Visual perception of spatial relations in depth

Michelle Doumen
Universiteit Utrecht

Helmholtz Instituut

Cover design: Mirjam van Tilburg
ISBN-10: 90-393-4298-9
ISBN-13: 978-90-393-4298-5

Visual perception of spatial relations in depth

Visuele waarneming van ruimtelijke relaties in de diepte

(met een samenvatting in het Nederlands)

Proefschrift
ter verkrijging van de graad van doctor aan de Universiteit Utrecht
op gezag van de rector magnificus, prof. dr. W.H. Gispen
ingevolge het besluit van het college voor promoties
in het openbaar te verdedigen
op woensdag 6 september 2006 des middags te 12.45 uur

door
Michelle Jacoba Adriana Doumen
geboren op 10 mei 1977, te Roosendaal en Nispen

Promotoren: prof dr. J.J. Koenderink
prof. dr. A.M.L. Kappers

The investigations described in this thesis were financially supported by the Netherlands Organization for Scientific Research (NWO, ALW).

Contents

Chapter 1
Introduction

Having thought about visual perception for five years, I have become more and more impressed by the power and ability of the visual system. It is amazing that this system is able to abstract so much information that is useful for the observer from an enormous amount of available data. In some respects the visual process is similar to the memory process: the basis of remembering is forgetting everything that is not useful and the basis of visual processing is perceiving important aspects of the environment by ignoring the rest. But how does the system work? How does it combine different features into an object? How does it integrate changes over time? And an important question for this thesis: how does the brain derive depth from two-dimensional projections on the retinae? My introduction is twofold: first of all I would like the reader to share my admiration for the processes involved in visual perception, and secondly I would like to explain why the work I have been doing is important.

I will begin by giving a short introduction to the human visual system. After looking at reference frames and depth cues, I will look at the traditional visual space research and go on to discuss the newer paradigms that have been proposed. Finally I will introduce my own work that is described in detail in the following chapters of this thesis.

1.1 The Visual System

The visual system consists of sensory organs (the eyes) and brain-structures that process the signals that have entered the eyes. However, visual processing already begins in the retinae. I have no intention of discussing every aspect of the visual system here, since the purpose of this introduction is merely to introduce the reader to the basics of the system.

Light entering the eyes

Light enters the eye through the lens. The shape of the lens causes the beam of light to bend. This bending of the light rays enables the eye to focus on objects at different distances from the observer. If an object is close to the observer, the ciliary muscle contracts to make the lens spherical. The light-rays are therefore bent to a considerable degree. If an object is far away, the lens has to be flat in order to have the object in focus. This happens when the ciliary muscles relax, leading to less bending of the light-rays.

The light hits the retina after entering the eye. The outer layer of the retina contains the photoreceptors, the rods and cones, which contain pigments that absorb photons. Bipolar cells transmit the signal from the photoreceptors to the ganglion cells. These ganglion cells have long axons that leave the eye at the blind spot. The blind spot got its name because this area of the retina has no photoreceptors and thus gives no output to the visual system. In the bipolar layer, there are neurons that contribute to the communication between adjacent bipolar cells. The initial processing of visual information (lateral inhibition) takes place in these cells. The light has to enter through this layer of cells to reach the photoreceptors.

The fovea is the spot on the retina on to which the part of the visual field that is in focus is projected. This area of the retina contains a large amount of cones. The cells that are used for further processing are not in front of the photoreceptors here. This is why we have a particularly sharp image of the part of the visual field that is projected on to the fovea. Figure 1.1 gives a schematic picture of the eye.

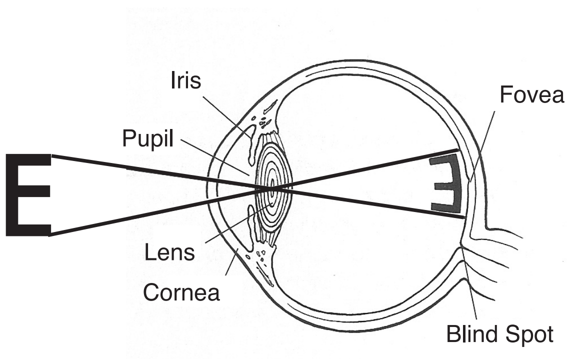

Figure 1.1 *Horizontal cross-section of the eye*

Visual processing

We can distinguish multiple types of ganglion cells. The most widely known ones are the M and P cells. The M cells have a larger cell-body and also have a larger receptive field. These cells are distributed fairly evenly throughout the retina. The P cells, on the other hand, are smaller and have small receptive fields. They are located in and close to the fovea. The axons of the ganglion cells together form the optic nerve as they leave the eye through the blind spot. The optic nerves of both eyes cross in the optic chiasm. Here the ganglion cells from the nasal areas of the retinae cross to the other side of the brain (see Figure 1.2). This causes the information from the left visual field to be projected on to the right side of the brain and vice versa. From the optic chiasm the optic tract passes through the lateral geniculate nucleus. From this point we can speak of two separate visual pathways: the magnocellular pathway and the parvocellular pathway; these are projected mainly by the M and P ganglion cells, respectively. The parvocellular visual pathway is responsible for color and form perception and the recognition of objects. It transmits detailed information that can only be derived from objects that are in focus, due to the presence of P ganglion cells primarily near the fovea. The magnocellular pathway is concerned with the perception of depth and movement, and derives information mainly from the M ganglion cells. These two pathways are not as strictly separated as is often assumed, as they interact at many stages of processing.

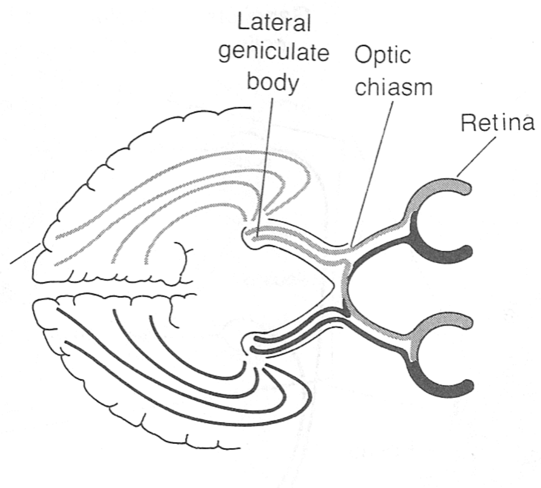

Figure 1.2
Horizontal cross-section of the brain showing the visual pathway from eyes to visual cortex

The visual pathways connect the geniculate nucleus to the primary visual cortex (V1) of the occipital lobe and then lead on to higher visual areas like the secondary visual cortex (V2). Some processing of depth-information takes place in V1, however spatial processing takes place primarily in the parietal lobe of the cortex, which is part of the magnocellular or dorsal visual pathway. There has been much discussion about whether this dorsal pathway is concerned mainly with visual control of action (Goodale, Milner, Jakobson, & Catey, 1991) or with a more general spatial representation of the world. However, according to Creem and Proffitt (2001), the discussion could be redundant if one thinks in terms of reference frames.

They suggest that in the parietal lobe different areas are involved in different reference frames. Since visual control of action and spatial representation of the world are based on different reference frames, this might explain the two theories of parietal lobe function (Creem, & Proffitt, 2001). However, others have suggested that the dorsal pathway is mainly concerned with egocentric representations (Neggers, van der Lubbe, Ramsey, & Postma, In press). Several authors suggest that it is the ventral pathway (the parvocellular pathway) rather than part of the dorsal pathway that is involved in the formation of allocentric representations.

1.2 Reference frames

The concept of reference frames is important in visual space research, and indeed for perception in general. By 'reference frame' we mean the locus (or set of loci) with respect to which the spatial position of an object is defined (Pick, & Lockman, 1981). Reference frames can be based on different aspects of the environment and are mainly divided into two types: egocentric and allocentric reference frames. By an egocentric reference frame we mean the relationship between objects and the observer. If observers use an egocentric frame of reference, they store information about positions of objects relative to their own position.

In an allocentric reference frame, on the other hand, one does not relate the spatial information to the observer, but to other objects in the environment. Of course, this frame of reference is dependent on the location of the observer. We will focus on the allocentric reference frame that could be used by the observers in our experiments. The frame consists mainly of the properties of the experimental room and the objects in the room. The experimental room has walls, a floor, a ceiling and some doors and windows. These elements together form several straight lines in the visual field of an observer. Based on a few assumptions about the usual shapes of objects, the visual input provides a basis for an allocentric reference frame. The assumptions people use are that windows and doors, for example, usually have rectangular shapes and that the walls are orthogonal to each other and to the floor and ceiling of the experimental room. If this is not the case, observers tend to make errors in judging the positions of objects. These reference frames provide a structure that an observer can use to make judgments about the positions of objects. Whether they actually use the available structure is a another question.

1.3 Depth cues

Humans have a fairly good sense of depth. However, the fact that we perceive depth is quite intriguing, since the image that is projected onto our retinae is two- instead of three-dimensional. Somehow, we derive the third dimension from these two-dimensional images. A number of possible sources of information (cues) have been described in the literature. However, the fact that these sources of information are present for observers does not mean that they actually use them to deduce depth in the scene. So we will have to be careful about assuming a certain degree of effectiveness when some cue is present. We can divide the various depth cues into three groups. I will discuss some physiological depth cues and pictorial depth cues below. Another important group of depth cues is derived from motion. I will not go into detail about these motion-related cues since they will not be discussed in the rest of the thesis.

Physiological depth cues

By physiological depth cues, I mean aspects of the physiology of the sensory organs that can contribute to the perception of depth. Ever since scientific interest in depth perception began, physiological depth cues have been regarded as very important. In fact, they are important for seeing depth in the range of 1 to 4 meters from the eyes. However, when distances from objects are larger, these cues are not as effective as people often assume.

> ***Box 1 Experiencing binocular disparity***
>
> Close your right eye and hold your right index finger at an arm's length from your nose. Hold the left index finger somewhere halfway between your eye and the right finger, so that your right index finger is blocked from view by the left one. Then close your left eye and open your right one. You'll notice that both fingers are visible now! This is due to binocular disparity.

An important physiological depth cue is binocular disparity. We have two retinae onto which two images of the world seen from slightly different positions are projected. These two images differ slightly from each other depending on the positions of the objects that are projected onto the retinae. Thus by combining the information of these images we obtain information about the relative distances between the observer and the objects in the scene. This difference between the two images of the retinae is called binocular disparity (see box 1 for a demonstration).

Another physiological depth cue is called accommodation. Accommodation occurs when the lenses of our eyes deform when the distance of the objects in focus changes. Accommodation of the lens is required to produce a sharp image of the object. This cue is derived from the muscle tension that is needed to have an object in focus; hence it can produce an absolute distance estimation as a function of muscle-tension.

A third physiological depth cue involves the vergence movement of the two eyes. If a person focuses on an object, then the two eyes will rotate so that the object is projected onto the fovea of the retinae. If this object is close to the observer, the two eyes will rotate inwards, but they will rotate outwards if the object moves further away from the observer. The tension of these muscles will therefore also give information about the absolute distance to the object that is in focus.

Pictorial depth cues

Although many disciplines already recognized the importance of pictorial depth cues, for many years the scientific community focused its attention mainly on physiological depth cues. However, pictorial depth cues are very important for perceiving depth. By pictorial depth cues we mean the cues that we can use to see depth in a picture. However, they are very useful for estimating depth in the actual world. There are lots of different pictorial depth cues, so I do not intend to name them all.

One of the oldest known depth cues is occlusion; when an object A is occluded by object B, object B is closer to the observer than object A. This is not an absolute cue to distance; thus one cannot give an accurate judgment of the distance from one of the objects solely on the basis of this cue. However, occlusion is a very reliable source of relative distance information for multiple objects.

Another well described pictorial depth cue is linear perspective. This depth cue makes use of the fact that parallel lines in the physical world around us seem to converge to one single point on the horizon(when not in a frontoparallel plane). A closely related cue is the relative size cue, which involves the fact that objects further away from an observer produce smaller retinal angle sizes than objects that are closer to the observer. Together with some knowledge of the sizes of familiar objects, this cue can even give absolute estimates of the distances of single objects. Gibson (1950) introduced the term texture gradient to refer to an effective depth cue in our environment. This cue is related to both the relative size cue and linear perspective. Texture gradients involve the fact that textures on surfaces change in size and compactness as the distance from the observer changes. Other pictorial depth cues are aerial perspective, height in the visual field etc.

Prior knowledge

There are numerous cues that contain information about depth. How an observer chooses which sources of information to use is also subject of debate. First, some scientists reasoned that observers weigh up the reliabilities of cues and will rely more on cues that have proved to be most useful in the past (Ames, 1953). In the Ames' room example (Box 2), linear perspective apparently is considered a more reliable cue than relative size (of the people in the room). Another possibility is that observers choose the possible lay-out of the scene that has been encountered most often before (Gibson, 1966). A rectangular shaped room is more familiar to most people than a trapezoidal shaped room; thus the Ames' room of Box 2 is seen as a rectangular shaped room. According to Yang and Purves (2003), our brains use Bayesian calculations to let us perceive the most probable scene that could produce the retinal image. These two views on the weighing up of cues or situations need not be contradictory, they can complement each other.

As well as finding out what kind of information sources we use to see depth, we can also investigate the accuracy or veridicality of our depth perception. I will refer to this work as visual space research. By visual space (short for visually perceived space) I mean the impression we have of our environment (physical space) on the basis of visual information.

1.4 Visual space research

Visual space research has a long tradition. The first systematic experiments were done under rather restricted conditions: they were done in dark rooms with small point-lights as stimuli and the heads of the observers were fixated. In this way scientists reduced the available cues to the physiological depth cues described above. This is why they talked about binocular visual space. Examples of these experiments are those performed by Hillebrand and Blumenfeld (Battro, di Pierro Nettro, & Rozestraten, 1976). They let observers make visual alleys consisting of points of light. These alleys were based either on parallelity (two rows of lights had to be placed parallel) or on equidistance (each pair of points is equidistant from the other pairs of points). In a Euclidean geometry these two tasks should result in the same settings, the only difference being the explanation given to the observers. However, differences were found between these two tasks. The difference between settings for these two tasks inspired Luneburg (1947, 1950) to formulate his conjecture that binocular visual space has a homogeneous non-Euclidean Riemannian metric. In other words there is a single visual space that is constant over viewing conditions, distances and tasks. If the curvature of visual space is zero, this would mean that visual space is Euclidean, which is the description

Box 2 Ames' room

One way of introducing the notion of depth cues is to explain the effects of Ames' room. Ames was an artist who was interested in depth perception. He wrote a book with lots of small experimental set-ups that clarified issues that had not been solved at the time (Ames, 1953). His most famous set-up is the "room" he created. To an observer, who views the room from a certain viewpoint, it seems to be a normal rectangular shaped room. However, the people in the room seem to be substantially different in size. When you look at the ground-plan of the room, you can see that the room is not rectangular-shaped at all and the people in it are approximately the same size but are at different distances from the observer. This effect is due to the strong effect of familiarity with rooms. Usually rooms are rectangular shaped, and windows and doors are also rectangular shaped. Thus a trapezoid that is extended in depth is easily mistaken for a square that is in another orientation.

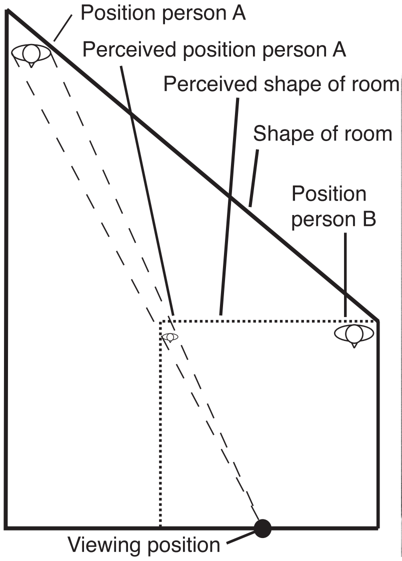

we give to the physical space. In a Euclidean metric, parallel lines never intersect and remain equidistant from each other. If the curvature is positive (the metric is said to be spherical) or negative (the metric is said to be hyperbolical) parallel lines do not remain at the same distance from each other. Blumenfeld, however, observed that equidistance alleys always lay outside parallel alleys and did not look parallel. This observation led Luneburg to conclude that the Riemannian metric of visual space had a negative (hyperbolic) curvature. Most of the early scientists agreed with Luneburg about the curvature of this metric (Blank, 1961; Zajaczkowska, 1956); however, there is still no consensus on this matter.

In order to generalize our knowledge to everyday vision, we need to take into account the other cues that contribute to our sense of depth in the visual world. Other cues like motion-related cues and pictorial cues provide a rich source of information about depth. For example, people with one eye can see depth, although they have no binocular disparity or convergence information available. Due to some behavioral adaptations, however, they have no difficulty in dealing with the world. Therefore we need to look at all depth cues in order to try and understand human depth perception. So far, most research into visual space that was done under less restricted conditions has taken place in outdoor settings. Larger distances were used than in the experiments discussed earlier, pictorial information was present (due to testing in daylight) and often observers could move their heads freely. Gilinsky (1951), for example, used a bisection task to study the deformation of visual space. In her experiments, a line was stretched across a lawn, starting a few inches from the observer. The task was to bisect the line in two equally long parts. She concluded that visual space can be described with a compressing distance function that is asymptotic to a constant (c in Equation 2.1). This distance function was derived from Luneburg's metric of visual space, the principles of linear perspective and the law of size constancy. The distance function has constant c that varies between observers and experimental conditions. Gilinsky theorized that c was also dependent on the availability of cues to distance.

Battro et al. (1976) did experiments with different tasks in large open fields. They compared the curvature of visual space using three tasks: visual alleys, horopters and triangles. Battro and colleagues (1976) concluded that visual space can be described by a Riemannian metric with a variable curvature. The curvature varied with scale and between observers. They demonstrated that visual space is not homogeneous in nature and thereby they falsified one of Luneburg's assumptions. Foley (1972) and Koenderink, van Doorn and Lappin (2000) came to the same conclusions.

The work described in this introduction mainly concerns measurements in horizontal planes. Indow and Watanabe (1984, 1988) examined not only the horizontal sub-space, but also the frontoparallel sub-space. They found no systematic deviations in this plane. Hence, they concluded that visual space is deformed in different ways for different sub-spaces: hyperbolic for the horizontal sub-space and Euclidean for the frontoparallel sub-space.

The aim of scientists nowadays is not to falsify Luneburg's conjecture, but to find new ways of describing visual space. For example, Kelly, Loomis and Beall (2004) used other tasks (a body pointing task and a collinearity task) to measure visual space in an open field. They concluded that systematic misjudgment in exocentric direction could not be caused by misperceived egocentric distances. However, Koenderink, van Doorn and Lappin (2003) were able to explain the misjudgments of egocentric direction, measured with an exocentric pointing task, in terms of a misperception of egocentric distances. Thus, we can conclude that not all the work done so far can be explained by one existing theory. Experiments have even produced conflicting results. Thus, one cannot speak of a single visual space. Instead, the structure of visual space seems to depend on numerous factors like the kind of task that is being used (Koenderink et al., 2000; Koenderink, van Doorn, Kappers, & Lappin, 2002), the viewing conditions under which the experiments are done (Wagner, 1985), the distances that are used and individual differences between observers (Battro, et al., 1976; Koenderink, et al. 2002).

Cuijpers and colleagues did laboratory research on visual space in an indoor setting. In a laboratory environment they tested whether observers showed systematic deviations from veridical settings under normal lighting conditions. Their observers were seated in an artificially illuminated room. To prevent the observers from deriving pictorial depth information from structures besides the actual objects used in their experiments Cuijpers and colleagues covered the walls of the experimental room with wrinkled plastic. They had observers seated in a small cabin that restricted the vertical field of view. These manipulations led to the situation where the observers could not see the floor, ceiling or walls of the experimental room. In addition to these restrictions, the heads of the observers were fixated with a chin-rest. And lastly, the objects that were used to do the task were scaled with distance from the observers, so that the observers always received an equally sized projection onto their retinae.

Cuijpers and colleagues did their research via three different tasks. The first task was an exocentric pointing task (Cuijpers, Kappers, & Koenderink, 2000A). In this task, a pointer could be rotated with a remote control. The task for the observer was to direct the pointer towards a target, which was a small sphere. The second task they used was a parallelity task (Cuijpers, Kappers, & Koenderink, 2000B). During this task, two rods are in the visual field of the observer. The experimenter placed one of the rods in a certain orientation. The observer's task was to place the other rod in the same orientation as the first one, i.e. to put it parallel. The last task was the collinearity task, in which also two rods were placed in the observer's visual field (Cuijpers, Kappers, & Koenderink, 2002). The task was to rotate them both so that they were in line, i.e. collinear. By means of these experiments Cuijpers and colleagues produced data that led them to the conclusion that the structure of visual space is dependent on the task that is given to a certain observer. Therefore, in their view one cannot speak of a geometry of visual space in general (Cuijpers, et al., 2002).

1.5 This thesis

Our first experiment was based on Cuijpers' work. Cuijpers and colleagues investigated whether a single visual space could also be defined in an illuminated environment. Although they found structural deviations, they found varying patterns for different tasks. We began by extending this work by introducing free viewing conditions (Chapter 2). Secondly, we changed the 2D exocentric pointing task into a 3D version, and tried to extrapolate to 3D scenes the knowledge we had gained from experiments in the horizontal plane (Chapters 3 and 4). Thereafter, we investigated in detail an interesting pattern we had found in the earlier chapters (Chapter 5). And finally, we compared our exocentric pointing task with a new task that we developed (Chapter 6).

Two lines of research can be distinguished in this thesis. One concerns spatial parameters and their effects on the perception of the positions of objects (Chapters 2, 3 and 6). The other line mainly concerns the effects of contextual information and reference frames (Chapters 4 and 5). This work should help us to understand which spatial and contextual parameters influence our perception of depth.

Chapter 2

In this chapter we discuss three experiments that we did to test whether systematic deviations of the visual perception of the positions of objects still occur under free viewing conditions. By free viewing conditions we mean that the observers had an unobstructed view of the walls, floor and ceiling of the experimental room. Furthermore, they could rotate their heads and upper-bodies freely. We used the exocentric pointing task, the parallelity task and the collinearity task as described above. All measurements were made in the horizontal plane on eye-height. We looked at the effect that relative distance and the horizontal separation angle had on the observers' settings and compared our findings to the results reported by Cuijpers and colleagues.

Chapter 3

In Chapter 3 the 3D exocentric pointing task is introduced in order to explore visual space in three dimensions. The pointer and ball could be positioned at different heights. In addition to rotating the pointer in the horizontal plane, the observer could also rotate it in the vertical plane. Thus, we used two dependent variables in this research: the slant and the tilt (the orientation in the horizontal and vertical orientation). We varied the horizontal and vertical separation angles and the relative distance. If visual space were isotropic, then the tilt would be dependent on the vertical separation angle as the slant is dependent on the horizontal separation angle. Moreover, both variables should be dependent on the relative distance.

Chapter 4

In this chapter we describe the experiments that were conducted to test whether external and internal references were used during a 3D exocentric pointing task. We compared a condition in which observers directed the pointer towards a ball while the pointing-direction was parallel to one of the walls with a condition that the pointing-direction was not parallel to one of the walls. Besides this, we compared conditions in which the pointing-direction was frontoparallel or not. By frontoparallel we mean in a plane perpendicular to the line of sight when one is looking straight ahead.

Chapter 5

In all our experiments, we found a difference for trials in which the pointer was positioned far from the observer (and the ball close by) and trials in which the pointer was positioned close to the observer (and the ball far away). The deviations we found were larger in the latter condition than in the first condition. In this chapter, we tried to find an explanation for this phenomenon. We tested two different explanations. First, we tested whether the position of the observer was used as a reference when the pointer was far from the observer. The position of the observer does restrict the pointing angle when the pointer is far from the observer. In our experiments we restricted the pointing angle to the same degree for the other condition (when the pointer is close to the observer) by means of poster-boards. In another experiment, we tested whether the different views of the pointer could explain the difference in deviations. If a pointer gives less information about its orientation when the pointer is close by and the ball far away, this could result in asymmetry in the amount of information an observer can use to orient the pointer.

Chapter 6

In the last chapter of this thesis we investigate visual space by means of an entirely different task. In this task, three red balls are hanging at different heights from the ceiling. A fourth ball can be adjusted in height by the observer. The task is to hang it in a plane that is defined by the three red balls. In this chapter, we describe the effect that the orientation in which the plane is tilted has on the deviations. Furthermore, we tested this with the red balls forming three different triangles; namely an acute, an obtuse and an equilateral triangle.

Chapter 2
Visual Space under Free Viewing Conditions

Abstract:
Most research on visual space has been done under restricted viewing conditions and in reduced environments. In our experiments, observers performed an exocentric pointing task, a collinearity task and a parallelity task in an entirely visible room. We varied the relative distances between the objects and the observer, and the separation angle between the two objects. We were able to compare our data directly with data from experiments in an environment with less monocular depth information present. We expected that in a richer environment and under less restrictive viewing conditions the settings would deviate less from veridical settings. However, large systematic deviations from veridical settings were found for all three tasks, but the structure of these deviations was task-dependent. The deviations were comparable to those obtained under more restricted circumstances. So the additional information was not used effectively by the observers.

Published as:
Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (2005). Visual space under free viewing conditions. Perception & Psychophysics, 67, 1177-1189.

2.1 Introduction

Humans are capable of interacting adequately with their surroundings on the basis of visual input. We can estimate the positions of objects well enough to interact with them effectively. However, sometimes we make large errors in estimating the positions of objects. For example, when someone standing next to you points to a person who is surrounded by other people, it may be hard for you to locate the person whom he or she is actually referring to.

People tend to make significant and systematic errors when estimating the positions of objects in a visual scene. It is as if "visual space", i.e. how we visually perceive the world around us, is deformed with respect to physical space, the physical lay-out of the world. Over the years many researchers have been interested in quantifying this deformation and in giving it a geometric description. Early experiments on visual space were generally done in dark rooms where observers viewed the stimuli binocularly while their heads were fixed. Such a setup effectively removes monocular distance cues except for accommodation. Thus, in these experiments the focus was on binocular depth cues. Helmholtz (1962), for example, measured apparent frontoparallel planes. Vertical threads that hung in a physically frontoparallel plane were not judged by observers to be in one plane. Therefore, Helmholtz concluded that the apparent frontoparallel plane is not the same as the physical one, but that it is curved. Other early research was carried out with luminous points as stimuli; observers had to do visual tasks which involved rearranging the points. The experiments of Blumenfeld and Hillebrand, who let observers make visual alleys based on parallelity or equidistance, inspired Luneburg to formulate a model for visual space (Luneburg, 1950). For these kinds of tasks under these conditions he suggested that visual space has a Riemannian geometry of constant hyperbolic curvature. Both Zajaczkowska (1956) and Blank (1961) confirmed this notion. Indow and Watanabe (1984, 1988) found that the metric of visual space varies over different planes in the visual world. They found a Euclidean metric for the frontoparallel plane (Indow & Watanabe, 1984, 1988) and Indow (1991) found a curved Riemannian metric for the horizontal plane at eye-height. This suggests that visual space is anisotropic, which contradicts Luneburg's assumption of isotropy. Due to the fact that these experiments were conducted in dark rooms, most monocular depth cues were not available to the observers. So if we want to generalize this knowledge to everyday vision, we should extend this research with experiments done under normal lighting conditions.

Some researchers concentrated on experiments in large open fields in normal daylight. In these open field experiments, distances between objects and observer are larger. Thus, different kinds of information (mainly monocular) become important when observers do tasks involving the estimation of depth in a scene. For example, at distances of more than four meters, binocular depth cues play a less important role. Testing in daylight, contrary to testing in the dark, provides depth-information from pictorial cues like linear perspective and texture. Gilinsky (1951), for example, did research aimed at obtaining insight into the relationship between the perceived distance and the perceived size of objects at distances of up to 22 m (70 feet). She developed a law to describe the compression of visual space perception she found in her experiments. She described perceived distance (P) with the following formula:

$$P = \frac{cr}{c + r} \tag{2.1}$$

where r is the physical distance and c is a constant that represents the distance at which an observer perceives objects that are an infinite distance away. In near space perceived distance is approximately equal to physical distance, but as the physical distance increases perceived distance increases less and less until it saturates at the distance c.

Some scientists concluded that there is no single geometry that can describe visual space under all conditions. For example, Battro, di Pierro Netto and Rozestraten (1976) and Koenderink, Van Doorn, & Lappin (2000) found that the curvature of visual space changed from elliptic to hyperbolic as the distance from the observer increases. Besides that, Koenderink, Van Doorn, Kappers and Lappin (2002) concluded from experiments that the structure of visual space varies over different tasks. Thus, studies that confirm theories with different geometries are not necessarily contradictory, they are merely complementary. According to Wagner (1985), visual space has an affine-transformed Euclidean geometry under full-cue conditions with free head-movements but restricted body-movements. The metric will get close to Euclidean when the perceptual information increases both quantitatively and qualitatively (Wagner, 1985).

Recently, Cuijpers, Kappers and Koenderink (2000a, 2000b, 2001, 2002) did indoor experiments in a room where most pictorial depth cues were eliminated from the visual field by means of wrinkled plastic that prevented observers from seeing the walls. The floor and ceiling of the rooms were not visible due to the fact that the observer was seated in a cabin that restricted the vertical visual field of view. The head of the observer was fixed using a chinrest. In the tasks they used, rods had to be made to point towards a target (exocentric pointing task), towards each other (collinearity task) or had to be put parallel to another rod (parallelity task). The angular deviations from veridical settings were measured. The pattern of these deviations was found to depend on the task. Cuijpers et al. (2002) claimed that there is no such thing as an invariant visual space because the form of the visual space is task dependent.

The experiments of Cuijpers et al. (2000a, 2000b, 2001, 2002) were done with artificial light, distances were less than 4.5 meters and most pictorial depth-information was eliminated from the scene. In this way the information present came mainly from physiological depth-cues. In everyday vision pictorial depth-information is available to an observer. Therefore, to be able to say something about everyday vision one needs to look at how visual space is deformed when contextual information is present. Normally people do not look at luminous points in a totally dark environment. Generally, they look at objects that are surrounded by other objects that can give a great deal of information about the relative positions of the objects. This is evident from the fact that one obtains spatial impressions from flat photographs where physiological cues are lacking or are inconsistent with monocular cues. According to Gibson (1950) visual space is dependent on what fills it; thus in studying visual space one should also look at contextual information. Another example that stresses the amount of information that can be provided by monocular depth cues comes from the perception of amblyopes, people that are unable to use binocular depth information.

Nevertheless they have no trouble perceiving depth in a normal environment. In fact very often they only discover their deficiency when they are subjected to stereo-tests.

Since monocular information provides a rich amount of information about depth in a scene, it seems logical to examine the perceived spatial relations between objects in an environment that provides both monocular and binocular depth cues. So the purpose of the present research is to increase our understanding of the structure of visual space as it occurs to us when both types of depth cues are present. To do this we studied the structure of visual space in an illuminated room under free viewing conditions, i.e. viewing without any restrictions on head movements or size of the visual field. We will test whether visual space is systematically different from physical space under free viewing conditions in a room where monocular depth information is available. Our hypothesis is that in a richer environment for similar tasks the settings will deviate less from veridical settings. We expected this because more pictorial depth cues are present and thereby one would expect more precise estimation of positions of objects (Wagner, 1985). Another issue in this research is whether the differences that Cuijpers et al. found between the different tasks, are also present in our setup.

