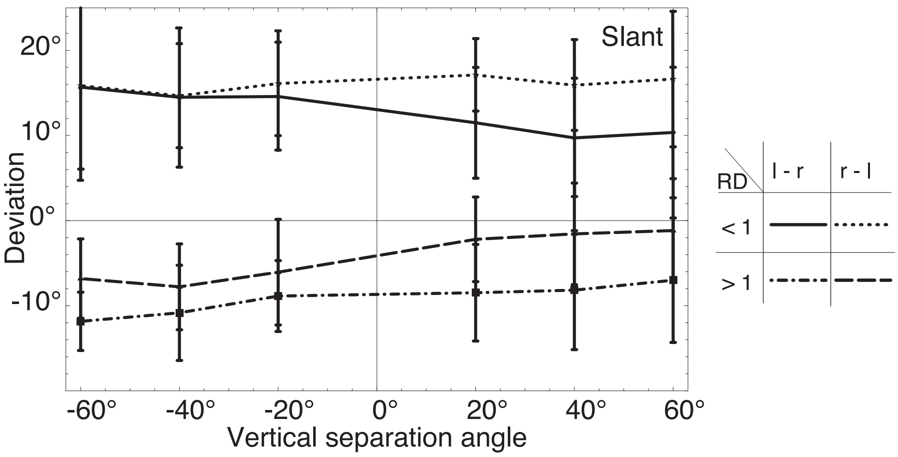

Figure 3.8
The data of the slant for experiment C. The deviation of the slant is plotted against the vertical separation angle in degrees. The four lines represent four different conditions: pointing from left to right with a relative distance smaller and larger than 1 and pointing from right to left with a relative distance smaller and larger than 1. Each data-point represents the mean of the values for the six observers. The error-bars represent the standard deviations.

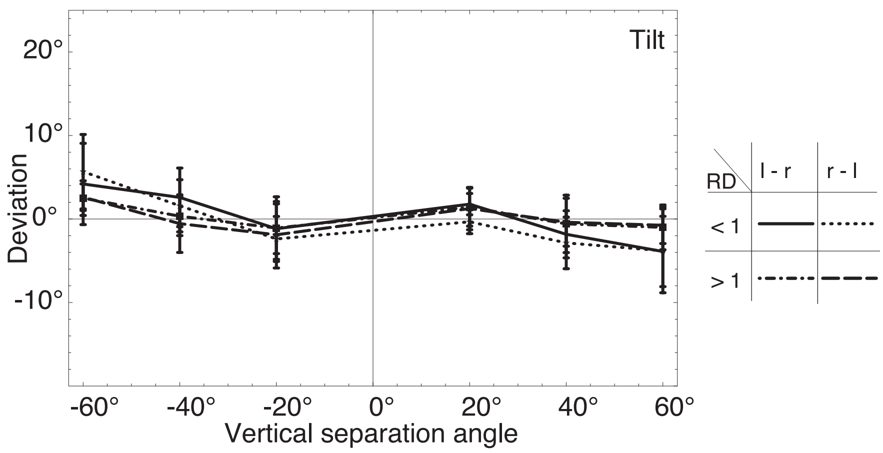

Figure 3.9
The data of the tilt for experiment C. The deviation of the tilt is plotted against the vertical separation angle. The four lines represent the same conditions as in Figure 3.8. Each data-point represents the mean of the values for the six observers. The error-bars represent the standard deviations.

values). The R^2 values are once again rather small, but this is due to individual differences that do not interfere with the trend that was found. With upward pointing and a pointer-position closer to the observer than the ball, the effect is also significant (see Table 3.2 for the values), whereas when the pointer is positioned further away from the observer and the pointing direction is upwards, no significant effect was found (see Table 3.2 for the values).

Although the deviations are rather small and possibly not significant, a striking feature of Figure 3.9 is the fact that with downward pointing with a large vertical separation angle (-60° / -40°) the deviations are mainly positive, whereas the deviations are negative when the vertical separation angle is -20°. In addition, with upward pointing, the deviations switch from positive to negative when the vertical separation angle increases.

