



**Dynamic neural representations
of human visuomotor space**

Stan van Pelt

DONDERS

series

Dynamic neural representations of human visuomotor space

Stan van Pelt

The research presented in this thesis was carried out at the Donders Institute for Brain, Cognition and Behaviour, Centre for Cognition of the Radboud University Nijmegen, the Netherlands, with financial support from the Netherlands Organisation for Scientific Research (NWO) VIDI program, grant number 452.03.307.

ISBN 978-90-9023945-3

Printed by PrintPartners Ipskamp, Enschede, The Netherlands

© Cover artwork: Jasper van der Graaf

© Stan van Pelt, 2009

Cover

Jasper van der Graaf

Muurschildering (zonder titel, 2004, 250 x 500 cm, SAFE Dalfsen)

[Mural, without title]

De schilderijen en muurschilderingen van Jasper van der Graaf bestaan uit verschillende structuren en patronen. De kunstenaar stelt zich tijdens het maakproces geen expliciet, vooraf bepaald eindresultaat ten doel, maar laat de vormen zichzelf creëren, waarbij elke nieuwe lijn als het ware vanzelf aansluit bij het aangrenzende patroon. Dit zelforganiserend vermogen kenmerkt ook onze hersenen. Er is geen centrale aansturing, maar connecties en representaties ontstaan dynamisch door lokale en lange-afstandsinteracties.

The paintings and murals of Jasper van der Graaf consist of various structures and patterns. The artist does not work up to an explicit, predefined end result, but rather lets the shapes create themselves, with each new line emanating as it were naturally from the adjacent pattern. This feature of self-organization also characterizes our brains. There is no central command, but connections and representations emerge dynamically through local and long-range interactions.

Dynamic neural representations of human visuomotor space

een wetenschappelijke proeve op het gebied van de Sociale Wetenschappen

Proefschrift

ter verkrijging van de graad van doctor
aan de Radboud Universiteit Nijmegen
op gezag van de rector magnificus prof. mr. S.C.J.J. Kortmann,
volgens besluit van het College van Decanen
in het openbaar te verdedigen op donderdag 19 maart 2009
om 13.00 uur precies

door

Stan van Pelt

geboren op 31 augustus 1978
te Zevenbergen

Promotor: Prof. dr. H. Bekkering

Co-promotor: Dr. W.P. Medendorp

Manuscriptcommissie: Prof. dr. A.J. van Opstal (voorzitter)

Prof. dr. J.D. Crawford (York University, Toronto, Canada)

Dr. S.F.W. Neggers (University Medical Center Utrecht)

Table of contents

Chapter 1	General introduction	7
Chapter 2	Gaze-centered updating of remembered visual space during active whole-body translations	29
Chapter 3	Updating target distance across eye movements in depth	59
Chapter 4	Visuospatial memory computations during whole-body rotations in roll	83
Chapter 5	Repetition suppression dissociates spatial frames of reference in human saccade planning	109
Chapter 6	Summary and discussion	127
Chapter 7	Nederlandstalige inleiding en samenvatting / Introduction and summary in Dutch	135
	References	153
	Publication list	165
	Dankwoord / Acknowledgements in Dutch	169
	Curriculum vitae	173
	Series Donders Institute for Brain, Cognition and Behaviour	175

Chapter 1

General introduction

Perhaps the most characteristic aspect of life and a powerful engine driving adaptation and evolution is the ability of organisms to interact with the world by responding adequately to sensory signals. In the animal kingdom, the development of a neural system that processes stimuli, learns and controls movement has proven to be a major evolutionary advantage in the struggle for existence. This allowed organisms to flee danger, actively search for food and inhabit new niches and habitats in a much faster pace than ever before in evolutionary history.

The more complex animals became, the more extensive and specialized became their nervous systems (Randall et al. 1997). Whereas some simple invertebrates like echinoderms lack a centralized brain and only have a ring of interconnected neurons to relay sensory signals, vertebrates such as mammals have developed a highly specialized neural network, consisting of a central and peripheral nervous system, in which each subunit has its own functional properties in controlling the body. While the spinal cord and brainstem are involved in controlling automated, internal vegetative processes such as heart beating, respiration and reflexes, the human prosencephalon (forebrain, containing the neocortex) has specialized in so-called higher-order functions, such as perception, action, learning, memory, language and cognition (Kandel et al. 2000). The specialization of the neural control of movement is a major feature that distinguishes primates from other animals. This has led to highly specialized and distinctive capabilities such as reaching, grasping, tool use and ocular foveation.

Sensorimotor transformations

Despite the apparent automation and effortlessness of goal-directed behavior like reaching and grasping, primate brains are presented with complex computational challenges that have to be solved to enable correct execution of such movements. For example, imagine you want to pick up a cup of espresso while reading the newspaper, as illustrated in Figure 1.1. To perform the appropriate grasping movement, you have to convert the visual information that your eyes provide about the cup's location into specific muscle contractions that move your hand to the espresso. How this process, called a *sensorimotor transformation*, comes about is one of the major research topics in neuroscience. To understand the brain's way of executing these transformations, neuroscientists often utilize the concepts of coordinate systems and *reference frames*, such as *eye-centered*, *head-centered*, *body-centered*, or *earth-centric* coordinate frames (Soechting and Flanders 1992). In the mathematical sense, a frame of reference is a set of rigid axes that are usually perpendicular to each other and intersect at one point, the origin. These axes allow the spatial location of an object to be defined by a set of



Figure 1.1. Sensorimotor control in a everyday situation: picking up a cup of espresso, while reading a newspaper. The cup's sensory location (in the right visual hemifield) needs to be transformed into the proper arm motor command ('move left'), which requires a reference frame transformation.

coordinates. This concept is used in a wide variety of applications, also outside of neuroscience. For example, Dutch GPS receivers (such as in car navigation systems) will provide you with your location according to the Dutch Rijksdriehoeksmeting (or RD-grid), in coordinates relative to the Onze Lieve Vrouwe tower in Amersfoort, in Cartesian (metric) units. However, internationally, the WGS84-grid is the GPS-standard, telling you where you are relative to the Greenwich meridian, in degrees longitude and latitude. So, one single location can be described by different coordinates, depending on the reference frame that is used to code this position (in this case, RD or WGS84).

The concept of reference frames can also be applied to spatial representations in the brain. For example, all visual information that falls on the retinas of our eyes is expressed in polar coordinates relative to the fovea, thus in an *eye-centered* (also called

retinal) frame of reference. In contrast, when making a reaching movement, we have to reposition our hands and arms relative to our body. So the required motor commands specify coordinates in a *body-centered* reference frame. In our example, the espresso cup is located in the right visual hemifield (so the retinal input coordinates are ‘to the right’), but it is to the left of the right hand (so the output command should be: ‘move to the left’). Using these terms, we can understand sensorimotor transformations by describing how the brain translates the position of the cup from the coordinates of the retinas into the coordinates of a reference frame linked to the hand.

Spatial constancy in motor control

It is essential for survival to be able to rely on a veridical representation of object locations in the outside world, e.g., for a predator to be able to catch its preys, or for humans to handle tools. The process of keeping track of the locations of objects around us is referred to as *spatial constancy* (Von Helmholtz 1867; Von Holst and Mittelstaedt 1950) and is an important component of goal-directed behavior. While maintaining spatial constancy seems relatively easy as long as objects are visible, it is essential that it can also be relied on in the absence of current spatial input, since we do not always immediately initiate actions upon objects that we see. In some instances, a movement to a previously specified location must be withheld, e.g. in Figure 1.1, seeing the espresso but not reaching for it assuming it is still too hot. In this case, the location of the cup should be stored in memory, in so-called *spatial memory* (Baddeley 1996; Goldman-Rakic 1996; Pierrot-Deseilligny et al. 2002), which can be used to guide a later movement. If the brain would store this information in retinal coordinates, as ‘to the right’, then these memorized coordinates will become immediately obsolete when the eyes look into another direction. In order to maintain spatial constancy, target locations that are important for programming future actions must be stored either in a form that is independent of eye movements, or internally updated to compensate for the eye movements (Crawford et al. 2004). The present thesis is an attempt to obtain insights in the mechanisms by which the brain achieves spatial constancy in motor control.

Over the last decades, a large number of psychophysical studies have shown that we can maintain spatial constancy fairly well. The first who systematically investigated this issue were Hallett and Lightstone (1976), using the now classic double-step saccade task (Figure 1.2). Subjects are shortly presented with two targets in the visual periphery (panel A; T1, T2), and are subsequently asked to make saccadic eye movements to both locations in sequence. In this task, the first saccade dissociates the




Figure 1.2. The double-step saccade task. **A.** While the subject foveates a visual fixation point (F), two peripheral targets (T1, T2) are briefly flashed in sequence. The subject's task is to make successive (memory-guided) saccades to both locations. **B.** After the saccade to T1, the subject has to account for the size and direction of this first saccade to accurately perform the second eye movement, to T2. In simple vector notation, this means that the initial retinal error of T1 ($\vec{T1}_i$) has to be subtracted from that of T2 ($\vec{T2}_i$), to generate the correct final motor error: $\vec{T2}_f = \vec{T2}_i - \vec{T1}_i$. The subject will make an erroneous saccade, towards T2', if the original retinal error of T2 is not updated.

retinal location of the second target ($\vec{T2}_i$) from the goal location of the second saccade (B). Thus, after the first saccade, subjects have to recompute or *update* the intended amplitude and direction of the second saccade to make it reach the goal, following $\vec{T2}_i - \vec{T1}_i$. In the experiment by Hallett and Lightstone, subjects could indeed perform the double-step task correctly, which was taken as evidence that spatial constancy can be maintained across saccadic eye movements. These findings were subsequently confirmed in a series of monkey neurophysiological experiments by Mays and Sparks (Mays and Sparks 1980; Sparks and Mays 1983). Since then, many other investigators have replicated these results. In more recent years, a large number of other behavioral experiments elaborated on these results by testing spatial constancy in different task conditions. By now, it has been shown that subjects can also keep track of memorized targets across smooth pursuit (McKenzie and Lisberger 1986; Gellman and Fletcher 1992; Blohm et al. 2003, 2006; Baker et al. 2003; Raffi et al. 2007) and vergence eye movements (Krommenhoek and Van Gisbergen 1994). Also head movements (Guitton 1992; Medendorp et al. 2002, 2003b; Vliegen et al. 2005) and whole-body movements (Baker et al. 2003; Li et al. 2005; Li and Angelaki 2005; Klier et al. 2008) do not disrupt spatial constancy to a great extent.

Reference frames in spatial constancy for motor control

From a computational perspective, there are many ways to achieve spatial constancy for motor control, depending on the underlying reference frame that is involved. One possibility is that spatial constancy is preserved in a body-centered reference frame. Spatial representations in body-centered coordinates are quite stable because they are not influenced by intermediate eye or head movements, but only by displacements of the body (Flanders et al. 1992; see Battaglia-Mayer et al. 2003 for review). Alternatively, the system could use a retinal frame of reference to code target locations. Employing an eye-centered representation (synonymously referred to in this thesis as *retinal*, *retinotopic*, *gaze-centered* and *gaze-dependent* representations) has as consequence that it needs to be updated every single time the eyes move, in order for spatial constancy to be maintained. For example, if the person in Figure 1.1 redirects gaze towards the watch, the retinal location of the espresso shifts from the right into the left visual hemifield. So internally, the spatial representation of the espresso's position has to be updated accordingly to maintain spatial constancy in eye-centered coordinates.

Another possibility is that the spatial reference lies outside the body. For example, locations can be coded relative to other objects (e.g., in Figure 1.1, relative to the table) or relative to gravity. These so-called allocentric representations do not require any updating for eye, head or body movements, so spatial constancy is always assured. Of course, allocentric representations need to be transformed into egocentric signals when guiding movements.

Which reference frame is the most advantageous in everyday behavior like grasping and looking at objects is not directly obvious and may not be the same for the various actions. For example, in the absence of allocentric cues, a body-centered




Figure 1.3. Parietal sensorimotor areas. **A.** Left hemisphere of the macaque cortex. Color-coded regions – with their corresponding names – indicate regions along the intraparietal sulcus (IPS) that are involved in spatial processing and motor preparation. **B.** Left hemisphere of the human cortex, with the putative homologues of the monkey regions. CS, central sulcus; AIP, anterior intraparietal area; LIP, lateral intraparietal area; MIP, medial intraparietal area; V6A, visual area 6A; VIP, ventral intraparietal area. Adapted from Culham and Valyear (2006).

representation seems optimal for reaching and pointing, because of the benefit of stability. But, notably, there are also arguments in favor of a common, eye-centered coding that is shared across effectors. First, the visual system is the dominant sensory system for spatial information and many brain regions are involved in visual processing. This would make it computationally and energetically beneficial to use a retinal frame of reference as much as possible. A second reason is related to the difference in spatial resolution of these coordinate frames. Retinotopic resolution might degrade when eye-centered information is transformed from into body-coordinates. Another advantage of an eye frame may be to simplify the orchestration of multiple effectors when they move to the same target (Andersen et al. 1997; Cohen and Andersen 2002). Finally, Henriques et al. (1998) proposed eye-centered coordinates to store and update multiple spatial targets, with only the target selected for motor execution being transformed into a body-centered frame ('conversion-on-demand').

Neural mechanisms for spatial constancy

The posterior parietal cortex (PPC) has been shown to play an important role in maintaining spatial constancy in the motor domain and processing sensorimotor transformations. In the primate brain, the PPC contains specialized subunits that are process spatial information for different kinds of movements, such as saccades, reaching and grasping (see Snyder et al. 2000; Colby and Goldberg 1999; Andersen and Buneo 2002; Culham and Valyear 2006; Jackson and Husain 2006, for reviews). Figure 1.3A displays a top view of a rendered representation of the left hemisphere of the macaque cortex, in which several of these parietal subregions are highlighted. The putative human homologues of these monkey areas are depicted in *B*, in a similar view. In both species, specialized areas are located within and around the intraparietal sulcus (IPS). For example, the anterior intraparietal area of the monkey (AIP) and its human functional equivalent are active during grasping, while the lateral intraparietal area (LIP) is involved in representing targets locations of saccades. MIP (medial intraparietal area) and V6A (an extrastriate visual area) constitute the so-called parietal reach region (PRR), coding targets of reaching movements.

Over the last decades, monkey neurophysiological studies have revealed that parietal regions such as LIP and PRR encode and store target locations in eye-centered maps. It was also shown that activity is shifted in these eye-centered maps to compensate for eye movements in order to maintain the correct coding of the target in retinal coordinates. A landmark study on the neural processes involved in spatial updating was performed by Duhamel et al. (1992), who recorded from macaque LIP

neurons while the animals were presented with visual targets (Figure 1.4). First, they showed that LIP neurons have retinotopic *receptive fields*, as is shown for one neuron in *A*. This implies that the cell fires whenever a stimulus is presented at a specific location that is fixed with respect to the fovea. Interestingly, the neuronal activity is sustained, even after the visual stimulus has disappeared. So, this cell continues to code the location of the stimulus in the absence of visual input, which can be regarded as the neural correlate of visuospatial working memory. Neural responses are similar when a saccade brings a stimulus into their receptive field (Panel *B*; different neuron than in *A*). This is a marker of retinotopic updating in LIP. Interestingly, the authors revealed that this updating also occurs when the saccade is made *after* stimulus offset (*C*), which proved that the memory trace is updated. Finally, many neurons in LIP were found to update their activity already *before* the saccade started, as is shown in *D* (same neuron as in *A*). This predictive updating implies that those neurons already “know” the size and direction of the upcoming eye movement, allowing for maintenance of spatial constancy. This is possible only if these neurons had received an *efferece copy* (also known as *corollary discharge*) of the motor command to the eye muscles (Von Holst and Mittelstaedt 1950; Sperry 1950; Sommer and Wurtz 2008). Recently, Nakamura and Colby (2002) observed a similar (predictive) updating of memorized targets at




Figure 1.4. Remapping of visual activity in monkey area LIP. **A.** Receptive field-related activity. The neuron responds to the onset of a visual stimulus in its receptive field, a specific region of space linked to the monkey’s fixation point (dashed circle). After stimulus offset, the firing rate remains at an elevated level, which indicates that the stimulus location is kept in memory. **B.** Spatial updating. This neuron responds when a saccade brings the eye-centered receptive field onto the location of a visual stimulus. **C.** Updating of spatial memory. The same neuron (as in *B*) responds when a saccade moves the receptive field onto a previously stimulated location. **D.** Predictive updating. The neuron (same as in *A*) responds prior to the saccade that brings the receptive field onto the stimulus location. Adapted from Duhamel et al. (1992).




Figure 1.5. **A.** A bilateral parietal region, shown on an inflated representation of the brain, mediates eye-centered spatial updating in a double-step saccade task. LH, left hemisphere; RH, right hemisphere. CS, central sulcus; IPS, intraparietal sulcus. **B.** Double-step saccade task. Subjects are presented with two targets, shortly flashed in sequence (T1, T2), to which they have to make saccadic responses after fixed delay periods. **C.** Responses within the areas highlighted in **A**. Each region shows lateralized activity corresponding to the location of the upcoming saccadic target (1st delay period). After the first saccade, memory activity for T2 is updated according to its new eye-centered location (second delay). Adapted from Medendorp et al. (2003a).

earlier stages within the visual hierarchy, in extrastriate visual areas V2, V3, and V3A. Furthermore, eye-centered remapping processes have been observed in the superior colliculus (Sparks and Nelson 1987, Walker et al. 1995) and the frontal eye fields (Umeno and Goldberg 1997). In addition, for reaching movements, Batista et al. (1999) showed eye-centered coding in PPR. In their experiment, monkeys had to reach for remembered targets, while initial hand and eye position were systematically varied. Across the population of neurons, a scheme that expressed reach vectors using eye-centered coordinates better explained the reach-related activity than a body-centered coding scheme. Notwithstanding all these findings, it should be emphasized that also non-retinal representations are found in other cortical areas such as VIP (head-centered coding; Zhang et al. 2004; Avillac et al. 2005) and AIP (hand-centered signals; Grefkes et al. 2002).