The research was done in a room in our laboratory. The observers were seated, could rotate head and upper-body if they liked and they had an unobstructed view of the floor, ceiling and walls of the experimental room. We used three different tasks. One task was an exocentric pointing task in which the observer had to direct a pointer towards a target. The second task was a parallelity task in which a rod had to be put parallel to another rod. The third task was a collinearity task in which two rods had to be placed in one line. We manipulated two different parameters for the three tasks. One of these parameters was the relative distance, which is the ratio of the distances between the two objects and the observer. The second parameter was the separation angle, i.e. the visual angle between the objects. For the parallelity task we had a third parameter, namely the orientation of the reference rod. The separation angle and the relative distance were chosen as parameters because together they can quite naturally give an indication of the positions of the objects with respect to the observer. Besides that, they were the major parameters in the experiments of Cuijpers et al. (2000a, 2000b, 2002). Since we want to test whether a room full of depth information will change the structure of visual space, it is important to be able to use Cuijpers' data as a baseline for our measurements.

2.2 General methods

Observers

The three tasks described in this paper involved the same four observers. The observers were undergraduates and were paid for their efforts. They all had normal or corrected to normal sight and were tested for binocular vision. All observers had stereovision with good acuity. The observers had no knowledge about the goal of the experiment and received no feedback regarding their performance during the experiment. They were tested individually.

Experimental setup

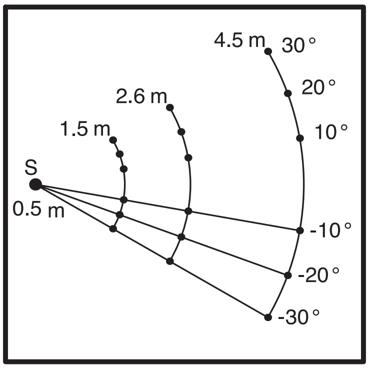

Figure 2.1
Schematic view of the experimental room. The black dots indicate the positions of the objects in the room. The larger black dot represents the position of the observer (S).

The experiment was set up in an empty room measuring 6m by 6m by 3.5m. On the left-hand wall blinded windows were visible. Under the windows were central heating radiators. Opposite the observer was an empty wall and on the right of the observer was a wall with two doors. On every wall, electric sockets were visible near the floor of the room. The ceiling was partly covered with oblong fluorescent lights and air-conditioning equipment. The room was illuminated with these artificial lights. On the floor, points were marked for the positioning of the objects. These points were marked at three different distances from the observer (1.5 m, 2.6 m and 4.5 m) at three different angular separations (20°, 40° and 60°) symmetrical around the line bisecting the room (see Figure 2.1). The objects used in the tasks consisted of yellow disks perpendicular to green rods. The rods were 25 cm long and 1.0 cm thick, and were sharp at each end. The disks had a diameter of 8.2 cm and a thickness of 1.0 cm. The rods were placed at eye-height and could be rotated around the vertical axis. The rods that were used as objects are depicted in Figure 2.2. The observer used a remote control to rotate the rods. The feet of the objects were square-shaped and contained a protractor from which the experimenter could read the pointing direction. A screen in front of the foot of an object prevented the observer from seeing the protractor and the square which was aligned with the walls. The observer's chair could be adjusted so that the objects were at eye-height. No chinrest was used and head movements were permitted.

Procedure

For every trial, two objects were placed on the marks in the room. We used three different separation angles (20°, 40° and 60°) and the objects were at three different distances from the observer (1.5 m, 2.6 m and 4.5 m). For every combination of positions on the floor, one object was always on the left of the observer and the other on the right. This setup gives a

Figure 2.2
A picture of the two rods that were used in the parallelity task and the collinearity task.

total of 27 possible combinations of positions on the floor (3 separation angles, 3 distances to one object and 3 distances to the second object).

For the analysis we used positive and negative separation angles. Positive separation angles (20°, 40° and 60°) were used for the trials in which one of the objects, the reference object, was positioned to the left of the observer and the other object, the test object, was positioned to the right. Negative separation angles (-20°, -40° and -60°) were used when the reference object was on the right of the observer, and the test object on the left.

The observers were allowed to move their heads, but were told to stay seated. In between the trials, the observers were asked to close their eyes so that they could not see the movements of the experimenter and the objects while the experimenter read the pointing direction of the test rod and changed the setup for the next trial.

All observers were tested with the three tasks. One task was completely finished before starting a second one. The order of the experiments was partially counterbalanced. This way, every observer had a unique order of experiments. The experiments were all conducted in sessions of approximately one hour. Mostly, the observers were tested for one hour each day, but sometimes we had two sessions a day with a break of at least 30 minutes in between the sessions.

2.3 Experiment 1: Exocentric pointing task

Methods

The exocentric pointing task involved the use of a pointer and a target. The target was an orange sphere with a diameter of 6.5 cm and was positioned at the same height as the pointer (at eye-height for the observer). For this task an object as described above was used as pointer, the only difference being that it had only one sharp conical end. The task was to rotate the pointer in such a way that it pointed towards the centre of the target. Each position on the floor was used as reference position (position of the target) and as test position (position of the pointer), so the number of combinations (27) has to be multiplied by two. Because we repeated every possible combination three times, the total number of trials

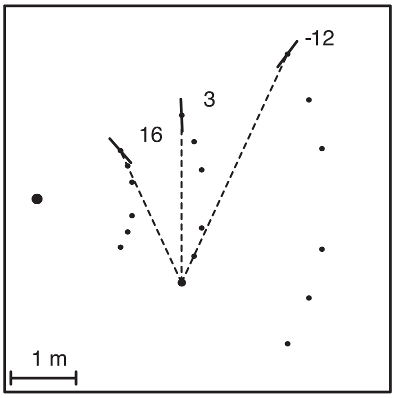

Figure 2.3
Schematic view of the experimental room with settings for the exocentric pointing task. The dashed lines represent veridical settings for the exocentric pointing task for conditions where the separation angle was -60°. The distance from the reference stimulus to the observer was 2.6 m. The solid lines represent the means of the settings for three trials for observer AW. They are shown for three different distances between the observer and the pointer. The numbers indicate the mean deviation from veridical settings in degrees.

needed for this task was 162 (27x2x3). It took each observer about three hours to complete the 162 trials.

Results

In the exocentric pointing task qualitatively similar systematic deviations were found for all observers, although the magnitude of the deviations was observer-dependent. In Figure 2.3 an example is given of the settings for observer AW, a separation angle of 60° and a distance of 2.6 m to the target. The lay-out of the floor is presented in this figure together with the three combinations of positions of target and pointer that are possible with a fixed target-position. The dotted lines represent the veridical pointing directions and the small thick lines represent the means of three settings of the observer. The numbers indicate the deviation from veridical directions in degrees. It is important to notice that the sign changes when the relative distance switches from larger than 1 to smaller than 1.

The influence of two parameters was analyzed, i.e. the separation angle and the relative distance. Figure 2.4 shows graphs for observer AW in which the deviation is plotted against the separation angle. A line is fitted through the data points in the figure using a least squares method. The five graphs represent the five relative distances that were used. For a relative distance of 0.3 and 0.6 the slopes are positive and for a relative distance of 1.7 and 3 the slopes are negative. For a relative distance of 1 the fit approaches a horizontal line. Table 2.1 gives the slopes in numbers for each observer and each relative distance. An asterisk indicates whether the slope deviates significantly from zero at a confidence level of .95. The same pattern was found for three observers, i.e. the slope deviated significantly from zero when the two rods were at different distances from the observer. The deviations for observer TL were very small and therefore the slopes were smaller. The results for BL, also, showed two slopes that did not deviate significantly from zero. However, in general the observed deviations could very well be approximated as a linear function of the separation angle.

Figure 2.5 is a graph in which the slopes of the lines in Figure 2.4 are plotted against the relative distance on a logarithmic scale. A large slope means that the deviations were very large for the larger separation angles, and vice versa. So the slopes in Figure 2.5 give indirect information about the size of the deviations for one relative distance. The data for the four observers are plotted on separate lines. As can be seen in the Figure 2.5, the lines have the same shape, which indicates that the observers show comparable behavior. We will address three points. First, an overshoot was present when the relative distance was smaller than 1 (the pointer was closer to the observer than the target). This means that the observers directed the pointer towards a point further away in depth than the target actually was. In contrast, an undershoot was present when the relative distance was larger than 1 (the pointer was further away from the observer than the target). This means that the observer directed the pointer towards a point closer to himself than the target actually was. Second, the size of the slopes approached 0 for a relative distance of 1. When the target and the pointer were at the same distance from the observer, the settings were almost veridical. Third, the slopes were larger for relative distances of 0.6 and 1.7 than for relative distances of 0.3 and 3.

Discussion

The deviations increase linearly with the separation angle. This means that when the two objects are further apart, i.e. the visual angle is large, the deviations are larger. An effect of relative distance was also found. There is an overshoot when the target is further away from the observer than the pointer and an undershoot when the target is closer to the observer. At a relative distance of 0.6 or 1.7 the deviations were larger than for the relative distances of 0.3 and 3. This could be due to the fact that when the relative distance approaches zero or infinity, the angle between the line connecting the observer with the more distant object and the line connecting the two objects approaches zero. In these cases, the task resembles an egocentric pointing task. Thus, the closer the relative distance approaches zero or infinity, the smaller the deviations are likely to be.

In the case of two observers (AW and TL) when the two objects were at the same distance from the observer (a relative distance of 1), the settings were close to veridical. For the other two observers a small overshoot was present at this relative distance.

Comparing these results with the data of Cuijpers and colleagues (2000a), we see that the same pattern of results was found for the relative distance. Cuijpers and colleagues (2000a) did not find any effect of the separation angle. However, they looked only at the separation angle for

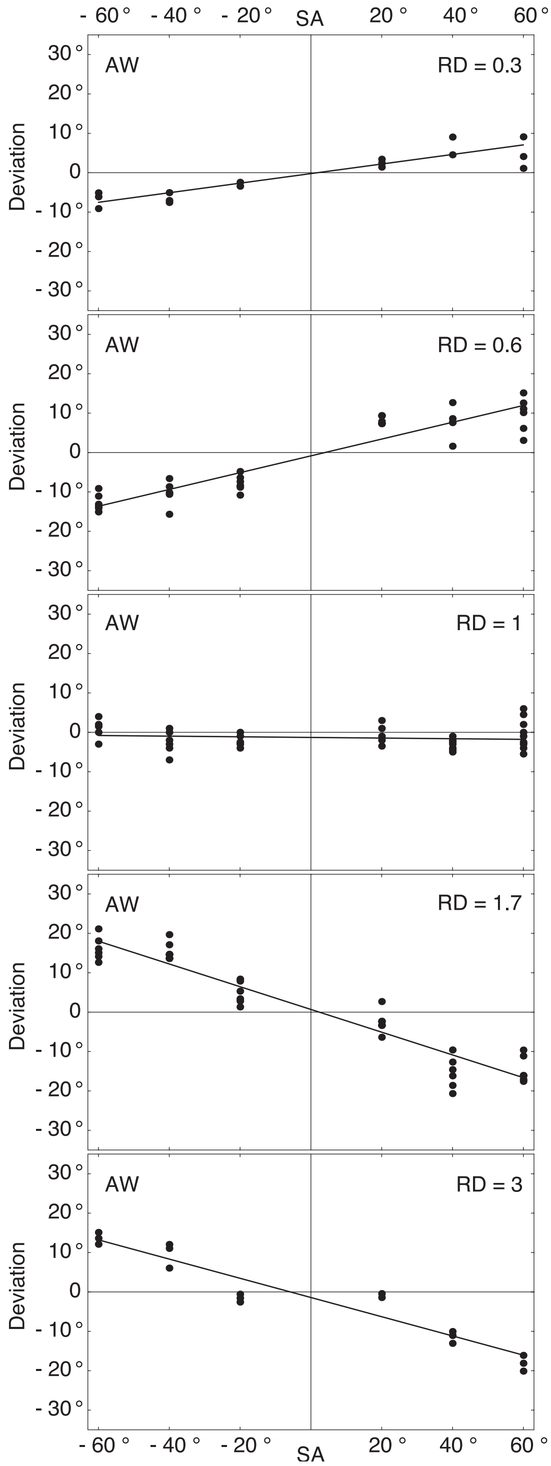

Figure 2.4
In each graph the deviations from veridical settings for the exocentric pointing task are shown as a function of separation angle for each relative distance. These are the data for observer AW. A line is fitted through the data points.

combinations of positions with a relative distance of 1 and for these trials we also found only minor deviations.

Table 2.1.
The Slopes of the Linear Fits through the Data Points as a Function of Relative Distance for all Observers in the Exocentric Pointing Task

	Relative Distance				
Observer	0.3	0.6	1.0	1.7	3.0
JP	0.11*	0.18*	-0.08*	-0.17*	-0.07*
AW	0.12*	0.21*	-0.01	-0.29*	-0.24*
BL	-0.02	0.04	-0.12*	-0.27*	-0.23*
TL	0.01	0.12*	-0.	-0.07*	-0.01

* *The slope deviates significantly from 0 (α = .05).*

2.4 Experiment 2: Parallelity task

Methods

For the parallelity task two rods were used, as described in the general methods section. One of the rods, the reference rod, was placed at a certain orientation by the experimenter. The task for the observer was to rotate the other rod, the test rod, so that the two rods were parallel. To clarify the word parallelity, we gave the observers an example of two parallel lines on paper. The reference rod could be either on the left or the right side of the room. Thus, as in the pointing experiment, the number of combinations of positions was doubled. Furthermore, four different orientations of the reference rod (22°, 67°, 112° and 157°) were used for every combination of the reference and test positions. We chose these orientations so that we could compare four oblique orientations with an even amount of rotation between them. We repeated all the measurements three times. As a result, the experiment consisted of 648 trials for each observer (27x2x4x3). Usually 54 trials were performed per session. Each session lasted an hour, so 12 hours were needed per observer.

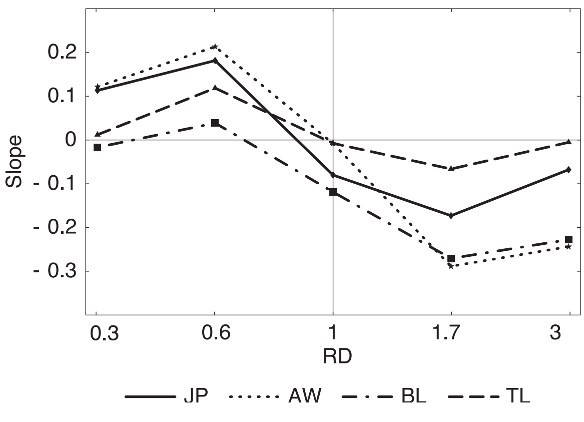

Figure 2.5
In this figure the slope of the lines from Figure 2.4 are plotted as a function of relative distance. The data for the different observers are given on separate lines.

Results

An example of the settings of observer BL with one distance to the reference rod and one orientation is given in Figure 2.6. The figure shows the layout of the floor of the experimental room. The point on the left represents the position of the observer, the other points the positions that were used for the rods. The figure gives both a graphical and a numerical view of the data. The lines and numbers (in degrees) represent the means of three trials for a reference orientation of 67° and a distance of 4.5 m between observer and reference rod. For this reference distance all possible combinations with the test distance are shown, as well as three different separation angles. The test rods on the outer line were tested with the reference rods on the outer line on the other side of the room (separation angle of 60°), the rods on the middle lines were also tested together (separation angle of 40°) and the same was done for the rods on the inner lines (separation angle of 20°).

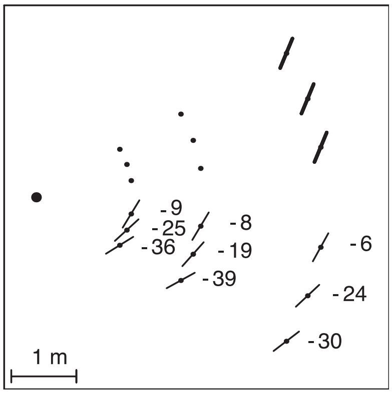

Figure 2.6
Schematic view of the experimental room with settings for the parallelity task. The thick lines represent the orientations of reference rods (distance 4.5 m, reference orientation 67°). The thin lines represent the means of three settings of observer BL. The numbers give the deviations from veridical settings in degrees.

Systematic deviations from veridical settings were found for all the observers. However, the size of the deviations was dependent on the observer. The size varied from 0° to 44°, observer AW produced the smallest deviations.

We will now look more closely at the three different parameters which may influence the pattern of deviations found in this experiment. These parameters are the separation angle, the relative distance and the reference orientation.

Figure 2.7 shows the graphs for observer BL and a reference orientation of 67° in which the deviations are plotted against the separation angle. Each graph represents the data for one relative distance. Each point in the graphs represents one trial. A line is fitted through these data points using a least squares method.

The slopes of the fits that were plotted in Figure 2.7 are plotted in Figure 2.8 against the relative distance. Each graph contains the data for one observer. The data for the four reference orientations are given on separate lines. The error bars indicate the confidence intervals for the slopes. The lines are nearly horizontal and the points for each reference orientation are all within the range of the confidence intervals of the other points on the line, so the relative distance has no effect on the slope. The slope-values represent the dependence of the deviations on the separation angle. This gives an indication of the range of the deviations that were measured. The more the slope deviates from zero, the wider the range of deviations. Thus, our method yields a pattern that resembles the one produced by plotting the deviations directly against the relative distance. We not only looked at the influence of the relative distance, we also looked at the effect of the absolute distance between the observer and the two rods. The absolute distance had no effect on the size of the deviations.

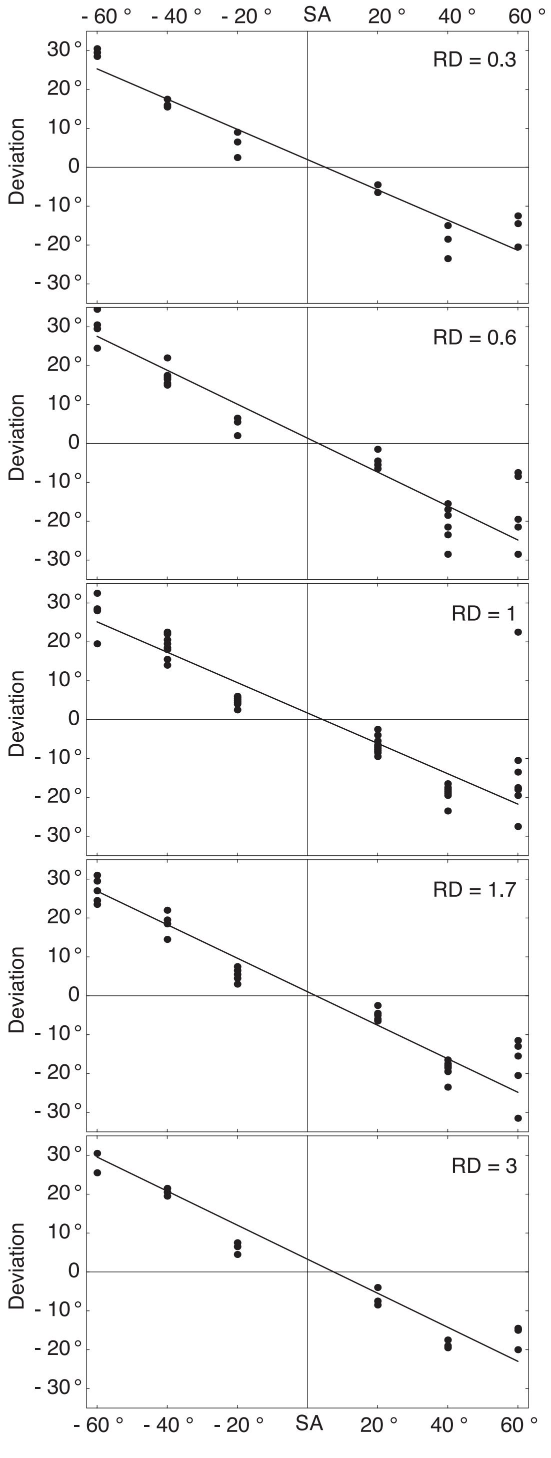

In Figure 2.8 it can be seen that the lines representing the different reference orientations are very similar for two observers (BL and TL). The results for the other two observers (AW and JP) show an effect of the reference orientation. For a reference orientation of 22° and 157°, the effect of the separation angle was small (the slopes in Figure 2.8 are small) but for the other two orientations, the effect of the separation angle was larger. For observers AW and JP a significant difference was found for reference orientations 22°/157° and 67°/112° (student's t-test, $p < 0.0001$ for both observers). This difference was not present for the other two observers ($p = 0.12$ for TL, $p = 0.15$ for BL).

The slopes of the linear fits of Figure 2.7 are shown in Table 2.2. The slopes are the means of all relative distances and two reference orientations (22° and 157°, 67° and 112°). All slopes deviated from zero significantly at a confidence level of .95.

Discussion

For the parallelity task the deviations increase linearly with separation angle. No effect of distance was found. So the distances between the two objects and the observer, and the ratio of these distances did not matter.

For two observers an effect of reference orientation was found. For these observers, the slopes of the fits were very small for two orientations (22° and 157°) and a bit larger for the other two orientations (67° and 112°). This means that, for these observers, there was an interaction of reference orientation with separation angle, the size of the effect of separation angle being dependent on the orientation of the reference rod. The difference

Figure 2.7
In each graph the deviations from veridical settings for the parallelity task are shown as a function of separation angle for each relative distance. These are the data for observer BL for a reference orientation of 67°. A line is fitted through the data points.

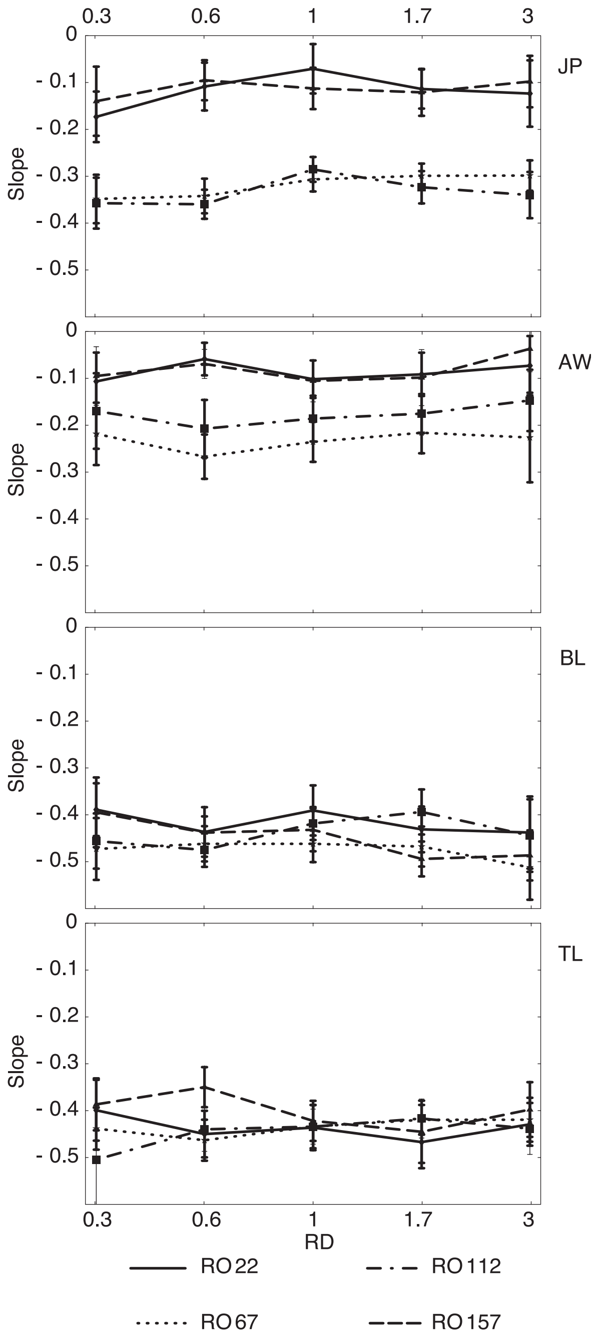

Figure 2.8
In each graph the slopes of the lines from Figure 2.7 are plotted as a function of relative distance for each observer. The data for the different reference orientations are given on separate lines.

Table 2.2 *The Slopes of the Linear Fits through the Data Points as a Function of Relative Distance for all Observers in the Parallelity Task*

Observer	Angle (deg)	Mean slope*
JP	22/157	-0.11
	67/112	-0.32
AW	22/157	-0.08
	67/112	-0.20
BL	22/157	-0.43
	67/112	-0.46
TL	22/157	-0.42
	67/112	-0.44

**The slopes are the means of the slopes found for all relative distances and two reference orientations.*

between the observers might be due to the different kinds of information that they abstracted from the scene in order to do the task.

The data were compared with the data of Cuijpers et al. (2000b). In their setup, a linear effect of separation angle was found without any effect of the relative distance. Observers differed greatly with regard to their dependence on reference orientation. However, they placed their reference rod at different orientations (0°, 30°, 60°, 90°, 120° and 150°). For most observers, they found that the slopes of the non-oblique orientations were negligible. As a possible explanation for this oblique effect they hypothesized that the observers were able to use some information about the 0° and 90° orientations from the walls of the room or the cabin in which they were seated, although an attempt had been made to conceal this information. They examined this further by varying the orientation of the observers, the cabin, in which the observers were seated, and the stimuli with respect to the walls. Some observers were dependent on the orientation of the walls, while some were mo dependent on the orientation of the cabin. A third group was somewhere in between (Cuijpers, R.H., Kappers, A.M.L., & Koenderink, J.J., 2001). In addition to this oblique-effect, they also found differences between the oblique orientations, for some observers. For these observers smaller deviations were found for trials in which a normal sized deviation would give a non-oblique setting. Since the perception of these non-oblique orientations is veridical, this is a conflicting situation. So the settings of these observers were somewhat in between veridical and non-oblique settings. Following this line of thinking one would not expect to find such good fits for our data as shown in Figure 2.7, because for all reference orientations one would have found smaller deviations for the positive or negative separation angles than for the other (dependent on the orientation). The difference could be due to the multiple sources of information about the orientation of the rod. If an observer, for example, is constantly looking at both the rod and the yellow disc perpendicular to it, he will have another pattern of deviations as an observer who looks primarily at the rod.

Thus, with minor exceptions the results of Cuijpers et al. (2000b) have the same size and follow the same pattern as the results presented in this paper. Therefore, for this task it can be concluded that the additional context did not make the settings of the observers more veridical.

2.5 Experiment 3: Collinearity task

Methods

For the collinearity task, the same two rods were used as in the parallelity task. The observer had two remote controls in his hands, enabling him to rotate the two rods. The task was to align the two rods, so they pointed towards one another. The instructions were that the observers had to rotate the rods so that they were in one line. Besides being given this verbal instruction, the observers were shown a picture of two collinear lines. The number of trials for this task was 81 (27x3 repetitions). It took the observer about two hours to perform this task.

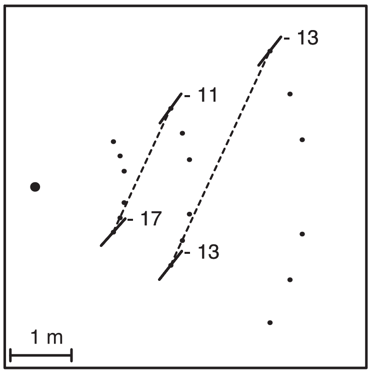

Figure 2.9
Schematic view of the experimental room with settings for the collinearity task. The dashed lines represent veridical settings for the collinearity task for conditions where the separation angle was -60° and the relative distance 0.6/1.7. The solid lines represent the means of the settings for three trials for observer AW for three different conditions. The numbers indicate the deviation from veridical settings in degrees.

Results

In the collinearity task systematic deviations were found for all observers. The size of the deviations was observer-dependent. Figure 2.9 gives an example of settings for observer AW. The means of the settings of three repetitions are shown graphically and numerically (deviations from veridical settings in degrees) for two different combinations of positions on the floor. The dotted lines are the veridical orientations of the rods.

We performed two kinds of analysis for this task. First, we looked at the two rods that were placed in the same

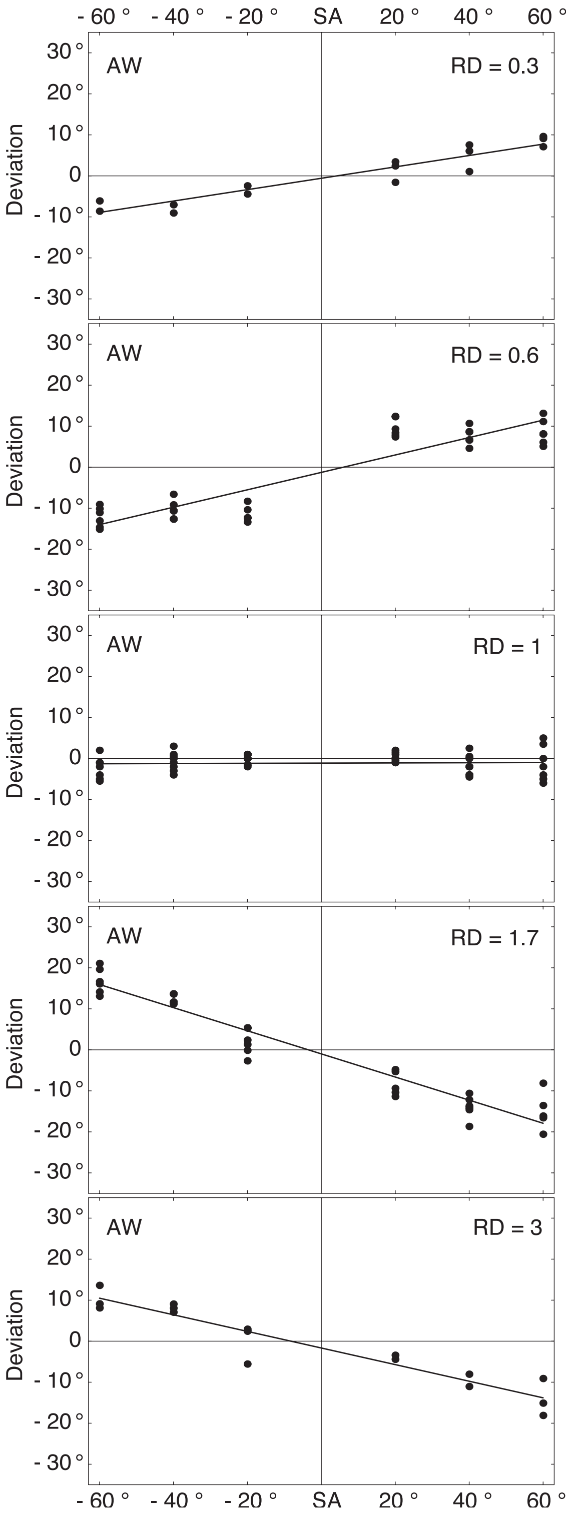

Figure 2.10
In each graph the deviations from veridical settings for the collinearity task are shown as a function of separation angle for each relative distance. These are the data for observer AW. A line is fitted through the data points.

trial separately. Second, we looked at the combination of the settings of the two rods.