Table 3.2 *The F, p and R^2 for the regression analyses of the data of the tilt of experiment C*

Relative distance	Pointer position	Height	$F_{(1,6)}$	p	R^2
< 1	Left	Down	6.92	.018*	.302
< 1	Left	Up	16.65	<.001*	.510
< 1	Right	Down	16.71	<.001*	.511
< 1	Right	Up	5.38	.034*	.252
> 1	Left	Down	5.65	.030*	.261
> 1	Left	Up	3.13	.096	.164
> 1	Right	Down	9.49	.007*	.372
> 1	Right	Up	2.16	.161	.119

* *The slope deviated significantly from 0*

Discussion

The vertical separation angle had no effect on the deviations in the case of the slant. However, the change in sign when the relative distance switches from smaller than 1 to larger than 1 is present in these results. The sizes of the deviations are different for each observer, but constant when varying the vertical separation angle. The size of the deviations of the slant was the same and for some observers even larger than the size of the deviations we found for experiments A and B. This is a point worth noting since the absolute distances from observer to objects were small in comparison to the other two experiments (80/160 cm versus 150/300/450 cm) whereas the horizontal separation angle was the same as the largest angle of experiment A, and the relative distances used were, among others, used in experiment B.

Concerning the tilt we found an effect of the vertical separation angle, except in the condition in which the observer had to direct the pointer upwards and away from himself. Although the variance explained by the regression model is rather small, the orientation of the pointer in the vertical plane is dependent on the vertical separation angle. When the vertical separation angle is ±20°, the observers overestimate the vertical separation angle slightly, i.e. they oriented the pointer a bit too high when the pointing-direction was upwards (vertical separation angle +20°) and a bit too low when the pointing-direction was downwards (vertical separation angle -20°). This is in agreement with the findings of experiments A and B in which the vertical separation angle was ±23°. Furthermore, an overestimation of the horizontal separation angle has been reported more often (Foley et al., 2004; Levin & Haber, 1993). However, when the vertical separation angle increases, the observers tend to direct the pointer too low when the pointing-direction is upwards and too high when the pointing-direction is downwards. One could say that for the larger vertical separation angles the observers underestimate the vertical separation angle. Although a linear increase with increasing vertical separation angle was found regarding the data of the tilt, just as in experiment A that investigated the effect of the horizontal separation angle on the settings of the slant, the deviations of the tilt were considerably smaller than we found for the data of the slant.

To investigate this effect of the tilt more thoroughly, we conducted a fourth experiment in which we varied the vertical separation angle in small steps between 7° and 45°.

3.6 Experiment D: Vertical separation angle 2

In order to investigate the effect of the vertical separation angle in more detail, we started this extra experiment to test whether we would still find a switch from positive to negative deviations with increasing vertical separation angle for upward pointing-directions (and vice versa, see Figure 3.3 B and C). To do this we used vertical separation angles between 7° and 45° in 6 steps.

Methods

In experiment D we used a configuration that resembles the configuration in Figure 3.1 C1. The distances, however, were different. The distances from the observer to the objects was 112 and 224 m and the horizontal separation angle 53°. Observers had to direct the pointer only from left to right to reduce the total number of trials. The pointer could be either closer to the observer than the ball or further away than the ball. The vertical separation angles we used were ± 7.7°, ± 15.3°, ± 22.8°, ± 30.0°, ± 37.5°, and ± 45.0°. As in the previous experiments, a positive deviation is defined as pointing upwards, whereas a negative deviation is defined as pointing downwards. Figure 3.1 C2 gives a side-view of experiment C. For experiment D a comparable configuration was used with six different vertical separation angles instead of three. Thus, for this experiment at both distances from the observer six dots below eye-height and six dots above eye-height would represent the positions of the ball and pointer correctly.