Recently, functional magnetic resonance imaging (fMRI) experiments have demonstrated eye-centered spatial updating in the human parietal cortex (Medendorp et al. 2003a, 2005; Merriam et al. 2003) as well as in other extrastriate visual areas (Merriam et al. 2007). In these human experiments, eye-centered updating was shown by demonstrating the dynamic exchange of activity between the two cortical hemispheres when an eye movement brings the representation of a stimulus into the opposite hemifield. The experiments by Medendorp et al. (2003a) were performed as follows. First a region in parietal cortex was located that showed lateralized responses for memory-guided eye movements (Figure 1.5A), analogous to observations by others (Serenio et al. 2001; Schluppeck et al. 2005). The activity in this region was then monitored using an updating task. Subjects fixated centrally and viewed two brief peripheral dots (Figure 1.5B), a ‘goal’ (T2) and ‘refixation’ target (T1), respectively. Both targets were presented either left or right of central fixation. After a delay, subjects performed a saccade to T1 (the refixation target), which made the remembered location of T2 (the goal target) switch hemifields. Crucially, in these trials, the region’s activation also shifted, as shown in Figure 1.5C. If T2 shifted into the contralateral hemifield after the first saccade, a high sustained activation was observed in the second delay period, but if it shifted to the ipsilateral hemifield the post-saccadic activity level decreased. In other words, this parietal region stored and updated a representation of the goal target relative to current gaze direction. Medendorp et al. (2003a) made these observations if the goal target served for a saccade, but also when it served for a reaching movement. A failure in this gaze-centered updating mechanism could explain the deficits that occur in visuomotor updating in patients with optic ataxia (Khan et al. 2005; Heide et al. 1995).

Additional evidence for the use of eye-centered representations in sensorimotor control comes from psychophysical studies. Behavioral experiments are valuable

complements to neurophysiological measurements, for a number of reasons. First, they can assess which reference frame dominates actual behavior, which is the resultant of all neural processing. Second, behavioral experiments allow for testing of many more motion conditions than simple eye rotations. For example, current fMRI technology does not allow body movements other than simple eye or small pointing movements. Behavioral experiments can test spatial updating across a plethora of dynamics such as head and body translations and rotations, and reaching and grasping movements (e.g., Baker et al. 2003; Karn et al. 1997; Li and Angelaki 2005; Klier et al. 2005, 2007, 2008; Medendorp et al. 1999, 2002, 2003b). A typical approach here is to evaluate the variable or systematic components of the errors subjects make under specific task constraints.

One behavioral experiment that contributed strongly to the current insights in spatial representations was performed by Henriques et al. in 1998, a few years after the classic Duhamel experiment described above. In this study, the experimenters investigated the reference frame for pointing across saccades, using the observation that subjects tend to overshoot (memorized) targets in the visual periphery when pointing (Bock 1986). Henriques and coworkers exploited this relationship by having subjects point to targets originally shortly presented at the fovea, but brought into the retinal periphery by an intermediate saccadic eye movement. Interestingly, the observed pointing errors corresponded to the new – updated – retinal target location, proving that the targets were kept in a dynamic, eye-centered map that was updated across eye movements. This corresponds to the neurophysiological findings described above.

Signals and models for spatial constancy

How is spatial constancy achieved? If the neural system is to maintain a correct representation of target locations in memory, it has to take the effect of all intervening body movements on that representation into account. The conceptual model shown in Figure 1.6 schematically describes the ways in which the amount of self-movement can be registered and how these signals are transformed and combined to guide the updating process. When movements are passive, such as when we ride the train or drive a car, the amount of motion has to be estimated by our internal sensors. E.g., the visual system can detect body movement on basis of changes in retinal input such as the absolute speed of objects moving across the retina (optic flow) and relative velocities due to differences in distance to the point of fixation (*motion parallax*) (Howard and Rogers, 1995). A second important sensory system for motion detection is the vestibular system, located in the inner ear, comprised of the otoliths and semi-circular




Figure 1.6. Conceptual model of sensorimotor processing for spatial updating. Movements (of parts) of the body could affect the the internal representation of a goal for a future action (Target Map). Both passive and active (self-initiated) movements are registered by multiple sensory organs, including the visual, somatosensory and vestibular system. These signals are subsequently processed and integrated by one or more internal models that compute the consequences of the movement on the representations in the target map, resulting in an updating signal to preserve spatial constancy. When the movement is self-initiated, an additional updating signal is computed based on the copy of the movement command, in a parallel, feed-forward process (in gray). This signal is predictive, rather than feedback-controlled, and may yield faster and more accurate updating.

canals. These detect the head's linear acceleration (such as the pull of gravity) and angular velocity, respectively, in all three dimensions (see Green and Angelaki 2007, for a review). Also the somatosensory system may detect differences in posture, including proprioceptive signals of the relative position of body, head, eyes and limbs. So when our car pulls up, together these sensors estimate how much we move, by combining the feel of acceleration, the optic flow caused by objects in the environment and the pressure on our backs by being pushed into the chair. By integrating these sensory signals in an appropriate fashion, a remapping signal can be generated that is used to update the target representation and maintain spatial constancy.

In our scheme, the computation of this remapping signal takes place in the *internal models* box. Although we will not go into detail here, it can easily be

understood that these computations are complex due to the differences in internal dynamics and intrinsic reference frames that each sensor uses (see e.g., Pouget and Snyder 2000; Laurens and Droulez 2007; Keith et al. 2007; Keith and Crawford 2008; Quaia et al. 1998; White and Snyder 2004, for computational approaches). E.g., the otoliths, hardly sensing very slow movements, detect movements in a head-centered frame of reference. In contrast, the visual system is very good at detecting even small changes but codes information in eye-centered coordinates. During the course of our lives, our internal models are calibrated such that they can correctly perform an optimal integration of sensory signals (Vaziri et al. 2006) and execute the required reference frame transformations to compute the proper updating signals in the coordinates used by the spatial memory map.

The relative importance of each sensor in updating for passive movements can be inferred when subjects are deprived of input in a specific sensory modality. This approach has revealed that vestibular signals are crucial in the computation of updating signals. For example, monkeys whose vestibular system was lesioned showed severe deficits in spatial updating during passive body motion (Li and Angelaki, 2005).

Importantly, when our movements are not passive, but self-initiated, for example when we refixate our gaze or during walking, our target representation has access to information from an additional updating system, parallel to the sensory feedback stream. This feed-forward process (shown in gray) exploits that we have knowledge about the kind of movement that is coming up. Here, the efference copy signals of the intended eye, head or body movement is fed into an internal model to generate updating signals. This model may be similar to or perhaps integrated in the one used in the sensory feedback stream, since it might predict or simulate the sensory information that the movement would generate. Because of the early initiation of this process – when there is only a movement plan but no actual movement yet (Mays and Sparks 1980) – the feed-forward updater is fast and may facilitate an optimal maintenance of spatial constancy. Neural correlates of the feed-forward updating signals include the predictive remapping activities observed in various cortical regions, some of which are described above (Duhamel et al. 1992; Walker et al. 1995; Heiser and Colby 2006). This way, updating for an active movement will be faster, and perhaps more accurate, than a similar, passive movement.

Different views exist on the computational operations involved in the updating process. One of the simplest models describes updating in terms of *vector subtraction* (Goldberg and Bruce 1990; Moschovakis et al. 1998; Cassanello and Ferrera 2007). In this model, the updating signal is a simple vector representing the intermediate movement. In the case of a saccadic eye movement, the novel retinal position of a remembered target is

computed by subtracting the size and direction of the saccade from the original target coordinates. In the double-saccade example of Figure 1.2, the location of T2 – expressed in vector coordinates with the origin in the fovea – after the saccade to T1 can be computed as $\vec{T2}_{updated} = \vec{T2}_{initial} - \vec{T1}$. Although a vector model of updating is fairly well able to describe updating for horizontal and vertical eye movements, its underlying geometry predicts inaccurate updating for eye movements between eccentric targets. However, when tested for this, subjects did not make the errors that this model predicted (Smith and Crawford 2001). Moreover, a vector model cannot explain the correct updating behavior that subjects show after rotating their heads (Medendorp et al. 2003b) or body (Klier et al. 2005) about a naso-occipital axis. This suggests that the updating system is able to deal with the nonlinearities of 3D rotations. This complexity in spatial updating was further examined in a study by Klier et al. (2007) who showed using a paradigm involving a sequence of rotations about different axes that spatial updating even accounts for the problem of non-commutativity of rotations. Together, these findings suggest that the brain is able to combine a memorized, two-dimensional retinal error vector with one or multiple 3D rotational vectors to generate a correct movement plan towards the remembered target location (Tweed and Vilis 1987).

Outline of this thesis

Although there is a large body of evidence for the dominance of eye-centered representations in sensorimotor control, most of this support is based on behavioral and neural signals derived during simple eye saccades with the head and body restrained (but see Baker et al. 2003). In everyday life, however, we do not only move our eyes, but continuously move our head and body around, in all directions. It is clear that spatial constancy can be maintained during more complex movements (see above), but the underlying spatial computations have not been identified. This thesis describes a number of experiments testing the mechanisms to preserve spatial constancy in more challenging updating conditions.

Translational updating

When the eyes translate through space – as occurs with head and whole-body translations, but also during natural head rotations, when the rotation axis lies in the neck – the updating system within the brain is faced with complex geometrical requirements. During translation, visual objects at different depths from the plane of fixation move with different velocities across the retina. This is known as motion parallax, and the same geometry needs to be accounted for in the eye-centered updating

of *remembered* targets during translational motion. Figure 1.7 illustrates this issue in more detail. The left-hand panels show the updating of two remembered visual stimuli T1 and T2 (A) across a leftward eye rotation (B). In this case, both targets locations simply have to be updated by the same (angular) amount to retain spatial constancy (C). In contrast, a translation of the eye requires the same two targets locations to be updated by different amounts (right-hand panels). This implies that the internal models in the brain must simulate the translational-depth geometry of motion parallax in order to instantiate spatial updating across translational motion. Medendorp et al. (2003b) showed that human subjects can maintain spatial constancy during self-initiated, sideward head translations. Similar results were found during passive whole-body translations of humans (Klier et al. 2008) and non-human primates (Li and Angelaki 2005). However, none of these studies explicitly tested the reference frames involved in this behavior. In Chapter 2, we took a further look at this issue by investigating the reference frame that is used to code and update target locations during active translational motion. Targets and movements were chosen such that retinal and non-retinal updating models made different predictions about the types of errors that subjects would make when reaching to the remembered targets.




Figure 1.7. Geometrical consequences of rotational (left-hand panels) and translational eye movements (right-hand panels) for spatial updating. Two targets (T1, T2) presented at the same retinal location (A) have to be updated by the same amount during an eye rotation (B, left-hand panel), rotating the stored locations through the inverse of the eye's rotation in space (C, left-hand panel). An eye translation (B, right-hand panel) requires target updating by different amounts, depending on each target's distance (C, right-hand panel). Adapted from Medendorp et al. (2008).

Updating in depth

Keeping track of the direction of a target may be adequate to direct saccadic or pointing movements, it will be insufficient to retain spatial stability for 3-D movements like reaching or grasping. For these types of actions, information about the distance (or depth) of the target is also required. It is generally accepted that depth and directional information are processed separately in the brain (Cumming and DeAngelis 2001; DeAngelis 2000; Flanders et al. 1992; Vindras et al. 2005), so it can be assumed that the spatial constancies of the respective dimensions are computed by different processes too. To date, not much is known about how spatial constancy in depth is preserved. Li and Angelaki (2005) showed that monkeys are able to update target depth after passive translation in depth, a finding which was replicated in humans (Klier et al. 2008). In addition, Krommenhoek and Van Gisbergen (1994) showed that vergence movements are taken into account when looking at remembered targets at different distances. However, none of these studies addressed the nature of computations that are involved in coding target depth. Especially in impoverished environments and without translational motion, the main cue for target distance is *retinal disparity* (Howard and Rogers 1995; Julesz 1971; Wei et al. 2003). As illustrated in Figure 1.8A, a target T lying behind fixation point F will appear at different locations on the retina of each of the two eyes, at retinal eccentricities α and β for the left and right eye, respectively. Together, these angles define the retinal disparity of the target (disparity = $\alpha - \beta$), which represents the distance of the target relative to the point of fixation. A representational system can store target depth in the form of its absolute body-centered distance (ϵ), which can be computed by combining target disparity with fixation distance information as is given by the binocular vergence angle (γ). Alternatively, target depth can also be kept in a retinal reference frame, in a dynamic retinal disparity map which is updated for vergence eye movements. There are neurophysiological indications for the latter. Gnadt and Mays (1995) described neurons in monkey LIP that had three-dimensional receptive fields that were fixed to the monkey's binocular point of fixation, as schematically illustrated in Figure 1.8B. Chapter 3 investigates how the human brain keeps track of target distance across vergence eye movements by exploiting the systematic errors in the reach towards targets in depth.

Gravitational signals in spatial constancy

Recent studies have shown that spatial constancy accounts for changes in head or body orientation relative to gravity, at least for small tilt angles $\leq 90^\circ$ (Medendorp et al. 2002; Klier et al. 2006). Is gravity used as anchor in the coding of spatial information during such body motion? In Chapter 4 we examined the role of gravity in spatial coding by testing saccades to remembered targets after large intervening body tilts about the naso-





Figure 1.8. Inference of target distance based on retinal disparity. **A.** Target T, lying behind fixation point F, is located at different locations relative to each eye's gaze line (α , β), and causes a retinal disparity equal to $\alpha - \beta$. Absolute target depth, expressed as the angle ε , can be computed by combining binocular fixation depth with retinal disparity, following $\varepsilon = \gamma - (\alpha - \beta)$. **B.** Schematic visualisation of a three-dimensional receptive field, based on single unit recordings in monkey LIP (from Gnadt and Mays 1995).

B 3D-receptive field



occipital axis. An eye- or body-centered coding principle requires that target locations are updated for body rotation, whereas a gravitocentric representation remains stable, not requiring any updating. To dissociate between these reference frames we exploited the well-known error in the perception of the Earth vertical when the body is tilted with respect to gravity. The error is termed the Aubert- or A-effect, after its discoverer (Aubert 1861). As Figure 1.9 shows, the A-error is quite substantial at large tilt angles $>90^\circ$ (Van Beuzekom and Van Gisbergen 2000; Kaptein and Van Gisbergen 2004). In Chapter 4 we tested whether the A-effect interferes with spatial constancy during body tilts, which is to be expected if targets are coded in allocentric, gravity-based coordinates, but not when coded in an egocentric reference frame.




Figure 1.9. Systematic response errors when subjects have to estimate the Earth vertical, as a function of body tilt angle. At large angles, the direction of gravity is strongly misestimated, called the A-effect. Adapted from Van Beuzekom and Van Gisbergen (2000).

Revealing spatial reference frames in the brain

As described above, most of the fMRI-experiments that investigated the reference frames of spatial representations exploited that many brain regions exhibit a lateralized internal organization (Medendorp et al. 2003a; Merriam et al. 2003; Gardner et al. 2008). This laterality, or spatial topography, is specified by an overrepresentation of the contralateral hemifield. For example, neurons in the primary visual cortex (V1) are orderly distributed both within and across the two hemispheres, according to the part of visual space they represent. But regions do not always have a topographic organization, which makes it difficult to identify their dominant reference frame based on topography measurements. This is illustrated in Figure 1.10. In a topographically organized brain region, neurons are distributed according to their receptive fields (*A*, left-hand panel). When stimuli are presented either in the left or in the right visual hemifield, this will primarily activate the neurons in the contralateral cortical hemisphere, as indicated by the gray rectangle and ellipse. This eye-centered activity is reflected in the BOLD signal measured with fMRI (*B*, left-hand panel). However, if the same spatially-selective neurons are randomly distributed across the hemispheres, there will be no bias in cortical activity (right-hand panels of *A* and *B*). So a stimulus in the right visual hemifield will yield similar BOLD signals in both hemispheres.

In Chapter 5, we used the method of *repetition suppression* to identify the spatial reference frame of cortical areas involved in saccade planning. This method might reveal the internal coding of spatial representations in regions that lack a clear topographic neuronal organization – such as in more downstream visual areas (Jack et al. 2007).