The first analysis we did is the same as the analysis we did for the other two tasks. We looked at the two rods separately and dealt with them in the same way as we dealt with the pointers in the pointing task. One rod was the test-rod and we took the centre of the other rod as the target. Again we looked at the effect of the separation angle and the relative distance on the size of the deviations. We therefore plotted the deviations against the separation angle in different plots for every relative distance (see Figure 2.10 for the data for observer AW). A line was fitted through these data using a least squares method. As can be seen in Figure 2.10, for a relative distance of 0.3 and 0.6 the slope is positive and for a relative distance of 1.7 and 3 the slope is negative. For a relative distance of 1 the slope for this observer is zero. Table 2.3 shows the slopes of the different fits for all observers and relative distances. The asterisks indicate whether the fit deviates significantly from zero. This was the case for most fits, with some exceptions. The slopes are plotted against the relative distance on a logarithmic scale in Figure 2.11. The data for the different observers are plotted on separate lines. As can be seen in this figure, when the relative distance is smaller than 1, the slopes tend to be positive. This means that when the rod is closer to the observer than the target (in this case the middle of the other rod) the observer tends to overshoot. On the other hand, when the target is closer to the observer than the rod, the observer tends to undershoot. The slopes are larger for a relative distance of 0.6 and 1.7 than for a relative distance of 0.3 and 3 respectively, which is the same pattern as we found for the pointing task. For two observers the slopes are zero at a relative distance of 1. For the other two observers, the slopes are negative at this relative distance.

We will discuss three different hypothetical situations that can occur if the two rods are viewed together. The first one is a veridical setting of rods. The second possibility is that the two rods are placed with deviations with an opposite sign, both overshooting or undershooting. The third possibility is that the rods are placed with deviations of a corresponding sign, so one of the rods is overshooting and the other undershooting. In the first situation both the sum of and the difference between the deviations of the two rods will be zero. If the rods are placed with deviations of the same size but with a different sign, the sum of the deviations will be zero. This is a special case of the second possibility, that is, deviations with opposite signs. If the two rods are placed parallel, i.e. with deviations of the same size (not equal to zero) and sign, the difference between the deviations will be zero, but the sum will not be zero. This is a special case of the third possibility.

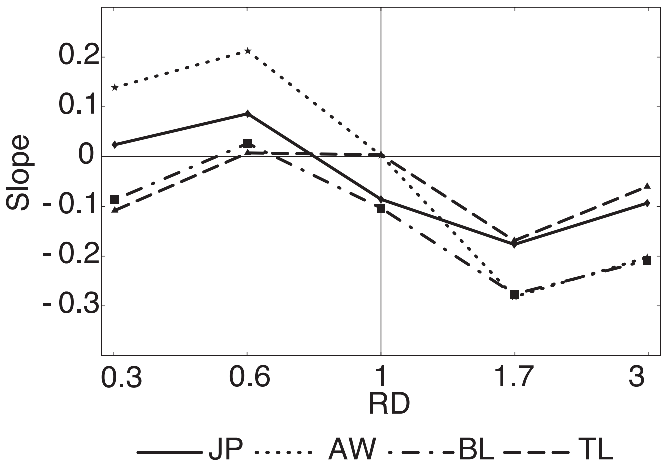

Figure 2.11
In this figure the slope of the lines from Figure 2.10 are plotted as a function of relative distance. The data for the different observers are given on separate lines.

Table 2.3. *The Slopes of the Linear Fits through the Data Points as Function of Relative Distance for all Observers in the Collinearity Task*

	Relative Distance				
Observer	0.3	0.6	1.0	1.7	3.0
JP	0.02	0.09*	-0.09*	-0.18*	-0.09*
AW	0.14*	0.21*	-0.	-0.28*	-0.20*
BL	-0.09*	0.26*	-0.10*	-0.28*	-0.21*
TL	-0.11*	0.01	-0.	-0.17*	-0.06

* *The slope deviates significantly from 0 (α = .05).*

In Figure 2.12 we have plotted the sum of the deviations of the corresponding rods as a function of the relative distance (on a logarithmic scale). Separate plots show the data for the four different observers. Each point represents the mean of 3 repetitions and for some relative distances a point represents a couple of combinations of points that have the same

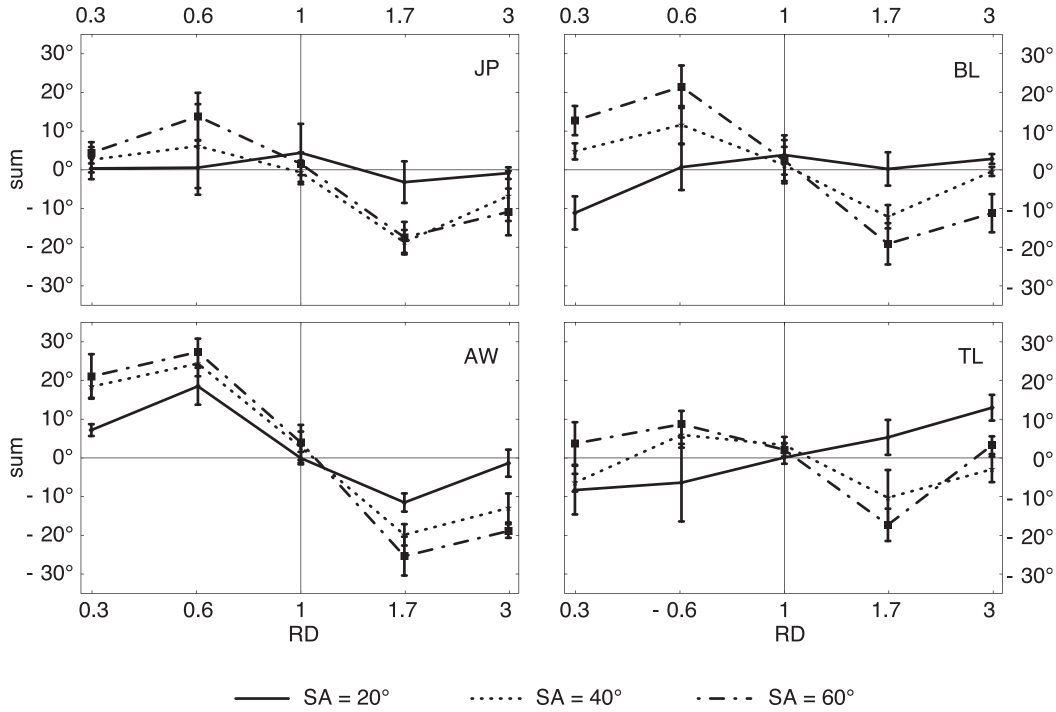

Figure 2.12
In each graph the sum of the deviations is plotted against the relative distance for one observer. The standard deviations are given via error-bars. The different lines represent different separation angles.

relative distance. The standard deviations are given via error-bars. What can be seen in the graphs is that the sign of the sum changes when the relative distance changes from smaller to larger than 1. If the relative distance is 1, then the sum is zero for all observers. The sum is larger for relative distances of 0.6 and 1.7 than for relative distances of 0.3 and 3 resp. For the smallest separation angle (20°) the relative distance had only a minor effect.

The difference between the deviations of the two bars is plotted in Figure 2.13 against the relative distance. Again the separate graphs represent the different observers, the relative distance is given on a logarithmic scale and the error-bars represent the standard deviations. Overall, the differences between the deviations are smaller than the sums of the deviations. The differences are not dependent on the relative distance. For observers AW and TL, for a relative distance of 1 the difference is zero. Since the sum is also zero, this means that the settings were veridical. For the other two observers, the difference was not equal to zero at a relative distance of 1. With a sum of zero, this means that the data were symmetrical. There was only a minor effect of separation angle on the difference between the deviations.

Discussion

First we will discuss the analysis in which we looked at the two rods separately. We performed this analysis mainly because we wanted to be able to compare it to the results of the exocentric pointing task. The same pattern was found for both tasks. As for the pointing task, the data of the collinearity task revealed a linear effect of the separation angle (see Figure 2.10). There was an effect of relative distance in that there was an undershoot when

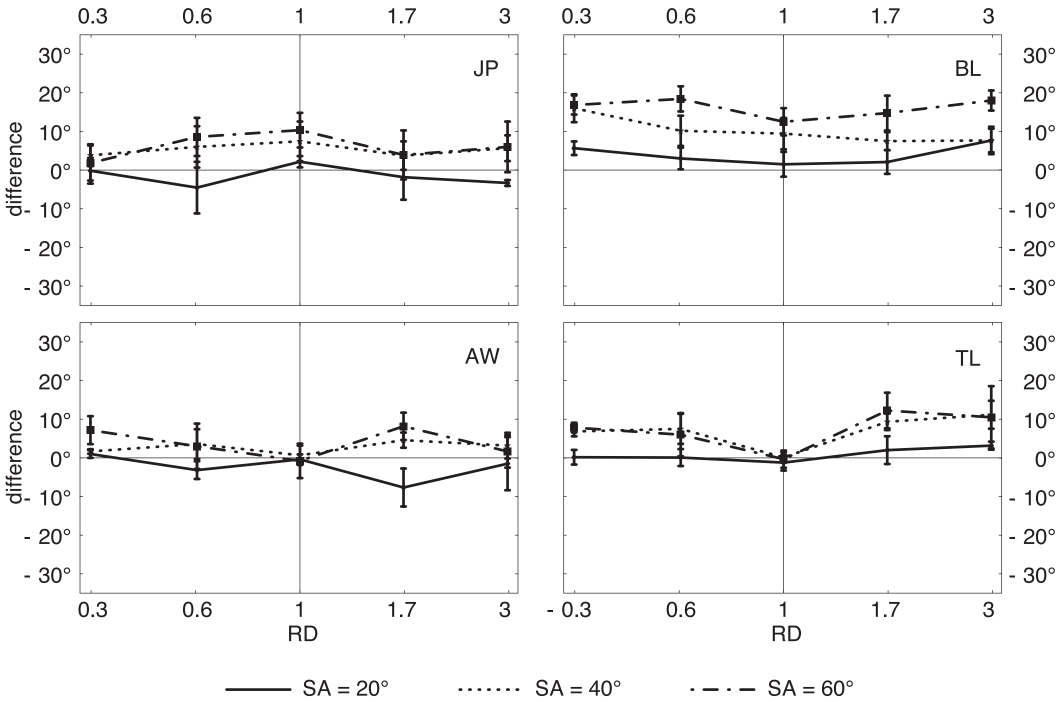

Figure 2.13

In each graph the difference between the deviations is plotted against the relative distance for one observer. The standard deviations are given via error-bars. The different lines represent different separation angles.

the relative distance was smaller than one and an overshoot when the relative distance was larger than one (see Figure 2.11). Again larger deviations were found for relative distances of 0.6 and 1.7 than for the relative distances of 0.3 and 3. A possible explanation is, as in the case of the pointing task, that if the relative distance deviates more from 1, the exocentric task will become more egocentric.

For two observers, AW and TL, we found (close to) veridical settings for a relative distance of 1. For the other two observers, we found settings with an overshoot. This pattern is the same as the one we found for the pointing task with the same two observers with veridical settings.

Looking at the results for two rods together we found that the sum of the deviations was dependent on the relative distance and separation angle. This is comparable to what Cuijpers et al. (2002) found. The difference was constant over different relative distances (with the exception of two observers who had a difference of zero when the relative distance was 1) and depended slightly on the separation angle. Cuijpers et al. also found a constant difference, even without the slight dependency on separation angle. Because (at least in geometry) collinearity is a special case of parallelity, one would expect the differences found between the settings of the two bars in the collinearity task to be of the same size as the deviations found for the parallelity task. This is not the case in the present experiments. Overall, the differences found for the collinearity task are smaller than the deviations found for the parallelity task. Next to that, the pattern of deviations is qualitatively different. The same discrepancy between the two tasks was found by Cuijpers et al. (2002). This shows that the geometrically similar parallelity task and the collinearity task are fundamentally different for the human observer.

2.6 General discussion and conclusions

As can be seen directly from the graphs (see Figures 2.5 and 2.11) the results of the exocentric pointing task are very similar to the results of the collinearity task. The same size and pattern of deviations was found. Even the distinction between two groups of observers, when one looks at the deviations found for a relative distance of 1, is the same. For this condition the same observers have veridical settings (AW and TL) for both tasks. Close to veridical settings for this condition were not found by Cuijpers and colleagues (2000a, 2002) as clearly as we did. An explanation for this difference could be as follows. When both objects were at the same distance from the observer, the veridical pointing direction was parallel to the back-wall. So for this condition there was direct information from the context available for the observer to do the setting.

If we compare these data to the data of the parallelity task (see Figure 2.8), we see a very different pattern. One of the differences is the size of the deviations: for the parallelity task the deviations are larger than for the other two tasks. In the parallelity task, the orientation of the reference rod and the test rod are misjudged. For the pointing task only the orientation of the test rod (the pointer) can be misjudged because there is no reference rod. For the collinearity task, one can split the task in two parts: pointing from one rod to the other and vice versa. This way, the orientation of one rod is not as largely dependent on the orientation of the other rod as it is in the parallelity task. Along this line of reasoning, one would expect the deviations of the parallelity task to be twice the size of the deviations in the other two tasks. This is exactly what we found in our experiments. The second difference is the effect of relative distance. Contrary to the other tasks, for the parallelity task no effect of relative distance was found. This distinction was also found by Cuijpers et al. (2000a, 2000b, 2002).

This quantitative and qualitative distinction between the parallelity task and the collinearity task seems to be quite strange when one considers the fact that in geometry, collinearity is a special case of parallelity. However, when one compares the two tasks one can distinguish between the way the tasks are performed. For the parallelity task, the observer does not have to look at the exact positions of the two objects. Instead, the view on the objects themselves is important. On the other hand, for the collinearity task this spatial relationship between the two objects is an essential part of the task next to the view on the objects. The collinearity task resembles a pointing task in which the task is to point with one pointer to the centre of the other pointer. So perhaps, it is not surprising to find a pattern and size of deviations comparable to that found for the pointing task. The comparison of the three tasks indicates that the concept of a single visual space is problematic. Apparently, “visual space” is deformed differently depending on the information in the environment necessary to do a certain task. One can describe the settings of the parallelity task geometrically by means of the separation angle and the reference orientation. The distance information is irrelevant for this task. This is totally different from the dependence on relative distance in the exocentric pointing and collinearity task. Thus, it is to be expected that human observers should perform differently in this task as compared to the other tasks. Furthermore, one would expect to find more veridical settings for the pointing and the collinearity task as compared to Cuijpers et al. data, since there is an increase in information about depth from monocular depth cues like linear perspective and texture segregation. In contrast, the performance for the parallelity task is less dependent on this kind of information about distances.

In the parallelity task a different degree of dependence on reference orientation was found for our observers. Cuijpers and colleagues (2001) discussed a difference between observers in dependence on references like the walls or the cabin the observers were seated in. But this was mainly a difference between oblique and non-oblique orientations. So this cannot explain our findings with differences between various oblique angles. Cuijpers et al. (2000b) noted a small difference between observers for oblique settings, but their explanation did not fit our data. An alternative explanation might be related to the degree of change in the view of an object when an object is rotated a small amount. For example, when an observer only looks at the rod and the rod is perpendicular to the line of sight, a rotation of 5° does not change the image of the rod on your retina as much as it would change the image of a rod collinear to the line of sight. On the other hand, if an observer looks both at the rod and the disc around it, the attended image on the retina will always have a rather large change. Thus, this difference in amount of change depends upon by the sources of information people use when performing a task.

For the pointing and collinearity task the veridical settings of the pointers can be described using the following formula:

$$\tan\beta = \frac{\sin\alpha}{\frac{r_1}{r_2} - \cos\alpha} \tag{2.2}$$

where β is the angle between the line between the pointer and the target, and the line between the pointer and the observer (Koenderink, Van Doorn, & Lappin, 2003). To use this formula for the collinearity task, we define one rod as the pointer. The middle of the other rod can be defined as the target. The variable α represents the separation angle between the two objects, and r_1 and r_2 represent the distance to the pointer and the target respectively. Traditionally, the focus has been on the perceived distances and how they are derived from the physical distances (r_1 and r_2). Different models have been fitted to different data-sets. For example, Wagner (1985) compared his data, acquired under full-cue conditions, to four models. Two of these were Riemannian models, one spherical, the other hyperbolic. These models did not fit his data: the spherical model produced such strange fits that it was rejected. The hyperbolic model did not produce good fits either, which was explained by noting that the model was made for reduced-cue conditions. Another model Wagner describes is an affine contraction model, which describes an affine transformation in depth (only in the direction straight ahead) using Cartesian metrics. Because humans are not thought to depend on Cartesian coordinates in dealing with depth, this model is refined into the vector compression model, which uses polar coordinates and fits very well to Wagner's data. In this model the physical distance is multiplied by a constant. In trying to explain our results with Equation 2.2, this constant will cancel out in the ratio that represents the relative distance. Thus, since multiplication of the distances with a constant has no effect on the pointing direction, the vector compression model cannot explain our results. We looked at Gilinsky's formula for perceived distance (Gilinsky, 1951) as can be seen in Equation 2.1. We replaced r_1 and r_2 from Equation 2.2 with $P(r_1)$ and $P(r_2)$. The equation did not fit our data well, but that is not surprising since Gilinsky formulated her theory on the basis of data obtained with larger distances. She described visual space as compressed. A compression of visual space does not suit our data, obtained with smaller distances. In fact, for two observers we found a (bad) fit with a negative constant, which is nonsense in the Gilinsky formula. For negative values of c, the formula is expanding, not compressing. Therefore, for the tasks in which spatial information was most important, the settings should be described by an expanding distance function (like a powerlaw) rather than a compressing distance function. This difference, possibly due to the varying distances used, is consistent with the ideas proposed by Battro, di Pierro Netto and Rozestraten (1976) and Koenderink, Van Doorn, & Lappin (2000) that the geometry varies with distance from observer to objects.

The data described above are quite similar to the data found by Cuijpers et al. (2000a, 2000b, 2002). This is not what we expected since the experiments were conducted in a very different environment. The walls, ceiling, floor, radiators, windows, doors, etc. were visible in our experiments in contrast to Cuijpers et al.. Despite this extra information provided by linear perspective, texture segregation, size constancy etc., the observers show a comparable magnitude and pattern of deviations. This can be explained in the following way. Perhaps the structure that was provided to the observers was not rich enough, so a richer structural context could make a difference. Wagner (1985) talked about the quantity and the quality of depth information and concluded that if both were maximal then visual space should be Euclidean. If one reduces both the quantity and the quality of the perceptual information, the deformation of visual space will increase as well. So a next step in this research should be to elaborate the context provided to the observers, and see whether the deformation of visual space will decrease. For example, we could put textures on the walls and floors or place extra objects in the room.

In summary, one cannot speak of a single visual space since the structure is dependent on the task that the observer is doing and the distance between the objects and the observer. The structure of visual space for two tasks that require spatial information from the objects (the exocentric pointing task and the colinearity task) and a distance of 1.5 to 4.5 meters between observers and objects, can be described by an expanding distance function like a powerlaw. For the parallelity task, a distance function is useless since the information about the exact positions of the objects is not necessary to do the task. Another conclusion that can be drawn from these data is that the settings of the observers in this environment full of monocular depth cues, were similar to the settings found for data obtained in a much poorer environment. Thus, the structure in this richer environment was not used effectively by the observers.

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

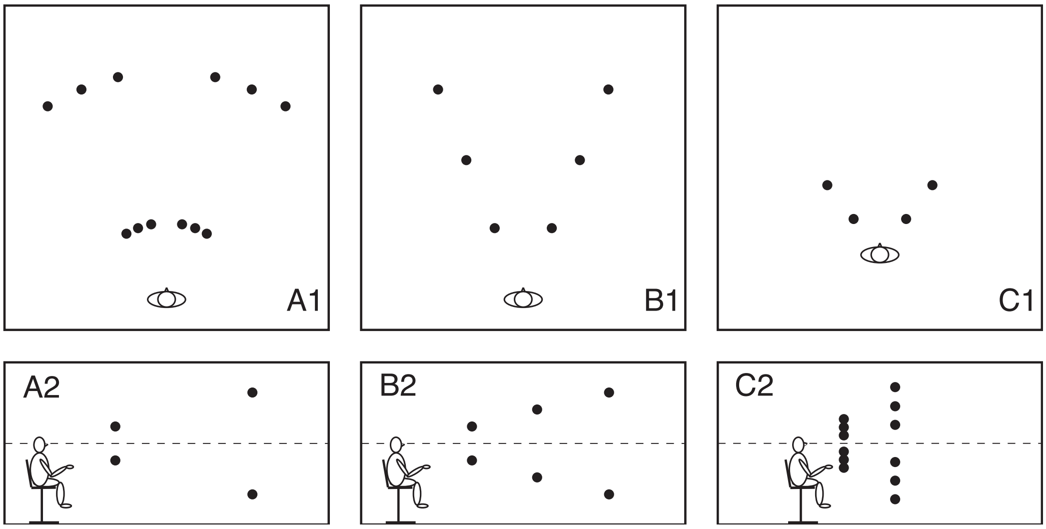

Figure 3.1
A1, B1 and C1 are top-views of the experimental room for experiments A, B and C, respectively. The squares represent the walls of the room, the black dots the positions of the objects in the azimuthal plane, and the small figure the position of the observer.
A2, B2 and C2 are side-views of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

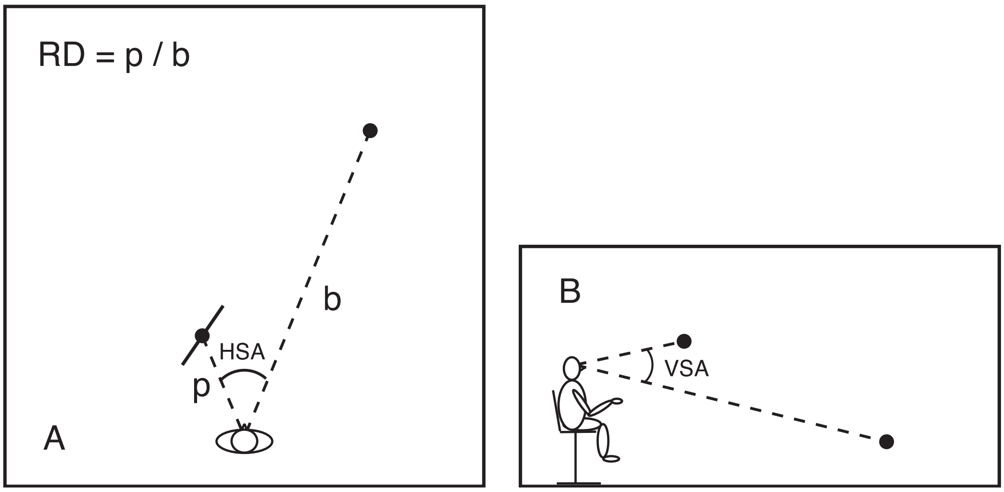

Figure 3.2
A graphical illustration of the parameters that were used. Figure A shows the horizontal separation angle (HSA) and the relative distance (RD). The relative distance is defined by the ratio of the distance from the pointer (p) and the distance from the ball (b). Figure B shows the vertical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

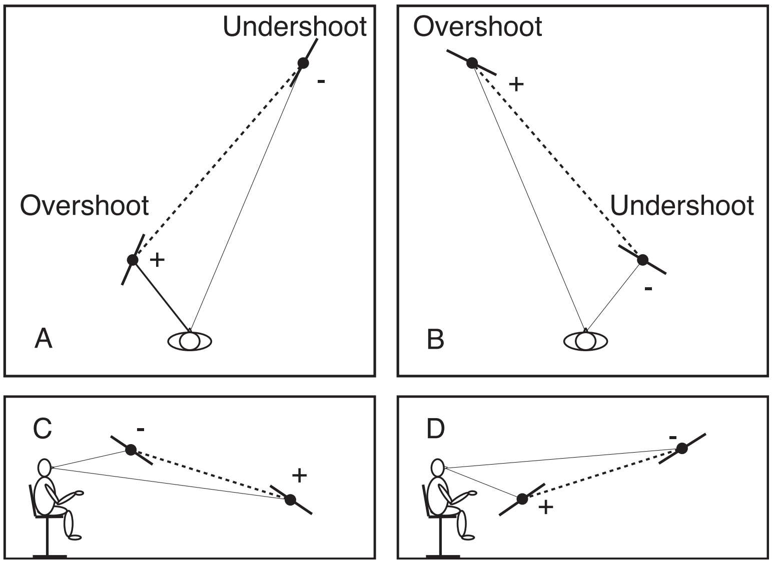

Figure 3.3
Examples of settings. The black dots represent the positions of the objects, the dashed lines the veridical pointing direction, and the small thick lines the pointing-orientation. The plus and minus signs represent positive and negative deviations. Figure A is a top-view of the room, with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

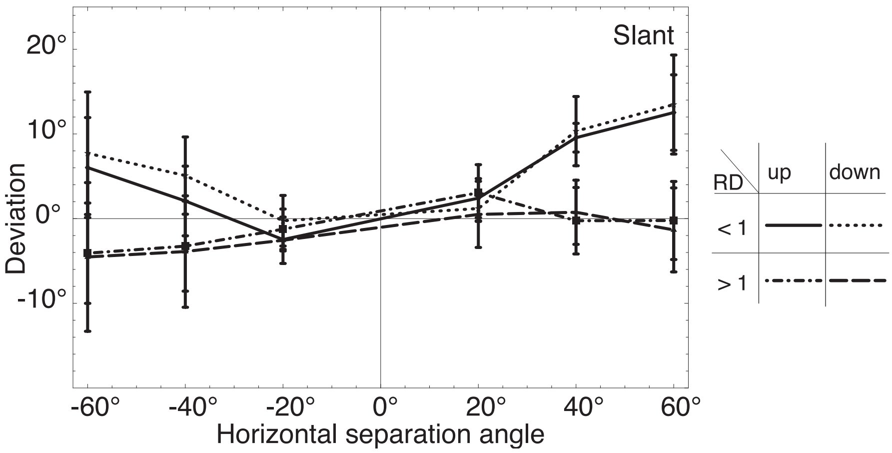

Figure 3.4
The data of the slant for experiment A. The deviation of the slant in degrees is plotted against the horizontal separation angle in degrees. The four lines represent four different conditions: pointing upwards with a relative distance smaller and larger than one and pointing downwards with 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).

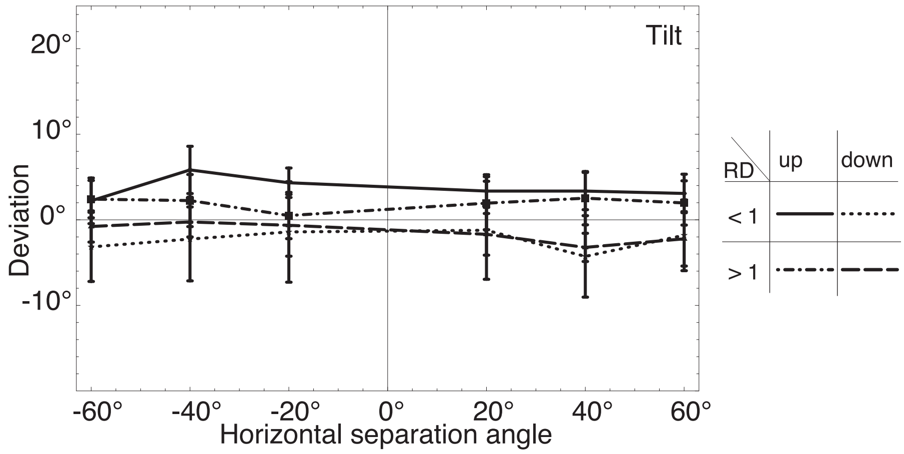

Figure 3.5
The data of the tilt for experiment A. The deviation of the tilt is plotted against the horizontal separation angle. The four lines represent the same conditions as in Figure 4. Each data-point represents the mean of the values for the six observers. The error-bars represent the standard deviations.

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.

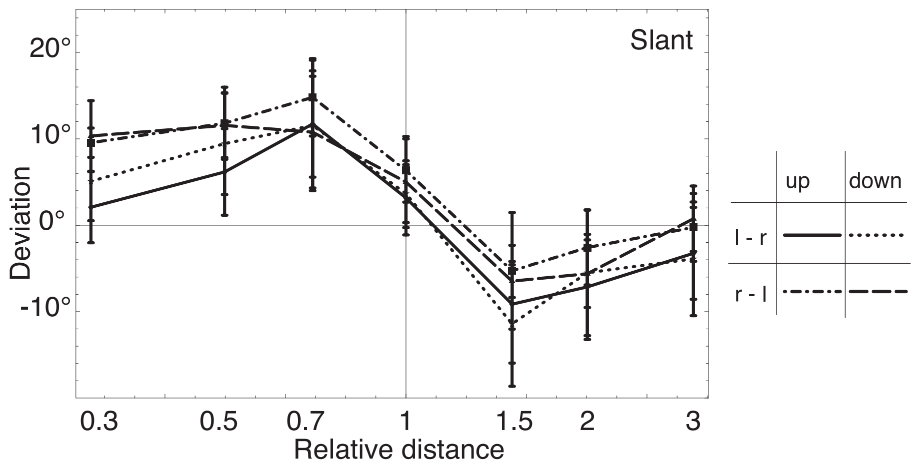

Figure 3.6
The data of the slant for experiment B. The deviation of the slant is plotted against the relative distance in logarithmic scale. The four lines represent four different conditions: pointing upwards and downwards from left to right and pointing upwards and downwards from right to left. Each 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.

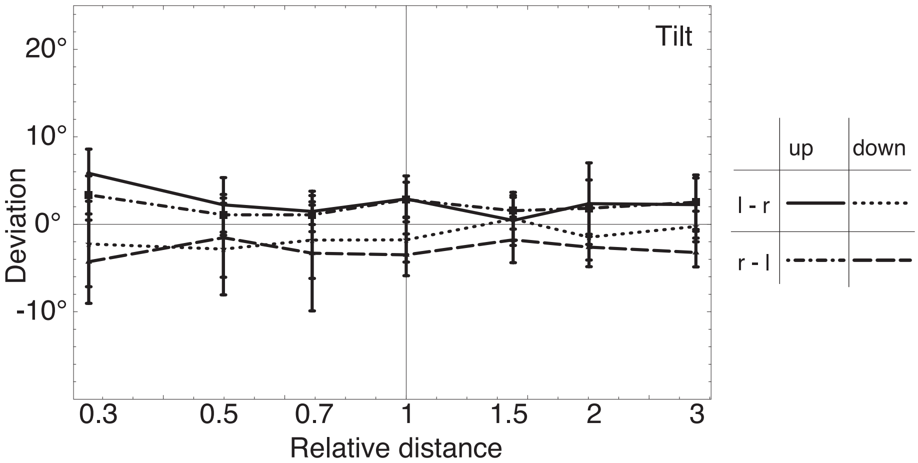

Figure 3.7
The data of the tilt for experiment B. The deviation of the tilt is plotted against the relative distance in logarithmic scale. The four lines represent four different conditions: pointing upwards and downwards from left to right and pointing upwards and downwards from right to left. Each 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

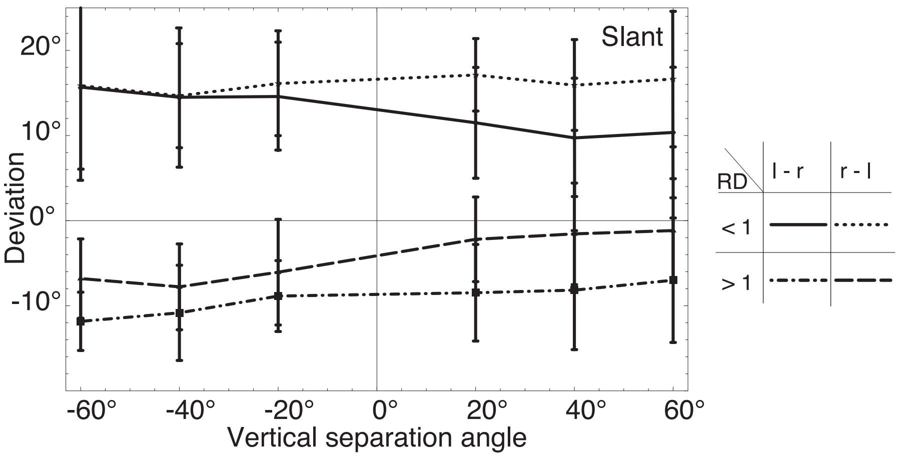

Figure 3.8
The data of the slant for experiment C. The deviation of the slant is plotted against the vertical separation angle in degrees. The four lines represent four different conditions: pointing from left to right with a relative distance smaller and larger than 1 and pointing from right to left 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.