Results

Figure 3.10 shows the data for the slant. The slant was plotted against the vertical separation angle. Each point represents the mean of the data of the four observers. The error-bars give the inter-observer standard deviations. The full line represents the data when the pointer was closer to the observer than the ball, the dashed line when the pointer was further away than the ball. No effect of the vertical separation angle on the deviations in the horizontal plane was found. Furthermore, we found an overshoot (positive deviations) when the pointer was closer to the observer and an undershoot when it was further away from the observer than the ball. The deviations were significantly different from zero, except for three points out of 24 that show a trend towards significance. The large error-bars show that the size of the deviations was observer-dependent. This pattern replicates our findings in experiment C. However, similar shapes of the graphs were found for all observers. Furthermore, the within observer variability (not plotted in the figures) was quite small.

Figure 3.11 shows the data for the tilt. This graph is similar to Figure 3.10, except that it depicts the deviation of the tilt instead of the slant. The deviations of the tilt are smaller than the deviations of the slant. When the pointing-direction is upwards, these observers showed the same pattern as was seen in experiment C. Particularly, when the pointer was closer to the observer than the ball, a change from positive to negative deviations is visible around a vertical separation angle of 30°. However, when the pointing-direction is downwards one can only see an increase of the deviations when the separation angle increases. As in experiment C, we did a weighted least squares regression analysis to test whether the tilt depended linearly on the vertical separation angle. We conducted four analyses: two for pointing towards the observer (pointing upwards and downwards) and two for pointing away from the observer (also pointing upwards and downwards). The slopes of the regression lines were all significantly different from zero (see Table 3.3 for F and p values). Again, the R^2 values were low (see Table 3.3), but the low values merely reflect the individual differences in the size of the deviations; they are not caused by differences in pattern.

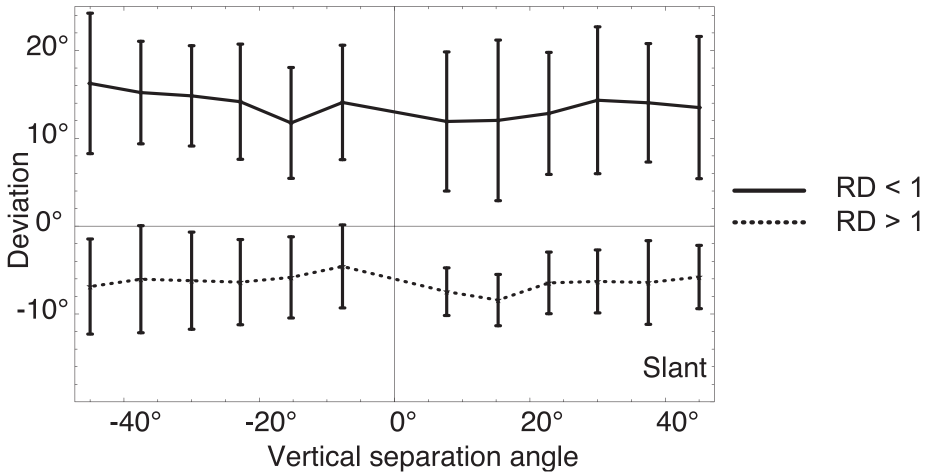

Figure 3.10
The data of the slant for experiment D. The deviation of the slant is plotted against the vertical separation angle. The solid line represents the data for a relative distance smaller than 1, whereas the dotted line represents the data for a relative distance larger than 1. Each data-point represents the mean of the values for the four observers. The error-bars represent the standard deviations.

Discussion

The data for the slant replicated earlier results. The tilt, however, was the main reason for adding experiment D to this paper. With regard to an upward pointing-direction, the results of experiment D replicated the results of experiment C. The deviations in the downward pointing-direction conditions were all positive, the size increasing with increasing separation angle. However, these deviations were very small, and probably did not differ from the deviations we found in experiment C. Thus, we must be cautious about speaking of a change in sign with increasing vertical separation angle.