Repetition suppression (RS) involves the decrease in cortical responses to repeated stimuli, which may reflect a ‘sharpening’ of cortical representations (Wiggs and Martin 1998; Desimone 1996) or shorter processing durations (Henson and Rugg 2001). RS has been observed in different fields, such as object perception (Grill-Spector and Malach 2001), number representation (Naccache and Dehaene 2001), and action observation (Hamilton and Grafton 2006, 2008). Interestingly, it has never been used




Figure 1.10. Two possible ways of organization of spatially-tuned neurons. With an orderly, topographic organization (e.g., in the primary visual cortex), neurons are distributed according to their response field locations (**A**, left-hand panel). Accordingly, a stimulus at $+30^\circ$ will activate the neurons highlighted by the rectangle in the left hemisphere, and a stimulus at -30° those marked by the gray ellipse in the right hemisphere. In fMRI, this lateralization is reflected in the BOLD signal, which will show higher levels in the left hemisphere for rightward than leftward targets, and vice versa for the right hemisphere (**B**, left-hand panel). Spatially-tuned neurons can also be distributed randomly across both hemispheres (right-hand panels). This has as a consequence that BOLD activity in such a region will not be lateralized, and thus that this region's spatial reference frame cannot be identified by exploiting laterality. An alternative approach is to exploit repetition suppression effects (**C**). Spatially-selective neurons will show an attenuated response when a stimulus location is repeated, independent of their distribution across the hemispheres.

to test between reference frames in spatial coding and movement planning tasks. It can be hypothesized that activity of neurons that code targets in a specific reference frame will be attenuated when the target location is repeated in that frame and not when it is repeated in another coordinate frame (Figure 1.10C).

Chapter 2

Gaze-centered updating of remembered visual space during active whole-body translations

This chapter has been published as:

Van Pelt S, and Medendorp WP. Gaze-centered updating of remembered visual space during active whole-body translations. *Journal of Neurophysiology* 97: 1209-1220, 2007.


In daily life, we appear to be perfectly aware of objects in our surroundings. Even when we move, we seem to have no difficulty in keeping track of objects and reach or look at their locations whenever necessary. This seemingly automatic behavior, called spatial updating, works even in darkness, and for targets that are otherwise no longer in view (Hallet and Lightstone 1976; Li and Angelaki 2005; Medendorp et al. 2002). But despite extensive investigations, the computational basis of spatial updating has remained controversial (Andersen et al. 1985; Duhamel et al. 1992; Baker et al. 2003).

A critical aspect of this issue is the reference frame with respect to which object locations for actions are encoded. A reference frame is characterized by a coordinate system, which represents locations using a set of coordinate axes fixed relative to some origin, like the eyes, head, body or earth. Obviously, in theoretical terms, spatial updating could work in any coordinate frame as long as the correct updating signals and computational operations are used (Medendorp et al. 2003b). Adding to this notion, various studies have argued that the reference frame used to encode a spatial memory is not fixed but depends on several factors, including the sensory inputs, task constraints, the visual background, memory interval, and the cognitive context (Battaglia-Mayer et al. 2003; Hayhoe et al. 2003; Bridgeman et al. 1997; Carrozzo et al. 2002; Snyder et al. 1998; Chapter 4 of this thesis). Within this view, psychophysical evidence obtained in neutral open-loop testing situations has suggested that the early feedforward mechanisms for internal spatial updating operate in gaze-centered coordinates (Henriques et al. 1998; Medendorp and Crawford 2002; Baker et al. 2003). In further support of this evidence, many brain regions in parietal and frontal cortex have been shown to update their activity patterns relative to the new gaze direction after an eye movement has occurred (Duhamel et al. 1992; Batista et al. 1999; Medendorp et al. 2003a; Merriam et al. 2003; Sommer and Wurtz 2002).

It is important to point out though that most of the actual evidence for gaze-centered updating was obtained using simple eye rotations only, with the head and body restrained, ignoring the fact that in natural situations our eyes also translate through space, as for example when we walk. When the body translates, correct updating in a gaze-centered frame seems computationally much more demanding because the required updating varies from object to object, depending nonlinearly on their depth and direction, as in motion parallax (Medendorp et al. 2003b; Li et al. 2005). In this respect, updating for translational motion seems much simpler if object locations were stored in, say, Cartesian body-centered coordinates because then the required updating would be the same for each object: the opposite of the amount of body displacement (Medendorp et al. 1999).

At present, it is unknown which reference frame is involved in the computations for the translational updating of remembered visual space. Here, we address this question by characterizing the pattern of errors in manual reaching movements toward briefly flashed targets, presented prior to a whole-body translation. Our goal is not to merely characterize a subject's ability to update spatial information for intervening translations. In fact, recent studies have already shown that humans and monkeys can look to remembered locations in near space, compensating for intervening eye translation induced by head or body motion (Israel et al. 1999; Medendorp et al. 2003b; Li et al. 2005). However, the computational principles underlying the spatial constancy in this behavior, whether gaze-related or not, remain to be revealed. We designed a novel experiment to discriminate between a gaze-dependent and a gaze-independent model of visuospatial memory updating during translations. In our test, subjects fixate centrally at fixation point FP, while a far or near target (T_f , T_n) is flashed onto the retinal periphery (Figure 2.1, middle column). Subjects then translate sideways (by making an active whole-body step displacement) while keeping their gaze at FP, and subsequently reach to the remembered target location. The logic behind the test is the following. Suppose that the targets were visible at all times, also when the

Figure 2.1. Predictions of the gaze-dependent and gaze-independent models of internal spatial updating during whole-body translations. The basic assumption in the test is that subjects generally misestimate the amount of self-motion when the body translates. **A.** The subject looks at a central fixation point (FP), while a target is flashed, either in front of (near $T = T_n$) or behind FP (far $T = T_f$). An internal representation of this target is coded in either a gaze-dependent frame (**A**, left-hand panel) or in a gaze-independent frame (right-hand panel). Thus, in the gaze-dependent frame, near and far targets are stored as memories coding opposite locations relative to the gaze-line. In the gaze-independent frame, say a body-frame, they are transformed and stored as memories reflecting positions at the same side from the body midline. **B.** After viewing and storing the target, the subject translates the body, e.g. in rightward direction, while keeping fixation at FP, and then reaches toward the remembered location of the target (**C**). If the target is stored in a gaze-dependent frame (**B**, left-hand panel), the subject should compensate for the induced change of gaze, by updating the gaze-dependent memory trace. That is, the near-target memory should be shifted to the left while the far-target memory should be shifted to the right. If compensation is only partial as a result of an erroneous estimation of step size, memory traces will be shifted partially and hence will not match the actual location of the targets. This will result in reach errors, denoted by E_f and E_n , which reverse in direction for remembered targets at opposite depths from fixation (**C**, left-hand panel). Alternatively, if targets are stored in a gaze-independent body frame, the subject should compensate for the induced changes of the body. In effect, when the body translates to the right, all memory traces should be shifted to the left, by equal amounts (**B**, right-hand panel). If then updating is only partial, this will result in reach errors in the same direction for all remembered target locations (**C**, right-hand panel).



body translates sideways. Then, parallax geometry dictates that targets in front and behind the eyes' fixation point (FP) shift in opposite directions on the retinas. Thus, if the brain were to simulate motion parallax also in the active updating of *memorized* targets (left column, black arrows), it can be predicted that if the body translation is not correctly taken into account (Glasauer et al. 1994; Medendorp et al. 1999), the updated locations (gray arrows) will deviate from the actual locations, leading to reach errors (E_f , E_n) in opposite directions for targets in front of and behind the FP (hypothesis A: gaze-dependent updating). Alternatively, parallax geometry plays no role if the brain codes locations in a gaze-independent reference frame, e.g., in a body-fixed frame (right column). If then translations are misjudged, the updated locations will also deviate from the actual locations, but with updating errors (as probed by the reach) in the same direction for all targets (hypothesis B: gaze-independent updating).

Our results demonstrate that translational updating follows the predictions of the gaze-dependent scheme. To obtain further insights in the putative computations in this process, we trained a simple three-layer recurrent neural network to perform gaze-centered updating in these translation conditions. The network learnt correctly the geometric computations involved, and preferred velocity, rather than position signals for updating remembered visual space during whole-body motion (White and Snyder 2004).

Methods

Subjects

Fifteen human subjects (four female, eleven male, mean age of 26 ± 4 years), were tested in four different task conditions, as described below. The main experiment involved ten naïve subjects, and the two authors. Each of the three additional control experiments tested five subjects (three naïve). All subjects signed informed consent to participate in the experiment. All subjects were right-handed, and all were free of any sensory, perceptual, or motor disorders. All pointing movements were made using the right arm.

Experimental setup

Subjects were standing in a completely darkened room, within a designated area of 60 cm width, which we will refer to as the 'translation zone'. A U-shaped ridge of 6 cm height was attached to the floor indicating the outer borders of the translation zone to the left, right and back of the subject. During the experiments, this ridge served as a reference for subjects to position their feet in order to accurately control their own

positions and self-induced translations. This configuration led to lateral body translations with an amplitude of 30 cm ($SD=7$ cm), averaged over all subjects. Within subjects, positions and translations were reproduced with an accuracy better than 3 cm.

We used an Optotrak 3020 digitizing and motion analysis system (Northern Digital) to record the position and orientation of various body parts in three dimensions (3D). This system tracks the 3D position of infrared-emitting diodes (ireds) with an accuracy better than 0.2 mm. We determined head position and orientation by means of four ireds attached to the eye tracking helmet worn by the subject (see below). Prior to the experiment, we calibrated the locations of the eyes and ears with respect to the ireds on the helmet. During this calibration procedure, the subject faced the Optotrak camera while wearing the helmet with three additional temporary ireds, one near the right auditory meatus and one on each closed eyelid. The 3-D locations of these ireds, which uniquely defined the location of the right ear and both eyes relative to the helmet, were recorded together with the ireds on the helmet. With this information, we were able, during the subsequent experiment, to compute the positions of the eyes and ear in space on the basis of the helmet ireds alone. The actual location of each eye, defined as its rotation center, was assumed to be 1.3 cm behind its cornea. In a similar fashion, we calibrated the position of the tip of the right index finger relative to four ireds attached to the middle phalanx of this finger. We further used the Optotrak system to record the position of the shoulder (acromion) as well as the positions of the stimulus targets. Optotrak data were sampled at 125 Hz. The ired coordinates were transformed to a right-handed space-fixed coordinate system. The x-y plane was aligned with the subject's horizontal plane. The positive x-axis was pointing forward, perpendicular to the subjects' shoulder line, and the positive y-axis was pointing leftward along the shoulder line, seen from the subject, and the z-axis pointing upward. The position of the central LED on the stimulus array (see below) served as the origin of the coordinate system. The orientation of the head was determined with respect to a reference position adopted when the subject faced straight ahead. Orientation and location measurements were accurate to within 0.2° and 0.2 mm.

We used an Eyelink II eyetracker (SR Research Ltd., Canada) to record binocular eye movements. We ensured that its camera system, which was mounted to the helmet, remained stable on the head during the entire experiment. Stable recording of eye position was further warranted by measuring corneal reflections in combination with pupil tracking, which reduces the errors caused by any helmet slip and vibration. As a further precaution, subjects were also instructed to minimize speaking during the experiments.

Eye movements were calibrated before the experiment, by having subjects face straight ahead and fixate the stimulus LEDs two times each, in complete darkness, both when standing left and right within in the translation zone. Eye recordings were calibrated in the head-fixed coordinate system of the eye-tracker. By combining the locations of the stimuli and the reconstructed locations of both eyes (using the helmet calibration data), as well as current head orientation, we computed the direction of the stimulus LEDs with respect to the subject's eyes in head-fixed coordinates. In this way, the eye-tracker data of both eyes could be matched to the corresponding vertical and horizontal stimulus directions, and expressed as eye-in-head orientation signals. During the actual experiments, eye-in-space orientation was calculated by combining head orientation and calibrated eye-in-head orientation signals. The eye calibration procedure resulted in a directional accuracy of the eye-in-head orientation better than 1.5° . Version and vergence positions were calculated from the left (L) and right (R) eye positions, as $(R+L)/2$ and $L-R$, respectively.

Two PCs controlled the experiment. A master PC was equipped with hardware for data acquisition of the Optotrak and Eyelink measurements, as well as visual stimulus control, while a slave PC contained the hardware from the Eyelink system.

Stimuli

Nine red light-emitting diodes (luminance < 20 mcd/m²) served as stimuli. They were attached to a frame in the shape of a cross that was mounted on a two-link robot arm. This robot arm, equipped with stepping motors (type Animatics SmartMotors; Servo Systems), could rapidly position the center of the frame to virtually any desired position within a hemisphere (radius 1 m) centered at its base. The frame was positioned with an accuracy of better than 0.2 mm, as confirmed by Optotrak recordings. During the experiment, the stimuli were presented at space-fixed locations, at eye-level in the subject's transverse plane (Figure 2.2A). The location of the central LED, which served as fixation point (FP), corresponded to the origin of our space-fixed coordinate system, which was straight in front of the center of the translation zone, at a distance of about 35 cm. Four other LEDs were lined up with the x-axis of the coordinate system and served as visual targets for task conditions described below. Two of these targets were behind the central LED (from the subject's perspective), at distances of 7 and 17 cm (T1, T2), and two were in front of the central LED, at distances of 6 and 10 cm (T3, T4). Using this configuration, we ensured that the target flashes stimulated both retinas during the experiments, at equal intervals of about 4° . We further positioned four other LEDs along the y-axis of the coordinate system, at either side of the central light at 6 and 12 cm (not shown in Figure 2.2A). These targets were used in catch trials to ensure that subjects did not simply make repeated




Figure 2.2. **A.** Sequence of stimuli and the subject's instructions during the translation trials. Subjects start by fixating a space-fixed target (FP) for 1.5 s. Next, a second space-fixed target was presented for 500 ms, either in front of or behind the fixation point. Subjects translated their body within a 2.3 s memory period while they maintained fixation on FP. After another 100 ms, an auditory cue signaled the subjects to reach toward the remembered location of the target. Stationary trials (not shown) differed from translation trials by the absence of subject translation during the memory period. Four space-fixed targets (1-4) served as potential target locations, presented such that their mutual distance in terms of retinal eccentricity was 4° , when the subject was standing at opposite ends within the translation zone. Figure not to scale. **B.** Typical performance of one subject (S1). Body position, eye position (version and vergence) and finger tip position (horizontal component) plotted against time for 16 stationary (left-hand panels) and 16 translation trials (right-hand panels). The target for memory was T1 (i.e., behind the fixation point). Black traces, leftward final position; gray traces, rightward final positions; dotted traces, geometrically-ideal signals; thin boxes, time intervals of the different trial stages (target presentation (T), memory period, reach interval).

stereotypic responses. Data for these catch targets were excluded from further analysis. We also made sure that subjects never saw the target configuration when the room lights were on, by positioning it to an elevated level using the robot.

Main task

The experiments were designed to test between a gaze-dependent and gaze-independent model of visuospatial updating for translational motion. In our test, subjects were instructed to perform memory-guided reaching movements under two conditions, which will be referred to as ‘stationary’ and ‘translation’ tasks. The experimental paradigm of the translation task is illustrated in detail in Figure 2.2A. Before the start of each trial, subjects positioned their feet on either the left or right end of the translation zone to certify a fixed starting position. A trial started with the onset of FP which was illuminated for 4.3 s and had to be fixated by the subject for its entire duration. At 1500 ms after the onset of FP, a target for memory (here T1), closer or farther than FP, appeared in the visual periphery for 500 ms. Then a 2.3 s time interval followed, in which subjects were instructed to either remain stationary (‘stationary task’) or to make a sideward step to the opposite side of the translation zone (‘translation task’), while still fixating FP. Then, FP was extinguished, the stimulus frame was retracted, and 100 ms later an auditory signal cued the subject to conjointly look and reach at the remembered location of T, keeping the body and head still. Subjects had to hold that reaching position until another auditory signal was presented 3.6 s later. Then, the next trial started. Targets were randomly chosen from the four locations. Each target location was tested 20 times for both starting positions, resulting in a total of 160 trials for each of the two tasks. Test trials were randomly interspersed with 32 catch trials. Subjects never received any tactile feedback during their reach. In all trials, subjects had to keep their head and body aligned, in the straight ahead direction. In the translation trials, the starting position of a trial was the end position of the previous trial, whereas in the stationary task the subject first moved to the other end within the translation zone before testing the next trial. Thus in the stationary task, response data was gathered at positions that also served either as initial or as final position in the translation task (F-test, $p > 0.05$). This allowed direct comparison of response behavior when updating was necessary (translation task) with that where no updating was needed (stationary task). For both test conditions, the total duration of each trial was 8.0 s.

During the reaching movement, visual feedback about hand position was provided by means of an LED attached to the fingertip. This way we tried to minimize the error attributable to an erroneous estimate of fingertip position during pointing (Beurze et al. 2006). We also allowed subjects to look where they were reaching in

order to eliminate contributions of errors occurring otherwise, i.e., when gaze would be off the reach location (Henriques et al. 2003, see control experiment II below).

The total experiment was divided into 2 sessions, tested on different days. In each experimental session, half of the translation trials were tested first, followed by half of the stationary trials. Subjects performed blocks of 12 consecutive trials, between which a brief rest was provided with the room lights on to prevent dark adaptation. During these periods, the stimulus frame was out of view. Each session lasted for about 60 minutes. One subject was tested over three sessions. During the experiments, subjects never received feedback about their performance. Before the actual experiment began, subjects practiced a few blocks to become familiar with the two task conditions.

Control tasks

We also performed three control experiments, in which we varied a number of task parameters to test their implications for updating behavior. All controls were performed with the same timing and stimulus durations as in the main experiment, unless indicated otherwise. First, we tested updating performance in the absence of visual feedback about fingertip position during the reaching movement (control I: reaching without feedback). This clarified whether the results in the main experiment were not critically dependent on a visually-monitored hand position during the reach. The next control experiment was inspired by the fact that reaching while looking where you reach is generally more accurate than reaching to a retinally peripheral location (Henriques et al. 2003). Therefore, in contrast with the main experiment, subjects performed the reaching movement in this task by keeping gaze fixed at the remembered location of FP (control II: reaching without looking). This tested whether the results of the main experiments were not mainly driven by one of the two motor systems (eye vs. arm). The final control was designed to test the effect of a visual fixation point (FP) during the updating task (control III: updating without FP). Therefore, in this task, FP was turned off immediately after the target flash, and subjects were instructed to make their body translation by keeping their gaze fixed on remembered FP. Reaching was performed under visual feedback of the fingertip, which had to be fixated. As the eyes may diverge from the remembered FP during the translation in darkness (Medendorp et al. 2003b), updating was tested for the two outermost targets only, since these were most discriminative in terms of the models outlined in Figure 2.1.