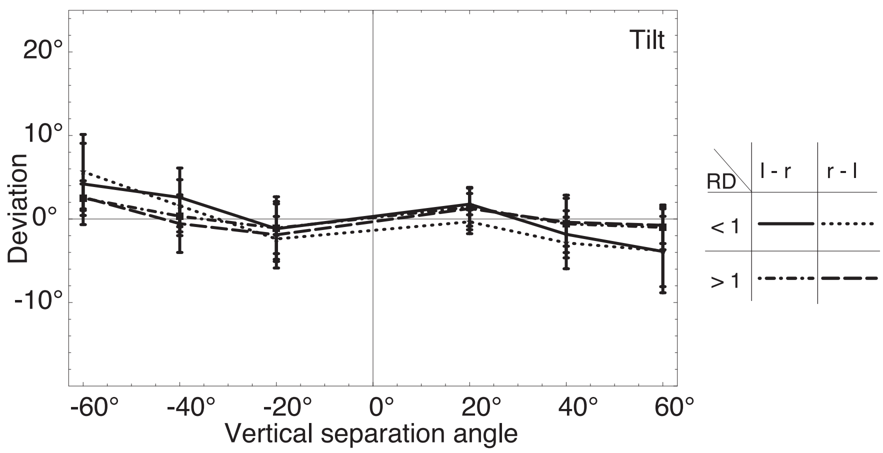

Figure 3.9
The data of the tilt for experiment C. The deviation of the tilt is plotted against the vertical separation angle. The four lines represent the same conditions as in Figure 3.8. Each data-point represents the mean of the values for the six observers. The error-bars represent the standard deviations.

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.

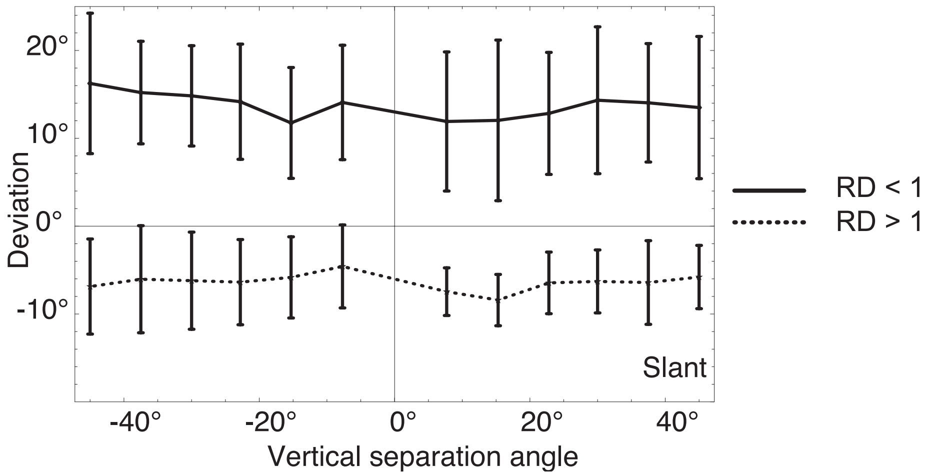

Figure 3.10
The data of the slant for experiment D. The deviation of the slant is plotted against the vertical separation angle. The solid line represents the data for a relative distance smaller than 1, whereas the dotted line represents the data for a relative distance larger than 1. Each data-point 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.

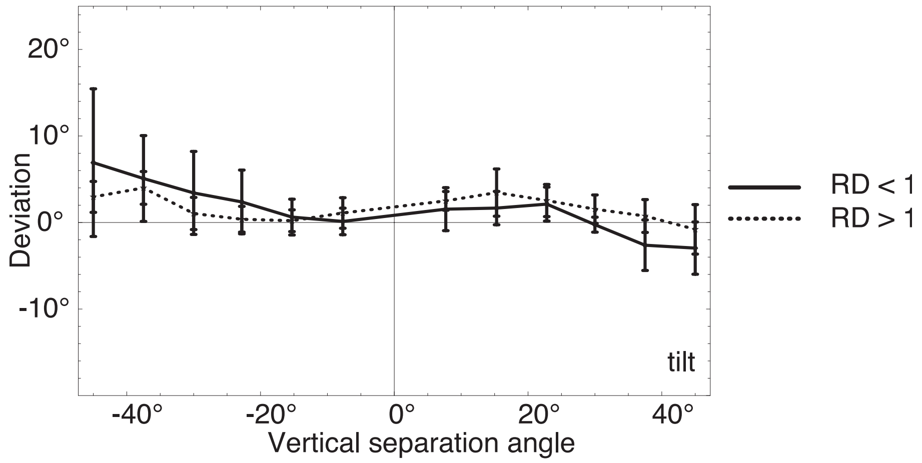

Figure 3.11
The data of the tilt for experiment D. The deviation of the tilt is plotted against the vertical separation angle. The solid line represents the data for a relative distance smaller than 1, whereas the dotted line represents the data for a relative distance larger than 1. Each data-point 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.

Chapter 4
Effects of context on a visual 3D pointing task

Abstract:
We examined the effects of egocentric and contextual references on a 3D exocentric pointing task. Large systematic deviations were found for the slant (angle in the horizontal plane). For most observers, the deviations were smaller when the veridical pointing direction was parallel to a wall. For some observers the size of the deviations was also dependent on whether the veridical pointing direction was frontoparallel or not. For the tilt (angle in the vertical plane), the deviations were smaller and less systematic. Hence, although observers show comparable systematic deviations, the way in which the presence of structure in an environment is used for judging positions of objects is observer-dependent.

In press as:
Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J., (2006). Effects of context in a visual 3D pointing task. *Perception, Pion Limited, London.*

4.1 Introduction

We know that there are multiple possible sources of information that we can use to see depth (Berkeley, 1732, Gibson, 1950). However, the fact that these cues are available to observers does not necessarily have to mean that they are actually using them. Hence we need to discriminate between the structure that is available and the information that observers extract from this structure, namely the actual cues for seeing depth. In the research described below, we examined the influence of prominent structure surrounding the set-up (e.g. walls of a room) and the information that observers can extract from internal references.

The term visual space is used in this paper as an operational entity that represents how we visually perceive locations of objects in the world around us. In the early research on visual space researchers largely ignored the fact that external cues are important. Instead they focused on egocentric references by using luminous stimuli in a completely dark context (Zajaczkowska (1956), Blank (1961)). By this method these researchers found large deviations from veridical settings. This is probably not surprising since the use of such a paradigm removes a considerable amount of informative structure that is normally present in the visual field.

Other researchers concentrated on contextual information. Ames pointed out that perceptual awareness is dependent on the weighting of the reliability of indications about depth (Ames, 1953). These reliabilities are learned through our past experience (Berkeley, 1732). With his distorted room demonstrations (Ittelson, 1952), Ames showed that prior knowledge and assumptions that, for instance, walls, windows and paintings are usually rectangular, can overrule physiological depth cues. Therefore, an irregularly shaped room could look like a regularly shaped one. If more information is present, i.e. under free viewing-conditions, this illusion is still present, but is somewhat weaker (Gehringen, & Engel, 1986).

Yang and Purves (2003) introduced a probabilistic approach that resembles the ideas put forward by Ames. The basis of this approach is that people are experienced observers and make use of prior knowledge (Gibson, 1966). Phenomena like the "specific distance tendency" (Owens, & Leibowitz, 1976), and the "equidistance tendency" (Gogel, 1965) can be explained by the fact that we are more likely to be surrounded by objects at a certain distance or combination of distances. According to Yang and Purves (2003) our brains use Bayesian calculations to let us perceive what is the most probable situation that is possible for a given visual input. Comparing this approach to Ames' theory, we see that both views incorporate prior knowledge. Ames talks about reliabilities of cues, Yang on the other hand does not talk about depth cues, but about the likelihood of certain situations. However, both theories incorporate the notion that people can make use of different sources of information, and that the choice of these sources is observer-dependent.

The "oblique effect" is an effect that is often found in research on visual space (Appelle, 1972). It refers to superior performance in tasks involving horizontally or vertically oriented stimuli than in tasks involving obliquely oriented stimuli. Chen and Levi (1996) found oblique effects for retinotopic coordinates with a parallelity discrimination task. Hence, perceptual phenomena like the "oblique effect " are influenced not only by allocentric references, but also by egocentric ones. For example, Darling and Bartelt (2005) found that an internally specified gravity reference is important for orienting objects visually. Cuijpers, Kappers and Koenderink (2000B) also found an oblique effect in research with a parallelity task in which the observers had to put a bar parallel to a reference bar. The oblique effect that Cuijpers et al. found is different from the former oblique effects. First, their research was done in a horizontal plane instead of a frontoparallel plane. Second, Cuijpers et al. found larger deviations for oblique orientations than for non-oblique orientations in a parallelity task. In contrast, the conventional literature on the oblique effects describes an increase in the variance of the data instead of a signed deviation. When we refer to “the oblique effect” in this paper, we mean the effect described by Cuijpers et al., i.e. an increase in the deviation from veridical settings.

In a second experiment Cuijpers, Kappers and Koenderink (2001) varied the orientation of the setup with respect to the room, the cabin the observer was seated in and the orientation of the observer himself. The goal was to measure which references in the environment were causing the oblique effect and thereby the veridicality of the settings of the observers. They concluded that people tend to make use of external references although they are not consciously aware that they are doing so. The structure that people use in their environment, is observer-dependent (Cuijpers et al., 2001). Since we found in a pilot experiment that observers use both egocentric and allocentric references for the exocentric pointing task, we examined this idea further.

The task we used in the present experiments was an exocentric pointing task. This task has been used in the horizontal plane at eye-height by several experimenters including Koenderink (Koenderink, Van Doorn and Lappin, 2000), Cuijpers (Cuijpers, Kappers and Koenderink, 2000A) and Schoumans (Schoumans, Kappers and Koenderink, 2002). During a 2D exocentric pointing task, an observer has to rotate a pointer with a remote control in the horizontal plane at eye-height. The task is to let it point towards a target. Cuijpers et al. (2000A) tried to conceal the information normally provided by walls, floor and ceiling of the experimental room by covering the walls with wrinkled plastic. In addition, they had the observers seated in a cabin with a chin-rest to prevent them from deriving information from head-movements. They found effects of the relative distance between the pointer and the ball and the visual angle. In our earlier experiments the same task and experimental room were used (Doumen, Kappers, and Koenderink, 2005 [Chapter 2]). The difference with Cuijpers' research was that the observers had an unobstructed view of the walls, ceiling and floor of the experimental room. Apart from this, the observers could move their heads freely. Our results were remarkably similar to the data of Cuijpers et al. (2000A) with regard to the size of the deviations and the structure of the deviations. An interesting aspect of the results of both experiments was that the deviations approached zero when pointer and ball were at the same distance from the observer. We wondered whether this decrement in deviations found for most observers was based solely on the fact that a wall was used as a reference, or whether it could be explained by the fact that these settings were also in the frontoparallel plane. Schoumans and colleagues (2002) showed that people use contextual information as a reference for doing an exocentric pointing task within grasping-distance. They concluded that observers used both egocentric and allocentric references to judge positions of objects at these distances.

Most research in visual space, including our own, has been concerned with horizontal planes. Therefore, it was a natural step to extend our knowledge to three-dimensional scenes. To do this, we used a pointer that could rotate both in the horizontal plane and vertical plane, and could be positioned at different heights. We found a comparable pattern for our 3D experiments for the settings in the horizontal plane as we found in the previously described 2D experiments (Doumen, Kappers, & Koenderink, in press A [Chapter 3]). The deviations in the vertical plane were smaller than in the horizontal plane and were differently dependent on spatial parameters. In the following experiment we were interested in whether egocentric references and contextual references could contribute to the accuracy and veridicality of the settings of the observers in a 3D task. Since the use of the structure in the environment in visual tasks is observer-dependent (Cuijpers et al., 2001), we expected that the observers would show differences in contextual influences. Furthermore, since oblique effects are common in visual perception, we hypothesized that in trials where the pointer and the ball were at different heights the size of the deviations would be larger than in trials where the ball and pointer were at the same height.

In the first experiment that we describe below, we compared two conditions: one in which the objects were in a plane that was both frontoparallel and parallel to the wall, and the other in which the objects were frontoparallel but not parallel to the wall. In the second experiment, the settings could be either frontoparallel or not and parallel to a wall or not. These two parameters were varied independently of each other.

4.2 General methods

Observers

Fourteen observers participated in our experiments, seven in each of the two experiments. They were paid for their efforts. They all had normal or corrected-to-normal sight and were tested for stereoacuity (all observers had an acuity of more than 60 arcseconds). The observers were tested individually, were naive as to the purpose of the experiments and had little or no experience of participating in psychophysical experiments.

Experimental set-up

The experimental room measured 6 m by 6 m by 3.5 m. See Figure 4.1 for a top-view of the experimental room. The square represents the walls of the room. The dotted line represents a wall that contained radiators below four blinded windows. The shorter dashed lines represent the light gray doors in the room. The four X's indicate the positions of four 3 m high wooden strips that were attached to the wall. The walls were plain white with visible texture.

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 the grid. These balls had a diameter of 6 cm and could be hung at two heights: 60 cm and 240 cm above the ground.

Metal strips were taped on the floor to position the pointer. The strips were visible to the observer. The pointer consisted of a 45 cm long orange rod that was connected to a device that contained the motor for rotating it in the vertical plane. The device was covered by a white sphere with a diameter of 14 cm. The pointer stood on a vertical iron rod that was attached to a circularly shaped foot with a motor to generate rotations in the horizontal plane. The motor was concealed by a cylindrical screen. The height of the center of the pointer was 60 cm. The pointer could be positioned on a pillar so that the total height was 240 cm. The pointer could therefore be positioned at two different heights: 60 and 240 cm above the ground. The observer could rotate the pointer using two small remote controls; one was for rotating in the horizontal plane, and the other was for rotating in the vertical plane. These remote controls consisted of small joysticks that by themselves gave no information about the orientation of the pointer.

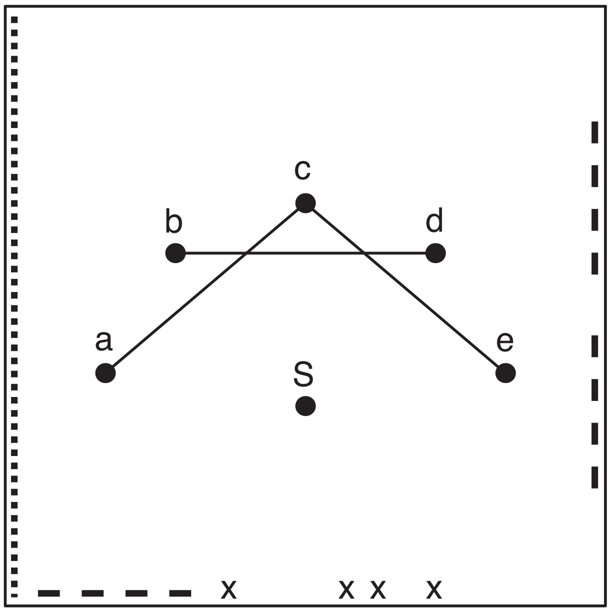

Figure 4.1
The square represents the walls of the room. The dotted line represents a wall with blinded windows and radiators along the base. The short dashed lines give the positions of the doors, and the crosses give the positions of four wooden strips that were attached to the wall. The observer was positioned at point S, the objects (pointer and ball) were at positions a through e. The lines between the points represent the pointing directions that were used.

The observer was seated on a revolving chair that could be adjusted in height so that every observer could have an eye-height of 150 cm. This was exactly halfway between the floor and the grid. The two objects, the ball and the pointer, could both be positioned either 90 cm above or below eye-height, independently of each other.

Procedure

The measurements were done in sessions of approximately one hour each. Before a session began, the height of the chair was adjusted so that the observer was seated with his eye-level at the correct height. The observers were instructed to stay seated during the measurements and not to move their upper bodies. They were allowed to rotate their heads or the chair in which they were seated. No references were made as to where they were allowed to look during the experiment. Furthermore, no feedback was given before the end of the entire experiment. The observer was asked to close his eyes when the exact settings of the pointer were read from the device and when the objects were being rearranged for the next trial. Before the observer was allowed to open his or her eyes again the observer was asked to rotate the pointer randomly in the two directions. This was done to make sure that the observer could not get feedback from the previous trial. At the end of each session the observer was asked to close his eyes and rotate the pointer again and the ball was removed from the scene. This way the observer only saw the pointer and the ball together from the position of the chair at an eye-height of 150 cm above the ground.

Analysis

Two values were collected for each setting: the slant and the tilt. These terms are used differently in various fields of research. Here slant is used as the angle in the horizontal plane (with a range of 360°) and tilt is used as the angle from the horizontal plane (with a range of 180°). We analyzed these two values separately. We had three reasons for analyzing the data in this way. The first reason was that this way we can compare the data to our previous work with a 2D task. Second, some authors assume that gravity (and with gravity a distinction between vertical and horizontal directions) is important for orienting (Paillard, 1991, and Darling & Bartelt, 2005). Third, the settings for the slant and the tilt were measured separately (one remote control for each of these two values).

For the slant, a positive deviation means that the observer pointed further away from himself than the position of the ball (overshoot), whereas a negative deviation means that the observer pointed more towards himself (undershoot). For the tilt, a positive deviation means that the observer pointed too high, a negative deviation means that the observer pointed too low. We conducted analyses of variance for the slant-data of the two experiments.

4.3 Experiment 1: Allocentric references

Methods

Seven observers participated in this experiment. Four of them were undergraduate students and three were new colleagues (one post-doc and two graduate students). One of the undergraduates was female, all the other observers were male.

Five points were marked on the floor (and the grid) to position the objects. These points were visible, although the points on the raster did not attract much attention. The letters A through E represent these positions in the azimuthal plane (see Figure 4.1). The distances between A and C, B and D, and C and E were 260 cm. The distance from the position of the cyclopean eye of the observer (point S in Figure 4.1) and these points in the azimuthal plane was 219 cm. Since the observer sat on a revolving chair and all possible positions of the objects were the same distance away from him, pointing from A to C, from B to D or from C to E could always be in the observer's frontoparallel plane. The difference between the pointing-direction B-D and the other two pointing-directions (A-C and C-E) was the presence of a wall parallel to the pointing-direction B-D. The distances from points B and D to the back-wall were approximately 2.5 m.

Since the objects were at two different heights, there were trials in which the observer had to point upwards, downwards or horizontally (low or high), see Figure 4.2.

For the three combinations of points we used the following parameters: the direction of pointing (pointing from left to right and vice versa), the height of the pointer and the height of the ball. We repeated every condition three times. Therefore for each observer we

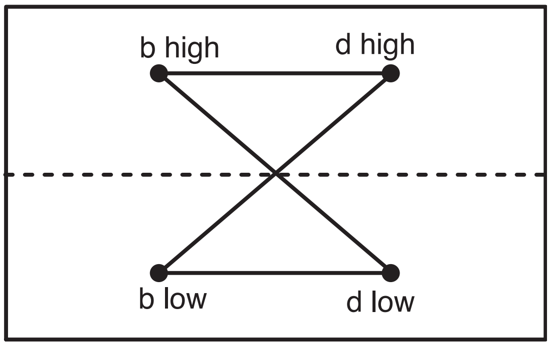

Figure 4.2
Frontal view of the set-up. The dashed line represents the eye-height of the observer (1.50 m above the ground). Of each point in the azimuthal plane (in the given example points a and b) two heights of the objects were used: 60 cm and 240 cm above the ground. The lines between the points represent the pointing directions that were used.

had 3x2x2x2x3 = 72 trials (# repetitions, # directions, # pointer-heights, # ball-heights, # combinations of azimuthal positions of the pointer and the ball). We completed the experiment in three sessions of one hour each. In each session, all combinations of pointer- and ball-positions were measured in a random order. Thus, each subsequent session contained the same trials in a different order.

Results of experiment 1

Slant

The data of the slant, the deviations in the horizontal plane, are plotted in three graphs in Figure 4.3. These graphs represent the same data sorted in different ways. The three figures show the deviations from the veridical pointing directions in degrees for the seven observers. The standard errors of the mean are represented in the error-bars. Quite large deviations from veridical settings were found. We found deviations with mean values of up to 17°. The sizes of the deviations were observer-dependent.

A second striking feature of Figure 4.3 is that most values for the slant are positive. This means that the observers almost always overshot the target. More specifically, the observers frequently pointed to positions beyond the target. One observer, BH, showed rather small deviations in the slant and these deviations were scattered around zero. He was the only observer who did not show a significant overshooting-bias.

Figure 4.3A shows two bars for each observer. The light gray bars represent the mean of all trials in which the veridical pointing direction was parallel to the wall. The dark gray bars represent the means of the trials in which the veridical pointing direction was not parallel to the wall. For all observers except BH the dark bars were larger than the light bars, which means that when the wall was parallel to the pointing direction, the deviations of the settings were smaller. In Figure 3B the light gray bars represent the trials in which the observer had to point from left to right, the dark gray bars represent the trials in which he had to point from right to left. There were differences between the bars for some observers, but these were not all in the same direction. In Figure 4.3C the light bars represent the trials in which the veridical pointing direction was horizontal (above and below eye-height respectively). The dark gray bars represent the trials in which the veridical pointing direction was downwards and upwards respectively. For four observers (JJ, JW, MW and TH) the dark bars are larger than the light bars, which means that the deviations from veridical settings were smaller when the observer had to direct the pointer horizontally.

We did an overall analysis of variance with repeated measures. The dependent variable was the deviation of the slant. We used three independent variables: context (pointing between A-B, B-D and C-E), direction (pointing from left to right and vice versa), and height (pointing

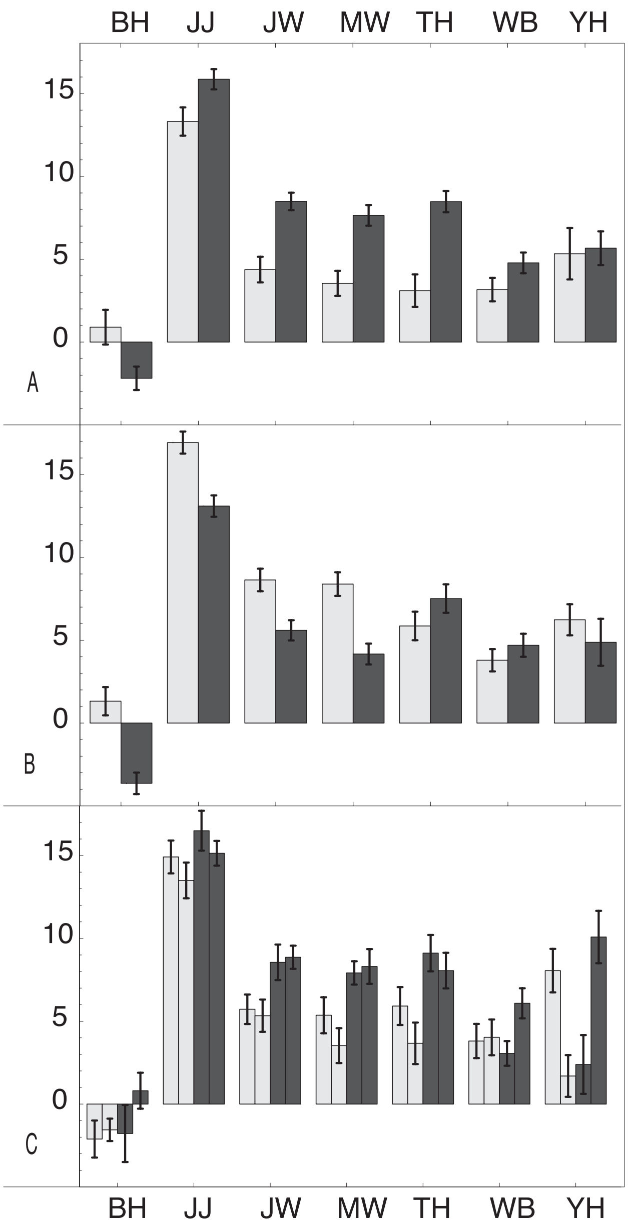

Figure 4.3
Bar-charts of the size of the deviations for the slant in degrees for each observer in experiment 1. (a) The light gray bars represent the means of the trials that were parallel to the wall, the dark gray bars show the means for the trials not parallel to the wall. (b) The light bars represent the means of the trials when the observer pointed from left to right, the dark bars when they pointed from right to left. (c) The light bars show the means for pointing horizontally (above eye-height and below eye-height respectively), the dark bars for pointing downwards and upwards.

from a high position to a high position, from high to low, from low to high and from low to low). We found a trend for both an effect of context and of direction of pointing ($F = 3.649$, $p = 0.09$ and $F = 4.708$, $p = 0.07$, respectively). The deviations were smaller when pointing parallel to the wall and when pointing from right to left. Furthermore we found an effect of height ($F = 5.800$, $p < 0.02$), with smaller deviations when the pointing-direction is horizontal.

Next to the analysis discussed in the previous paragraph, we did an analysis of variance on the slant-data for each observer separately. For four observers we found an effect of context, as was expected (p-values <.005). Post-hoc tests reveal that for these observers pointing between A-C and C-E (neither parallel to the back-wall) did not differ from each other, but there was a difference between these two combinations and B-D (parallel to the wall). These observers, JJ, JW, MW and TH, thus showed more veridical settings when a wall was positioned parallel to the pointing-direction. One observer, BH, showed an effect of context, but it was contrary to the former trend, i.e. he showed larger deviations when pointing from B to D or vice versa. However, his deviations were rather small and scattered around zero (see Figure 4.3 A). Hence, we do not feel justified in concluding that the observer's results were affected by context. The other two observers, WB and YH, showed no effect of context. See Table 4.1 for F and p values for all observers.

The direction of pointing did have an effect for four observers, BH, JJ, JW and MW. All showed the same bias: pointing from left to right resulted in

larger deviations of the slant. For BH however, this was arbitrary because of the positive and negative values of his data.
The height also affected the data for four observers, JW, MW, TH and YH. Post-hoc tests reveal that when pointing from a high position to a high position and from a low position to a high position, YH had larger deviations than revealed by the data for pointing from high to low and from low to low. The post-hoc tests also showed that there was a difference between the conditions where the veridical pointing direction was horizontal (pointing from high to high and from low to low) and the conditions where it was not horizontal (for JW, MW and TH). For the other three observers, no effect of height was found.

Table 4.1. *The F and p values for the ANOVA of the data of the slant in experiment 1*

	Context		Direction		Height		Interactions
	$F_{(2,48)}$	p	$F_{(1,48)}$	p	$F_{(3,48)}$	p	
BH	4.795	.013*	1.974	.130	27.444	.000*	
JJ	4.911	.011*	2.427	.077	23.642	.000*	H*D, $F_{(3,48)}$= 4.887. p = .005 D*P, $F_{(2,48)}$= 3.716. p = .032
JW	16.864	.000*	7.649	.000*	20.735	.000*	D*P, $F_{(2,48)}$= 5.475, p = .007
MW	15.943	.000*	10.652	.000*	37.476	.000*	D*P, $F_{(2,48)}$= 4.741, p = .013
TH	17.808	.000*	8.065	.000*	3.784	.058	H*D, $F_{(3,48)}$= 2.949, p = .042 D*P, $F_{(2,48)}$= 4.095, p = .023
WB	2.280	.113	2.775	.051	1.348	.251	D*P, $F_{(2,48)}$= 6.044, p = .005
YH	1.051	.357	10.812	.000*	1.163	.286	H*D, $F_{(3,48)}$= 5.947, p = .002 D*P, $F_{(2,48)}$= 5.865, p = .005

For the parameters parallel and height α = .05 (one-sided).
For the parameter direction α = .025 (two-sided).
** the slope deviated significantly from 0.*

Tilt

For the tilt, the deviation from veridical settings in the vertical plane, we found very small deviations (see Figure 4.4), most deviations being less than 5°. Figure 4.4 shows in each plot four bars for each observer. In both plots, the light gray bars show the deviations for trials in which the pointing direction was parallel to the wall, and the dark gray bars show the deviations for the trials that were not parallel to the wall. Figure 4.4 A shows the data for the trials when the observer had to point horizontally. For each observer the first two bars give the mean values for pointing horizontally above eye-height, whereas the second two bars give the means for trials in which the observer had to point horizontally below eye-height. Figure 4.4 B gives the data for pointing downwards (first two bars for each observer) and pointing upwards (second two bars). No clear pattern emerged from these data. Furthermore, the results did not deviate systematically in one direction as they did for the slant.

Discussion of experiment 1

In the horizontal plane, observers mainly overshot when pointing to the ball. This was reported earlier by Cuijpers et al. (2000A) and Doumen et al. (2005) [Chapter 2]. In those experiments, most deviations for frontoparallel settings were negligible, but the deviations that were found were mostly positive.

We did both an overall analysis and separate analyses for all observers. With the overall analysis we found a trend for an effect of context. This is because some observers show an effect while others do not. Since observer-dependency is quite common with respect to dependence on allocentric references (Cuijpers et al., 2001), we think it is more fruitful to look at the individual analyses. In experiment 1, four out of seven observers showed significantly larger deviations for the slant when pointing parallel to a wall. Three of our observers showed a non-significant difference between the two conditions, while one observer (BH) showed results that differed from those of the other observers (rather small negative deviations). Hence, these results suggest that people differ in the extent to which they use allocentric sources of information to estimate the positions of objects in a scene.

The deviations in the vertical orientation were rather small. This is in line with our previous work (Doumen et al., in press A [Chapter 3]). In those experiments, we found the deviations of the tilt to be smaller than the deviations of the slant, independent of the relative distance. Thus, our first idea that the deviations of the tilt were small due to the fact that the

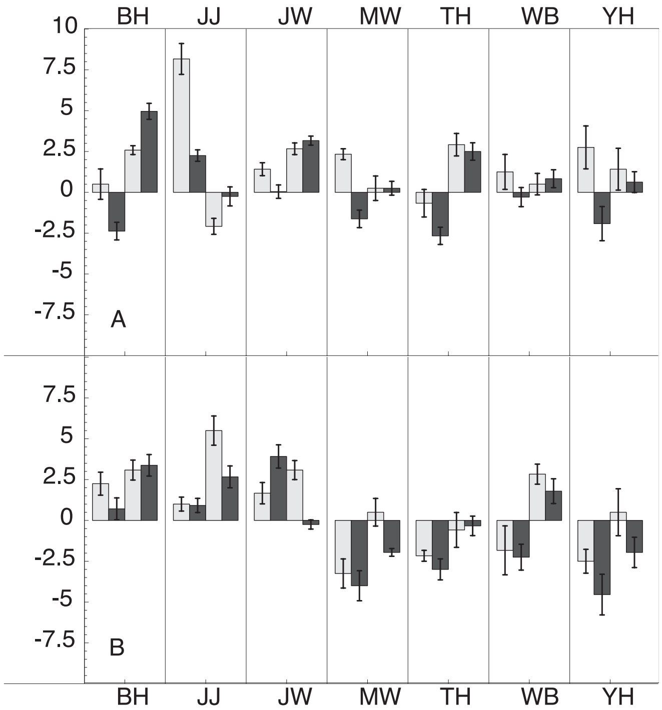

Figure 4.4
Bar-charts of the size of the deviations for the tilt in degrees in experiment 1. The light bars give the deviations for pointing parallel to the wall, the dark bars give the deviations for pointing not parallel to the wall. (a) This bar-chart gives the data for pointing horizontally: the first two bars for each observer give the data when pointing above eye-height, the second two bars when pointing below eye-height. (b) This bar-chart gives the data for pointing downwards (the first two bars) and upwards (second two bars).

task could be solved in the visual field, does not hold since the deviations are also small for settings outside the frontoparallel plane.