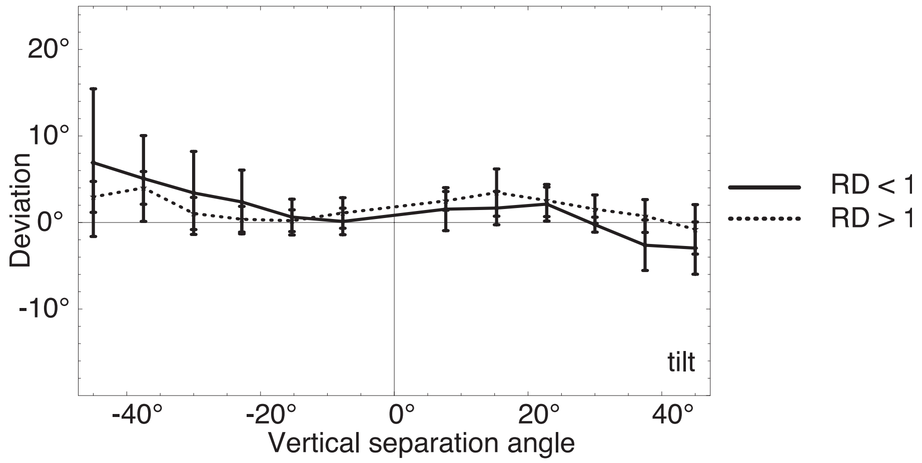

Figure 3.11
The data of the tilt for experiment D. The deviation of the tilt is plotted against the vertical separation angle. The solid line represents the data for a relative distance smaller than 1, whereas the dotted line represents the data for a relative distance larger than 1. Each data-point represents the mean of the values for the four observers. The error-bars represent the standard deviations.

3.7 General Discussion and Conclusions

First we will discuss the data for the slant. The deviations of the slant were not affected by the separation angle in the vertical plane. However, we did find a linear dependence on the horizontal separation angle when the relative distance was smaller than 1. When the relative distance was larger than 1, we found no significant dependence. However, we did find a dependence on the relative distance: the settings change from overshooting to undershooting when the relative distance changes from smaller to larger than 1. This is largely in agreement with the results we obtained with a 2D exocentric pointing task (Doumen et al., 2005 [Chapter 2]) and with the results of Cuijpers and colleagues (2000). This indicates that the extension to a 3D task did not change the demands of the task or the performance of the observers. This finding is not in agreement with the work of Schoumans and Denier van der Gon (1999) though the two studies are not quantitatively comparable because they used different definitions for slant and tilt. Furthermore, they used smaller distances and a different kind of setup (virtual instead of real-life).

The fact that the sign of the deviations of the slant changes with relative distance is interesting in the light of the traditional type of research on visual space. If visual space has a Riemannian geometry with a constant curvature, one would expect to find an overshoot or undershoot for all relative distances. Visual space would be elliptic if overshoots had been found for all relative distances and would be hyperbolic if undershoots had been found. Clearly this was not the case in the present work and in the work described in the introduction (Cuijpers et al., 2000; Doumen et al., 2005 [Chapter 2]). This means that visual space does not have a Riemannian geometry. Since in the horizontal plane, the observers overshot the position of the ball when the pointer was closer to themselves than the ball and

Table 3.3 *The F, p and R^2 for the regression analyses of the data of the tilt of experiment D*

Relative distance	Pointer position	Height	$F_{(1,6)}$	p	R^2
< 1	Left	Up	12.37	.002*	.360
< 1	Left	Down	8.65	.008*	.282
> 1	Left	Up	6.30	.020*	.223
> 1	Left	Down	8.49	.008*	.279

** The slope deviated significantly from 0*

vice versa, we can best describe the space in the horizontal plane with an expanding distance function (for a more elaborate discussion see Doumen et al., 2005 [Chapter 2]).

The curvature of visual space varies not only with relative distance but its structure seems to be dependent on a multiplicity of factors. Previously, we found that visual space differed for different tasks (Doumen et al., 2005 [Chapter 2]). Furthermore, factors like task-demands (Koenderink et al., 2000, 2002), distances of objects (Battro et al., 1976; Koenderink et al., 2002), observers (Battro et al., 1976) and viewing conditions (Wagner, 1985) can also influence the visual perception of spatial relations. Thus, visual space cannot be regarded as a well-structured geometrical entity.