Data analysis

Data were analyzed off-line using Matlab (The Mathworks). We excluded trials in which subjects did not keep their eyes directed at FP within a 3° interval or made a

saccade during target presentation. We also discarded trials in which the subject had not correctly followed other instructions of the paradigm, e.g., when stepping or reaching too early, or not making a step when this was required. Typically, 23 ± 11 trials ($\sim 7\%$) were discarded based on the arm and eye movement criteria. For each of the remaining trials, final reaching positions were selected manually at the time when the arm had the greatest degree of stability within the last 2 s of the response interval. For each trial, an average position was computed over a 6 sample interval (48 ms) centered at this point in time. After categorizing the stationary and translation trials by starting position and translation direction, respectively, we computed the mean reach endpoint separately for each of the targets within these categories. Starting and final body positions were defined by the location of the center of the two eyes at the time of target presentation and reach response, respectively. The difference between these two positions determined the amplitude of the translation (step size). We tested between gaze-dependent and gaze-independent updating models by comparing the horizontal components of the updating errors of reaches toward the targets flashed in front of and behind FP in the translation trials. Since both variables are subject to natural variation and measurement error, a Model II regression (also referred to as a major-axis regression) was used to determine their relationship, with slope and confidence limits estimated by the bootstrap method (Press et al. 1992). We used the results of the stationary paradigm as a measure for errors attributable to perception or motor effects assuming that both contributed equally. A further 2-D vectorial analysis was performed to entail how the interaction between initial target position, translational motion and reach response can be described in both gaze-dependent and gaze-independent coordinate frames (see later). Statistical tests were performed at the 0.05 level ($p < 0.05$).

Neural network model

In order to understand our findings in neurophysiological terms, we trained a simple recurrent three-layer Elman-type neural network using backpropagation to perform gaze-centered updating for both intervening rotations and translations of the eye. We used a similar type of network architecture as White and Snyder (2004) who modeled the updating process for (conjugate) eye rotations only. The predictions of this model will be discussed in the discussion. In the present model, the input layer of the network includes a map of neurons with similar spatial tuning properties as those observed in parietal region LIP: Gaussian-like receptive fields for the eye-centered direction of a stimulus and its relative depth from the plane of fixation (retinal disparity) (Gnadt and Mays 1995). For simplicity we used a 2-D horizontal-disparity map of 121 units (11×11 units; horizontal range -50 to 50° disparity range -25 to 25°). Each unit within the map had a 2-D Gaussian tuning curve, with a $10 \times 5^\circ$ horizontal-disparity receptive

field ($1/e^2$ width), so that receptive fields of units at neighboring locations overlap considerably. Stimulus direction and disparity input to the network were limited to $< 20^\circ$ and $< 9^\circ$, respectively. The network also received 4 eye position units: one pair of units represented binocular gaze (version); another pair encoded binocular depth (vergence). For each unit, the activity was linearly scaled within the range -1 to +1, corresponding to -40 to $+40^\circ$ version angle and 0 to $+10^\circ$ vergence angle, respectively. In each pair, the second unit had the opposite activity of the first (push-pull arrangement). Another two pairs of push-pull input units coded for version velocity between -250 and $250^\circ/\text{s}$ and vergence velocity between -10 and $10^\circ/\text{s}$, respectively. Finally, two push-pull units encoded translation velocity of the eye between -250 and 250 cm/s ; another unit pair represented the integrated velocity between -50 and 50 cm (translational path) of the eyes. The output layer was modeled corresponding to the input map. All units in the network were fully connected, with each input unit connected to all hidden units, and each hidden unit connected to all output units. The hidden layer had recurrent connections to enable the network to remember past events. Both the hidden layer units and the output neurons were characterized by a logarithmic sigmoid activation function of the form $A(x) = 1/(1+\exp(-x))$. We simulated a trial as a series of 11 consecutive time steps, with each step defined as a 200 ms interval. We tested the network with different numbers of units (25, 50 and 100) in the hidden layer. Each type of network was trained four times with random initial weights to validate reproducibility of behavior. The analysis presented in this paper was performed with 50 hidden units.

During training, targets were presented at one of five locations in space, at 25, 29, 35, 42, and 52 cm in front of the subject, when viewing them from straight ahead (translation position 0). The other translational positions of the eyes at the start of the trial were 5, 10, 15, and 18 cm to the left or right from position 0. The binocular point of fixation was at the location of either the 25, or 35, or 52 cm target. The simulated translational motion was 0 (no translation), ± 10 , ± 20 , ± 30 , and $\pm 36 \text{ cm}$. To simulate trial conditions with only rotational motion of the eyes (without translational motion), the fixation spot was moved by either 0, 5, 10, 15, or 18 cm to the left or right. Targets were presented for one time step, i.e. 200 ms, at the onset of a trial. Translation of the subject, or translation of binocular fixation point, which followed a bell-shaped velocity profile, was initiated 400 ms after the target disappeared, and lasted for 1 s. The network's output, the direction and disparity of the target in eye-centered coordinates, was read at the final time step of the trial. Trial types which moved the horizontal target direction $> 20^\circ$ in the output map were excluded to minimize edge effects at the boundaries of the workspace. Together, this led to 1129 different types of trials in the training set.

Network testing included all combinations that can be comprised with the binocular fixation position at 33 cm, targets presented at either 27, 35, or 48 cm, the translational offset of the eyes either -16, -6, 0, 3 or 9 cm, translation motion of 25, 12, 8, 0, and 14 cm, and movements of the fixation point of -13, -7, 0, 4 and 15 cm. The network was built, trained, and tested using the Matlab Neural Network Toolbox, with a training function that updates weight and bias values according to gradient descent momentum and an adaptive learning rate. For training, individual weights were initially set to random values between -0.1 and +0.1.

Results

We exploited the geometry of motion parallax to address the question whether the location of a space-fixed target, briefly presented before an intervening whole-body translation, is stored and updated in a gaze-dependent or gaze-independent coordinate frame (see Figure 2.1). A gaze-dependent coding predicts that if the translation is not correctly taken into account, the updated locations will deviate from the actual locations, with updating errors in opposite directions for targets in front and behind the FP. Alternatively, updating within a gaze-independent framework requires the readouts of the memories of such targets after the body translation to be affected by errors in the same direction. We tested between these hypotheses using memory-guided reaching movements in stationary and translation trials.

Task performance

Twelve subjects participated in the main experiment, outlined in Figure 2.2A. Using the stationary trials, we first tested the ability of stationary subjects to look and reach to memorized locations of space-fixed targets flashed at different distances from the fixation point. The left column of Figure 2.2B shows the performance of a typical subject over the time course of sixteen trials, either when standing at the leftward position (black traces) or at the rightward position (gray traces) within the translation zone, with a target that was flashed 17 cm behind the eyes' fixation point (T1, see Figure 2.2A). The top panel depicts the horizontal component of the subject's body position during the entire trial. Both within and across trials this position remained constant, as instructed, also during reaching, at about 17 cm left or right of the center of the translation zone. The second panel displays binocular gaze direction, superimposed on the average signals for ideal performance (dotted lines), which were computed on the basis of the Optotrak data. Binocular gaze showed steady fixation when the target was presented and during the memory interval (as required to meet the

3° accuracy range of the trial inclusion criteria, see Methods), and small saccades at the time of pointing. These saccades direct the eyes toward the finger tip, which is to point at the remembered location of the stimulus flash. The third panel shows a similar pattern for binocular fixation depth (in degrees, as indicated by the vergence component of the eye positions). The decline in vergence during the reach seems to match the requirements (dotted lines) to look at the remembered location of the flash, which is farther away than the fixation point. Finally, the bottom panel demonstrates the horizontal position of the finger tip (in cm), showing that the subject reached fairly accurately to the remembered location of the stimulus flash, with errors < 3 cm. These few trials are exemplary for the performance of all subjects in the stationary trials, showing that they can localize a non-foveated flashed target fairly well.

The question is how well are these subjects able to localize these flashed targets when they have translated after viewing the flash? This was tested using the translation task. Recall that a whole-body translation effectively disturbs the spatial registry of the location of the flash relative to any reference frame attached to the body. Hence, in any egocentric reference frame, whether gaze-dependent or gaze-independent, the location of the reach goal after the translation is different from the location of the flash before the translation.

The right column of Figure 2.2*B* shows the typical performance of the same subject over the time course of sixteen translation trials, in which the translation was either rightward (gray traces) or leftward (black traces). As in the stationary examples, illustrated on the left-hand side, the target for updating was T1, flashed 17 cm behind the fixation point. As instructed, the subject only began moving after the target had flashed, and reached his final position before FP offset (upper panel). Kinematics of the self-induced translation were highly reproducible across trials, with a mean (\pm SD) displacement of 32 ± 2 cm. During the translation, changes in binocular fixation direction and depth matched the geometrically required modulations (dotted lines) to keep gaze fixed at FP quite well (second and third panel). In other words, the body translation had negligible influence on the ability to keep fixation at a lit fixation target. In accordance with the instructions, the changes of these signals during the reach period indicate a change in the binocular fixation point toward the remembered location of the target. The accuracy of the respective reaching movement reflects the accuracy of the spatial memory update, as well as the perceptual and motor deficits involved. The reaching movements here show clearly larger errors than in the stationary condition, ranging up to about 7 cm.

To demonstrate the differences in performance in both tasks more clearly, Figure 2.3 compares the reach endpoints in the stationary (left column) and the translation task (right column), in separate top-view panels for the four targets, ordered by their




Figure 2.3. Reaching positions (circles) of one subject in the stationary (left column) and translation task (right column). Data from subject S1. Data presented in separate top-view panels for the four targets (squares), ordered by their location from FP. Errors in the translation trials appear to depend on the direction of the intervened translation and on the depth of the target from fixation, which is most consistent with the predictions of the gaze-dependent updating model.

location from FP, for one subject. In both conditions, a general underestimation of target distance seems to be present. In the stationary task, errors are only small, with a slight dependence on the subject's body position. Undeniably, errors in the translation trials exceed those in the stationary trials, irrespective of step direction. Both size and horizontal direction of this error seem to depend on the direction of the intervened translation and on the location of the target. For rightward translations, the subject reached too far to the right for the farthest target, while there was a leftward bias for the nearest target. The opposite pattern is observed for a leftward translation. There is also a tendency for errors to increase for the targets flashed at farther distances from the fixation point, despite the same amount of intervened translation. Thus, for this one subject, the pattern of errors in the translation trials seem to follow the prediction by the gaze-dependent updating model: pointing positions deviate in opposite directions for targets in front and behind the FP, with a nearly mirror-symmetric pattern of errors for leftward and rightward translations.

Error analysis

To analyze these findings quantitatively, we assumed that the reach errors in the static trials reflect a sensorimotor deficit while the reach errors in the translation trials reflect sensorimotor deficits as well as deficits in the spatial memory update (see Methods). Therefore, to compute the latter, i.e., the updating errors, we subtracted the mean horizontal reach error observed in the static trials from the horizontal reach errors that occur in the translation trials, for each target separately. Figure 2.4A plots these horizontal updating errors for targets behind FP versus the errors for their corresponding equiangular counterparts in front of FP, for each translation direction. Thus, updating errors of target T1 were plotted versus the updating errors of target T4, and errors from target T2 with target T3. This pair-wise comparison was performed by picking, without return, the errors randomly from the respective trials, yielding a maximum of 80 data points. The gaze-dependent updating hypothesis predicts that these errors have equal size but opposite signs (Figure 2.1). Accordingly, data points should fall in the even quadrants, ideally along the dashed line with slope -1. In contrast, the gaze-independent updating hypothesis predicts that these errors have equal size and signs, which would be indicated by data points along the positive diagonal (slope +1). Any other slope values, whether 0 (the data scatter around the x-axis), infinity (the data scatter about the y-axis) or any other value reflect a measure intermediate these two models. To deal with this in further analysis, we converted all slope values to a reference frame index (RFI) between -1 (perfect gaze-dependent coding) and +1 (perfect gaze-independent coding). For example, slopes of +2 and -2 correspond to a reference frame index of 0.5 and -0.5 respectively. Figure 2.4A presents




Figure 2.4. Reach errors to targets at opposite, but equiangular distances from fixation, plotted versus each other. Data would fall along the negative diagonal if subjects had updated remembered target locations in a gaze-dependent frame (reference frame index, $\text{RFI} = -1$). Data would scatter along the positive diagonal if subjects had employed a gaze-independent updating mechanism ($\text{RFI} = 1$). **A.** Subject (S1) favoring the gaze-dependent model. The best-fit line (in gray) which characterizes the distribution of the data points, has a clear orientation to the negative diagonal. (gray dashed lines: $\pm 95\%$ confidence intervals), **B.** Best-fit lines from all subjects. **C.** RFI values (with bootstrap confidence intervals) from all subjects, with $\text{RFI} -1$ supporting the gaze-dependent model and $\text{RFI} +1$ the gaze-independent model. Subjects typically support the gaze-dependent updating model.

the results of this analysis for the same subject as in Figure 2.3, showing that the majority of the data points fall in the even quadrants. According to a Model II regression, the best-fit line that characterized the direction of the data point clustering was closely directed along the line with slope -1 . The reference frame index of this subject had a value of -0.93 ± 0.06 (mean \pm SD), which is illustrative for a data distribution that best supports the gaze-dependent updating model. The best-fit lines

of all twelve subjects are superimposed in Figure 2.4*B*, generally indicating an orientation in the direction predicted by the gaze-dependent model. Figure 2.4*C* summarizes the corresponding reference frame indices (\pm SD) for all subjects (black bars), showing a clear bias toward the gaze-dependent model. Averaged across subjects, the reference frame index was -0.68 ± 0.23 , which was significantly different from zero (t-test, $p < 0.05$), indicating that our data is most supportive for a gaze-centered coding and updating of spatial memory.

For completeness, Table 2.1 provides further statistical information about the data distribution of each subject, showing the mean correlation coefficient (r), reference frame index (RFI), and a variance ratio (VR), defined as the ratio between variance of the data along the main axis of the distribution and the variance in the direction orthogonal to it.

Table 2.1. Results of the horizontal error analysis in each subject. r , correlation coefficient of Model II regression; RFI, mean reference frame index; VR, variance ratio. All values, bootstrap estimates (mean \pm SD). N , number of data points.

Subject	r	RFI	VR	N
S1	-0.65 ± 0.05	-0.93 ± 0.06	4.17 ± 1.22	80
S2	-0.49 ± 0.06	-0.89 ± 0.08	3.03 ± 1.29	78
S3	-0.22 ± 0.04	-0.26 ± 0.28	1.61 ± 0.91	75
S4	-0.35 ± 0.08	-0.76 ± 0.17	2.04 ± 0.79	63
S5	-0.12 ± 0.09	-0.42 ± 0.32	1.56 ± 1.22	61
S6	-0.41 ± 0.08	-0.78 ± 0.13	2.33 ± 0.76	69
S7	-0.50 ± 0.14	-0.83 ± 0.13	2.30 ± 0.75	53
S8	-0.45 ± 0.04	-0.84 ± 0.11	2.56 ± 0.99	77
S9	-0.43 ± 0.08	-0.81 ± 0.13	2.37 ± 0.81	63
S10	-0.20 ± 0.08	-0.49 ± 0.30	1.59 ± 1.13	64
S11	-0.20 ± 0.06	-0.38 ± 0.31	1.56 ± 1.17	80
S12	-0.30 ± 0.10	-0.76 ± 0.20	1.85 ± 0.62	59
Mean \pm SD	-0.37 ± 0.19	-0.68 ± 0.23	2.25 ± 0.73	12

Vectorial analysis

Although the data of most of our subjects lend support for the gaze-dependent updating hypothesis, it should be pointed out that this conclusion is based on an (1-D) analysis of the horizontal reach errors. Since subjects also make updating errors in depth (see Figure 2.2), it is desirable to validate this conclusion in a 2-D analysis. Therefore, we investigated how the position of the target before the translation (\vec{T}_i , estimated by the average response in the stationary task), the position of the same target after the translation (\vec{T}_f), the actual translational motion ($\vec{T}_f - \vec{T}_i$) and reach response (\vec{R}), expressed as Cartesian 2-D vectors, are related in the coordinate frames of the two updating models (see Figure 2.5A). The two coordinate axes of the gaze-dependent model were chosen to be aligned with and orthogonal to the gaze line, respectively, with the origin at the center of the two eyes (cyclopean eye). At the same origin, the coordinate axes of the gaze-independent model were arranged to be aligned with and orthogonal to the shoulder line, respectively. Note that the same (space-fixed) target \vec{T}_i in this example is described by quite different vectors in each coordinate system. In both coordinate frames, the following updating relationship can be specified,

$$\vec{R} - \vec{T}_i = a \cdot (\vec{T}_f - \vec{T}_i) + \vec{b} \quad (\text{Eq. 2.1})$$

in which $\vec{T}_f - \vec{T}_i$ represents the ideal updating vector, $\vec{R} - \vec{T}_i$ the actual updating vector, fit parameter a the updating gain, and vector \vec{b} the bias in the updating process. If a subject had a correct percept of \vec{T}_i , but did not account for the intervening translation, reach vector \vec{R} would be equivalent to target vector \vec{T}_i , and hence the internal updating vector $\vec{R} - \vec{T}_i$ would equal zero, thus $a = 0$, $\vec{b} = \vec{0}$. In contrast, if translational updating were flawless, \vec{R} would be identical to the new target vector \vec{T}_f , and thus $a = 1$, and $\vec{b} = \vec{0}$.