For some observers, the deviations in the horizontal orientation are smaller when an external reference is present directly behind the objects. We would like to know whether these observers also depended on internal references. Therefore we arranged a second experiment in which the set-up could be both parallel to a wall or not and frontoparallel or not. This way, we can investigate the possibility of interaction between internal and external references.

4.4 Experiment 2: Egocentric and allocentric references

Methods

The seven observers who participated in this experiment were all undergraduate students. Four of them were female, the other three male.

Points were marked on the floor and the raster, which lay on a circle with a radius of 130 cm around the middle of the room (see Figure 4.5). These were the positions of the ball and the pointer in the azimuthal plane. Four positions on the floor were marked for the observer (S1 through S4). The azimuthal distance between the observer and the middle of the circle was always 200 cm. For each position of the observer there was always a frontoparallel combination of points on the circle and two different oblique orientations that were at an angle of 45° to the frontoparallel plane. For each position of the observer, we used two pointing directions in the frontoparallel plane (from left to right and vice versa), and two oblique orientations with the pointer closer to the observer than the ball (from left to right and vice versa). The positions were chosen so that for two positions of the observer (S1 and S3), the frontoparallel combinations were also parallel to one of the walls, and so that for the other two positions of the observer (S2 and S4) the oblique orientations were parallel to one of the walls.

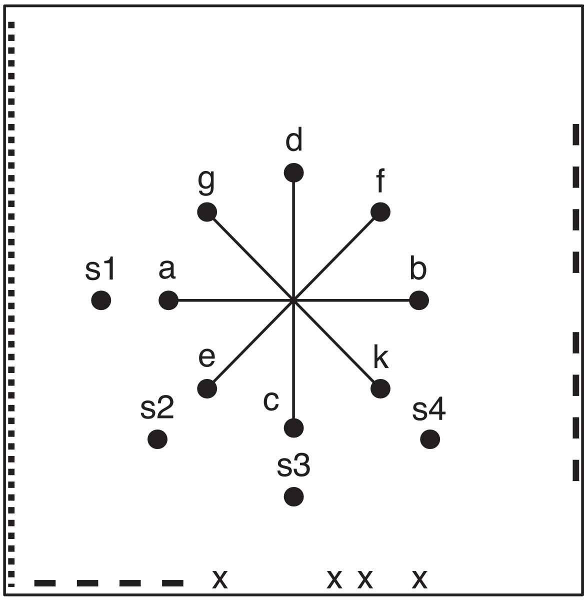

Figure 4.5
The square represents the walls of the room. The dotted line represents a wall with blinded windows and radiators along the base. The short dashed lines give the positions of the doors, and the crosses give the positions of four wooden strips that were attached to the wall. The observer was positioned at points S1 through S4, the objects (pointer and ball) at positions a through k. The lines between the points represent the pointing directions that were used.

The combinations of heights (60 and 240 cm above the ground) were the same as used in the previous experiment (See Figure 4.2). To do this experiment we therefore needed 4x2x2x4x3 = 192 trials (# positions of observers, # pointer heights, # ball heights, # combinations of azimuthal positions of pointer and ball, # repetitions). The randomization of the trials was the same as in Experiment 1. The observer finished all 64 different trials in a random order before starting with the first repetition. After the first repetition, all trials were presented for the third time. Each observer needed approximately seven hours to complete this experiment, which was done in sessions of an hour each.

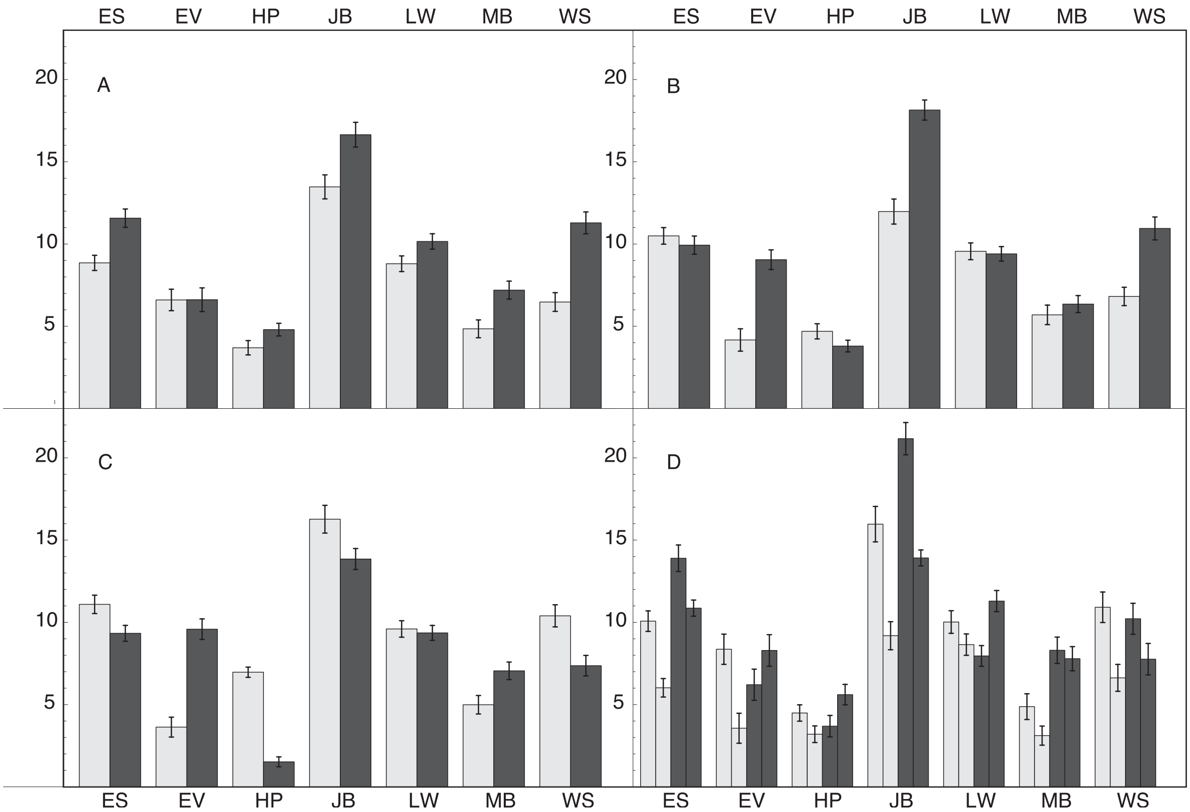

Figure 4.6
Bar-charts of the size of the deviations for the slant in degrees for each observer in experiment 2. (a) The light gray bars represent the means of the trials that were parallel to the wall, the dark gray bars show the means for the trials not parallel to the wall. (b) The light gray bars show the means of the trials that were frontoparallel, the dark gray bars represent the means for the trials not frontoparallel. (c) The light bars give the means of the trials when pointing from left to right, the dark bars when pointing from right to left. (d) The light bars represent the means for pointing horizontally (above eye-height and below eye-height respectively), the dark bars for pointing downwards and upwards.

Results of experiment 2

Slant

As in experiment 1 we found large deviations from veridical settings in the horizontal plane. All deviations (with only a few minor exceptions) were in the same direction, i.e. the observer pointed further away than the actual position of the target (overshooting). Figure 4.6 consists of four bar-charts in which the length of the bars represents the size of the deviations

***Table 4.2**. The F and p values for the ANOVA of the data of the slant in experiment 2*

	Walls		Frontoparallel		Direction		Height	
	$F_{(1,160)}$	p	$F_{(1,160)}$	p	$F_{(1,160)}$	p	$F_{(3,160)}$	p
ES	23.361	.000*	1.000	.319	7.614	.006*	33.271	.000*
EV	.000	.984	39.028	.000*	19.895	.000*	8.388	.000*
HP	6.300	.013*	4.147	.043*	78.852	.000*	5.709	.001*
JB	23.947	.000*	90.667	.000*	1.433	.233	58.641	.000*
LW	4.962	.027*	.066	.797	.019	.891	5.940	.001*
MB	13.043	.000*	1.014	.316	9.512	.002*	14.198	.000*
WS	44.943	.000*	33.186	.000*	4.801	.030	7.942	.000*

For the parameters height, walls and frontoparallel $\alpha = .05$ (one-sided).
For the parameter direction $\alpha = .025$ (two-sided).
**the slope deviated significantly from 0.*

in degrees. A positive value means an overshoot, a negative value an undershoot. The error-bars give the standard errors of the mean.

Figure 4.6 A shows the data for each observer sorted into two groups: parallel to a wall (light bars) and not parallel to a wall (dark bars). Most light bars are smaller than the dark ones, which means that the deviations were smaller when pointing was parallel to a wall. Figure 4.6 B shows the same data sorted in a different way, the light bars representing the trials in which the pointing direction was frontoparallel and the dark bars representing the trials in which the pointing direction was not frontoparallel. For three observers, EV, JB and WS, the light bars are smaller than the dark ones; thus for these observers pointing in a frontoparallel plane resulted in smaller deviations. Figure 4.6 C gives the data sorted for the pointing direction: pointing from left to right (light bars) and from right to left (dark bars). For some observers pointing from left to right yielded more veridical settings, whereas for other observers these settings were achieved by pointing from right to left. Figure 4.6 D gives four bars for each observer. The two light bars represent pointing horizontally (above and below eye-height respectively) and the two dark bars represent pointing from a high position to a low position and vice versa. The deviations for pointing horizontally are smaller for three observers (ES, JB, and MB).

A repeated measures analysis over all observers shows an effect of both parallelity to a wall ($F = 8.4$, $p = 0.027$) and frontoparallelity ($F = 14.0$, $p = 0.010$) but not of the direction of pointing (pointing from left to right and vice versa, $F = 3.4$, $p = 0.117$) and height differences (pointing horizontally or not horizontally, $F = 0.3$, $p = 0.615$). However, since we can see in Figure 4.6 that there are large differences between observers, we believed it to be more fruitful to analyse the data for each observer separately.

We conducted an analysis of variance on these data for each observer separately. We looked at the effect of parallelity to one of the walls and the effect of frontoparallelity, the effect of pointing direction (from left to right and vice versa) and the difference in height (pointing from a high to a high position, from high to low, from low to high, and from low to low). The combined effects of frontoparallelity and parallelity to a wall suggest four possible results: (1) only frontoparallelity has an influence, (2) only parallelity to a wall has an influence, (3) frontoparallelity and parallelity to a wall both have an influence and (4) neither frontoparallelity nor parallelity to a wall has an influence. Three of these possibilities, but not

Table 4.3. *The F and p values for the interactions found for the ANOVA of experiment 2*

	Effect 1		Effect 2		Effect 3	
ES	H*A	$F_{(3,160)}$=4.419, p=.005	A*E	$F_{(1,160)}$=6.343, p=.013		
EV	H*E	$F_{(3,160)}$=4.469, p=.005	D*E	$F_{(1,160)}$=9.447, p=.002	D*A*E	$F_{(1,160)}$=7.220, p=.008
HP	H*E	$F_{(3,160)}$=2.913, p=.036	D*A	$F_{(1,160)}$=21.541, p=.000	D*A*E	$F_{(1,160)}$=20.887, p=.000
JB	H*E	$F_{(3,160)}$=11.950, p=.000	A*E	$F_{(1,160)}$=4.265, p=.041	D*E	$F_{(1,160)}$=5.772, p=.017
LW	H*E	$F_{(3,160)}$=6.288, p=.000	A*E	$F_{(1,160)}$=8.189, p=.005		
MB	H*E	$F_{(3,160)}$=10.743, p=.000				
WS	D*E	$F_{(1,160)}$=6.157, p=.014	A*E	$F_{(1,160)}$=27.899, p=.000	H*A*E	$F_{(3,160)}$=3.307, p=.022

Only the interaction with an $\alpha < .05$ are given.

no. 4, applied to our group of seven observers. EV showed only an effect of frontoparallelity. This means that her deviations from veridical settings were smaller when the pointing-direction was frontoparallel than when this was not the case. ES and MB were not affected by this frontoparallelity but showed an effect of the presence of a wall parallel to the pointing direction. For these observers the deviations were smaller when a wall was present parallel to the pointing direction than when there was no wall parallel to this direction. Two observers, JB and WS, showed a statistically significant effect of both parallelity to a wall and frontoparallelity. The final observer, HP, showed an effect of both parallelity to a wall and frontoparallelity, but the weak effect of frontoparallelity was opposite to the trend described above. Hence, we can conclude that for this observer the deviations were smaller when a wall was present parallel to the pointing direction but not when the pointing direction was frontoparallel. See Table 4.2 for the F and p values of all observers. Figures 4.6 A and 4.6 B give a graphical view of these results.

EV, ES, HP and MB were all affected by the pointing-direction. HP and ES had larger deviations from veridical settings when pointing from right to left than vice versa, while EV and MB had larger deviations when pointing from left to right (see Figure 4.6 C). All observers were affected by the height- differences. We hypothesized that a difference would be found between trials where the pointer and ball were at different heights and trials where they were at the same height. To test this hypothesis, we added a planned contrast comparison to the analysis. For three observers (ES, JB, and MB), this comparison revealed an overall difference between pointing to a different height and pointing horizontally. Post-hoc tests revealed some other differences, but no clear pattern was found among observers. See Figure 4.6 D for the details.

There were many significant interactions. Five of our observers showed an effect of the interaction between height and frontoparallelity, four observers showed an effect of the interaction between parallelity to a wall and frontoparallelity. For three observers we found an interaction effect between direction of pointing and frontoparallelity. See Table 4.3 for the exact values. We will not give further details about these interactions because they do not contribute much to the overall discussion in this paper.

Tilt

With regard to the data of the tilt, the deviation in the vertical direction, the deviations showed positive and negative signs (see Figure 4.7). A positive value means that the observer pointed too high, a negative value means that the setting was below the veridical direction.

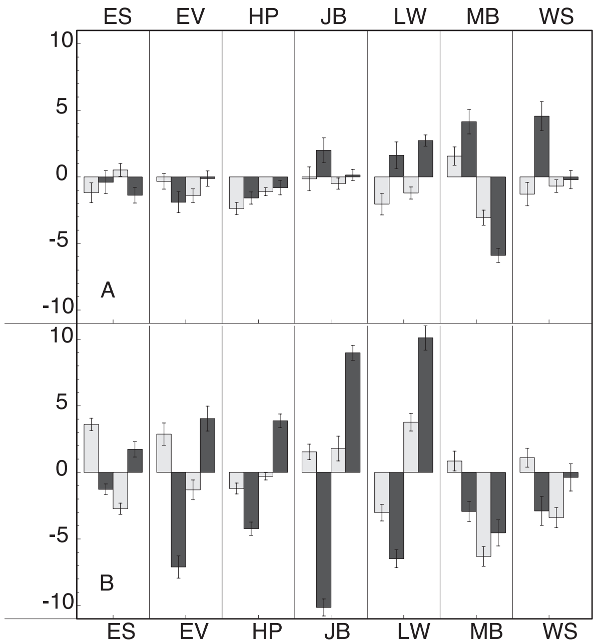

Figure 4.7
Bar-charts of the size of the deviations of the tilt in degrees for each observer in experiment 2. The light bars show the deviations for pointing frontoparallel, the dark bars show the deviations for not pointing frontoparallel. (a) This bar-chart represents the data for pointing horizontally: the first two bars for each observer represents the data when pointing above eye-height, the second two bars when pointing below eye-height. (b) This bar-chart shows the data for pointing downwards (the first two bars) and upwards (second two bars).

An analysis of variance is useless for these data because of the signed deviations. Hence, we will look at the data and give an extensive description of our findings.

Figure 4.7 gives bar-charts for each observer. In both graphs each group of four bars represents the data for all observers. The light bars represent the combinations of points that were frontoparallel, the dark bars those combinations that were not frontoparallel. Figure 4.7 A gives the data for pointing horizontally. The first two bars of each group represent the data for pointing above eye-height, the second two bars give the data for below eye-height. No clear pattern was found in these data. In Figure 4.7 B, the first two bars represent pointing from a high position to a low position, and the second two bars represent pointing from a low position to a high position. The most striking feature of this pattern is the change in sign that is evident in the first and second dark bars of the observers EV, HP, JB and LW. This means that the deviations were large and opposite in sign when pointing was neither horizontal nor frontoparallel to the observer. When the observers had to point upwards, the deviations were positive, meaning they pointed too high, whereas when the observers had to point downwards, they pointed too low. The other observers do not show such a clear pattern.

Discussion of experiment 2

Regarding the slant, the observers' results revealed large deviations. Positive deviations were found for both pointing in a frontoparallel plane and pointing away from the observer (the pointer was closer to the observer than the ball). The sign of the deviations when observers pointed away from themselves corresponded to the results of our 2D experiments with an exocentric pointing task (Doumen et al., 2005 [Chapter 2]). In those experiments we found that the deviations from veridical settings depended on the ratio of the distances of the two objects (ball and pointer) from the observer. When the pointer was further away from the observer than the ball, the deviations reflected an undershot, while when the pointer was closer to the observer, the deviations reflected an overshot. This effect was also found by Cuijpers and colleagues (2000A). However, the positive deviations for the frontoparallel settings, where the ratio between the distances from observer to objects was 1, were negligible in our previous work but in the present experiments we also found large deviations for frontoparallel settings.

Furthermore, we examined whether there was an effect of height-difference in the sense that we hypothesized that the deviations from veridical settings would be smaller when the veridical pointing direction was horizontal. For three observers, we found smaller deviations for the slant when the veridical pointing direction was horizontal. For the other observers we found no difference between pointing horizontally and not horizontally for the data of the slant.

For four observers we found that the pointing direction affected the slant-data, i.e. a difference in pointing from left to right and vice versa. This effect was found for both directions (pointing from left to right and from right to left) and could not be explained by the structure provided by the walls. Each condition was measured in two different orientations in the room, thus with two different backgrounds. This way, the structure on the walls surrounding the setup varied between the conditions. For position S1 (see Figure 4.5), the wall directly in front of the observer contained two light gray doors, whereas for position S3 the wall in front of the observer contained no doors or windows. For the positions S2 and S4 one of the walls that was visible when the observer looked straight ahead was the white wall, and the other was a wall with doors or windows. For position S2 the wall with more structure was on the right, while for position S4 the wall with more structure was on the left. Since the data for the two corresponding positions (S1 and S3, S2 and S4) were combined and the effect of pointing direction is present in both directions, we believe these effects of pointing direction are due to personal biases.

We found some effects of interaction between parameters in the analysis. One of the most salient interaction-effects that we found was the interaction between frontoparallelity and parallelity to a wall in the case of observers WS, LW, ES and JB. We can conclude from the data that for these observers, the deviations were smallest when the settings were both frontoparallel and parallel to one of the walls.

We found smaller deviations for the tilt than for the slant. For the frontoparallel settings, the same random pattern was found as in experiment 1. In the conditions oblique with respect to the observer, we found larger deviations for some observers. For four of our observers the size of these deviations was also dependent on the height-differences. Thus if the settings are frontoparallel and horizontal, the deviations of the tilt are small, but if the settings are oblique with respect to the observer and not horizontal, than the deviations are larger. So in these trials we found an “oblique effect” with regard to the size of the deviations. This is in contrast with the oblique effect that is most often discussed in the literature (Appelle, 1972), concerning the variance of the data. The “oblique effect” we found was especially prominent with regard to egocentric references, which is in agreement with the results reported in the work by Chen and Levi (1996) and Darling and Bartelt (2005).

4.5 General discussion and conclusions

The first important point we want to make is that large deviations from veridical settings were found with the 3D exocentric pointing paradigm. For the slant, deviations of up to 20 ° were found. When comparing these results to results obtained with a 2D exocentric pointing task (Doumen et al., 2005 [Chapter 2]), we can conclude that the spread of the data is larger for the 3D task. This might reflect the increased level of difficulty in this task that involved an extra degree of freedom, namely the orientation in the vertical plane. The sizes of the devations are in agreement with our previous work with the 3D exocentric pointing task (Doumen et al., in press A [Chapter 3]).

Looking at the data for the slant, we see that most observers overshoot when the pointing-direction is frontoparallel. When we compare this result to the results of research by Koenderink & van Doorn (1998), we find the same pattern even though that research was done outdoors while ours was done indoors. At least for frontoparallels that are closer than 4 m from the observers, the perceived frontoparallels were all concave towards the observer. For this distance this effect is irrespective of the task that is given to the observer. Koenderink, van Doorn, Kappers and Lappin (2002), for example, used a line bisection task and found comparable results at these distances.

Overshooting was also found in the conditions where the pointer was closer to the observer than the ball (experiment 2). This replicates the results of our 2D experiments in which we used an exocentric pointing task, and varied the relative distance of pointer and ball to the observer (Doumen et al., 2005 [Chapter 2]).

We found much smaller deviations from veridical settings for the tilt than for the slant. The deviations rarely exceeded 5°, particularly when the veridical settings were in the frontoparallel plane. In the frontoparallel trials, the task can be solved in the visual field, i.e. the coordinates of the retina correspond to the coordinates in this plane, which makes the task less difficult. When the veridical orientation was oriented obliquely towards the observer (exp 2) and not horizontally, the deviations were larger for most observers. For four observers we found a pattern that involved pointing too high when pointing downwards and pointing too low when pointing upwards. Hence, for these observers visual space seems to be compressed in the vertical orientation. We could link this to the idea that visual space seems "stretched" in the horizontal plane (Doumen et al., 2005 [Chapter 2]). When visual space is "stretched" in the horizontal direction, the perceived tilt will be smaller than the physically veridical tilt.

The most salient question posed in this paper concerns the types of information people use when doing an exocentric pointing task. To obtain some insight into this issue, we let observers point parallel to a physical (allocentric) reference, a wall parallel to the pointing direction, and to an egocentric reference, a frontoparallel plane. In the first experiment, all trials were in the frontoparallel plane, but only one condition had pointing orientations parallel to a wall. In this experiment we found that four out of seven observers were influenced by the presence of the wall. This means that the deviations from veridical settings were smaller when the observers pointed parallel to the wall. We could say that these people were able to use the structure provided in order to perform the task more veridically. In the second experiment, where we varied both the allocentric and the egocentric references, six out of seven observers showed smaller deviations when there was a wall parallel to the pointing direction. This experiment gives us information not only about the use of allocentric references but also about the egocentric references we discussed above. The observer who was not influenced by the walls used this egocentric information. Two other observers were also influenced by this egocentric reference, so these observers were influenced by both egocentric and allocentric references provided during the task. This result partly contradicts the results of Cuijpers et al (2001). All their observers were influenced by allocentric references, although which structure of the visual scene was used (the cabin or the wall of the room) was observer-dependent. However, they found none of their observers to be dependent on egocentric references. This difference could be due to the fact that the observers used by Cuijpers et al. were restricted in their head-movements by a chin-rest. But it could also be due to the different task used (a parallelity task instead of an exocentric pointing task). In our earlier work, we also found differences between settings for a parallelity task and an exocentric pointing task. We believe that this could be due to the fact that for an exocentric pointing task an observer needs to make a judgment about the positions of both the pointer and the ball whereas for a parallelity task the observer does no have to know the actual positions of the two rods in order to do the task. If the positions of the two objects are judged by their relative distances from the observer, it is not surprising to find that observers depended on egocentric references. Furthermore, following this line of reasoning, it is not surprising that no such dependence was found in the experiments of Cuijpers and colleagues, that were done with a parallelity task.

From our experiments we can conclude that people make a personal selection of the sources of information they use for judging depth in a scene. This set of useful sources is probably based on the prior experiences and knowledge of each individual. We cannot yet say whether this involves a weighting up or a selection of these different sources, or whether the knowledge is based on a statistical process or on the weighting up of reliabilities of cues. Of course, we should be careful with making generalizations to other situations and experimental set-ups. However, in our experiments, we found effects of contextual information, and we believe that our set-up is not a special case. Thus, we assume that our results can be generalized to other conditions. Therefore, we know that there are differences between observers and that both information from the context (allocentric references) and information from the observer himself (egocentric information) are important for judging the positions of objects.

Chapter 5
Forward-backward asymmetry: the effect of body position and pointer shapes

Abstract:
In previous experiments with an exocentric pointing task we found that deviations from veridical settings were larger when the pointer was closer to the observer than the ball but smaller when the pointer was further from the observer than the ball (the forward-backward asymmetry). We investigated possible origins of this asymmetry. First we tested whether, when the pointer is further from the observer than the ball, the observer can use his own position as a reference to orient the pointer. The pointing angle is restricted in this condition. We therefore tested whether restricting the pointing angle by poster-boards when the pointer was closer to the observer than the ball could cause a reduction in the size of the deviations. This turned out to be the case. Furthermore, we tested whether the observer's view of the pointer could also account for the differences we found. Therefore we tested whether a pointer providing the same amount of structure for the two combinations of pointer- and ball-positions that we used would induce settings that differed from those induced by a pointer not providing the same amount of structure in both conditions. However, the results show that the view of the pointer cannot explain the forward-backward asymmetry we found in our earlier experiments. From these studies we can conclude that people can use their own position as an egocentric reference for directing pointers to targets.

In preparation for publication

5.1 Introduction

Much of the research into visual space is concerned with describing the spatial structure. The structure of visual space varies with the viewing conditions. Different tasks often yield different results. For example, Koenderink and colleagues found conflicting results for an exocentric pointing task and an apparent frontoparallelity task (Koenderink, van Doorn, & Lappin, 2000; Koenderink, van Doorn, Kappers & Lappin 2002). It is probable that people need different kinds of information in order to perform adequately in experimental tasks (Cuijpers, Kappers, & Koenderink, 2002; Doumen, Kappers, & Koenderink, 2005 [Chapter 2]). This might be reflected in the results obtained in different experiments. Furthermore, the environment in which experiments are conducted also influences the data (Battro, di Pierro Nettro, Rozestraten, 1976). To what extent are the conditions constrained and do they simulate natural viewing conditions? By constrained conditions we mean that the stimulus is reduced to the point where some depth cues are no longer present. Some scientists investigate visual space by removing these depth cues. This way, other cues for the perception of positions of objects can be investigated directly (for example Cuijpers, Kappers & Koenderink, 2000). In another approach natural viewing conditions are used to investigate the perception of observers who have all the information that is present in normal situations (e.g. Kelly, Loomis, & Beall (2004) and Koenderink et al. (2002)). These two lines of research are both important for our understanding of visual space. Not only are there differences in viewing conditions, the distances used also vary considerably: from within arm's reach to distances of up to 25 meters. Since the deformation of visual space was found to vary with distance from the observer (Battro, et al., 1976; Koenderink et al., 2002), it is not surprising to find that similar experiments produce conflicting results.

Even with the same experimental conditions, we found some discrepancies that require an explanation. We have been working with a two-dimensional exocentric pointing task. During this task, observers had to orient a pointer in the horizontal plane at eye-height towards a ball in the same horizontal plane. We found a difference between the condition in which the pointer was further away from the observer than the ball and the condition in which the pointer was closer to the observer than the ball (i.e. when we vary the relative distance) (Doumen, et al., 2005 [Chapter 2]). The observers overshoot the target when the pointer is closer to the observer than the ball (situation A from Figure 5.1), whereas they point between themselves and the ball when the pointer is further away from the observer than the ball (situation B from Figure 5.1). However, the sizes of the deviations also differ for these two conditions. In situation B (when pointing backward), the deviation is smaller than in situation A (when pointing forward). For example, typical settings for an exocentric pointing task are deviations of 3° in situation B and 10° in situation A (mean over all observers in Doumen et al. (2005) [Chapter 2] with a separation angle of 60° and a relative distance of 3 or 1/3). In this paper we will refer to the difference in size of the deviations for the two situations as "forward-backward asymmetry".

We can think of a couple of explanations for this asymmetry. These explanations do not concern the sign of the deviations but they concern their size. Why the sign of the deviations changes when the relative positions of the pointer and the ball change, is another question that needs to be investigated. However, we will not address this question here. In situation B, the observer could be using his own body position as a reference. In this case the

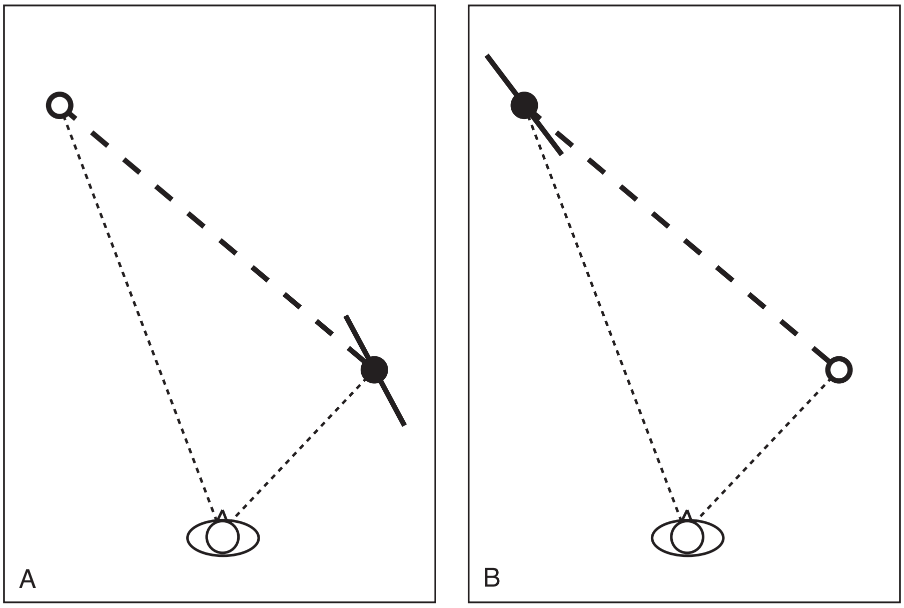

Figure 5.1
An example of typical settings in an exocentric pointing task. The thick dashed line is the veridical orientation between the pointer and the ball, the short lines the settings of the observer. Both conditions are plotted here: the pointer further away from the observer than the ball and vice versa.

angle between the lines connecting the pointer to the ball and the observer serve as a restriction for pointing to the ball. An observer, for example, will not point towards himself while trying to point at the ball. Another possible explanation concerns the view of the pointer. When the pointer consists only of a single rod, the observer can experience a difference in the amount of information when the pointer is at different positions in the room. For example, in situation B, a slight rotation of the pointer will induce a large change in the observer's retinal image of the pointer. This is in contrast with situation A in which a small rotation of the pointer will not change the observer's retinal image very much. In the latter condition, less information is supplied to in the retinal image of the observer.

In the present paper, we investigate the reasons for the difference that we found in the size of the deviations for different pointer and ball positions. To do this, we conducted two experiments with the exocentric pointing task. To address the first possible explanation, we restricted the space in situation A in the same way as in situation B. We did this by placing poster-boards in the room as depicted in Figure 5.2 (at the thick black and gray lines). We measured two conditions: the first with the poster-boards as described above and the second without boards. If the observer used himself as a reference and thereby pointed more veridically when the pointer was far away from him (situation B), then placing poster-boards on the other side of the room would give him the same restriction in pointing angle and would therefore result in smaller deviations than if the poster-boards were not present when the pointer was close to the observer (situation A).