The tilt was neither affected by the horizontal separation angle nor by the relative distance. It was, however, affected by the vertical separation angle. The signs of the deviations switch with increasing vertical separation angle in Experiment C. In experiments A and B, the vertical separation angle was held constant at 23°. In these experiments, we found that observers pointed too high when pointing upwards, and too low when pointing downwards. In experiment C, we see that there is a tendency towards the same conclusions when the vertical separation angle is only 20°, whereas the pattern reverses (pointing too high when pointing downwards and pointing too low when pointing upwards) when the vertical separation angle gets larger. One could therefore conclude from these observations that small angles are perceived to be larger than they are, and large angles are perceived to be smaller than they are. In the horizontal plane, a known tendency is the specific distance tendency (Gogel, 1965; Owens & Leibowitz, 1976; Yang & Purves, 2003). This is a tendency to see an object at a certain distance, when all information about distance is absent from the visual scene. It could be that a comparable tendency exists in the vertical dimension: a tendency to see an object at a certain angle from the horizontal plane at eye-height. In experiment D the effect we found in experiment C was replicated for the trials in which the veridical pointing-direction was upwards. For pointing downwards we did not find the switch from pointing too high to pointing too low. Here we found only a decrease in positive deviations. For the small negative vertical separation angles, the positive deviations were negligible, this is quite predictable since with a vertical separation angle of 7.7° the objects are almost in the same horizontal plane. Thus, we must be cautious about drawing conclusions based on the data for the tilt when we are dealing with small vertical separation angles.

We kept two dependent variables separate in this paper: the slant and the tilt. We did this because the way in which the data were collected. However, we did look at the total deviations. By total deviations we mean the total angular difference between the veridical pointing-direction and the observer's pointing-direction. We found approximately the same pattern as we found for the slant. Due to the difference in the size of the deviations the tilt did not contribute much to the deviations of the slant.

Nevertheless, we should relate the deviations of the slant and the tilt to each other. For example, if the visual space is expanded in the horizontal plane, the perceived vertical separation angle between two objects will decrease automatically. And this is in fact what we found for the large vertical separation angles in experiments C and D. Thus, if one assumes that there is no distortion in the vertical dimension, one would expect an underestimation of the vertical separation angle when the visual space is expanded. We need to be able to explain our findings for the small vertical separation angles. Observers seem to overestimate these angles when they are smaller than 30°. The sign of the deviation of the slant does not change in these conditions. Thus, we cannot assume that there is no distortion of the structure of visual space in the vertical orientation. However, this deformation in the vertical orientation is not fully understood so more research is needed to clarify this point. From the results it is clear that the slant and the tilt behave differently with varying spatial parameters. The deviations for the tilt are smaller than the deviations for the slant. Although we kept the dimensions of the set-up equal (both the vertical and horizontal separation angles varied from 20° to 60°), the dimensions of the experimental room were not equal in the horizontal and vertical direction. Thus we cannot discard possible contextual effects. More important, however, is the fact that the slant depends on the relative distance (see Figure 3.6) whereas the tilt shows no dependence at all on the relative distance. This means that the structure of visual space varies with direction. Visual space therefore is anisotropic; this conclusion contradicts Luneburg's assumption of homogeneity (Luneburg, 1950).

In summary, we can conclude that visual space is expanded in the horizontal direction. This is in agreement with the findings of Cuijpers et al. (2000), Koenderink et al. (2002, for distances up to 5 m from the observer), and Kelly, Loomis, & Beall (2004). In addition, the distortion increases when the distance between the objects increases horizontally. We cannot say much about the distortions in the vertical dimension. However, for large visual angles, we could describe the distortion-pattern as a decrease in the perceived visual angle, whereas for the small visual angles we could describe it as an increase in the perceived visual angle. Some of these findings correspond to our findings for previous work done in a horizontal plane. We conclude that the tasks make similar demands on the observer. The extra parameter, namely the height differences, did not influence the structure we found for the horizontal settings, but did influence the vertical settings. Thus, we can conclude that the structure of visual space is distorted in both the horizontal and vertical direction. The deformation, however, is not isotropic.