We fitted Eq. 2.1 in terms of the predicted updating error \vec{E} (see dashed gray vector in Figure 2.5A). The results of this analysis are shown in Figure 2.5B for one subject, for the rightward translation trials. The actual average endpoints (left) are compared with those predicted by each of the two models, on basis of the fit parameters of Eq. 2.1. Close scrutiny indicates that the predictions of gaze-dependent model (middle) better match the observed reach endpoints than the gaze-independent model (right). The gaze-dependent model seems to capture the observed pattern of opposite errors for targets behind and in front of the fixation point whereas the gaze-independent model shows only a small rightward shift of each of the reach endpoints. On a population level (Figure 2.5C), Eq. 2.1 gave a better description (higher correlation coefficients) of the updating errors when expressed in gaze-dependent coordinates than in gaze-independent coordinates (t-test, $p < 0.01$), which is consistent




Figure 2.5. 2-D vectorial analysis of updating performance. **A.** Target location before (\vec{T}_i) and after (\vec{T}_f) the body translation, and reach location (\vec{R}) expressed as 2-D vectors in a coordinate frames fixed to either the gaze line, G, (left) or the line perpendicular to the shoulder line, B, (right). Eq.1 was fitted in the predicted updating error \vec{E} , in both coordinate frames, with $(\vec{R} - \vec{T}_i)$ representing the terms of actual and $(\vec{T}_f - \vec{T}_i)$ the ideal amount of updating. **B.** Actual reach endpoints (open circles) of one subject (subject S1) flanked by reach endpoints based on model fits (filled circles) for the four targets (squares) for rightward translation trials. The gaze-dependent model predicts the actual pattern of endpoints best. **C.** Correlation coefficient for the fit in gaze-dependent versus in gaze-independent coordinates, for all subjects. The gaze-dependent model made the best description of the data in nine out of twelve subjects. Two subjects showed very low correlations for both models.

with the 1-D analysis described above. Within individual subjects, the gaze-dependent model produced the best description for nine out of twelve subjects. The gaze-independent model performed slightly better in three subjects, although its performance remained rather low in two of them (see Table 2.2). Table 2.2 lists the best-fit coefficients of Eq. 2.1, showing the updating gain, a , and bias vector, \vec{b} , for both models, for each subject separately. Across the population, the bias vector was not significantly different from a zero vector (t-test, $p > 0.05$ for all components), for both of the two models. In the gaze-dependent model, the updating gain, a , specifies how well the translational-depth geometry is taken into account in the updating of remembered visual space. Averaged across subjects, its value was 1.16 ± 0.15 (SD), which was significantly different from 1 (t-test, $p < 0.05$). This suggests that this model takes the systematic reach errors into account in fitting the data, or in other words, that subjects generally overestimated the amount of self-motion when updating targets in 3-D space during active whole-body translations. In contrast, the gaze-independent model yielded

Table 2.2. Fit performance of Eq. 2.1 ($\vec{R} - \vec{T}_i = a(\vec{T}_f - \vec{T}_i) + \vec{b}$) in gaze-dependent and gaze-independent coordinates, in each subject. r , correlation coefficients (values also shown in Figure 2.5C). Best-fit values of a and \vec{b} refer to updating gain and bias vector, respectively.

Subject	Gaze-Dependent Model			Gaze Independent Model		
	r	a	\vec{b} (cm)	r	a	\vec{b} (cm)
S1	0.83	1.38	[-0.84 1.90]	0.06	0.99	[-0.48 2.34]
S2	0.65	1.26	[0.15 0.87]	0.32	0.97	[0.46 1.52]
S3	0.47	1.28	[0.06 1.00]	0.61	0.95	[0.29 2.37]
S4	0.58	1.11	[0.73 -0.88]	0.07	1.00	[0.87 -0.88]
S5	0.45	0.89	[1.09 0.08]	0.07	1.00	[0.77 -0.20]
S6	0.45	1.12	[0.28 1.10]	0.25	1.02	[0.22 0.69]
S7	0.65	1.19	[0.31 2.13]	0.04	1.00	[0.98 2.24]
S8	0.93	1.10	[0.22 1.88]	0.4	1.06	[0.21 1.05]
S9	0.74	1.37	[0.60 2.39]	0.35	0.95	[0.75 2.10]
S10	0.09	1.03	[0.36 -1.25]	0.33	0.97	[0.66 -0.96]
S11	0.01	1.00	[-0.02 0.57]	0.09	0.99	[0.15 0.65]
S12	0.53	1.15	[-0.26 1.27]	0.31	1.03	[-0.13 1.31]
Mean ± SD	0.59 ±0.42	1.16 ±0.15	[0.23 ±0.49 0.92 ±1.15]	0.25 ±0.20	1.00 ±0.03	[0.39 ±0.44 1.02 ±1.20]

an average updating gain that was statistically not distinguishable from 1 (t-test, $p=0.62$), which essentially indicates that this model has no provision to account for the systematic errors observed in the data.

Control experiments

To determine the robustness of these findings we performed three control experiments (see Methods). The task designs of these controls were kept identical to that of main experiment as much as possible. In the analysis of these experiments, each performed on five subjects, we focused on the horizontal reaching errors, investigating the relationship between the errors for targets in front of FP and errors to targets behind FP. As above (see Figure 2.4), a negative relationship would confirm gaze-dependent coding (ideal slope -1); a positive relationship would be suggestive of a gaze-independent coding scheme (ideal slope +1). We first asked whether the same results would be obtained if the reaching movement toward the updated target locations were not accompanied by any visual feedback about hand position movement (control I: reaching without feedback). The results show that the absence of hand feedback does not alter our main conclusion. All subjects performing the task without hand feedback produced data consistent with the gaze-centered updating hypothesis (see Figure 2.6A). This is reflected by the average reference frame index, which was -0.70 (SD=0.18) and significantly different from a value of 0 (t-test, $p>0.05$).

Next, we investigated if the effects were mainly specific to moving the eyes to the updated target locations, rather than to moving the hand (control II: reaching without looking). For eye movements, the sensory frame of reference imposed by the retina and oculomotor reference frames for the eyes are quite similar (Snyder 2000). Hence, for the eyes to look at the remembered target locations, saccadic amplitude must depend nonlinearly on target depth and direction. If saccadic amplitude was not scaled appropriately (Medendorp et al. 2003b), and the eyes lead the arm, the errors that appeared reflected an eye-centered motor representation, rather than information about the spatial representation that codes the target. Arm movements do not suffer from this drawback: the sensory frame of the retina is quite distinct from the motor frame of reference imposed by the joints and muscles of the arm (Snyder 2000). Therefore, in this control experiment, subjects were instructed to keep gaze fixed at FP at all times during the trials, also when probing the remembered target by the reach. These results show once again clear evidence for the gaze-dependent coding scheme (Figure 2.6B). All subjects had reference frame indices significantly smaller than zero. Moreover, the average RFI across subjects was -0.87 (SD=0.11), which was significantly different from zero (t-test, $p<0.05$) but not from -1 (t-test, $p>0.05$).




Figure 2.6. Results of three control experiments, each performed on five subjects. Reference frame indices according to Fig 2.4C: +1 reflects a gaze-independent scheme; -1 the gaze-dependent scheme. **A.** Control I: Reaching without visual feedback of the fingertip. All subjects support the gaze-dependent model. **B.** Control II: Reaching without looking at the finger tip provides unanimous support for the gaze-dependent model. **C.** Control III: updating without FP. Subject translated their body keeping gaze fixed on a remembered fixation point. Although the reference frame test may be less discriminative due the vergence drift, clear support for the gaze-dependent model can be seen in most subjects. Error bars, bootstrap confidence intervals.

Finally, we asked whether the visual FP, available during the main experiments, was a biasing factor for the gaze-centered updating hypothesis. To test this, we conducted an experiment in which subjects had to keep their eyes fixated on the remembered FP during the self-motion and then looked and reached to the remembered location of the flashed target (control III: updating without FP, Figure 2.6C). It is important to realize that in this situation, our test has less discriminative capabilities. Because of possible vergence drift caused by the absence of a visual FP during translation in this paradigm, updating vectors in gaze-coordinate will not be of equal size for targets in front and behind FP (compare Figure 2.1). In spite of that, across the five subjects that participated here, three followed the gaze-dependent model. The RFIs in the other two subjects had values around zero. Averaged across subjects, we found a RFI of -0.48 ($SD=0.47$) – a clear bias in favor of the gaze-dependent updating model.

Taken together, the results of all our experiments lead to the conclusion that the brain uses a gaze-dependent reference frame to store and update visuospatial memories during self-generated whole-body translations.

Discussion

Inspired by the work of Von Helmholtz, investigators have made abundantly clear over the last decades that humans can remember visual direction across rotary eye and head

movements (Von Helmholtz 1867; Hallet and Lightstone 1976; Herter and Guitton 1998; Blouin et al. 1998; Schlag et al. 1990; Medendorp et al. 2002; Wexler 2005). Since Gibson, vision scientists have also become aware of the complexity of motion parallax for seeing in depth when the eyes translate through space (Gibson et al. 1955; Rogers and Graham 1979). Here, we have exploited a paradigm based on the conjunction of these two challenges for visual stability, testing if the brain internally simulates motion parallax when updating remembered visual space during active whole-body translations. We called this the gaze-dependent hypothesis, as it predicts a systematic pattern of updating errors depending on gaze fixation if the intervening translation is not correctly taken into account, with the errors reversing in direction for targets at opposite depths from gaze fixation. As a contrasting hypothesis we set up the predictions of a gaze-independent coding scheme. According to this hypothesis, the brain codes remembered space irrespective of gaze fixation, and therefore predicts no such reversal of updating errors if translations are misjudged. We emphasize that the central premise behind our test was that subjects misestimate their traveled distance during self-generation motion, as shown by many studies (Medendorp et al. 1999; Israel 1993; Glasauer et al. 1994; Philbeck and Loomis 1997; Kudoh 2005), although the exact explanation for why this occurs is not directly relevant (but see later). Our results show that target updating for translational motion is compromised by small errors, which increase with depth from fixation and reverse in direction for opposite depths from fixation. This is consistent with the gaze-dependent prediction, so we conclude that the brain employs a gaze-centered mechanism to internally update remembered visual space during whole-body translations.

We will now list a number of observations that further support this conclusion. First, reaching errors were larger in translation trials (with intervening body translation) than in the stationary trials (without body translation), suggesting that the differences indeed arose during the updating of spatial information (Figure 2.3). Second, a quantitative analysis of these errors revealed that they were opposite for targets in front of and behind FP (Figure 2.4). Third, a two-dimensional vectorial analysis of the translational-depth geometry in the transversal plane showed that the interaction between target location, translational motion and reaching response is much better described in a gaze-centered than in a gaze-independent coordinate system (Figure 2.5). Fourth, the gaze-centered updating errors were quite robust and invariable among various task constraints (Figure 2.6). More specifically, the same error pattern was found irrespective of whether the eyes and hand moved to the memorized target location or the hand alone. Neither did the pattern of errors change when subjects performed the reaching movement with or without visual feedback of hand position.

Even the presence or absence of a visual fixation point during the translations was not essential for a gaze-centered description of updating errors.

Although our data provides support for the gaze-dependent model across subjects, it is important not to overstate this. The results are not perfect, and our conclusions follow from relatively small systematic errors. As a matter of fact, three of our subjects did not show support for the gaze-dependent hypothesis in all conditions and analyses (see Figures 2.5C and 2.6C). It is also important to note that our test was based on relative simple geometry, whereas the brain may actually represent visual space in a more complex manner (Cuijpers et al. 2002). Furthermore, we should emphasize that we have focused on only one important signal, the central representation of body translation, as an underlying basis for the updating errors, which is but one of a myriad of variables which might lead to errors. In this respect, further experiments are needed to isolate the various signals related to overall performance of the present task. Nevertheless, despite these reservations, we think that our behavioral tests provide evidence that the brain possesses a geometrically complete, dynamic map of remembered space, whose spatial accuracy is maintained by internally simulating motion parallax during volitional translational body movements.

It is true that even when you walk around normally in the environment, it is difficult to experience motion parallax, even if you try (Palmer 1999). And without doubt it is even harder to imagine motion parallax with locations of remembered objects, or objects that are out of view. Nevertheless, this cannot be taken to imply that the neural mechanism for spatial coding cannot act by simulating the parallax geometry to maintain spatial constancy, as we have shown here.

Recently, various studies have shown that both human and nonhuman primates can adjust the amplitude of memory-guided eye movements after intervening translation, taking into account the amount of translation and distance of the memorized target (Li et al. 2005; Li and Angelaki 2005; Medendorp et al. 2003b; Israel and Berthoz 1989). None of these studies, however, explicitly assessed the exact nature of the representation of remembered visual space during these tasks. Here, for the first time, we were able to establish that targets in such tasks are stored in a gaze-centered reference frame, an inference based on the assessment of the operational errors in the system.

Our evidence for gaze-centered updating during translational motion agrees well with recent studies showing gaze-centered updating for rotational motion (Henriques et al. 1998; Medendorp and Crawford 2002; Pouget et al. 2002; Baker et al. 2003). The first three showed that subjects overshoot the direction of a previously seen, but foveally-viewed target when reaching toward it after an intervening eye rotation. Interestingly, here we show a similar type of overshoot for translation-induced changes

of gaze, corroborating these gaze-centered results. Baker et al. (2003) investigated updating behavior during horizontal whole-body rotations using a memory-guided saccade task. Based on the assumption of noise propagation at various processing stages in the brain, they found their results most consistent with a gaze-centered representational system for storing the spatial locations of memorized objects.

Which signals are needed in the updating process? In the present study, the updating mechanism may have received information about the self-motion through efference copy and proprioceptive signals (available in the context of active motion), and by vestibular inputs (Medendorp et al. 2003b; Klier et al. 2005; Li et al. 2005; Li and Angelaki 2005; Chapter 4 of this thesis). Li et al (2005) found updating during passive translation to be compromised after bilateral labyrinthectomy, attributing an important role of the vestibular system. Also, Israel and Berthoz (1989) have provided evidence for spatial updating with the vestibular system as the main extraretinal source of motion-related information. Furthermore, in the present study, the changes in eye position to keep the eyes fixed at FP during the translation – the version and vergence eye movements – are essential for a well-functioning updating system. All of this information must be must be integrated at a central level within the brain and unified with retinal information about target direction and depth to mediate the computations for gaze-centered spatial updating, as outlined in detail in Medendorp et al. (2003b).

In line with our findings, many brain regions have been demonstrated to store and update target locations within an eye-fixed, gaze-centered reference frame (Gnadt and Andersen 1988; Duhamel et al. 1992; Batista et al. 1999; Medendorp et al. 2003a; Merriam et al. 2003; Sommer and Wurtz 2002). However, the majority of these studies have focused on directional updating of target location in the frontal parallel plane. For example, the lateral intraparietal area and superior colliculus have been shown to update its retinotopic map of target directions for each eye movement (Duhamel et al. 1992; Walker et al. 1995). On the other hand, it is also known that the activity of LIP neurons is modulated by retinal disparity information, providing them with three-dimensional receptive fields (Gnadt and Mays 1995; Genovesio and Ferraina 2004). Moreover, Cumming and DeAngelis (2001) indicated that the updating of target distance may be expressed by changes in retinal disparity representations.

To obtain further insights in the interactions between self-motion information and retinal signals at the level of the parietal cortex we designed a simple recurrent neural network performing gaze-centered target updating during translations and rotations (see Figure 2.7A, and Methods). The input to the network was a transient distributed representation of target direction and disparity in a 2-D retinotopic map (as a hill of activity), as well as a variety of extraretinal signals, including angular gaze




Figure 2.7. A. Diagram of the three-layer recurrent network model, trained to perform gaze-centered updating during both eye rotations and translations. Inputs: target location as a two-dimensional Gaussian hill of activity within a 11x11 units horizontal-disparity (retinal eccentricity-depth) map and extraretinal signals, including gaze position (version/vergence), gaze velocity signals (version/vergence), and the eyes' translation velocity and path signals. Hidden layer contained 25, 50 or 100 units. Output layer encodes the memory of the target in terms of direction and disparity relative to the binocular fixation point, i.e., in gaze-centered coordinates. **B.** Performance of the network ($n=50$) when particular input signals are removed. Updating error in the direction of the target (in deg) is shown for the intact network, the network with gaze position inputs removed, with gaze velocity inputs eliminated, and with translational inputs removed. The network has a strong preference for gaze velocity inputs over gaze position inputs. Similar results were obtained for the networks trained with 25 or 100 hidden units. Error bars denote SD for 4 networks.

position and velocity signals (version / vergence), and translational velocity and path signals of the eyes. The network was trained to store the memory of the target for successive time intervals and update its representation for any intervening rotational or translational eye motion..