In the second experiment, we varied the shapes of the pointer. One of the pointers consisted of two rods perpendicular to each other. This double-rod pointer contained the same amount of information for the various pointer positions that are used in experiment 2.

Another pointer consisted only of a rod (the single-rod pointer) and provided completely different information depending on whether pointing towards and away from the observer. If observers use the kind of information that is provided by the pointer, we would be unable to see any difference in the size of the deviations when the double-rod pointer was used in situation A or in situation B. However, we would see forward-backward asymmetry when the single-rod pointer was used.

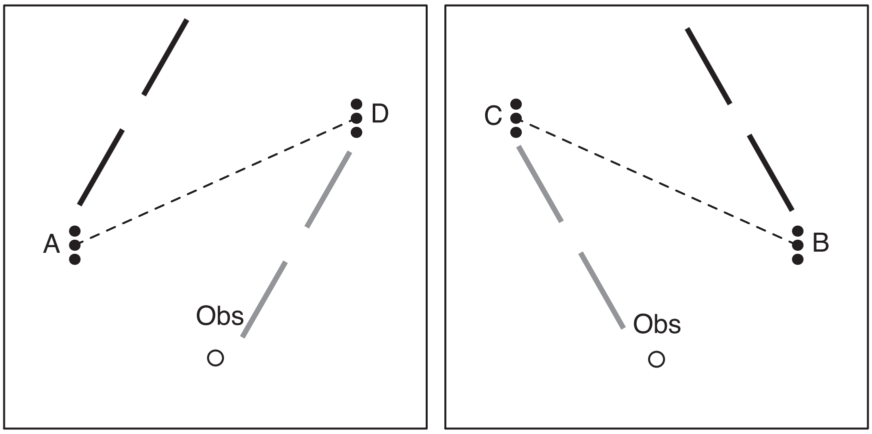

Figure 5.2
Two top-views of the experimental room for experiments 1. The squares represent the walls of the room, the circles the positions of the observer and the dots the positions of the ball and pointer. The pointer was always positioned on the middle dot, the ball could be positioned on all dots. The dashed lines connect the positions that are used together as pointer and ball-positions. The two panels give the situation for pointing between A and D and between C and B respectively. The black lines in these panels represent the positions of the poster-boards far away from the observer. The gray lines represent the poster-boards close to the observer.

5.2 General methods

Observers

Undergraduate students, who were paid for their efforts, participated in the experiments described below. They had little or no experience in psychophysical experiments and were naive as to the goals and purposes of the experiments. They all had normal or corrected-to-normal visual acuity. Before the start of the experiment, each observer was tested for stereovision. They all had a stereo-acuity of more than 60".

Experimental set-up

The experimental room measured 6 m by 6 m. The wall opposite to the observer was white, with some electrical sockets near the floor. The wall on the left-hand side of the observer contained four blinded windows with radiators underneath them. The wall on the right-hand side contained two grey doors. Markers were placed on the floor to mark the positions of the objects that were used in the experiments. From the ceiling a horizontal iron grid was suspended (3 m above the floor). Attached to this grid were some white cubes that were used for other experiments. The cubes measured 20 cm x 30 cm x 15 cm and did not interfere with the stimuli in the visual field of the observer. From the iron grid hung a green ball 1.5 meters above the ground. The ball had a diameter of 6 cm and could be placed on different positions. The pointer we used for the first experiment consisted of a horizontal green rod (25 cm long and 2 cm thick, with one sharp endpoint) bisected by a yellow disk (diameter 8.2 cm and thickness 1 cm) that was perpendicular to the rod. The rod passed through the center of the disk and the disk was attached to a vertical iron rod so that the height of the rod was 1.5 meters. The pointer is depicted in the third panel of Figure 5.5. The foot of the pointer contained the motor that rotated the pointer and a protractor that was used to read the settings of the observer. A screen in front of the foot prevented the observer from seeing the protractor and the square shape of the foot. The observer could rotate the pointer using a remote control. The observer was seated on a revolving chair that could be adjusted in height so that the eye-height of each observer was at 1.5 meters.

Procedure

In both experiments described below, the same procedure was followed. The observer sat on the revolving chair and the chair-height was adjusted. The experimenter explained the task: with the sharp endpoint of the rod point to the ball as accurately as possible. The observer was instructed to remain seated and not to move his upper-body. He was allowed to rotate his head and the chair. The observer always had a practice trial before starting the real experiment. After each trial, the observer had to close his eyes so that he could not see the experimenter reading the protractor and rearranging the objects for the next trial. The observer was also asked to rotate the pointer a little before opening his eyes again. This ensured that he did not have information about his previous setting.

Analysis

The dependent variable of the experiments discussed below is defined as follows: a positive deviation means that the observer overestimates the distance between the two objects in the y-direction. The deviation is negative when this distance is underestimated. This definition differs from the definition used in our previous experiments! We changed our definition because we wanted to be able to quantify differences in the size of the deviations for all relative distances.

For each experiment we did a number of paired t-tests in order to analyze the data. We chose this type of analysis because the more standard analyses could not answer our questions directly.

5.3 Experiment 1: The effect of body position

In experiment 1 we investigated the question whether we could influence the observer's settings when the pointer was positioned closer to the observer than the ball by placing poster-boards to restrict the pointing angle. The idea was to restrict the pointing angle in situation A to the same degree as was the case in situation B. Figure 5.2 depicts the lay-out of the floor of the experimental room. The observers always had to point from A to D, from B to C, from C to B or from D to A. In the figure, the angle between the black lines (representing the positions of the poster-boards) and the line connecting A and D in the first panel (B and C in the second panel) is the same size as the angle between the line between the observer and D (or C in the second panel) and the line connecting A and D (B and C in second panel). We used the poster-boards to restrict the angle of pointing just as the observer's body restricts the angle of pointing when the pointing-direction is towards the observer. Furthermore, we placed poster-boards on the gray lines in figure 5.2. We did this because the mere presence of the poster-boards close to the pointer might affect the data. Thus, in the experimental condition the pointing angle was restricted by poster-boards for both situations in Figure 5.1. In the control condition no poster-boards were visible to the observer.

Methods

Six male and two female observers participated in the experiment. They were all undergraduate students.

The experimental set-up consisted of a pointer, a ball, four poster-boards and a revolving chair. The poster-boards measured 200x90 cm (hxw) each board was covered with equally sized posters with a brick wall pattern on them. The posters were used to help the observer to see the orientation of the poster-boards clearly. The coordinates (in cm) of the positions of these objects were as follows: the observer was placed at (0, 0), the pointer and ball could be either at position A (-200, 160), B (200, 160), C (-200, 340) or (200, 340) as depicted in the two panels of Figure 5.2. We varied the positions of the ball to prevent the observer from using the information provided by his view of the pointer to guide his settings. The difference between the positions was 20 cm. For example, the coordinates of the ball at position A could be either (-200, 140), (-200,160) or (-200,180). The difference was large enough for the observer to notice the difference between the positions, but small enough not to change the relative distance too much. In addition to the experimental condition with the poster-boards we had a control condition without the boards. This resulted in 4x2x3 = 24 trials (# pointer & ball combinations, # experimental conditions, # repetitions). The experiment was measured in approximately half an hour for each observer.

Results

The data of experiment 1 are plotted in Figure 5.3. Each bar represents the mean of the values for all eight observers. The black bars represent the data for the control condition, i.e. without the poster-boards. The gray bars represent the data for the experimental condition, i.e. with the poster-boards. Each group of bars represents the results for one position of the pointer. The first two groups are the cases in which the pointing direction is forward, the second two groups of bars represent the cases in which the pointing-direction is backward.

The first thing to notice is that in the condition without poster-boards all deviations are positive and the deviations for positions A and B are larger than for C and D. Furthermore, the deviations are larger in the conditions without poster-boards than in the conditions with poster-boards.

We conducted three paired t-tests on these data ($\alpha = 0.0167$ (0.05 / 3)). We found a difference between pointer-positions far away and close by for both experimental conditions (without boards: $T = 6.925$, $p < 0.001$, with boards: $T = 5.411$, $p = 0.001$). Furthermore, we found a significant difference between the two experimental conditions when the pointer was closer to the observer than the ball ($T = 4.825$, $p = 0.002$).

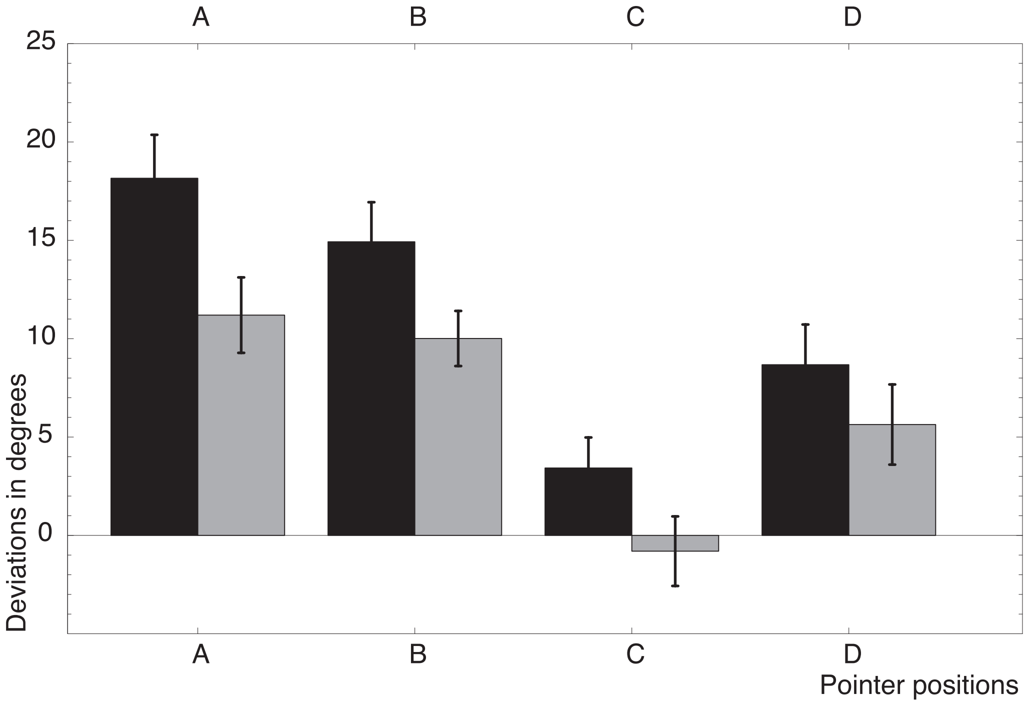

Figure 5.3
The data of experiment 1 for all observers together. The deviation in degrees is plotted against the pointer positions (A through D). The black bars represent the condition without poster-boards, the gray ones the condition with poster-boards. The error-bars represent the standard error of the means.

Discussion

The effect of the difference in the sign of the deviations when the pointer is close to the observer or far away, which we found in our earlier work and in the work of Cuijpers and colleagues (Cuijpers, et al., 2000; Doumen, et al. 2005 [Chapter 2], in press A [Chapter 3]]), is replicated in the present work. The observer overshoots the ball. We also replicated the forward-backward asymmetry.

The presence of the poster-boards seems to decrease the size of the deviations when the pointer is closer to the observer. However, if the restriction of the pointing angle were to be the complete explanation for the relative distance effect, there would be no difference in the size of the deviations when the pointer is at position A or B with boards and when the pointer is at position C or D. However, the deviations are larger for pointer positions A/B with poster-boards than for pointer-positions C/D. Hence, the restriction of the possible pointing-directions is not the whole story; possibly other factors play a role or the presence of the poster-boards is not as strong a cue as the presence of the observer himself.

5.4 Experiment 2: The effect of pointer shapes

In the introduction we suggested that the difference in the size of the deviations for different pointer-positions might be due to the different views that observers have of the pointer. In all our experiments conducted so far we used a pointer consisting of a rod and a circular disk perpendicular to the rod. Since the shape of a rod differs from the shape of a disk, one could say that, depending on the observer's view of the pointer, an observer can have different amounts of information about the orientation of the pointer. Consider a very simple pointer, a rod, positioned close to the observer and pointing towards a ball further away from the observer. Rotation of the pointer by a few degrees does not change the observer's retinal image of the pointer as much as it would if the pointer further away from the observer than the ball were rotated by the same amount. For example, in the experiment described below, the pointer has a visual angle of 7.38° when it is positioned at A or B and pointing towards D or C, respectively (see Figure 5.4). Rotating the pointer 2° to the left and 2° to the right results in visual angles that differ in size by 0.25°. This value divided by the total visual angle gives a ratio of 0.034. When the pointer is at positions C or D (and the ball at positions B or A), however, the total visual angle is 1.18°, the difference in visual angle is 0.21 and the ratio is 0.17. The ratio for situation B is five times larger than it is for situation A. Thus in situation A, the observer obtains less information from the pointer itself than in situation B. In the present experiment, we tested whether the difference in the amount of information given by the pointer can explain the forward-backward asymmetry. To do this, we constructed two new pointers. One pointer consisted of two rods, one perpendicular to the other. Since the angles between observer, pointer-position and ball-position for the two relative distances differ by exactly 90°, the pointer contains the same information for the two pointer-positions. The other pointer consisted of only one rod. If our hypothesis is correct, we would expect to find smaller deviations for situation B than for situation A with the single-rod pointer but no difference between the size of the deviations with a double-rod pointer. Furthermore, we compared the settings of these two pointers with the settings of the pointer with the disk that was used in our earlier experiments.

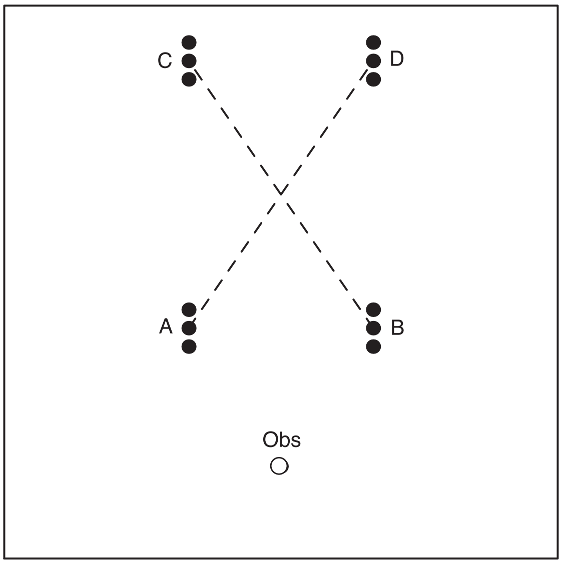

Figure 5.4
A top-view of the experimental room for experiment 2. The lines represent the walls of the room, the circle represents the position of the observer, and the black dots the positions of the ball and pointer. The pointer could be positioned only at the middle dots, the ball on all dots. The dashed lines connect the points that are used together as pointer and ball-positions.

Methods

In this experiment, eight undergraduate students (five male and three female) participated as observers. We did the experiment in the same experimental room as the previous experiments. However, the positions of the observer and the pointers and ball differed from the first experiment. We had to change the positions of the ball and pointer in order to obtain a difference in angle between the observer, pointer and ball of 90° for the two pointer-positions. The ball and pointer could be in four different positions (with the following coordinates in cm with the position of the observer at the origin): A(-100, 150), B(100, 150), C(-100, 440) or D (100, 440) (see Figure 5.4). The balls could be either at one of the given positions or 20 cm in front of beyond that point; thus the ball in position A was hanging at position (-100, 130), (-100, 150) or (-100, 170).

As mentioned above, we used three different pointers that could be positioned on the vertical iron rod at a fixed orientation. One pointer, the double-rod pointer, consisted of two blue rods 25 cm long and 1 cm thick, one rod having a sharp end-point 5 mm above and perpendicular to a rod without a sharp end-point. The upper-rod was at a height of 1.5 meters; this was the rod that the observer had to use to point to the ball. The difference in height was such that the observer had an (almost) unobstructed view of both rods at all times. The other pointer, the single-rod pointer, consisted of one blue rod 25 cm long and 1 cm thick with a sharp end-point in one direction. This rod was placed on a height of 1.5 meters. The third pointer was the pointer that was used in experiment 1. This pointer consisted of a green rod, with one sharp point, perpendicular to a yellow disk. Figure 5.5 depicts of the three pointers that were used.

We had three different pointers, four different pointer- and ball-positions, and three repetitions, which resulted in 3x4x3 = 36 trials.

Figure 5.5
Photographs of the three pointers used: and the double-rod pointer, the single-rod pointer, and the rod with disk pointer respectively.

Results

In figure 5.6 the data of experiment 2 are presented in a bar-chart. The groups of bars represent the four different pointer-positions, whereas the three bars in each group represent the three pointer-conditions, namely the double rod (dark bar), the single rod (gray bar) and the rod with disk (light bar). Just as in the previous plots, the deviations are positive when the position of the ball is overshot and negative when this position is undershot. Each bar represents the mean of the data for all eight observers who participated in this experiment. The error-bars represent the standard error of the means. The first thing to notice when looking at this bar-chart is that the deviations for the double-rod pointer are generally smaller than for the other two conditions. Furthermore, the deviations for pointer-positions A and B are slightly larger than for pointer-positions C and D.

For experiment 2 we did six paired t-tests ($\alpha = 0.0083$ (0.05 / 6)). We tested whether for each of the three experimental conditions there was a difference in the deviations for pointer-positions A/B and C/D. Furthermore, we tested whether there was a difference between the three experimental conditions for the pointer-positions A/B. The only trend we could find was a difference between the single-rod pointer and the double-rod pointer (for pointer-positions A/B: $T = 3.410$, $p = 0.011$). This difference, however, does not reach significance with the Bonferroni-correction. The results of all comparisons are given in Table 5.1.

Besides looking at the mean deviations, we investigated the effects of the pointer shapes and positions on the variance of the data. More specifically, we again conducted six paired t-tests. We tested whether the size of the standard deviations differed for the same groups of trials that were tested for the mean deviations. As can be seen in Table 5.1, there is a significant difference between situations A and B for the single-rod pointer ($T = 3.632$, $p = 0.008$) and a small trend towards a difference between the double-rod pointer and the single-rod pointer in situation A ($T = 3.163$, $p = 0.016$).

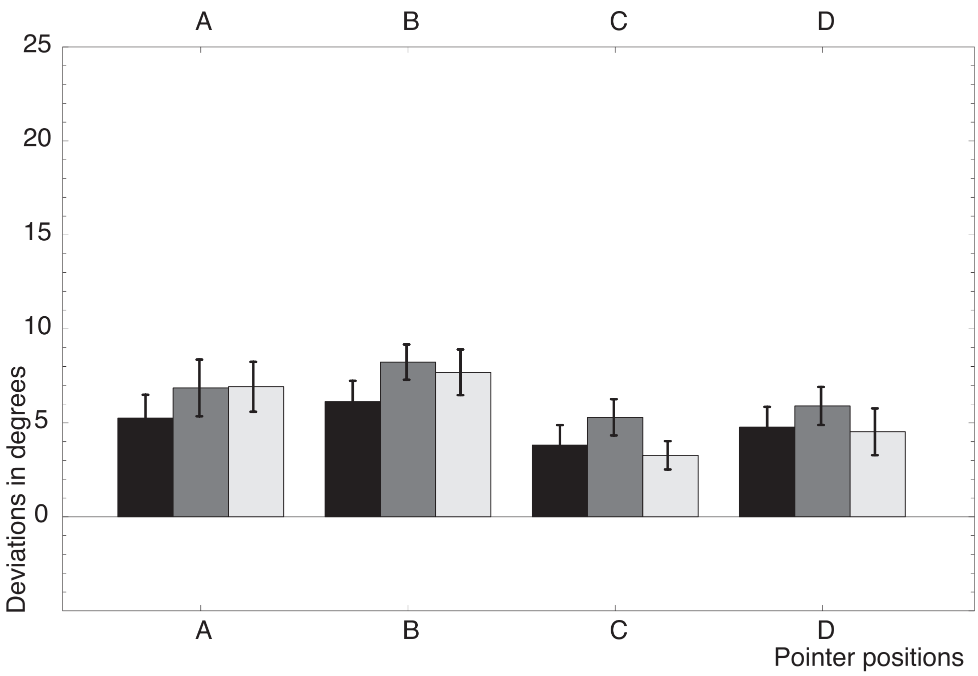

Figure 5.6

The data for experiment 2 for all observers together. The deviation in degrees is plotted against the pointer positions (A through D). The black bars represent the condition with the double rod pointer, the dark gray ones the condition with the single rod pointer, and the light gray ones with the rod with disk pointer. The error-bars represent the standard error of the means.

Discussion

We found no significant effect of the mean deviations for the different pointer positions for the three experimental conditions. Although the differences between the pointer positions A/B and C/D are largest for the pointer that we have been using so far, the differences do not reach significance in this experiment. When looking at the data for each observer separately, we see that there are differences between observers. Most observers show an effect of the relative distance of the pointer and ball, but some observers show a rather incoherent pattern. This might be due to the relatively small visual angle that was used in this experiment together with the short distance to the pointer- and ball-positions close to the observer.
However, we did find a significant difference between the standard deviations of the data for situations A and B for the single-rod pointer only. Thus, the difference in the amount of information of the single-rod pointer and the double-rod pointer reveals itself in the accuracy of the settings rather than in the mean deviations.

Table 5.1: *The results of the paired t-tests of experiment 2 for the mean values and for the standard deviations*

Group 1	Group 2	T mean	p* mean	T sd	p* sd
Double rod AB	Double rod CD	0.838	0.430	-1.259	0.248
Single rod AB	Single rod CD	1.145	0.290	3.632	0.008
Rod with disk AB	Rod with disk CD	1.985	0.088	-0.789	0.456
Double rod AB	Single rod AB	-3.410	0.011	-3.163	0.016
Double rod AB	Rod with disk AB	-2.055	0.079	-0.249	0.810
Single rod AB	Rod with disk AB	0.384	0.712	1.694	0.134

**The difference is significant at $\alpha = 0.0083$ (with Bonferroni correction)*

5.5 General discussion and conclusions

In this paper we presented two explanations for a difference we found in deviations in an exocentric pointing task with varying relative distances. We tried to explain why deviations are smaller for backward pointing than for forward pointing. We gave two possible explanations: one concerning the position of the observer that can be used as a reference point when the pointer is far away from the observer. The second explanation was the difference in the view of the pointer in the two conditions. The position of the observer restricts the pointing direction when the pointer is far away from the observer. Restricting the pointing direction to a similar degree in the other direction resulted in smaller deviations in the condition where the pointer is close to the observer and the ball far away. However, the size of the deviations was not reduced to the size of the deviations in the condition when the pointer is far away from the observer. From this we can conclude that the position of the observer restricts the pointing direction in the conditions where the pointer is further away from the observer than the ball more effectively than do the poster-boards when the pointer is closer to the observer than the ball. However, it could be that restriction of the pointing angle is not the whole story: another factor may be involved as well.

Our second explanation concerned the observer's view of the pointer. We found that although there is a trend towards a difference in the mean deviations between two extreme pointers (single rod or double rod) there was no difference in mean deviations for the double-rod pointer and the pointer used in the previous experiments. However, we did find a difference in standard deviations for the two situations for the single-rod pointer but not for the other two pointers. This means that the amount of information provided by the shape of the pointer does influence the accuracy of the settings in an exocentric pointing task.

There are a few other differences between pointing forward and backward. For example, in the backward pointing condition, the visual angle of the pointer is much smaller than for forward pointing. Furthermore, when pointing backward with the rod-with-disk pointer, a large part of the rod is occluded by the disk. This is not the case in the forward pointing condition. However, these differences would predict more veridical settings in the forward pointing condition than in the backward pointing condition. This is contradictory to our findings.

Although the shape of the pointer does affect the accuracy of the settings, it cannot explain the forward-backward asymmetry. However, the body position of the observer does seem to restrict the angle of pointing for conditions in which the pointer is further away from the observer than the ball. Apparently, people make effective use of an egocentric reference such as body position.

Chapter 6
Visual perception of tilted planes

Abstract:
The ball-in-plane task tests whether observers can accurately place a ball by eye in a plane defined by three reference balls located in three-dimensional space. We varied the direction in which the plane was tilted and the shapes of the triangles formed by the reference balls. We found that the observers placed the test-ball too low when they looked at the plane from above. The observers had small negative or positive deviations when they look at the plane from below, with a downward trend when the azimuthal angle diverges from 0°. These data suggest that, at least when the observers looked at the plane from above, they perceived the plane as concave.

In preparation for publication

6.1 Introduction

Visual space has been investigated by means of various tasks. The first experiments were done with parallel and equidistance alleys in the horizontal plane (Hillebrand, 1902; Blumenfeld, 1913). The main conclusion from these experiments was that parallel and equidistance alleys differed from each other: the equidistance alleys were found to lie outside the parallel alleys. Battro, di Pierro Nettro and Rozestraten (1976) used comparable alley experiments on a larger scale in outdoor settings. They found no uniform deformation of visual space across scales and observers. Furthermore, Indow and Watanabe (1984, 1988) investigated the frontoparallel and the horizontal subspaces with alley-experiments. They concluded that the deformation of visual space varies with scale, observers and subspaces (Indow & Watanabe, 1988).

Another task that is often used is a horopter task (an 'apparent frontoparallelity task' is a better name). In this task, observers have to form with light-points or stakes (depending on the experimental set-up) a horizontal line that is frontoparallel. One or more points are fixed and the other points have to be placed in the same plane. Some scientists found that most observers form an apparent frontoparallel line that is concave towards the observer for smaller distances (up to 5 meters) whereas the apparent frontoparallel is convex towards the observer when the distances are larger (Foley, 1991). Other scientists found the apparent frontoparallel plane to be concave for both far and near distances (Koenderink, van Doorn, Kappers, & Lappin 2002); while Battro, et al. (1976) found no consistent pattern for their observers. Nevertheless, most scientists agree that the apparent frontoparallel plane is concave towards the observers at distances of up to 5 meters.

Other methods of studying the relative distances of objects from an observer consist of exocentric pointing, parallelity, body-pointing or collinearity tasks. These tasks were done mainly in the horizontal plane at eye-height. Results were dependent on experimental conditions. For example, Kelly, Loomis and Beall (2004) compared the results of a body-pointing task, in which the observer had to direct his body in the same orientation as a line-segment (defined by two stakes), with a collinearity task, in which the observer had to place a point collinear to two other points. They found similar results for the tasks, although there was a bias in the body-pointing task that was not present in the collinearity task. Furthermore, one should be careful in comparing tasks that may rely on different sources of information. For example, different results can be obtained with a parallelity task (putting a horizontal rod parallel to another rod) and an exocentric pointing task (directing a pointer at a ball). We found that the different results can be explained by the fact that for the parallelity task an observer does not need to know the exact positions of the two rods (Doumen, Kappers, & Koenderink, 2005), whereas this information is essential for the exocentric pointing task. Thus, particular constraints of the tasks do influence the results.

Recently, we developed a three-dimensional exocentric pointing task (Doumen, Kappers and Koenderink, in press A [Chapter 3], in press B [Chapter 4]) to enable us to progress from two-dimensional experiments to three-dimensional versions. The deviations in the horizontal plane found with the three-dimensional exocentric pointing task were comparable to the deviations found with two-dimensional tasks. Furthermore, we found that the relative distance of the ball and the pointer to the observer had no effect on the deviations in the vertical plane. Since the deviations we found for the horizontal plane were dependent on the relative distance, we conclude that visual space is anisotropic, as already mentioned by Indow and Watanabe (1988).

Just as we can expand our knowledge by proceeding from a two-dimensional exocentric pointing task to a three-dimensional one, we can also expand the apparent frontoparallelity task to a three-dimensional version. In the light of the anisotropic data that were produced for different sub-spaces (Indow & Watanabe, 1988), it is a logical step to develop tasks for measuring three-dimensional space. We wanted to investigate three-dimensional space instead of various sub-spaces. In the ball-in-plane-task, a plane is defined by three balls that are suspended from the ceiling. The observer can adjust the height of the fourth ball. The task is to hang the ball in the plane defined by the other three balls. In the experiments described in this paper we investigated the orientation of the plane with respect to the position of the observer. By varying the orientation of the plane, one automatically varies the observer's view of the plane: the observer can look at the plane from above, from below or the plane might pass through the cyclopean eye of the observer. These different views of the pointer may elicit different settings. If we assume that flat planes are perceived as concave (Foley, 1991; Koenderink et al., 2002), we can hypothesize that the observer will hang the test-ball too high when he views the plane from below and too low when he looks at the plane from above.

In this paper we will discuss the results of three experimental conditions in which we varied the angle at which the plane, defined by the three balls, was tilted (the azimuthal angle). Furthermore, we varied the shape of the triangle that was formed by the three balls. We compared the results for conditions where the balls formed an acute, an obtuse and an equilateral triangle. Thus, the goal of the experiment was to obtain insight into the effect of the azimuthal angle of the plane and to find out whether the shape of the triangle influenced this effect.

6.2 Methods

Observers

Twenty-four undergraduate students, who were paid for their efforts, participated as observers in the experiments described below; there were eight observers per experimental condition. They were naive as to the purposes of the experiments and had little or no experience as observers in psychophysical experiments. Before starting these experiments, they had been observers in another experiment that consisted of an exocentric pointing task. However, they were not given any feedback about the purpose and their own performance in that experiment before the end of the following experiments. Furthermore, they all had normal or corrected-to-normal sight and they were tested for stereo-acuity. Each observer had a stereo-acuity of more than 60".

Set-up

The experiment was conducted in a laboratory room measuring 6 m by 6 m by 3.5 m. The wall opposite the observer was white, with some electrical sockets near the floor. The wall on the left-hand side of the observer contained four blinded windows with radiators underneath them. The wall on the right-hand side of the observer contained two gray doors. The floor was empty except for some markers that were used to position the objects that were used for the former task. A horizontal iron grid was suspended below the ceiling at a height of 3 m above the ground. From this grid, white cubes were hung that contained the balls that were used for the experiments. The balls used were snooker-balls (with a diameter of 6 cm). Three red balls were hung in a triangle (the reference-balls), while another ball (brown or yellow) was hung between the other balls (the test-ball). The height of the balls could be adjusted with a PC with an initial speed of 70 mm/sec. The observer could adjust the height of the test-ball with a remote control. The observer could move the test-ball up and down with two jog-speeds (26 and 62 mm/sec). The observer could press a “ready-button” when he was satisfied with the height of the test-ball. The position of the test-ball was read at that moment and the value was written in an output-file by the PC. The balls were calibrated each morning: balls on the floor were at height 0. The observers sat on a revolving chair that was adjustable in height. All observers were seated at an eye-height of 150 cm.

For each experimental condition, the balls and observers were placed at different positions in the experimental room. The reference balls were hung in different triangles from the grid, as can be seen in Figure 6.1. We performed measurements with an acute, an obtuse and an equilateral triangle (in the azimuthal plane). In Figure 6.1, the circle corresponds to the positions of the observer, and the black dots correspond to the positions of the reference balls in the azimuthal plane. Furthermore, the black lines represent the direction of the steepest upward vector in the plane that was constructed with the reference balls. The center of the lines gives the position of the test-ball in the azimuthal plane. In Table 6.1, the exact coordinates of the positions of the balls and observers are given in cm. Table 6.2 gives the values used for the elevation and the azimuthal angles used for the construction of the planes

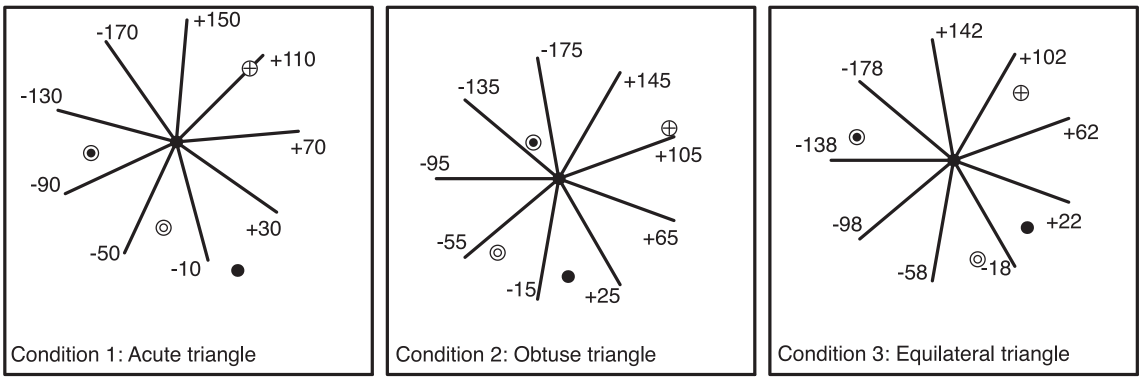

Figure 6.1

Top-views of the room for the three different experimental conditions. The black dot represents the position of the observer. The three differently filled circles give the positions of the reference balls in the azimuthal plane. The point where the lines meet is the azimuthal position of the test ball. The lines represent the directions in which the plane is tilted (the azimuthal angle); the plane is defined by the heights of the reference balls.

Table 6.1 *The azimuthal positions of the balls and observer in cm*

	Ref ball A		Ref ball B		Ref ball C		Test-ball		Observer	
	x	y	x	y	x	y	x	y	x	y
Exp 1	120	120	-20	-120	-140	-20	0	0	100	-210
Exp 2	180	80	-100	-100	-40	60	0	0	15	-160
Exp 3	110	110	40	-160	-160	40	0	0	120	-110

Table 6.2 *Definitions of the planes. The heights are given in mm from the ground. The azimuthal angles of the plane are given as the angle between the line direction of the steepest upward vector of the plane (in the horizontal plane) and the line from test-ball to observer.*

	Elevation	Heights	Azimuthal angles*
Exp 1	25°	1150 / 1325 / 1500	-170° + i 40°
Exp 2	28°	1150 / 1325 / 1500	-175° + i 40°
Exp 3	30°	1500 / 1650 / 1800	-178° + i 40°

* $i \in \{0,8\}$

for each experiment. The elevations were 25°, 28° and 30° respectively, We used nine azimuthal angles, distributed evenly from -180° to +180°. The height of the plane was varied because otherwise the observer would always have had to put the test-ball at the same height. Thus, for each azimuthal angle, we had three different heights. These heights are defined by the veridical heights of the test-ball and are given in Table 6.2. The results obtained for these three different heights were averaged in the analysis. Thus, each experimental condition consisted of 9x3 = 27 trials (# azimuthal angles, # veridical heights).

Procedure

The procedure was the same for each experimental condition. The observers were informed about the remote control and the task at hand. They were told that the task was to put the brown or yellow ball in the plane defined by the red balls. In experimental condition 3, the experiment with the equilateral triangle, we used a brown test-ball, whereas in the other experiments a yellow ball was used. In a pilot experiment we found no differences between settings of observers when differently colored balls were used. The observers were allowed to rotate the chair and their heads, but were not allowed to make upper-body-movements to the side or forwards (backwards was not possible). Before starting the experiment, we conducted a test-series of three trials so that the observer could get used to the remote control and the task he had to do. If there were no questions, we started the real experiment. First the computer placed all balls at a specified height. The reference balls were hung at the heights that defined the first plane. The test-ball was hung at a random height that was between 75 cm below and above the veridical value. A light appeared on the remote control when the observer could move the test-ball. The observer had 50 seconds time to adjust the height of the test-ball. After 40 seconds, the light on the remote control flickered to draw the observer's attention to the fact that he had only 10 seconds left to finish the trial. The observer pressed the "ready-button" when the correct position was reached before the end of the 50 s period. After the trial was terminated (by a button-press or when the end-time was reached) the balls were hung at different heights and everything started afresh. Most observers finished each trial well within 50 seconds. On the rare occasions when this was not the case, the observer was adjusting the height only a few mm up or down. The observers were told that if the 50 second period was over and the position of the test-ball was really wrong, then they should say so. These trials were retested after the entire experiment had finished. This only happened in the case of five trials in all three experiments with all observers.

Analysis

We calculated the shortest distance from the measured position of the test-ball to the plane defined by the reference-balls. In this way, we ensured that the size of the deviations was not dependent on the elevation of the plane.

6.3 Results

Figures 6.2 A, B and C give an impression of the observer's visual field for the different experimental conditions. Each figure gives nine panels for the nine different azimuthal angles used at a height (of the test-ball) of 150 cm. The circles with the cross, circle and disk in the middle represent the positions of the balls A, B and C in the visual field. The gray triangles represent the surface areas in the visual field limited by the reference balls. The surface is dark gray when the observer looks at it from below, and light gray when he looks at it from above. The open circle (sometimes not visible) represents the veridical position of the test-ball, whereas the black disk represents the mean position in which the test-ball was hung by the observers.

The results of the three experimental conditions are shown in Figure 6.3. Each graph represents the data for one condition. The deviation from the plane is plotted against the azimuthal angle. The deviations are given in cm; a positive deviation means that the ball was positioned above the plane, whereas a negative deviation means that the ball was positioned below the plane. The error-bars give the standard errors of the mean. The azimuthal angle is the number of degrees separating the projection of the steepest upward vector (of the plane) on the horizontal plane from the orientation of the line between observer and test-ball. A positive azimuthal angle of 30° means that the steepest upward vector is oriented 30° rotated counterclockwise from the observer. Thus, the observer is looking from below at the plane when the azimuthal angle is between –90° and 90°; otherwise he is looking at the plane from above. What can be seen in all graphs is that there is a downward trend when the azimuthal angle diverges from 0°. Most deviations were negative which indicates that the observers usually place the ball too low. For the acute and equilateral triangles the trials with azimuthal angles close to 0 have positive deviations. Thus, the ball is placed closer to veridical, or sometimes even above the plane, when the observer looks at the plane from below. However, the downward trend is also present for these trials.

Figure 6.4 shows the standard deviations plotted against the azimuthal angles. Each azimuthal angle was measured three times for each observer. The standard deviations were calculated for these trials for each observer. The points represent the means of the standard deviations for all observers. The error-bars give the standard error of the mean of these values. The three graphs represent the data for the three experimental conditions. However, we did not find that the azimuthal angle had an overall effect on the standard deviations.

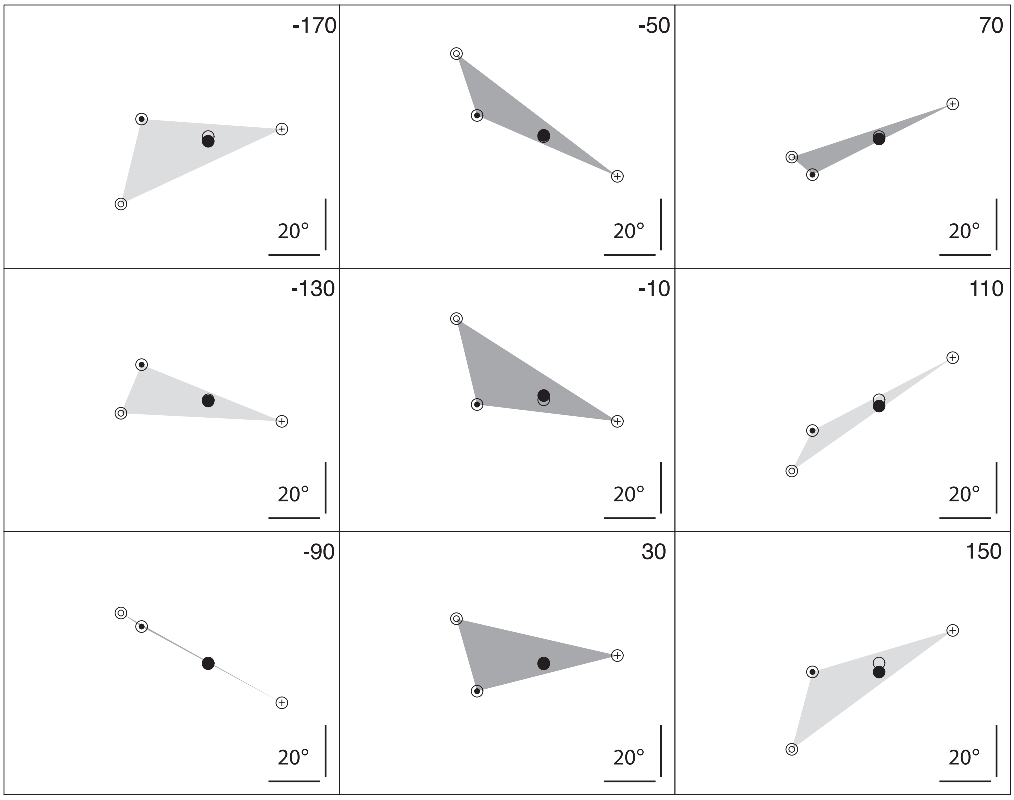

Figure 6.2 A: *The visual field in the acute triangle condition*

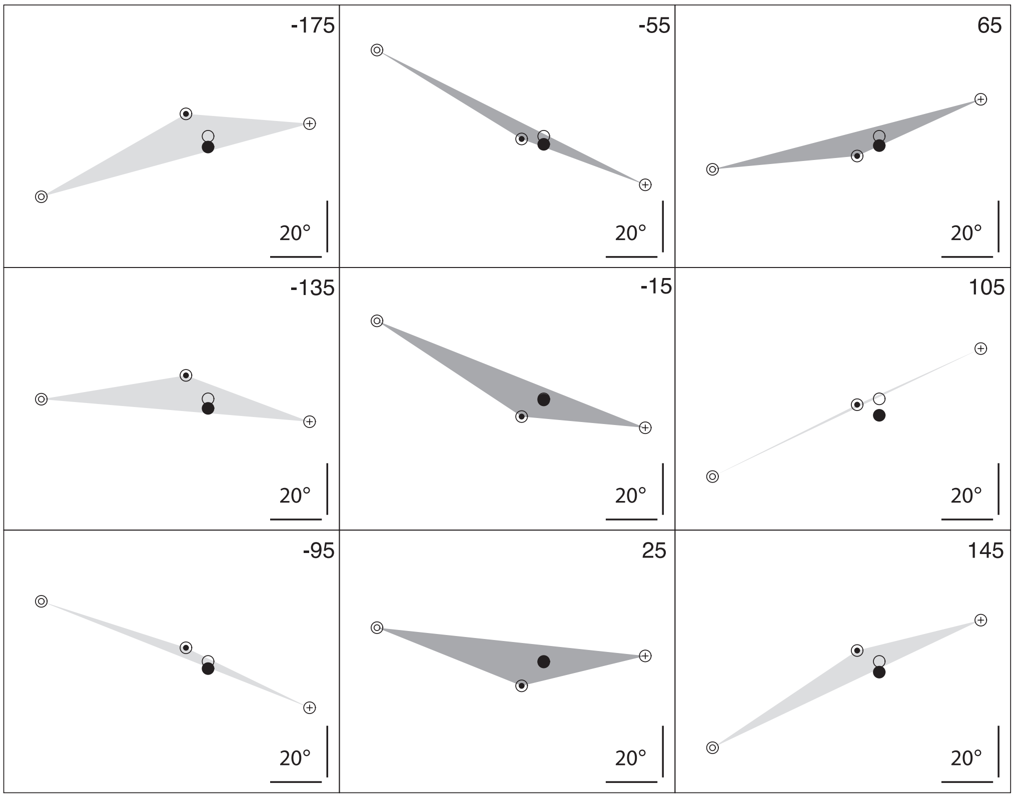

Figure 6.2 B: *The visual field in the obtuse triangle condition*

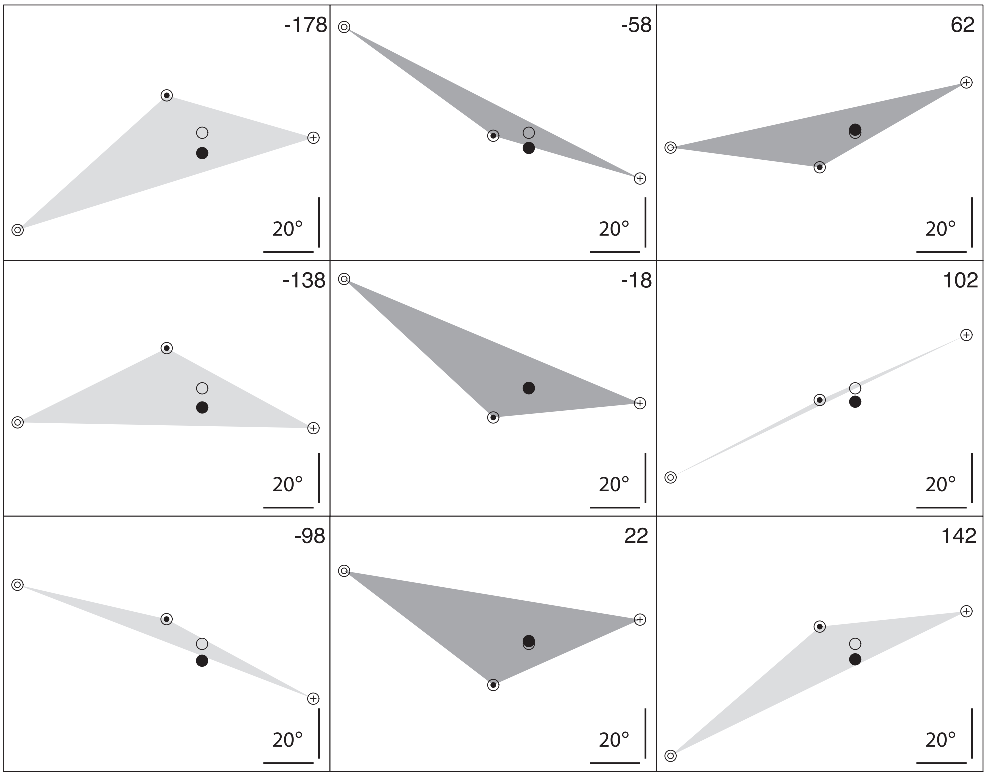

Figure 6.2 C: *The visual field in the equilateral triangle condition*
Figures A, B and C show nine panels with the observer's visual fields for the different azimuthal angles at a plane height of 150 cm. The circles containing a cross, a disk or a circle represent the positions of the reference-balls A, B and C in the visual field, respectively. The gray triangle represents the triangle in the visual field formed by the reference-balls. The triangle is dark gray when the observer sees the plane from below, and light gray when seen from above. The circle represents the veridical position of the test-ball, the disk the mean value of the settings of all observers.

Besides plotting the deviations against the azimuthal angle, we looked at other variables that varied with the azimuthal angle. Among these are the vertical visual angle, the total visual angle and the surface area of the triangle. Looking at Figure 6.2, one can see that the surface area in the visual field is not correlated to the size of the deviations. We plotted the deviations of all conditions against the three-dimensional surface area, the vertical visual angle and the total visual angle and we fitted a line through these points. We found no dependence on the three-dimensional surface area ($p = 0.318$), a trend for the vertical visual angle ($p = 0.055$) and a significant effect of the total visual angle ($p = 0.005$). However, the R^2 values were too low to be regarded as a good fit (0.001, 0.105, and 0.242 respectively).

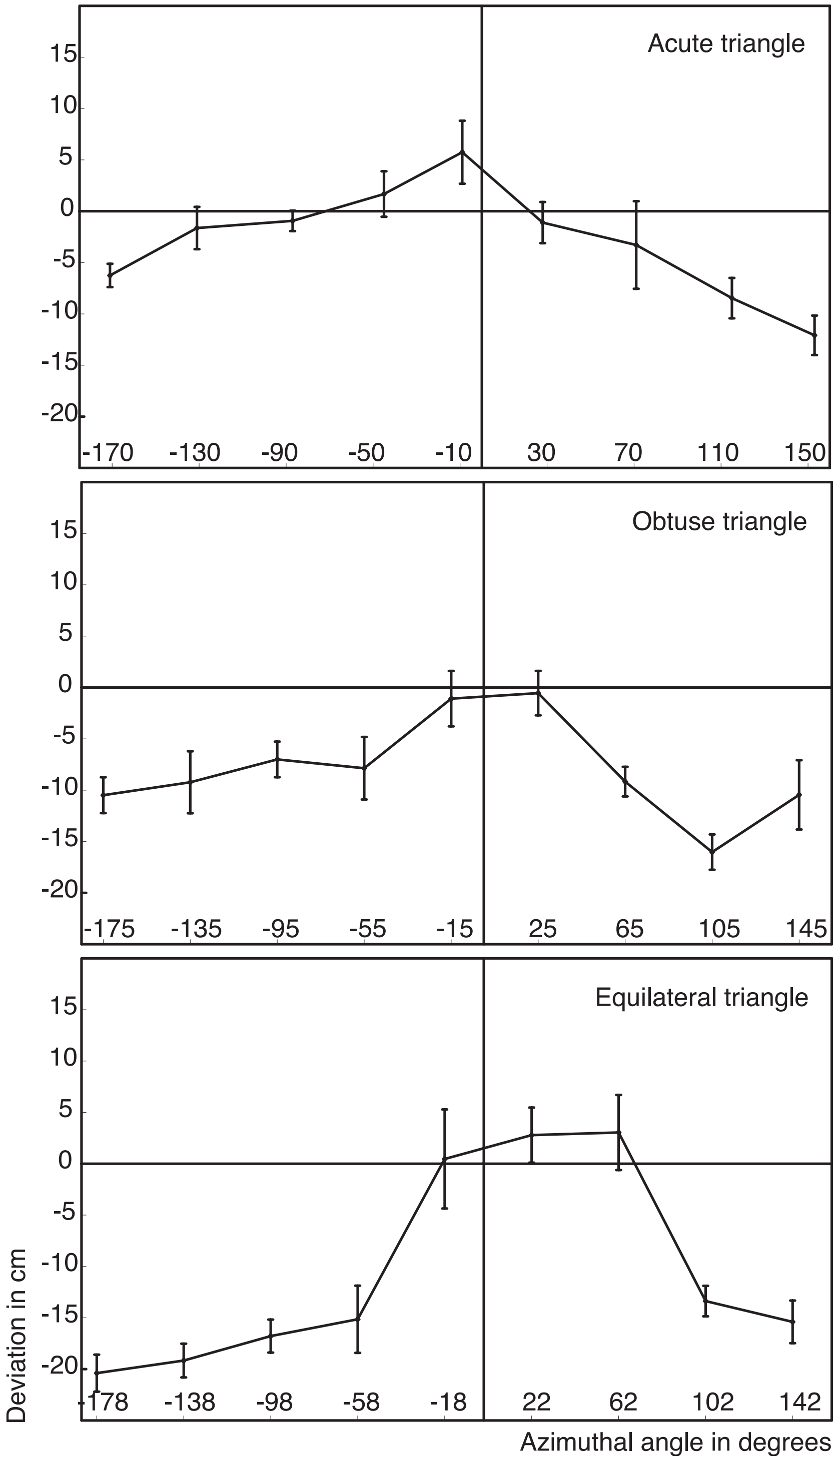

Figure 6.3
The three graphs give data for the three experimental conditions: for the acute, the obtuse and the equilateral triangle respectively. In each graph, the deviation from veridical settings (given in cm) is plotted against the azimuthal angle. A positive deviation means that the observer hung the test-ball above the veridical position, a negative deviation means that he hung the test-ball lower than the veridical position. The data-points give the means of all observers, the error-bars the standard error of the means.

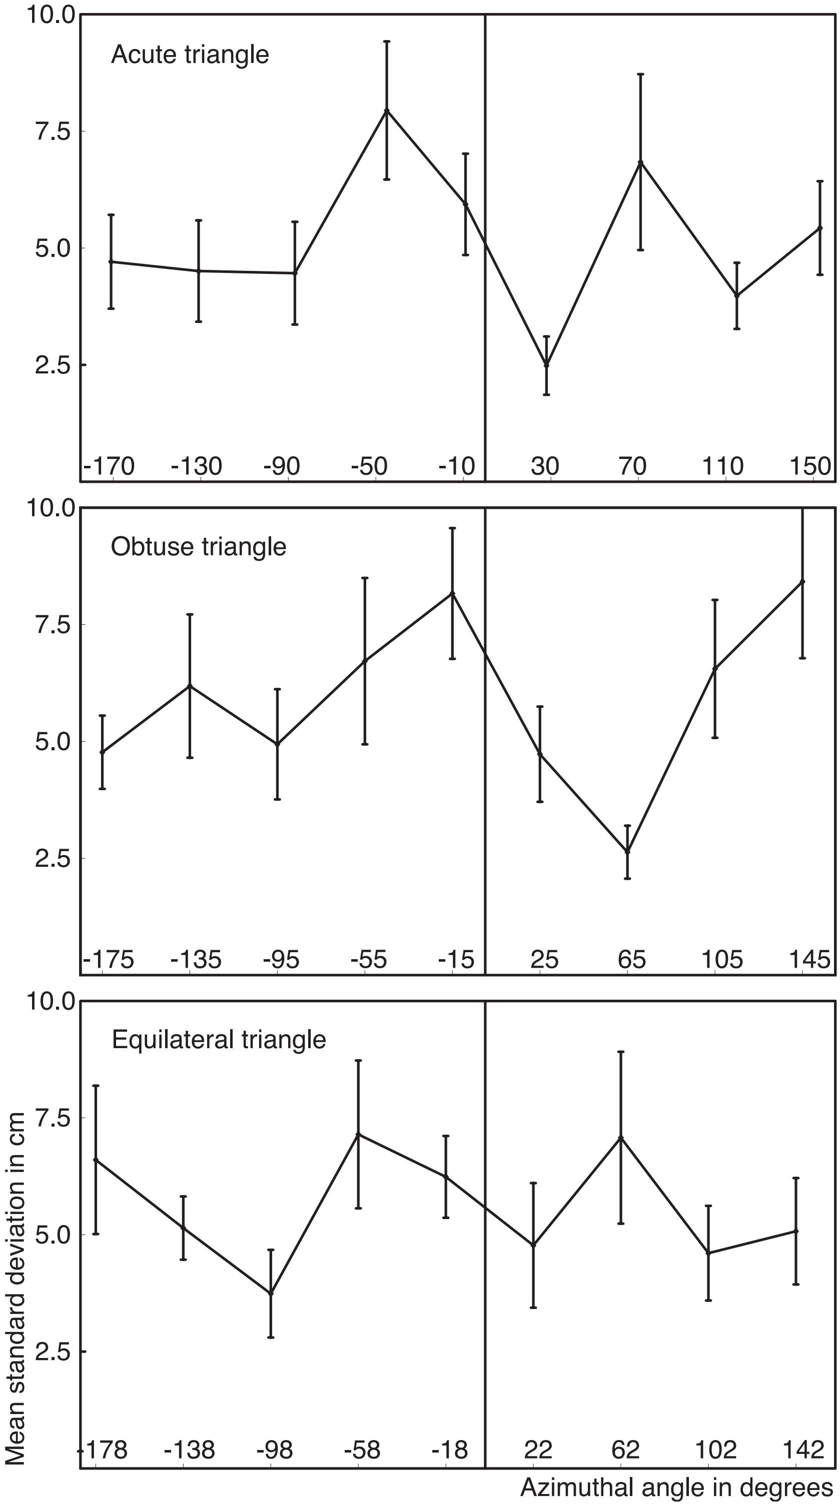

Figure 6.4

The three graphs give data for the three experimental conditions: for the acute, the obtuse and the equilateral triangle respectively. In each graph, the intra-observer standard deviations (given in cm) are plotted against the azimuthal angle. Each data-point gives the mean of the standard deviations of all observers, the error-bars give the (inter-observer) standard error of the means.

6.4 Discussion and conclusions

In this chapter we looked at the effect that a few parameters had on the deviations we found with the ball-in-plane-task. Parameters involving the surface area of the triangle did not affect our data and from these data we cannot conclude that the sizes of the deviations are dependent on the visual angle.

The most important result, however, is that the deviations decrease (become (more) negative) when the azimuthal angle diverges from 0°. Usually, the observers hang the ball too low when they look at the plane from above. Thus from the data we can conclude that at this scale flat planes are seen as concave when the observers look at the plane from above. This is in agreement with the results obtained with an exocentric pointing task (when pointing frontoparallel) and an apparent frontoparallelity task (Foley, 1991; Koenderink, et al., 2000, 2002). However, the sign of the deviations is less clear when the observers look at the plane from below. Small negative and positive deviations were found in these trials with a downward trend when the azimuthal angle diverges from 0°. We did not expect to find small deviations in these trials since most observers said they found these trials more difficult. In the trials in which the observers looked at the plane from below or above the background behind the balls differed. When the observer looked at the plane from above, he could see the walls and floor behind the balls. However, when he looked at the plane from below, he could see the ceiling behind the balls as well as the walls. Since the ceiling was covered with an iron grid, the visual field in these trials is much more crowded. This is probably why observers thought that the trials in which they had to look at the plane from below were more difficult.

From these data we can conclude that the best description we can give for the perception of flat planes is that they are perceived as concave at distances of up to 5 m from the observer. Although most deviations are close to veridical when the observers look at the plane from below, the description fits very well when the observers look at the plane from above.

Summary

In this thesis we report a series of experiments designed to discover factors that are responsible for the deformation of visual space. We looked at spatial and contextual parameters.

Spatial parameters

In chapter 2 we described three experiments in which we investigated whether the deformation of visual space was dependent on the visual angle of the two objects that were used as stimuli and on the relative distance of these two objects from the observer. We used three tasks. In the exocentric pointing task, the observer had to direct a pointer with a remote control towards a ball. In the second task, the collinearity task, the observer had to rotate two rods in such a way that they were in one line. The third task was the parallelity task; in this task the observer had to place a rod parallel to another rod. For the exocentric pointing task and the collinearity task we found that the visual angle had a linear effect and the relative distance had a non-linear effect on the deviations. For the parallelity task, however, we found no effect of relative distance. The dependence on the visual angle was linear just as in the other two experiments. Besides these two variables, we looked at the orientation of the reference rod in the parallelity task. For two out of four observers we found an effect of reference orientation, which led us to conclude that the observers probably differ in their use of contextual information.

In chapter 3 we extrapolated the results of a 2D exocentric pointing task to three dimensions. In the 3D exocentric pointing task the observers could rotate a pointer in the horizontal plane (slant) and in the vertical plane (tilt). This meant that we could place the pointer and the ball at various heights. In this experiment we varied the horizontal visual angle, the vertical visual angle and the relative distance. If Luneburg's conjecture is correct, visual space should be isotropic. This would mean that the deviation of the slant would depend on the horizontal visual angle in the same way as the tilt depends on the vertical visual angle. Furthermore, both dependent variables should depend in a similar way on the relative distance. This is not what we found in the experiments described in chapter 3. The inevitable conclusion is that visual space is anisotropic.

Contextual parameters

In addition to studying the effects of spatial parameters, we investigated the effects of contextual parameters. In previous experiments, we found small deviations when both the ball and the pointer were at the same distance from the observer. From these experiments, however, we cannot be certain whether this is due to the fact that the pointing-direction was parallel to one of the walls of the experimental room or to the fact that the pointing-direction was frontoparallel. Thus, in the experiments described in chapter 3, we varied the positions of pointer, ball and observer in such a way that we could discriminate between the effects of frontoparallelity, parallelity to a wall and possible interactions between these effects. We found differences between observers in the way their results were dependent on an egocentric factor like frontoparallelity or an allocentric factor like the walls of the experimental room.

In a 2D exocentric pointing task we investigated two other possible contextual effects on the deformation of visual space. We performed these experiments because in all the experiments with an exocentric pointing task performed so far we found an effect of relative distance. We wanted to find out why the deviations were larger when the pointer was closer to the observer than the ball, than when the pointer was further away from the observer than the ball. We examined two possible explanations for this observation: one concerning the possible effect of restricting the pointing angle, the other concerning the shape of the pointer. When the pointer is further away from the observer than the ball, the position of the observer restricts the pointing angle. In contrast, when the pointer is closer to the observer than the ball, there is no extra reference such as one's body position. This difference could explain the difference in the size of the deviations we found for the exocentric pointing task. We tested this by using poster-boards. We placed these boards between the pointer-positions close to the observer and a position on the other side of the room. In this way we restricted the pointing-direction by the same angle as the observer' body position restricted the pointing angle for the condition in which the pointer is further away from the observer than the ball. We found that the placing of the poster-boards did result in smaller deviations for the condition in which the pointer is closer to the observer than the ball, although the size of this deviation is not as small as in the other condition.

Our second possible explanation for the effect of relative distance on the size of the deviations was that the view of the pointer was different for the two conditions. When the pointer is further away from the observer than the ball and the pointer is rotated slightly, the observer's retinal image changes more than when the pointer is rotated when it is closer to the observer than the ball. Thus the two conditions differ in the amount of information that is available about the exact orientation of the pointer. Whether observers are influenced by this difference in the amount of information can be tested easily by using two pointers that differ in shape. This is exactly what we did. We used two extra pointers (in addition to the pointer we used in previous experiments): one with a single rod and another with two rods perpendicular to each other. This double-rod pointer contained the same amount of information in all experimental conditions, in contrast to the other two pointers. Although we found a difference in the standard deviations when we used the single-rod pointer and the double-rod pointer, we did not find an effect for the size of the deviations themselves. Thus, the only explanation we can find for the difference in the deviations we found by varying the relative distance is that the pointing angle is restricted by the position of the observer.

Ball-in-plane task

The last chapter of this thesis is about the ball-in-plane task. This task enables us to gain more insight into all three dimensions of visual space. Chapter 6 describes an experiment to test whether we can visualize planes tilted in visual space. The plane is defined by three red balls that are hung at different heights in the experimental room. The observer can adjust the height of a fourth ball that is suspended from the ceiling somewhere among the other balls. The task is to hang the fourth ball in the plane defined by the red balls. In particular, we investigated whether the deviations are dependent on the direction in which the plane is tilted. We tested this with three different configurations of red balls forming an acute, an obtuse and an equilateral triangle. For the three triangles we found small negative or positive deviations when the azimuthal angle was close to 0. When the azimuthal angle increases in size, the deviations are increasingly negative. This pattern can best be described as concave settings towards the observer. This has also been found for bisection tasks in the horizontal plane.

Conclusions and further directions

Looking at spatial parameters, we found similar patterns for the observers. For the ball-in-plane-task, we found a pattern with concave settings towards the observer. For the exocentric pointing task however, we found the deviations to be dependent on the relative distance in an unexpected way that cannot be easily described with a metric function. Possible descriptions are that positions of objects are overestimated proportional to the distance from an object or that the deviations are dependent on the visual angle. An effect of visual angle could be explained by the uneven distributions of photoreceptors in the retinae and the organization of the visual cortex. However, these ideas need to be elaborated further. Furthermore, we found that visual space is anisotropic, which is in contrast with Luneburg's conjecture. We can therefore conclude that visual space is not homogeneous in different sub-spaces.

For the experiments concerning contextual parameters, the results were less clear. Large differences occurred between observers for cues that were not crucial for doing the task, namely ones such as the presence of the walls and the frontoparallel plane of the observer. It seems that each observer has his own preferred information sources that he uses to solve a task. However, some sources of information are so prominent that observers cannot ignore them. An example is the restriction of pointing angles that is discussed in chapter 5. Furthermore, varying the shape of the pointer did not result in any differences between the settings. Thus, there was no difference in the amount of information provided by the pointers that could explain the forward-backward asymmetry. This is an interesting field of research that needs more attention than it has received so far. It fits into current visual space research since attention is gradually shifting from really sober environments to more elaborate ones. Thus a straightforward route for future research to follow is to investigate the contributions made by various information sources.