Figure 2.8 shows the simulation results for updating a target in front of ('near' target) and target behind ('far' target) the eyes' fixation point during a translational motion of the eyes. The extraretinal signals involved are the same in both situations (Figure 2.8A), for which the geometrical relationships are depicted in Figure 2.8B. As shown in Figure 2.8C, the near target appeared at 8° to the right of the center of gaze, at -2° disparity, and shifted to an 8° deg leftward, -2° disparity position after the rightward translational motion. The activity pattern of the far target evolves in the opposite direction of the map during the translation (Figure 2.8D). In other words, the updating network must have used information about target depth in order to determine how the hill of activity should move over the map. The exact location of the target was decoded from the map by means of a weighted average of the activity of all neurons (see open circles in the bottom panels of Figures 2.8C, D), which closely follow the geometrically-required changes for ideal updating over time (thin lines). Likewise, the network also incorporated the geometrically-required properties of updating targets in the same direction on the map, irrespective of their depth, when the eyes rotate only (not shown). Using 25 neurons in the hidden layer was already sufficient to learn the task acceptably, but performance improved for the 50 and 100 hidden units networks.

Since the network was trained to perform these tasks under the provision of extraretinal position and velocity signals, an interesting question to ask is whether one input is more relied on than another (White and Snyder 2004). To this end, we removed one of the inputs after training ('artificial lesion') and looked at the performance of the network in terms of its updating errors (Figure 2.7B). As the figure shows, the network has a clear preference for gaze velocity over gaze position inputs, which is consistent with findings by White and Snyder (2004) for rotational updating. The use of velocity signals may give the network a benefit to update continuously, irrespective of initial or final gaze position. Thus, our simulation results provide good evidence for the idea that the brain synthesizes ego-velocity signals and stereoscopic depth and direction information to update the internal representation of 3-D space during self-motion. This integration may occur in parietal area LIP using the computations that we have described, or in any other cortical or subcortical structures involved in updating, as long as they have the necessary signals at their disposal.




Figure 2.8. Network performance for target updating during a rightward translation trial. **A.** Activation of the units representing the eyes' rotational and translational kinematics at each time step. **B.** Geometry that has been simulated. **C.** Updating of a target flashed in front of the eyes' fixation point. The hill of activity coding the target memory shifts across the horizontal-disparity map. The bottom panels show the target representation encoded by the output layer, showing that a near target shifts from right to left relative to the gaze line (open circles), as geometrically required (lines). The network also matches the required changes in disparity. **D.** A target farther away than the fixation point shifts in the opposite direction on the map. Thus, activity patterns evolve during translation in a way that depends on target depth.

Chapter 3

Updating target distance across eye movements
in depth

This chapter has been published as:

Van Pelt S, and Medendorp WP. Updating target distance across eye movements in depth. *Journal of Neurophysiology* 99: 2281-2290, 2008.

Maintaining spatial constancy across self-generated movements is crucial for veridical perception of the world and for accurate control of goal-directed actions. Over the past few decades, the quality of spatial constancy has been investigated systematically across various types of self-motion, including eye, head and body movements. As a result, it is now well established that spatial constancy is preserved across intervening saccadic (Hallet and Lightstone 1976; Sparks and Mays 1983) and smooth pursuit eye movements (Schlag et al. 1990; Baker et al. 2003). Also a reorientation of the head or displacement of the body does not compromise spatial stability to a great extent (Mergner et al. 2001; Li and Angelaki 2005; Medendorp et al. 1999, 2002, 2003b; Israel et al. 1993; Chapters 2 and 4 of this thesis).

From a mechanistic perspective, there has been considerable debate over how the brain solves the spatial constancy problem. In the absence of allocentric cues, it seems that an egocentric, gaze-centered reference frame dominates in the mechanisms of spatial stability for simple saccade or reaching tasks (Henriques et al. 1998; Medendorp and Crawford 2002; Klier et al. 2005; Chapter 2 of this thesis). In support, cells in monkey extrastriate visual areas (Nakamura and Colby 2002), posterior parietal cortex (Duhamel et al. 1992; Batista et al. 1999; Colby and Goldberg 1999), frontal cortex (Goldberg and Bruce 1990) and superior colliculus (Walker et al. 1995), as well as in the human posterior parietal cortex (Medendorp et al. 2003a; Merriam et al. 2003, 2007) have been shown to update the gaze-centered coordinates of remembered stimuli to maintain an accurate representation of visual space across saccades.

Notwithstanding these convincing observations and clear insights, it should be emphasized that nearly all these studies were limited by only examining the directional aspect of spatial constancy. For many spatially-guided actions, however, directional constancy is not the only spatial requirement; the constancy of target depth (or distance) is another essential component that should be mediated by the signals and mechanisms for spatial stability.

Since it is generally assumed that target depth and direction are processed in functionally distinct visuomotor channels (Flanders et al. 1992; DeAngelis 2000; Cumming and DeAngelis 2001; Vindras et al. 2005), the mechanisms to preserve their constancy may also operate independently, at least to some extent. To date, only few studies have explicitly assessed the constancy of target depth during self-motion (Krommenhoek and Van Gisbergen 1994; Medendorp et al. 1999, 2003b; Philbeck and Loomis 1997; Li and Angelaki 2005). Krommenhoek and Van Gisbergen (1994) showed that subjects can look at a remembered position of a target in depth after a vergence eye movement. Li and Angelaki (2005) reported that nonhuman primates can keep track of changes in the distance of nearby objects when their body moved toward or away from them. Despite these quantitative observations, the computational

mechanisms underlying depth constancy have not been addressed. The objective of the present study is to fill this lacuna by testing between two models for depth coding in the visuomotor system.

While a variety of cues to depth can be used by the visual system, binocular disparity dominates in the creation of a cohesive, three dimensional depth percept (Julesz, 1971; Howard and Rogers, 1995; Wei et al. 2003). Binocular disparity is caused by the slight difference in viewpoint of the two eyes, due to their differential location in the head. Objects at different distances from the eyes' fixation distance project onto different positions on each retina, and thus cause different horizontal binocular disparities. Likewise, a single object at a fixed position from the eyes will have different horizontal disparities for different viewing distances.

In this study, we investigated how the brain codes the distance of a remembered space-fixed target during intervening changes of the binocular fixation point (i.e., vergence eye movements). We reasoned that if the brain were to encode a binocular disparity representation, i.e., target depth relative to the eyes' fixation point (Shadmehr and Wise 2005), each disjunctive change of gaze will require an active update of this representation to maintain spatial constancy. Alternatively, if the brain were to store a non-retinal depth representation of the target by integrating binocular disparity and vergence signals at the moment of target presentation (Genovesio and Ferraina 2004; Genovesio et al. 2007), this representation should remain stable for subsequent vergence eye movements.

To test between these hypotheses, we employed a memory-guided reach paradigm adopted from Henriques et al. (1998), who originally developed it to examine the computations for directional spatial constancy. We expanded this test by examining actual versus predicted localization errors in depth when vergence eye movements intervene between viewing a target and reaching toward its remembered location, as will be further outlined in Figure 3.1.

The assumption behind our test was that subjects make systematic distance errors in their reach toward memorized targets, depending on their fixation depth (static reaching – Figure 3.1A) – as they have shown to make directional errors depending on their gaze direction (Figure 3.1B, Henriques et al. 1998). In the latter case, the phenomenon has been termed the 'retinal exaggeration effect' since subjects tend to overshoot the target relative to current gaze direction, although individual subjects show considerable variations in this pattern (Bock et al. 1986; Henriques and Crawford 2000). It is not known if a similar overshoot effect occurs for depth; however, as long as distance errors depend on fixation depth, even if only in a complex and idiosyncratic manner, this relationship can be exploited to distinguish between retinal and non-retinal target representations.




Figure 3.1. Testing between nonretinal and retinal models for memory encoding of spatial depth and direction. **A, B.** Static condition: it is supposed that in the reach towards a space-fixed target will be erred depending on gaze fixation position in depth (**A**) and direction (**B**). The exact relationship is not of importance for the test. **C, D.** Dynamic condition: a gaze shift intervenes between target presentation and reaching. The nonretinal model predicts no effect of the gaze shift on reaching. The retinal scheme requires target updating relative to the new gaze position, predicting reach errors as in the static situation with the eyes at the same final position.

The critical part of the test is based on the errors that occur when subjects reach after an intervening eye movement toward the location of a target that was viewed only before this eye movement (dynamic situation, Figure 3.1C and D). Figure 3.1C depicts the situation for the depth dimension, when subjects changed gaze from far to near fixation after initial target perception. Reaching in depth as in the static case without an intervening vergence eye movement (Figure 3.1A, static FP far condition) would argue in favor of the use of a non-retinal depth representation. However, if the intervening eye movement leads to a depth error like that observed when the same target was viewed from the final eye position (Figure 3.1A, static FP near condition), this would provide evidence for the use of an updated eye-centered binocular disparity representation. Following its original design, the test can likewise discriminate between a retinal and a non-retinal representation of target direction across saccadic eye movements (as shown by the panels in Figure 3.1D).

Our results suggest that the brain codes dynamic disparity and direction representations to store target locations for reaching across eye movements in depth and direction. Regression analyses revealed that these representations are modulated by using both eye position and eye displacement signals, consistent with recent observations in monkey neurophysiology (Genovesio et al. 2007).

Methods

Subjects

Fourteen human subjects (four female, ten male; mean age 26 ± 4 years) signed informed consent to participate in this study. All were free of any known sensory, perceptual, or motor disorders. Twelve participants were right-handed; two were left-handed; reaching movements were made using the preferred arm. Two subjects (the authors) were aware of the purpose of the experiments, while the others were naïve.

Experimental setup

Subjects were seated in a completely darkened room, with their torso securely strapped into a custom-made chair by means of two safety belts across both the torso and pelvis to minimize body movement. Their head was mechanically stabilized using a chin rest and a helmet, which was fixed to the chair by means of a frame that was adjustable in height. This ensured that only the preferred arm and the eyes could move, while the rest of the body remained stationary.

The stimulus array (see Figure 3.2B, left panel) consisted of nine LEDs, each 3 mm in diameter, and each could be flashed in two different colors, either as a green or a red light (luminance < 20 mcd/m²). The LEDs had fixed positions on a frame, which could be moved by a robotic arm. Stimuli were presented in front of the subject, in a horizontal plane, slightly below the eyes, at the intersections of three imaginary horopters (equal vergence lines: 8°, 13° and 18° vergence, i.e., 46.5, 28.5, and 20.5 cm from the subjects' eyes) and three equidirection lines (-10°, 0°, and +10° version), based on an average interocular distance of 6.5 cm. The robotic arm was equipped with stepping motors (type Animatics SmartMotors, Servo Systems) and could rapidly move the stimulus array to various positions within the workspace, bringing it within 200 ms out of touch during the reaching task performed by the subject (see below and Van Pelt and Medendorp 2007). During the experiments, the total movement time of the robot was always 2.3s. Also between trials, when the room lights were on, the stimulus array was close to the ceiling of the experimental room. This way, only the frame's rear side could be viewed, which gave no information about the stimuli's spatial configuration.




Figure 3.2. Experimental paradigm. **A.** Sequence of stimuli and the subject's instructions. A trial started with the illumination of a red fixation light (FP1). Then, after a delay of 1.5 s, a green target (T) was cued for 0.5 s. After a further 0.5 s, the subject had to change fixation to fixation light 2 (FP2) in dynamic trials. In static trials, fixation was to be kept at FP1, thus FP2=FP1. Next, 1.5s later, FP2 was extinguished, and an auditory cue instructed the subject to reach toward the remembered location of the target, while keeping fixation at the remembered location of FP2. Open circles: possible stimulus locations. Filled circles, exemplar stimuli presented. **B.** Potential locations of the stimuli, which served either as (initial or final) fixation point or memory target, or both, were on the intersections of three isoversion (-10° , 0° , and 10°) and three isovergence (8° , 13° , and 18°) lines. In dynamic trials, gaze displacements could consist of a pure version movement (second panel), a pure vergence change (third panel), or a combination of both (rightmost panel).

Prior to the experiments, we measured the location of the eyes in space and the locations of the space-fixed stimulus LEDs using an Optotrak Certus system (Northern Digital Inc., Canada). With this information, we were able to compute the direction

and distance of the stimulus LEDs with respect to the subject's eyes. During the experiment, the Optotrak continuously recorded the location of the tip of the index finger. We ensured that the fingertip was at least always visible during the last part of the reaching movement (see Van Pelt and Medendorp 2007). Optotrak data were sampled at 125 Hz with an accuracy of better than 0.2 mm and saved on a PC for offline analysis.

We recorded the subjects' binocular eye movements using an Eyelink II eyetracker (SR Research, Canada) mounted to the chair-fixed helmet. This system tracks the pupils' positions using infrared light reflection at a sampling rate of 250 Hz. Before the experiment began, eye movements were calibrated by fixating the stimulus LEDs three times each, in complete darkness. This resulted in a calibration accuracy $< 0.5^\circ$. Calibration was checked offline, to allow for drift correction due to headband slippage or other factors. Since the head and body stayed fixed during the experiment, the orientation of the eyes within the head, as measured by the tracker, was equivalent to the orientation of the eyes in space (gaze). Rightward rotations were taken as positive.

Two PCs in a master-slave arrangement controlled the experiment. The master PC contained hardware for data acquisition of the Optotrak measurements and visual stimulus control. The slave PC was equipped with hardware and software from the Eyelink system.

Experimental paradigm

The main focus of this study is to reveal the reference frame employed by the brain to maintain spatial constancy for depth. To allow for comparison with previous studies (Henriques et al. 1998; Medendorp and Crawford 2002; Beurze et al. 2006), we employed a paradigm that also tested the mechanisms for directional constancy.

Figure 3.2*A* illustrates the paradigm. A trial started with the onset of a red fixation LED, which we refer to as FP1 (fixation point 1), to be fixated for its entire illumination duration of 2.5 s. FP1 could be any of the nine stimulus locations on the stimulus array (Figure 3.2*B* – leftmost panel). At 1.5 s after the onset of FP1, a target for memory (T, a green LED) was flashed for 0.5 s, while the subject kept gaze fixed at FP1. Thus, T was on the fovea when presented at the same location of FP1, but on the peripheral retina for any of the eight other possible locations. Next, 0.5 s after the flash, a time interval of 1.5 s followed during which the subject either changed gaze fixation to a second illuminated fixation light (FP2 – dynamic paradigm) or maintained fixation of the first fixation point when FP2=FP1 (static paradigm). Subsequently, at FP2 offset (4.0 s after trial onset), the stimulus array was retracted, followed 100 ms later by an auditory signal that cued the subject to reach to T, while keeping gaze fixed

at (the remembered location of) FP2. The subject had to hold the reaching position until the end of a 2.4 s interval, indicated by a second auditory signal. Then the next trial started, with FP1 at a different location than the location of T in the preceding trial, to avoid any visual feedback about performance in the previous trial. Between trials, subjects had their reaching arms resting unencumbered on their lap, with the hand close to their knees. FP1, T and FP2 were pseudo-randomly selected from the stimulus array, such that all combinations of FP1, T and FP2 were tested once. This yielded a total of 729 unique trials: 81 trials were pure static trials (FP1=FP2) and the other trials were dynamic trials. Of the dynamic trials, 162 trials had a pure conjugate change in eye position (version eye movement) while vergence (disjunctive part) remained constant (Figure 3.2*B*, second panel); 162 trials had a vergence change but constant version (ignoring the small vertical version eye movements due to the vertical offset in the positioning of near and far LEDs; third panel); and 324 trials had a combined vergence-version change (rightmost panel).

The total experiment was divided into three sessions, each of which lasted for about 60 minutes each, and were tested on different days. In each session, subjects performed blocks of 15 or 16 consecutive trials, between which a brief rest was provided with the room lights on to avoid dark adaptation. During the experiments, subjects never received feedback about their performance. Before the actual experiments, subjects practiced a few blocks to become familiar with the task.

Data analysis

Data were analyzed off-line using Matlab (The Mathworks). Optotrak data were first transformed to a right-handed Cartesian coordinate system, referenced to the position of the cyclopean eye. In this coordinate system, the positive y-axis pointed leftward along the shoulder-line (from the subject's perspective), the x-axis pointed forward and the z-axis upward.

Horizontal gaze direction was computed for each eye separately; binocular version and vergence angles were calculated from the left (L) and right (R) gaze directions as $(R+L)/2$ and $L-R$, respectively. Rightward rotations were taken as positive. Cartesian positions of FP1, FP2, T and the fingertip were also expressed in binocular coordinates, in terms of depth and direction (in degrees) from the cyclopean eye. This allowed for the computation of reach and target depth relative to the plane of fixation, expressed in terms of angular disparity (Howard and Rogers 1995). By convention, crossed disparities were taken as positive.

We discarded trials in which subjects did not maintain fixation within a $5^\circ \times 4^\circ$ (version \times vergence) interval around the fixation points or made a saccade during target presentation. For the remaining trials, eye fixation accuracy was $2.60 \pm 0.83^\circ$ (mean \pm

SD). We also excluded trials in which subjects had not correctly followed the reaching instructions of the paradigm, i.e. when they started their reaching movement too early or did not adopt a stable reach position during the response intervals (fingertip velocity > 5 cm/s based on Optotrak data). Overall, $< 3\%$ of the trials was discarded on the basis of these arm and eye movement criteria.

The endpoint of each reaching movement was selected at the time at which the velocity of the fingertip first dropped below 5 cm/s within the 2.4-s reaching interval, under the requirement that the arm had correctly followed the instructions of the paradigm. An average position was computed over an 8 sample interval (64 ms) centered at this time point.