Nederlandse samenvatting

Mensen hebben het gevoel alsof ze heel goed afstanden tot objecten in kunnen schatten. Dit valt echter, als je het onderzoekt, heel erg tegen.

Er zijn meerdere bronnen waar mensen informatie uit kunnen halen om diepte te kunnen zien. Men weet echter nog niet goed hoe mensen precies bepalen welke bronnen ze wel of niet gebruiken om de afstand tot een object in te schatten. Dit soort onderzoek wordt veel gedaan onder zeer gecontroleerde omstandigheden: vaak zodanig dat er van de vele informatie-bronnen er maar eentje aanwezig is. Daarbij wordt er gekeken of de proefpersoon de afstanden goed inschat. Wat daarbij vergeten wordt, is dat mensen nu eenmaal al die informatie tot hun beschikking hebben en dat ze waarschijnlijk van meerdere bronnen gebruik maken. Dus is er ook onderzoek nodig waarbij meerdere bronnen aangeboden worden om te kijken hoe mensen het dan doen. Echter, bij dit type onderzoek worden grote verschillen gevonden. Deze zijn aan meerdere factoren toe te schrijven: de taak die is aangeboden, de afstanden die gebruikt zijn en de verdere condities waaronder gemeten wordt. Met onze experimenten hebben we geprobeerd een bijdrage te leveren aan de kennis die bestaat over de visuele ruimte. Eerst hebben we gekeken wat er gebeurt als we mensen met een pijl naar een bal laten wijzen als ze in een normaal verlichte kamer zitten waarvan de muren, plafond en vloer zichtbaar zijn. Daarna hebben we de taak van een twee-dimensionale naar een drie-dimensionale taak uitgebreid. Vervolgens wilden we weten of vergelijkbare resultaten zichtbaar werden als we een andere taak gebruikten.

Ruimtelijke parameters

In hoofstuk 2 wordt het onderzoek beschreven dat gedaan is met drie verschillende taken: een exocentrische aanwijstaak, een parallel-zet-taak en een collineair-zet-taak. Bij de exocentrische aanwijstaak moet de proefpersoon met een pijl naar een bal wijzen. Bij de parallel-zet-taak moet de proefpersoon een staaf parallel zetten aan een andere staaf. En bij de collineair-zet-taak moet de proefpersoon twee staven in elkaars verlengde zetten door beide staven te draaien. Al deze pijlen, ballen en staven waren altijd in het horizontale vlak op ooghoogte. Uit de data bleek dat de instellingen van de proefpersoon bij de aanwijstaak en de collineair-zet-taak afhangen van de afstand tussen de twee objecten (de horizontale visuele hoek) en de relatieve afstand tussen de twee objecten en de proefpersoon. Bij de parallel-zet-taak, daarentegen, zijn de instellingen niet afhankelijk van de relatieve afstand van de objecten tot de proefpersoon. Wel zijn de instellingen bij deze taak afhankelijk van de visuele hoek tussen de objecten en de hoek waarin de referentie-staaf wordt gezet. Het patroon van fouten dat mensen maken bij een dergelijke opstelling is dus afhankelijk van de taak die aan hen wordt opgelegd. Deze resultaten komen in grote mate overeen met de resultaten van Cuijpers en collega's (2000A, 2000B, 2002). Zij hebben dezelfde taken uitgevoerd terwijl proefpersonen niet de muren, vloer en plafond van de experimentele ruimte konden zien. Het lijkt er dus op dat de proefpersonen in ons experiment geen effectief gebruik maken van de extra structuur die beschikbaar was in het visuele veld.
De exocentrische aanwijstaak die beschreven is in hoofdstuk 2 betreft alleen het horizontale vlak op ooghoogte. In hoofstuk 3 worden echter de resultaten van een drie dimensionale exocentrische aanwijstaak beschreven. De proefpersonen konden een pijl draaien in het horizontale vlak en in het vertikale vlak. De pijl en de bal konden ook op verschillende hoogtes geplaatst worden. Met dit onderzoek wilden we verifiëren of de eerder gevonden afwijkingen van een twee dimensionale taak hier ook gevonden konden worden en wilden we bekijken wat de nieuwe afhankelijke variabele (de afwijking in het vertikale vlak, m.a.w. de tilt) voor patroon laat zien. Allereerst lijken de afwijkingen in het horizontale vlak (de slant) op het patroon dat beschreven wordt in hoofdstuk 2. Daarnaast vinden we een heel ander patroon voor de afwijkingen van de tilt. De tilt is niet afhankelijk van de horizontale visuele hoek of de relatieve afstand. Daarnaast is de tilt wel afhankelijk van de verticale visuele hoek. Het feit dat de tilt niet afhankelijk is van de relatieve afstand geeft aan dat er sprake is van een anisotropie van de visuele ruimte. Dit wil zeggen dat de ruimte niet in alle richtingen hetzelfde vervormd is.

Contextuele parameters

In hoofdstuk 4 wordt een ander onderzoek besproken met dezelfde 3D aanwijstaak. De nadruk ligt bij dit hoofdstuk niet zozeer op de invloed van ruimtelijke parameters op de afwijkingen van de slant en de tilt, maar juist op de invloed van contextuele parameters. Je kan een onderscheid maken tussen referentiekaders die vanuit de persoon zelf komen, zoals de waarneming van afstanden van objecten tot jezelf (egocentrisch referentiekader), en referentiekaders die door de omgeving worden gevormd (allocentrisch referentiekader), als je in een kamer staat is dat dus de kamer om je heen. In voorgaand onderzoek vonden we dat als objecten op dezelfde afstand van de proefpersoon stonden, de afwijkingen zeer klein waren.

Maar met de opstelling zoals we gebruikt hebben in die experimenten is niet te zeggen of dit komt doordat de afstanden tot de proefpersonen hetzelfde zijn (een egocentrische maat) of dat het komt doordat de juiste wijsrichting ook parallel aan een van de muren van de kamer is (een allocentrische maat). Dit hebben we onderzocht in de experimenten die in hoofdstuk 4 besproken worden. Het blijkt dat er grote verschillen bestaan tussen proefpersonen in het gebruik van dit soort informatie. De ene persoon vertrouwt meer op de allocentrische informatie, terwijl een ander juist let op de egocentrische informatie. Er zijn ook mensen die van beide factoren afhankelijk zijn. Deze effecten werden gevonden voor de afwijkingen van de slant, maar minder in de afwijkingen van de tilt. Dit is misschien te verklaren doordat de tilt de richting in het vertikale vlak is en de referenties die gevarieerd zijn vooral betrekking hebben op het horizontale vlak. Dus om hier echt uitsluitsel over te geven zouden we andere structuren van de experimentele ruimte moeten variëren.

In de voorgaande onderzoeken werd steeds gevonden dat de afwijkingen groter zijn als de pijl dichter bij de proefpersoon staat dan de bal dan als de pijl ver weg staat en de bal dichtbij. Twee verklaringen werden getest met de experimenten die in Hoofdstuk 5 beschreven staan. Ten eerste zou het kunnen zijn dat proefpersonen hun eigen lichaam als referentiepunt kunnen gebruiken voor de instellingen als de pijl verder van hen vandaan staat dan de bal. Hierdoor wordt de hoek waarin gewezen kan worden namelijk in grote mate beperkt. Ten tweede is de zichtbaarheid van de pijl anders voor de twee condities. Als de pijl een stukje gedraaid wordt als hij verder van de proefpersoon af staat (en wijst in de richting van een bal die dichtbij hangt), verandert het beeld op het netvlies bij een kleine rotatie van de pijl meer dan wanneer de pijl dicht bij de proefpersoon staat en de bal ver weg. Om de eerste verklaring te testen hebben we een experiment gedaan waarin we de wijshoek in dezelfde mate beperkten voor een ver weg staande pijl als voor een dichtbij staande pijl. Dit hebben we gedaan door poster-borden te plaatsen die dezelfde hoek afbakenden. Dit heeft zeker invloed op de wijsrichting, maar de afwijkingen voor de conditie waarbij de pijl dichtbij staat en de bal ver weg, waren nog steeds groter dan we vonden voor de conditie waarbij de pijl ver weg stond en de bal dichtbij. Dus zou het zo kunnen zijn dat het plaatsen van de poster-borden niet zo'n sterke invloed heeft op de proefpersoon als zijn eigen lichaams-positie, of er zijn meerdere factoren die voor dit effect zorgen.

De tweede verklaring die we hadden betreft de vorm van de pijl. Om te testen of dit invloed heeft, hebben we twee nieuwe pijlen gemaakt: een pijl die bestaat uit een enkele staaf en een pijl met twee staven die loodrecht op elkaar staan. De enkele staaf is als mogelijke informatie bron zeer ongelijk voor de twee condities terwijl de dubbele staaf in beide condities evenveel mogelijke informatie geeft. Er werd een klein verschil gevonden in grootte van de afwijking tussen de enkele en dubbele pijl. De pijl die in voorgaande experimenten gebruikt werd, geeft echter geen verschillende afwijkingen met de enkele of de dubbele pijl. Het lijkt er dus op dat de vorm van de pijl in voorgaande experimenten geen invloed heeft gehad op het verschil in afwijkingen die we vonden voor de hierboven beschreven condities.

Bal-in-vlak-taak

In hoofdstuk 6 wordt een nieuwe taak geïntroduceerd, de bal-in-vlak-taak. In deze taak moet de proefpersoon een bal in een vlak hangen dat gedefinieerd wordt door drie andere ballen. In dit onderzoek varieerden we de richting waarin het vlak gekanteld was. Dit deden we voor drie verschillende driehoeken: een scherpe, een platte en een gelijkzijdige driehoek. Het belangrijkste resultaat van dit onderzoek is het feit dat de proefpersonen de bal te laag hingen als ze van boven tegen het vlak aankeken. Als de proefpersonen van onder tegen het vlak aan keken, hingen ze de bal dicht bij de veridicale hoogte. Met name als de steilste vector van het vlak in de richting van de proefpersoon wees, waren de afwijkingen positief of dicht bij 0. Het lijkt erop dat een plat vlak als hol wordt waargenomen door mensen.

Algemene conclusies

Voor de ruimtelijke parameters geldt hetzelfde voor alle proefpersonen. Dit is echter wel afhankelijk van de taak die aangeboden wordt. Zo zijn de afwijkingen in een exocentrische aanwijstaak en een collineair-zet-taak afhankelijk van de relatieve afstanden van de twee objecten tot de proefpersonen en de afstand tussen deze twee objecten (de visuele hoek). Dit terwijl de afwijkingen van een parallel-zet-taak niet afhankelijk zijn van de relatieve afstand, maar wel van de visuele hoek. De vervorming van de visuele ruimte is dus afhankelijk van de karakteristieken van de taak die gebruikt wordt voor het experiment. Daarnaast vonden we dat de vervorming van de ruimte ook afhankelijk is van het vlak waarin men meet: in het horizontale vlak is de vervorming anders dan in het verticale vlak. Hiermee hebben we aangetoond dat de visuele ruimte niet homogeen is en dus geen constante kromming kan hebben zoals Luneburg aannam. Aan de hand van de bal-in-vlak-taak hebben we hetzelfde kunnen concluderen als aan de hand van de voorgaande taken, nl dat de ruimte uitgestrekt is of hol ten opzichte van de proefpersoon.
Voor contextuele parameters is dit echter een heel ander verhaal, er zijn grote verschillen te vinden tussen proefpersonen in de effecten die het variëren van de context met zich meebrengt. Zo maken sommige mensen kleinere afwijkingen als de wijsrichting van een exocentrische aanwijstaak parallel is aan een muur terwijl andere mensen juist kleinere afwijkingen hebben als de wijsrichting frontoparallel is. Hier staat echter tegenover dat als de contextuele informatie zeer prominent aanwezig is, zoals het inperken van de wijshoek door posterborden, de proefpersonen die informatie niet kunnen negeren en dus een vergelijkbaar patroon van afwijkingen laten zien. Dit is een onderzoeksgebied dat naar mijn idee meer aandacht verdient dan het tot nu toe gekregen heeft. De aandacht is wel die kant op aan het verschuiven: van donkere kamers tot omgevingen waar een heleboel potentiële bronnen van informatie aanwezig zijn. Nu wordt het dus zaak om precies te onderzoeken welke structuren in de omgeving ook daadwerkelijk als bron van informatie kunnen dienen.

Reference list

Ames, A. (1953). Reconsideration of the origin and nature of perception in situations involving only inorganic phenomena. In Ratner, S. (Ed.), *Vision and Action (p 251-274).* New Brunswick: Rutgers University Press.

Appelle, S. (1972). Perception and discrimination as a function of stimulus orientation: the "oblique effect" in human and animals. *Psychological Bulletin, 78,* 266-278.

Battro, A. M., di Pierro Netto, S., & Rozestraten, R. J. A. (1976). Riemannian geometries of variable curvature in visual space: visual alleys, horopters, and triangles in big open fields. *Perception, 5,* 9-23.

Berkeley, G. (1732). An essay towards a new theory of vision (4th ed.). http://psychclassics.yorku.ca/Berkeley/vision.htm.

Blank, A. A. (1961). Curvature of binocular visual space. An experiment. *Journal of the Optical Society of America, 51,* 335-339.

Blumenfeld, W. (1913). Untersuchungen über die scheinbare Grösse im Sehraume [studies on the apparent size in visual space]. *Zeitschrift für Psychologie, 65*, 241-404.

Chen, S., & Levi, D.M. (1996). Meridional anisotropy in the discrimination of parallel and perpendicular lines – effect of body tilt. *Perception, 25,* 633-649.

Cline, M.G. (1967). The perception of where a person is looking. *American Journal of Psychology, 80*, 41-50.

Creem, S.H., Proffitt, D.R. (2001). Defining the cortical visual systems: "what", "where", and "how". *Acta Psychologica, 107,* 43-68.

Cuijpers, R. H., Kappers, A. M. L., & Koenderink, J. J. (2000a). Investigation of visual space using an exocentric pointing task. *Perception & Psychophysics, 62*, 1556-1571

Cuijpers, R. H., Kappers, A. M. L., & Koenderink, J. J. (2000b). Large systematic deviations in visual parallelism. *Perception, 29,* 1467-1482

Cuijpers, R. H., Kappers, A. M. L., & Koenderink, J. J. (2001). On the role of external reference frames on visual judgements of parallelity. *Acta Psychologica, 108,* 283-302

Cuijpers, R. H., Kappers, A. M. L., & Koenderink, J. J. (2002). Visual perception of collinearity. *Perception & Psychophysics, 64*, 392-404

Cutting, J.E., & Vishton, P.M. (1995). Perceiving layout and knowing distances: the integration, relative potency, and contextual use o different information about depth. In Epstein, W, & Rogers, S.J. (Eds.). *Perception of space and motion* (pp. 69-117). San Diego: Academic Press, Inc.

Darling, W.G., & Bartelt, R. (2005). The visual perception coordinate system uses axes defined by the earth, trunk, and vision. *Perception, 34,* 17-30.

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (2005). Visual space under free viewing conditions. *Perception & Psychophysics, 67*, 1177-1189.

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (in press A). Horizontal-vertical anisotropy in visual space. *Acta Psychologica.*

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (in press B). Effects of context on a visual 3D pointing task. *Perception.*

Foley,J.M. (1972). The size-distance relation and intrinsic geometry of visual space: implications for processing. *Vision Research, 12*, 323-332.

Foley, J.M. (1991). Binocular space perception. In D. Regan (Ed.), *Vision and visual dysfunction: Binocular vision* (Vol. 9, pp. 75-92). London: MacMillan Press.

Foley, J.M., Ribeiro-Filho, N.P., & Da Silva, J.A. (2004). Visual perception of extent and the geometry of visual space. *Vision Research, 44,* 147-156.

Gehringen, W.L., & Engel, E. (1986). Effect of ecological viewing conditions on the Ames' distorted room illusion. *Journal of Experimental Psychology, 12,* 181-185.

Gibson, J.J. (1950). *The perception of the visual world.* The Riverside Press, Cambridge.

Gibson, J.J. (1966). *The senses considered as perceptual systems*. Boston: Houghton Mifflin Company.

Gilinsky. A. S. (1951). Perceived size and distance in visual space. *Psychological review, 58,* 460-482

Gillam, B. (1995). The perception of spatial layout from static optical information. In Epstein, W, & Rogers, S.J. (Eds.). *Perception of space and motion* (pp. 23-67). San Diego: Academic Press, Inc.

Gogel, W.C. (1965). Equidistance tendency and its consequences. *Psychological Bulletin, 64,* 153-163.

Goodale, M.A., Milner, A.D., Jakobson, L.S., & Carey, D.P. (1991). A neurological dissociation between perceiving objects and grasping them. *Nature, 349,* 154-156.

Helmholtz, H. L. F. (1962). *Helmholtz's treatise on physiological optics* (J. P. C. Southall, Trans.). New York: New Dover ed. (Original work published 1866).

Hillebrand, F. (1902). Theorie der scheinbaren Grösse bei binocularem Sehen [Theory of apparent size in binocular vision]. *Denkschriften der Wiener Akademie, Mathematisch-Naturwissenschaft Klasse, 72,* 255-307.

Indow, T., & Watanabe, T. (1984). Parallel-alleys and distance-alleys on horopter plane in the dark. *Perception, 13,* 165-182.

Indow, T., & Watanabe, T. (1988). Alleys on an extensive frontoparallel plane: a second experiment. *Perception, 17,* 647-666.

Indow, T. (1991). A critical review of Luneburg's model with regard to global structure of visual space. *Psychological review, 98,* 430-453

Ittelson, W.H. (1952). *The Ames demonstrations in perception*. Boston: Houghton Mifflin.

Kelly, J.W., Beall, A.C., & Loomis, J.M. (2004). Perception of shared visual space: establishing common ground in real and virtual environments. *Presence, 13,* 442-450.

Kelly, J.W., Loomis, J.M., & Beall, A.C. (2004). Judgments of exocentric direction in large-scale space. *Perception, 33,* 443-454.

Koenderink, J.J., van Doorn, A.J. (1998). Exocentric pointing. In Harris, L.R. , & Jenkin, M. (Eds.), *Vision and Action* (pp. 295-313). Cambridge: Cambridge University Press.

Koenderink, J. J., Van Doorn, A. J., & Lappin, J. S. (2000). Direct measurement of the curvature of visual space. *Perception, 29,* 69-79.

Koenderink, J. J., Van Doorn, A. J., Kappers, A. M. L., & Lappin, J. S. (2002). Large-scale visual frontoparallels under full-cue conditions. *Perception, 31,* 1467-1475.

Koenderink, J. J., Van Doorn, A. J., & Lappin, J. S. (2003). Exocentric pointing to opposite targets. *Acta psychologica, 112,* 71-87.

Levin, C.A., & Haber, R.N. (1993). Visual angle as a determinant of perceived interobject distance. *Perception & Psychophysics, 54,* 250-259.

Luneburg, R. K. (1950). The metric of binocular visual space. *Journal of the Optical Society of America, 40*, 627-640.

Luneburg, R.K. (1947). *Mathematical analysis of binocular vision*. Princeton, NJ: Princeton University Press.

Neggers, S.F.W, Van der Lubbe, R.H.J., Ramsey, N.F., & Postma, A. (in press). Interactions between ego- and allocentric neuronal representations of space. *NeuroImage*.

Owens, D.A., & Leibowitz, H.W. (1976). Oculomotor adjustments in darkness and the specific distance tendency. *Perception & Psychophysics, 20*, 2-9.

Paillard, J. (1991). Motor and representational framing of space. In Paillard, J. (Ed.) *Brain and Space*. Oxford: Oxford University Press.

Pick, H.L. jr., Lockman, J.J. (1981). From frames of reference to spatial representations. In Liben, L.S. (Ed.), *Spatial Representation and behavior across the life span (p 39-61)*. New York: Academic Press.

Schoumans, N., & Denier van der Gon, J.J., (1999). Exocentric pointing in three-dimensional space. *Perception, 28,* 1265-1280.

Schoumans, N., Kappers, A.M.L., & Koenderink, J.J. (2002). Scale invariance in near space: pointing under influence of context. *Acta Psychologica, 110*, 63-81.

Sinai, M.J., Ooi, T.L., & He, Z.J. (1998). Terrain influences the accurate judgment of distance. *Nature, 395,* 497-500.

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

Yang, Z., & Purves, D. (2003). A statistical explanation of visual space. *Nature Neuroscience, 6,* 632-640.

Zajaczkowska, A. (1956). Experimental determination of Luneburg's constants σ and K. *Quarterly Journal of Experimental Psychology, 8,* 66-78.

Publication list

Publications

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (2005). Visual space under free viewing conditions. *Perception & Psychophysics, 67*, 1177-1189.

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (in press A). Horizontal-vertical anisotropy in visual space. *Acta Psychologica*.

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (in press B). Effects of context on a visual 3D pointing task. *Perception*.

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. Forward-backward asymmetry in visual space (in preparation).

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. Visual perception of tilted planes (in preparation).

Refereed Abstracts and conference proceedings

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (2003). Collinearity and exocentric pointing under free viewing conditions. *Perception, 32 (supplement),* 84c. (ECVP 2003)

Kappers, A.M.L., Doumen, M.J.A., & Koenderink, J.J. (2003). Parallelity under free viewing conditions. *Perception, 32,* 84b. (ECVP 2003)

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (2003). Visual perception of spatial relations under free viewing conditions. *Proceedings NVP wintercongres, 9*, 58.

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (2004). The effect of visual angle on the perception of planes in depth. *Perception, 33 (supplement),* 95a. (ECVP 2004)

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (2005). Systematic deviations in a 3-D exocentric pointing task. *Perception, 34 (supplement),* 111. (ECVP 2005)

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (2005). Effects of context on a visual 3D pointing task. *Proceedings NVP wintercongres, 10*, 46.

Doumen, M.J.A., Kappers, A.M.L., & Koenderink, J.J. (2006). Effects of context on a 3D pointing task. *Journal of vision 6(6),* 728. (VSS 2006).

Dankwoord

Promoveren is goed te vergelijken met het scoren van een doelpunt in een hockey-wedstrijd. Een promovendus moet natuurlijk zelf het doelpunt zetten. Het is echter wel belangrijk om een goed team om je heen te hebben met de juiste mensen op de juiste posities.

De voorhoede

Om het doelpunt te zetten, is het handig om zelf in de *spits* te staan.

In mijn promotie-team staat Jan *links-voor*, een plek voor zeer behendige spelers die vanuit een moeilijke hoek de bal toch nog goed voor kunnen geven. Jan, het is duidelijk dat dit jouw plek was tijdens mijn promotie. Aan de hand van jouw eigenzinnige ideeën wat betreft visuele waarneming kon ik zelf mijn richting bepalen in mijn onderzoek.

Astrid zou ik *rechts-voor* zetten, een plek voor iemand die de bal mooi voor kan geven met een verfijnde techniek. Astrid, jouw strak omlijnde ideeen hebben zeker geleid tot een soepel verloop van mijn AiO-traject.

De midden-linie

Normaal staan daar drie spelers. Ik heb er wel wat meer mensen staan: JP, AW, BL, TL, IL, MR, MB, LB, BB, JL, LB, PB, MF, ET, BH, JJ, JW, MW, TH, WB, YH, ES, EV, HP, JB, LW, MB, WS, KK, ST, AB, LW, BB, MS, TC, MK, MP, DB, VJ, JB, GG, AK, TV, BV, MR, MH, BK, MB, MP, LD, TW, SB, CT, SL, MD, NR, KM, SS, MV. Dit zijn alle proefpersonen die mee hebben gewerkt aan de in dit proefschrift beschreven onderzoeken. Daarnaast zijn er ook nog veel proefpersonen geweest die mee hebben gedaan in de pilot-experimenten. Deze mensen hebben het ren-werk gedaan door grote hoeveelheden data voor me te produceren!

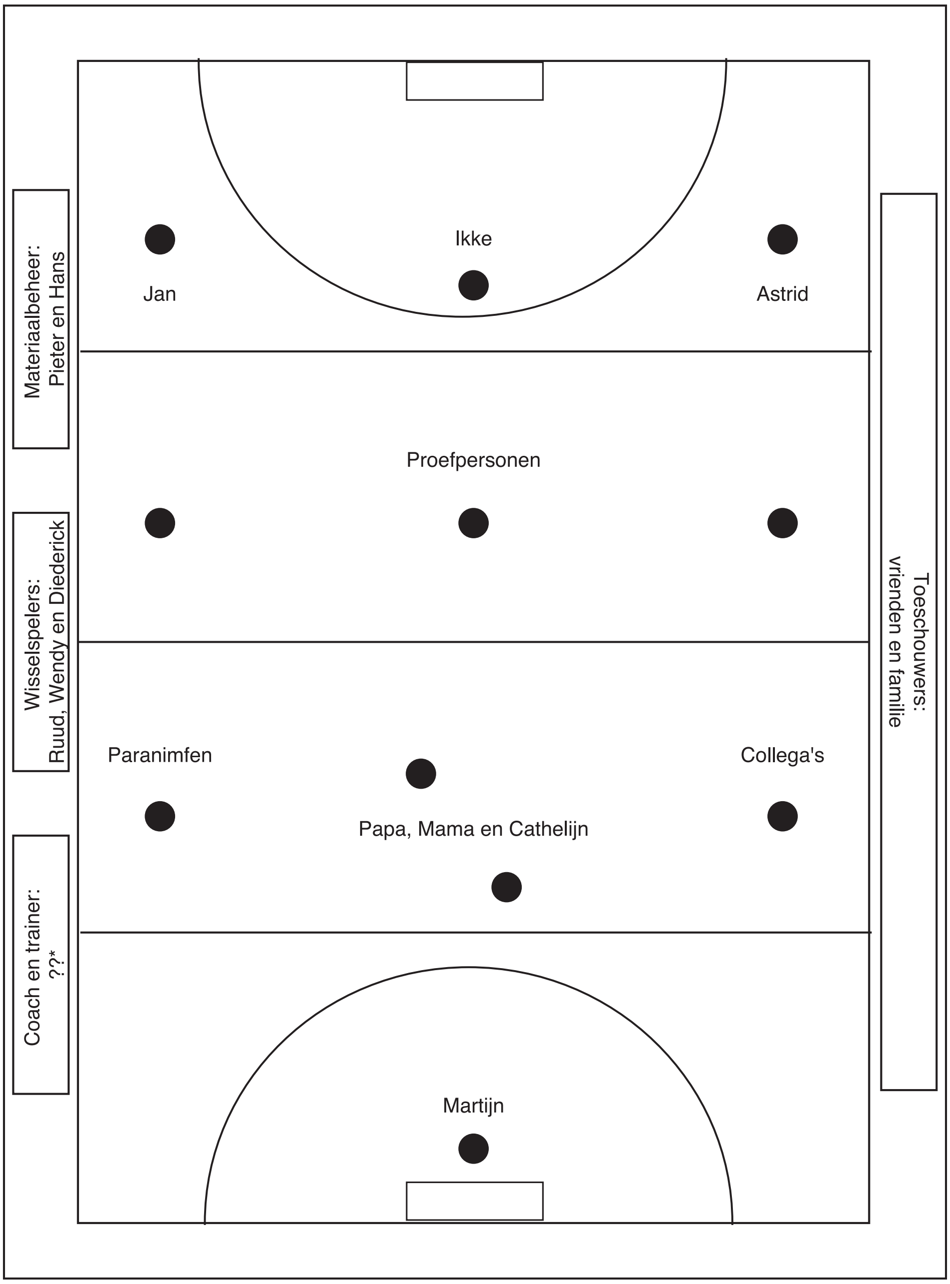
Ikke
Jan
Astrid
Proefpersonen
Paranimfen
Collega's
Papa, Mama en Cathelijn
Martijn
Materiaalbeheer:
Pieter en Hans
Wisselspelers:
Ruud, Wendy en Diederick
Coach en trainer:
??*
Toeschouwers:
vrienden en familie

De achterhoede

De *rechts-achter* speler is ook wat in aantal toegenomen: al mijn collega's die er waren voor leuke discussies (voor het opbouwen van de aanval) maar ook voor ondersteuning als het even wat minder ging (opvangen van een tegenaanval). Alex, Ans, Ansgar, Bala, Bernard, Bob, Bram, Casper, Dagmar, Eelco, Elena, Ellen, Evelijn, Evert-Jan, Gijs, Harold, Jan Joris, Jeroen, Loes, Maaike, Maarten, Martin, Myrthe, Oscar, Raymond, Rob, Robert, Stefan K., Stefan L., Sylvia, Tomas and Wouter: thank you for the good times in our lab!

De *links-achter* speler is ook zeer belangrijk voor het opvangen van een tegenaanval als het allemaal wat tegen zit: mijn paranimfen Els en Mirjam. Els, bedankt voor het delen van alle AiO-perikelen van de afgelopen jaren! Mirjam, bedankt dat ik altijd mezelf bij jou kan zijn en voor het ontwerpen van de kaft!

Als *centrale verdediging* staat daar natuurlijk mijn familie. Papa, mama en Cathelijn, door jullie onvoorwaardelijke steun lukt het me altijd weer om te bereiken wat ik bereiken wil.

En als *keeper* staat er Martijn: je geeft me rust als ik zelf niet meer zo rustig ben.

Een team kan natuurlijk niet draaien zonder mensen buiten het veld:

Ten eerste de *wisselspelers*: wachtend totdat ze het veld in mogen om ook hun kunsten te vertonen. De studenten die ik begeleid heb, nemen deze positie in: Ruud, Wendy en Diederick, het was fijn om samen met jullie onderzoek te doen (training).

Dan is er natuurlijk het *materiaal beheer*: deze taak wordt met grote glorie uitgevoerd door Hans en Pieter. Jullie houden het balletje in onze vakgroep zeer vakkundig aan het rollen!

Ook de *PR* was in goede handen: Sheila, bedankt voor alle verbeteringen van de Engelse teksten.

Zeer belangrijk zijn natuurlijk ook de *supporters*. Zonder supporters zal een team nooit presteren: Hiermee bedoel ik natuurlijk mijn vrienden die me toegejuichd hebben om door te gaan, door lekker te kletsen, te hockeyen, te schaatsen, te carnavallen of een bordspelletje te spelen.

*Dit team moet het zonder *coach* en *trainer* doen aangezien wij wetenschappers natuurlijk veel te koppig zijn om aanwijzingen op te volgen ...

Curriculum Vitae

Michelle was born in Roosendaal, the Netherlands, on the 10th of May 1977. After attending the primary school in Teteringen, she got her VWO diploma in 1995 at the Newman College in Breda. She moved to Maastricht to study psychology at the Maastricht University, specializing into neuropsychology. In 2000, she did an internship at the Biology-department of the Utrecht University on a project within the Behavioural biology group. Here she investigated the relationship between age, social rank and cognitive performance on a discrimination-reversal learning task with long-tailed macaques (Macaca fascicularis).

After her graduation Michelle worked for 4 months at the Maastricht University and the "Expertise Centrum Actief Leren" (ECAL). During this time she developed an ethological practical for the second year students of the Psychology faculty and for fifth year VWO-students. After finishing this project in january 2001, she worked for half a year as a "toegevoegd docent" at the Psychology-faculty in Maastricht where she teached/supervised work-groups of first to third year psychology students.

In September 2001 Michelle started as a PhD-student in the Physics of Man group, at the Physics and Astronomy department of the Utrecht University, which resulted in the present PhD-thesis.