Analyses were performed separately for the directional and depth dimensions. We assessed performance by quantifying reach errors in both dimensions, for each trial. Using multiple linear regression analysis, we investigated the effects of eye displacement and eye position on the reach errors that were observed. Statistical tests were performed at the 0.05 level ($p < 0.05$).

Results

The experiments were designed to test between retinal and non-retinal models of the coding of target depth and direction. The basic premise for this test is a difference in accuracy of reaching movements toward remembered space-fixed but non-foveally viewed targets for different eye fixation positions (Henriques et al. 1998; Beurze et al. 2006). The two models make clearly different predictions about the reach errors that would arise when a gaze displacement intervenes between seeing the target and reaching to its remembered location (dynamic condition). The non-retinal model predicts an error similar to that observed in the static condition without the intervening gaze displacement, whereas the retinal model predicts an error similar to that observed for a target viewed from the same final fixation position (see Figure 3.1).

Task performance

Figure 3.3 illustrates task performance of a typical subject, demonstrating static and dynamic trials testing either for depth constancy (left hand panels) or for direction constancy (right hand panels). The left panels show binocular vergence (gray traces) and measured fingertip depth (black traces) over the time course of eight trials, with a target flashed at the middle depth (13° vergence when foveated, dotted line, see also Figure 3.2B). Following instructions, the eyes fixated at either one of the far targets (Figure 3.3A, upper panel – requiring a smaller vergence angle), or at one of the near




Figure 3.3. Typical performance of one subject. Eye position (version, vergence, in gray), target (dotted), fingertip position (in black), expressed in binocular depth and direction coordinates (in deg), plotted as function of time. Dashed trace, geometrically-ideal eye position. Fingertip traces are only shown for the last 1.5s of the reach period. Left panels investigate depth coding of a target at 13° vergence angle, with in **A** two static conditions (far vs. near fixation), and in **C** the corresponding dynamic condition (far to near). Right panels investigate directional coding of a craniotopically central target, with in **B** two static conditions (left vs. right fixation), and in **D** the corresponding dynamic condition. Thin boxes, time intervals of the different trial stages (Target presentation (T), F1 and F2 periods, reach interval (R)).

targets (Figure 3.3A, lower panel – requiring a larger vergence angle), or reoriented from far to near fixation after stimulus presentation (Figure 3.3C). In all conditions, final eye fixation, which was to be maintained during the reach, showed a small decline in vergence after the offset of the fixation point, at the go cue for the reach, but note that reach responses were performed in complete darkness. In the static trials, reaches (black) showed small overshoots (smaller response angle than required) depending on the eyes' fixation depth, with errors of about $-2.2^\circ (\pm 0.8^\circ)$ and $-4.5^\circ (\pm 0.9^\circ)$ for far and near fixation. In the dynamic trials, in which the target is viewed with far fixation and the reach is performed with the eyes fixating near (Figure 3C), the errors seem qualitatively indistinguishable from those in the static near situation (Figure 3A, lower panel), with a mean error of $-4.7^\circ (\pm 0.7^\circ)$. Thus, a change of gaze in depth affects the reaching responses to previously seen targets, making them look like those with gaze stationary at the same final depth.

Figure 3.3B and D show eight typical time courses in each of three trial types serving to illustrate how the directional coding of a craniotopically central target depends on (changes in) gaze direction. Again, in all trials, the eyes act according to instructions, showing steady fixation during target presentation and only small saccades at the moment of reach. Reaching behavior in the dynamic trials, in which gaze changes from 10° leftward to 10° rightward direction (Figure 3.3D), matches more closely the observations made in static trials with gaze in the 10° rightward direction than with gaze at the same direction of stimulus presentation (Figure 3.3B). For the static trials, mean horizontal reach error was $2.5^\circ (\pm 2.8^\circ)$ and $-12.1^\circ (\pm 3.5^\circ)$ in the static left and right conditions, and $-11.5^\circ (\pm 3.4^\circ)$ in the dynamic condition. In other words, the change in eye position has had a marked effect on reaching behavior.

Reach patterns

To demonstrate performance in the static and dynamic conditions more clearly, Figure 3.4 shows spatial plots of reach endpoints (filled circles) for a single subject (RV). The size of each circle represents the corresponding confidence limit (see legend for computation). In the static conditions with far or near fixation (Figure 3.4A), the mean reach endpoints towards the 9 targets are interconnected with gray lines, and superimposed on the spatial structure defined by the stimulus locations (thin black lines). Perfect behavior would require that reach responses align with the stimulus matrix. This is clearly not the case: the subject makes substantial errors with regard to both depth and direction of nearly all target locations. Reaching movements of this subject undershoot the distance of some targets, whereas they are more accurate in others. More importantly, reach patterns in the two static conditions seem noticeably different, depending on the eyes' fixation depth.




Figure 3.4. Reach patterns across all targets from one subject (RV). **A.** Depth coding with reach patterns (in gray) in two static conditions (Far vs. Near), superimposed on the target grid (in thin black). **C, E.** Reach pattern in the corresponding dynamic conditions (in black), superimposed on the patterns (in gray) from **A**, which serve as predictions of the nonretinal and retinal models. **B.** Directional coding with reach patterns in two static conditions superimposed on target grid. **D, F.** Reach patterns in the corresponding dynamic conditions, superimposed on the predicted patterns by the two models. RMSE (in deg), a measure of the deviation of the dynamic data to the model prediction. Each data point represents the mean of three pure trials, with version the free parameter for the depth trials (**A, C, E**) and vergence for the direction trials (**B, D, F**) dimension, respectively. Circle size, SE of each data point.

The question is which of these reach patterns is observed in a dynamic condition in which the subject viewed the target with far fixation, then changed gaze toward near fixation, and subsequently reached toward the remembered target location. If the subject had stored absolute target depth relative to the body (non-retinal model),

computed at the time of seeing the target, the intervening gaze deviation should have no systematic effect on the reach responses. Thus, the non-retinal model predicts reach errors as in the static condition with gaze in far space. However, if the subject had stored target depth relative to the plane of fixation, as a retinal disparity signal (retinal model), this signal must be updated for the gaze change, predicting an error pattern similar to that observed in the static condition with near fixation.

Figure 3.4C shows the systematic reach patterns obtained in the related dynamic trials (thick black lines), superimposed on the predictions of either of the two models. Note that we based this illustration on ‘pure’ trials only, i.e., trials with a change in vergence (far-to-near fixation) but constant version (see Methods), in order to demonstrate the effect of the vergence change in the clearest possible fashion. It is important to realize that using ‘combined’ trials here (trials with a change in both vergence and version, see Methods) could easily obscure the main effects in either dimension. Clearly, for this subject and when gaze was displaced from far to near (Figure 3.4C), the data seem more consistent with the predictions of the retinal model than with those of the non-retinal model. Likewise, we can ask the question of which reach pattern is observed when targets are presented with gaze near, but the reaches to them executed with gaze far. Also in this case, as shown in Figure 3.4E, the reach pattern is more similar to that predicted by the retinal model (now being the static far pattern).

To quantify these observations, we took the root of the average of the squared difference between the error in each dynamic trial and its predicted value by either the non-retinal or retinal model. Consistent with the qualitative observations, this yielded root-mean-squared error (RMSE) values lower for the retinal model than for the non-retinal model, for both dynamic conditions (0.8° vs 2.2° and 0.7° vs 2.0° , respectively). So, for these pure vergence displacements, the subject’s average reach patterns seem to correspond best with the retinal coding scheme.

The right panels of Figure 3.4 present data of reaching movements across two directional gaze changes. Again, the two static conditions show different patterns of reach endpoints for different gaze directions (10° left and 10° right; left- and right hand panels of Figure 3.4B, respectively). Based on these results and using the same arguments described above, our two models make two different predictions about the errors in the dynamic paradigm. Evidently, as shown by Figure 3.4D and F, which again are based on pure version trials (see Methods), the observed reach patterns in the dynamic situations (left-to-right and right-to-left gaze changes, respectively) have a much greater similarity with the corresponding patterns predicted by the retinal model than with the endpoint distributions predicted by the non-retinal model. RMSE values, computed as above, yielded 3.1° vs 12.7° and 4.7° vs 14.3° (retinal vs non-

retinal model) for the rightward and the leftward dynamic gaze changes, respectively. This corroborates findings in previous literature that also made a clear case for the retinal coding and updating of target direction (Henriques et al. 1998; Medendorp and Crawford 2002; Poljac and Van den Berg 2003).

To further quantify the results of this subject (RV), we computed RMSE values related to the retinal and non-retinal model across all possible gaze displacement manipulations, thus also including the trials with a combined vergence-version change (see Methods). In this analysis, performed in either dimension (depth / direction), we included in one manipulation the three reach patterns that were obtained with the eyes always starting at the same fixation point and ending at points that have the same vergence (or version) difference from this point, irrespective of the version (or vergence) component. The RMSE was then computed based on the average reach pattern from three refixations (always one pure change and two combined version-vergence changes). Since we used nine initial fixation points, and two vergence (or version) differences relative to each point, this makes 18 manipulations in total per dimension. We plotted these values versus each other in Figure 3.5, separately for the depth (A) and direction (B) dimension. Data points above the diagonal would indicate a preference for the non-retinal model; data below the diagonal would be more consistent with the retinal model. For this subject, this quantitative comparison clearly




Figure 3.5. RMS errors in relation to the two models plotted versus each other. Data pooled across trials that have the same gaze displacement (same dynamic situation), in one subject (RV, same as in Fig 3.4). For both depth (A) and direction (B), most data points fall below the diagonal, indicating that the retinal model gives a better description of the data in this subject.

supported the retinal model, showing a better provision to account for the systematic depth and direction errors observed in the data.

Next, before we proceed further, recall that the efficacy of our test is based on the premise that the reach error in the static condition depends on eye position. Figures 3.4 and 3.5 confirm this for a single subject. To test this assumption in an analysis across subjects, we performed a 3x3 repeated-measures ANOVA on the static reach errors with eye position and craniotopic target location as within-subject factors. This analysis revealed a significant effect of either factor, in both the depth and direction dimension (in all tests $F(2,12) > 6.3$, $p < 0.05$), which validates the basic premise of our test.

Subsequently, under this confirmed assumption, Figure 3.6 quantifies the results of all 14 subjects in two typical dynamic conditions (far-to-near, near-to-far and left-to-right, right-to-left). Analogous to the observations made in Figure 3.4, the figure demonstrates that the retinal model fits the mean pattern of reach errors across subjects better than the non-retinal model (compare the corresponding RMSE values). The mean RMSE values across all testing conditions are given in Figure 3.7A and B, for each subject separately. Across subjects, the retinal model produced the best description for the coding of both target distance and direction (paired t-test, $p < 0.001$).




Figure 3.6. Population data. **A, C.** Depth updating: reach patterns for two typical dynamic conditions (same as Figure 3.4, in black), averaged across subjects, each superimposed on the predictions of the two models (in gray). Circle size, SE of each data point. The retinal model makes the best match to the data, for both conditions. **B, D.** Direction updating: data in same format. Data are most consistent with predictions of the retinal model.




Figure 3.7. Comparison of the mean RMS error value (across all trials) for the retinal and non-retinal model, for all subjects separately. In each subject, the retinal model performed best for both depth (A) and direction (B) updating, with all respective data points lying below the diagonal.

Model analysis

Although the data of our subjects seem to lend support for the retinal model, this interpretation may be flawed if reach errors were to depend non-retinally on final eye position, instead of being caused by an updated retinal representation. To examine this, in the following analysis, we further quantified reaching behavior by performing a multiple linear regression to investigate how the reach error relates to either eye displacement or final eye position. We fitted the following relationship, separately for the depth and direction dimension,

$$Err = a_0 + a_1 \cdot (T_{ret} - u \cdot \Delta E) + a_2 \cdot E_f \quad (\text{Eq. 3.1})$$

to the data of each subject, with Err the reach error in degrees, T_{ret} the retinal location (eccentricity or disparity) of the target, ΔE the amount of eye displacement (version or vergence, in deg), E_f the final eye position in craniotopic coordinates (version or vergence, in deg), and a_0 , a_1 , u and a_2 free parameters in the fit. Parameter a_0 quantifies the bias in the reach error, irrespective of target location or eye position. Parameters a_1 and u characterize the error term related to the processing of an (updated) target representation relative to the eyes, with a_1 a scaling term and u the updating gain. If errors were to arise solely at the level of target presentation, the effect of eye displacement would be zero, thus the updating gain would be zero; $u=0$. If the errors

depend on the location of the target relative to the new eye position, this means that the system has taken possible eye displacements into account, which ideally requires the updating gain to be 1, thus $u=1$. Finally, fit parameter a_2 in Eq. 3.1 quantifies the dependence of the errors on final eye position per se.

For all subjects we found significant correlations: $0.2 < r < 0.9$ ($p < 0.05$ for all subjects) for the depth dimension and $0.3 < r < 0.7$ ($p < 0.01$ for all subjects) for the direction dimension. Parameter a_0 had a mean value (\pm SD) that was significantly different from zero for depth (2.69 ± 1.68 , $p < 0.001$), but not for direction (-1.33 ± 5.11 , $p = 0.35$). The histograms in left-hand panels of Figure 3.8 show the distribution of updating gains u across subjects, for both depth and directional updating. The data give no clear sign of differences between depth and directional updating. Across subjects, the mean updating gain (mean \pm SD) is not significantly different from 1 either for the depth component ($p = 0.33$; $u_{\text{depth}} = 0.91 \pm 0.32$) or for the directional component, ($p = 0.18$; $u_{\text{direction}} = 1.13 \pm 0.36$). The fact that the updating gain u is close to 1 indicates that the reach errors arise in relation to an updated retinal representation, derived by correctly compensating for intervening eye displacements.

To characterize the relative contribution of a retinally shifted eye-centered target representation and final eye position to the reach error, we computed the ratio $(a_1 - a_2)/(a_1 + a_2)$. This ratio would be one when the reach error depends only on the updated target representation ($a_2 = 0$) and would be minus one when the error depends solely on final eye position ($a_1 = 0$). Figure 3.8C and D depict this ratio, separately for the depth and direction dimensions. As shown, for target depth, the ratio settles between the two extremes, showing a mean value of $-0.10 (\pm 0.27)$, indicating that both final eye position and an updated disparity representation contributed about equally to the reach errors. For the direction dimension, in contrast, the ratio had a value of $0.65 (\pm 0.39)$, indicating that the reach errors seem to arise primarily in relation to an updated retinal representation.

Binocular versus monocular updating

As a final note, throughout our analyses, we have assumed that target direction and depth are processed as separate signals. Theoretically, it may be possible that depth information is not processed in the form of an explicit disparity signal, but rather is computed on demand on the basis of two monocular direction representations that are stored and updated in separate maps. To investigate this possibility, we fitted Eq. 3.1 in terms of directional components only, for each eye separately. On basis of the fitted parameters, we inferred reach depth by computing the point of intersection of the reach directions predicted based on monocular processing. We then quantified how well the resulting depth errors correlated with the actual, observed errors. Performance




Figure 3.8. Updating gains and final eye position effects in depth and directional reaching. **A, B.** Updating gains (u) are distributed around a value of one, for both depth ($u = 0.91 \pm 0.32$ (SD)) and $u = 1.13 \pm 0.36$, respectively). **C, D.** Contributions of the updated target position and final eye position to the reach error, expressed by the ratio $(a_1 - a_2)/(a_1 + a_2)$. For depth, both contributed about equally, mean ratio: -0.10 ± 0.27 . For direction, mean ratio: 0.65 ± 0.39 , meaning the error arises for the most part in relation to an updated direction representation.

of this description in monocular coordinates was very poor, with $0.01 < r < 0.37$, and significantly lower (paired t-test, $p < 0.001$ using Fisher z-transformation for comparing correlation coefficients) than predicted by the binocular coding scheme, as shown above. This warrants our assumption of binocular processing in terms of depth and direction components in the exploitation of the test described above.

Discussion

Over the last few decades, many studies have investigated how the constancy of spatial direction for motor actions is achieved across conjugate eye-movements. In contrast,

the mechanisms involved in the maintenance of spatial depth across disjunctive eye movements have remained largely unexplored. Here we have addressed this issue using the accuracy of memory-guided reaching movements to visual targets, briefly presented at different depths prior to a shift of gaze.

We tested between two models of the implementation of depth constancy: a retinal vs. a non-retinal model. To make this distinction, we exploited the fact that the accuracy of human reaching movements toward remembered space-fixed, but nonfoveally viewed, targets depends on the eyes' fixation distance, analogous to the utilization of systematic reach errors for testing directional coding (Henriques et al. 1998). We hypothesized that, if spatial depth is stored non-retinally, the intervening vergence shifts in the dynamic trials should have no effect on reaching. This model predicts that reaches in these trials should be similar to those made in static trials at the same initial eye position, but without the intervening vergence shift. Alternatively, if depth coding is retinal, updating for the gaze shift becomes essential and reaches should match those of static trials performed at the final eye position, under the assumption that the sensory consequences of the gaze shift have been perfectly taken into account (perfect updating).

With intervening gaze-shifts, the memory-guided reaches showed an error pattern that was based on the new eye position and on the depth of the remembered target relative to that position (Figures 3.3-3.8). This suggests that target depth is recomputed after the gaze shift, as would be required if the brain encoded depth in retinal coordinates. We found the values of the updating gain near unity (Figure 3.8), demonstrating the persistence of a correct representation of target depth relative to fixation across vergence eye movements. This is in line with perceptual observations by Gonzalez et al. (1998), who reported that perceived depth of random-dot stereograms is not affected by changes in vergence. We deduce that the systematic reach errors in the present study must therefore arise after the updating stage, in the subsequent visuomotor transformation from this updated retinal representation to the arm-centered representation for reaching (Khan et al. 2005). The latter is further emphasized by the influence of final eye position on the reach error (Figure 3.8).

As far as we are aware, no other studies have investigated the reference frame in depth constancy of across vergence eye movements. Krommenhoek and Van Gisbergen (1994), who tested human subjects in double step eye movement experiments with combined version-vergence movements, reported that the saccadic and vergence system can use non-retinal feedback about a prior eye movement in direction and depth. They did not, however, address the spatial representation that underlies spatial constancy in this behavior.

All subjects tested ($n=14$) favored the retinal model of depth coding. This result is warranted by the fact that our control results on spatial direction, obtained in the same experiments, provide a strong confirmation of the earlier literature. Several behavioral studies on directional constancy have reported evidence for retinal updating of target direction (Henriques et al. 1998; Medendorp and Crawford 2002; Chapter 2 of this thesis). Corroborating these findings, we found the updating gain for directional updating to be close to one (see Figure 3.8). Moreover, several monkey and human brain areas show activity related to the retinal coding and updating of the direction of remembered targets, including frontal (Goldberg and Bruce 1990) and parietal areas (Duhamel et al. 1992; Batista et al. 1999; Medendorp et al. 2003a). The present study suggests that target depth is also coded and updated in retinal maps, presumably in the form of absolute disparity coordinates (Cumming and DeAngelis 2001). Changes of vergence alter the values of absolute disparities, so they must be updated to maintain spatial constancy.

Previous neurophysiological work suggests that depth representations may be constructed in areas within occipital, frontal and parietal cortex (Sakata et al. 1997; Dobbins et al. 1998; Ferraina et al. 2000, 2002; Fukushima et al. 2002; Gnadt and Beyer 1998; Gnadt and Mays 1995; Rosenbluth and Allman 2002; Genovesio and Ferraina 2004; Genovesio et al. 2007). For example, Gnadt and Mays (1995) described neurons in the lateral intraparietal area (LIP) of the macaque that have three-dimensional receptive fields. Activity of these neurons is expressed as a function of spatial parameters in the frontoparallel plane (horizontal and vertical eccentricity) and the relative depth from the plane of fixation (retinal disparity). Also, Genovesio and Ferraina (2004) found LIP neurons that were sensitive to the retinal disparity of a target, but further showed that this disparity tuning is modulated by fixation distance. A brief report from Bhattacharyya et al. (2005) on reaching in depth also suggest that neural activity in the parietal cortex reflects distance to target and vergence angle. Given these signals, it has been argued that the parietal cortex plays a role in the integration of retinal and extraretinal information to determine the egocentric distance of a target located in three-dimensional (3-D) space (Genovesio et al. 2004). It remains to be investigated whether the computation of egocentric depth is an automated process or is enforced on demand only when a (reach) action is prepared. Cumming and DeAngelis (2001) indicated that the updating of target distance may be expressed by changes in retinal disparity representations. Recently Genovesio et al. (2007) recorded neural activity in LIP while monkeys performed saccades between targets in different depths. They showed that in the postsaccadic period, neural activity is influenced conjunctively by both the eye displacement and the new eye position. In this respect, the behavioral observations made in the present study indicate a striking

correspondence by showing effects of the same types of signals on the depth component of the reaching errors. It can be argued that these signals play a role in the dynamic retinal representation of visual space and in the further transformation of spatial information in other coordinates systems (Genovesio et al. 2007; Andersen et al. 1985).

In support of this dynamic spatial representation, it has also been shown that neurons in LIP actually begin to respond before the eye movement to stimuli that will enter the receptive field after the eye movement (Duhamel et al. 1992; Nakamura and Colby 2002). In other words, LIP neurons anticipate the sensory consequences of the future eye position before the saccade is executed, which suggest that the updating mechanisms relies on a copy of the eye motor command (Sommer and Wurtz 2002), rather than on sensory feedback that arrives much later. A useful experiment to be performed in this context would be to investigate if predictive updating also occurs in relation to vergence eye movements. Along these lines, Kaiser and Lappe (2004) reported recently that visual objects flashed shortly before or during a saccade are mislocalized, resembling a compression of space around the saccade target. They attributed this distortion to the remapping process in parietal cortex, and it would be interesting to see whether similar spatial distortions occur across vergence eye movements, and if so, whether they have a similar time course. Viewed from a different perspective, a recent hypothesis here, put forward by Vaziri et al. (2006), is that the brain integrates the predicted sensory consequences of motor commands with the actual sensory (feedback) information to produce an estimate of sensory space that is better than possible from either source alone. More experiments are required to see if this hypothesis is upheld across combined saccade-vergence movements, or in further extended conditions involving movements of the head and body in space.

Finally, it is important to emphasize that the present study considered direction and depth as independent spatial variables, processed and updated separately during updating across saccades. We made this assumption based on results of various reaching studies, performed in static conditions, showing that the spatial distributions of movement endpoints of reaches toward remembered targets were elliptical in shape with a tendency of the major axis to be directed to the subject's eyes (McIntyre et al. 1997; Henriques et al. 2003; Baud-Bovy and Viviani 1998). This implies that noise in the reach is larger for the depth than for the directional component in these cases, suggesting that both dimensions are controlled separately. The present analyses supported this assumption by demonstrating that an alternative, implicit depth representation emerging from two monocular signals is less consistent with our data, demonstrated by low correlations.

That being said, in more complex updating conditions, depth and directional signals must interact to preserve spatial constancy in retinal coordinates (Medendorp et al. 2002, 2003b; Chapter 2 of this thesis; Li and Angelaki 2005; Ferraina et al. 2000). For example, when the body translates, correct updating in a retinal frame requires updating to vary from object to object, depending nonlinearly on their depth and direction (Medendorp et al. 2003b; Li and Angelaki 2005). Recently, we showed, using memory-guided reaching movements, that the updating of target direction for translational motion is compromised by small errors, which increase with depth from fixation and reverse in direction for opposite depths from fixation (Chapter 2 of this thesis), consistent with translational updating in retinal coordinates. Li and Angelaki (2005) reported that monkeys can update target distance during body motion in depth, using extraretinal vestibular information. Borne out by the present results, we propose that these vestibular signals interact with retinal disparity and eccentricity information to retain 3-D stability during body motion in space.

Chapter 4

Visuospatial memory computations during whole-body rotations in roll

This chapter has been published as:

Van Pelt S, Van Gisbergen JAM, and Medendorp WP. Visuospatial memory computations during whole-body rotations in roll. *Journal of Neurophysiology* 94: 1432-1442, 2005.

How the brain represents space and how this information is used to generate goal-directed behavior has been subject of longstanding debate (Howard 1982; Von Helmholtz 1867; Andersen et al. 1985; Duhamel et al. 1992). While currently viewed targets impinge on the retina and are always available, locations of previously viewed targets must be stored in memory if needed for actions at later time. It is known that these memories remain quite accurate over long times, even after we have moved around (see e.g., Israel and Berthoz 1989; Hallett and Lightstone 1976; Sparks and Mays 1983; Mergner et al. 2001; Baker et al. 2003; Medendorp et al. 2002, 2003b).

What are the computational strategies that underlie this behavior? The brain could store spatial memories in egocentric as well as allocentric coordinates (see Battaglia-Mayer et al. 2003, for review). Target representations computed and stored in an allocentric frame of reference (or Earth or inertial frame) are most stable, since they remain correct for any type of intervening self-motion. In turn, for motor planning, they must be converted backwards to eye-, head- or limb-related coordinates, depending on the motor system that is being employed. In contrast, target locations stored within an egocentric framework must be continuously recomputed, or updated, whenever the axes of the specific ego-frame (e.g. limb, eye, head or torso) move, if they are to remain useful for guiding motor action.

There is neurophysiological evidence for either view. For example, hippocampal place cells code self-position in allocentric coordinates (see Burgess et al. 2002; Best et al. 2001, for reviews). Snyder et al. (1998) reported evidence for separate allocentric and egocentric representations within the posterior parietal cortex. In addition, many areas in parietal and other cortical and subcortical brain regions have been shown to update their egocentric information when the egocentric reference frame moves, using extraretinal information about self-motion (Gnadt and Andersen 1988; Duhamel et al. 1992; Medendorp et al. 2003a; Sommer and Wurtz 2002; Nakamura and Colby 2002; Goldberg and Bruce 1990; Walker et al. 1995).

In which reference frame are locations of targets for saccades maintained? Since neurophysiological data suggest that several coding schemes co-exist, they cannot distinguish which of these dominates behavior. However, behavioral performance can provide important insights. By assessing the operational errors in the system during various task conditions, one may be able to make inferences about the nature of the computations.

Several studies have used this strategy by investigating the variability in the endpoints of saccades to memorized targets, after intervening eye, head and body movements. On this basis, Baker et al. (2003) found evidence for eye-centered target representations, while Skavenski and Steinman (1970) and Karn et al. (1997) suggested

the use of an extraretinal, possibly space-centered frame of reference. These different suggestions may simply imply that no single frame of reference is being employed in these various conditions, or that the reference frame question cannot be addressed by looking at the variable errors alone.

In the present paper, we have investigated the internal mechanisms underlying spatial memory by exploiting a robust systematic error that occurs in human external space perception. More specifically, we have designed a paradigm to assess whether human subjects store target locations for saccades, presented prior to a whole-body rotation, in egocentric or in allocentric coordinates. Our test is based upon the observation that subjects, when tilted sideways in darkness, make systematic errors when indicating the direction of gravity, an allocentric variable. For tilts $> 60^\circ$, these errors are found in the same direction of the body tilt, mounting up to 50° when the body is tilted at about 130° (Kaptein and Van Gisbergen 2004). This type of error is known as the Aubert effect, or A-effect (Aubert 1861), and has been observed in various test paradigms, such as classical visual-line tests and oculomotor paradigms relying on saccadic pointing (Mast and Jarchow 1996; Van Beuzekom and Van Gisbergen 2000). Noteworthy, for small tilt angles ($< 30^\circ$), errors in the subjective direction of gravity are much smaller, and commonly observed in a direction opposite to the body tilt. Such errors are known as the E-effect (Howard 1982). In the present study, for clarity, both the A- and E-effects will be collapsed and referred to as A-errors, unless indicated otherwise. It is important to note that errors are virtually absent when subjects are asked to estimate their body-tilt in space (Kaptein and Van Gisbergen 2004). This indicates that the A-error is not merely caused by inaccuracies in the underlying head orientation in space signal, but rather reflects a property of the central computation involved in external, world-centered space perception (Mittelstaedt 1983; Eggert 1998; Kaptein and Van Gisbergen 2004).

Figure 4.1 illustrates how the robust relationship between the A-error and body tilt can be exploited to design a test to distinguish between egocentric and allocentric coding of spatial memory. As Figure 4.1A shows, if a subject, tilted sideways at a given angle (ρ_1) stores the direction of a target T in an allocentric, world-centered frame of reference, the corresponding memory (ϕ) will be affected by the perceived distortion of this frame at this tilt angle, (A-error, A1). In other words, the direction of the target will be stored in a distorted world frame, i.e., relative to \hat{V} , rather than to the actual world-centered coordinates (V). If then, this subject is rotated to a final position (ρ_2), and a read-out of this allocentric memory representation is obtained, the response will also incorporate the error (A2) in the subject's representation of the world-centered axes at the new position (see Figure 4.1B). Thus the allocentric, world-centered model makes a very precise prediction: the directional error of a saccade S towards a target,




Figure 4.1. Allocentric versus egocentric spatial memory computations during whole-body rotations in roll. In our tests, a subject stores a world-fixed target to memory at initial tilt position (ρ_1). After being rotated to a new final position (ρ_2), the subject directs a saccade to the memorized location of the target. A,B: Allocentric coding. **A.** If the direction of a target T is stored in an allocentric, world-centered frame of reference, its memory representation (ϕ) will be affected by the perceived distortion of this frame (A-error, A1), and thus be encoded relative to the internal representation of the physical world-frame \hat{V} . **B.** A similar error (A2) will be incorporated in the read-out (S) of this allocentric memory, depending on subject's representation of the world-centered axes at the new position (\hat{V}). C, D: Egocentric coding. **C.** The memory (ϕ_1) is encoded relative to the axes of an egocentric frame of reference (head/body/eyes). **D.** This representation must be updated during the intervening rotation ($\Delta\rho = \rho_2 - \rho_1$) in order correctly represent the direction of a space-fixed target location at the new body position (ρ_2), resulting in an updated memory (ϕ_2), that is ideally equal to $\phi_1 - \Delta\rho$. Note, it suffices here to refer to this operation by a subtraction, but strictly speaking, rotations do not commute and therefore must be described by multiplicative operations. The dashed arrow represents the previous body tilt position in the world (ρ_1). The static tilt position of the head (ρ) is defined as the angle between the direction of the physical vertical (V) and the subject's positive long-body axis (Z), as seen from behind the subject (positive: rightward tilt; negative: leftward tilt).

briefly presented before a whole-body rotation, should be equal to the difference in subjective distortion of the earth-frame when probing the memory and when storing the memory ($A_2 - A_1$). Accordingly, this model predicts zero response errors in absence of an intervening body rotation ($\rho_2 = \rho_1$), since A_2 will then be equal to A_1 . This also means that any response error found in this condition would merely imply the involvement of other independent processes, affecting memory preservation.

However, if a spatial memory is stored in egocentric coordinates (Figure 4.1C), such allocentric distortions, as expressed by the A-error, are supposed to play no role. In that case, the memory (ϕ_1) is stored relative to the axes of a given egocentric frame of reference, e.g. relative to the subject's positive Z-axis, Z. While this reference frame can be the eyes, head or body, these can be treated equivalently in the present study, ignoring the small effects of eye countertorsion (but see later). Within an egocentric framework, however, a spatial memory about a target location must be updated when the body rotates to a new position to keep correct registry with its true spatial location, represented by ϕ_2 (Figure 4.1D). Therefore, if the egocentric model is correct, we would expect the readout of this memory after the body rotation only to be affected by errors, if any, related to the amount of intervening rotation $\Delta\rho$ ($=\rho_2 - \rho_1$). Note that, for the sake of argument, we have assumed here that the brain can calculate the change in angle perfectly, irrespective of initial and final tilt position. We will take up this issue in more detail in the Results section to test whether this assumption is correct. As with allocentric coding, the egocentric scheme would predict zero updating error in static conditions, when intervening motion is absent (excluding the errors caused by other deterioration processes of memory).

In the present study, we asked human subjects to make eye saccades to remembered targets after whole-body rotations in roll. We show that the saccadic response is systematically affected by the distorted percept of the external world. We will interpret these results in terms of their computational and physiological significance for the brain.

Methods

Subjects

Six subjects (1 female, 5 male), aged between 24 and 60 years, gave their informed consent to participate in the experiment. All subjects were without any known visual, vestibular or other neurological disorders. Three of them were naïve with respect to the purpose of the experiments. No systematic differences in performance were found between naïve and non-naïve subjects.

Setup

The subjects were seated in a computer-controlled vestibular chair. They were secured tightly into the chair using seat belts, trunk and hip supports, a foot rest and straps around the feet and legs. A padded adjustable helmet firmly stabilized the head in the normal upright position. For each subject, seat adjustments were made so that the eye of which the orientation was measured coincided with the roll axis. During the experiment, subjects were rotated around the roll (naso-occipital) axis in complete darkness. The chair rotated with a constant velocity of $45^\circ/\text{s}$, with equal values of acceleration and deceleration of $30^\circ/\text{s}^2$. Chair orientation was measured using a digital position encoder with a resolution of 0.04° , and recorded on disk.

Two-dimensional eye orientation of either the left or the right eye was measured with the scleral search coil technique, using oscillating magnetic fields generated by two sets of orthogonal coils inside the chair ($0.77 \times 0.77 \text{ m}$). The signals were amplified, low-pass filtered at 200 Hz, and recorded at 500 Hz per channel.

Targets (red light emitting diodes, LEDs, luminance 8 mcd) were presented on a chair-fixed screen at a distance of 115 cm in front of the subject's measured eye; the central LED on the screen coincided with the axis of rotation. Peripheral LEDs ($n=36$) were positioned in an array on this screen, at the intersections of three circles at 11 , 22 and 31° of visual angle and 12 equally-spaced meridians. Prior to the experiment, the subject fixated each of the LEDs in random order, in complete darkness, to calibrate the search coil.

Experimental paradigm

We used a memory-guided saccade task to test a subject's ability to memorize locations of targets, briefly presented before he or she underwent a whole-body rotation.

ROTATION PARADIGM

This experimental paradigm is illustrated in detail in Figure 4.2. Roll angle (ρ) is defined as the angle of the longitudinal body axis with the earth vertical, taken positive for right-ear-down rotations. Each trial began by turning off all lights after which the subject was rotated to an initial body tilt position ρ_1 ($t=0-4\text{s}$). Next, a central fixation LED was presented (at $t=4\text{s}$), which cued the subject to make self-paced saccades into the four Earth-centric cardinal directions to indicate the perceived world horizontal and vertical directions (Van Beuzekom and Van Gisbergen 2000). These responses, to be completed within six seconds, were used to determine the subject's A-error at this position. Next, the central LED had to be fixated again and a peripheral target was flashed for 1000 ms, at a retinal eccentricity of 22° . The subject was explicitly instructed to memorize the location of this target as if it were world-fixed, for as long as




Figure 4.2. Experimental paradigm. **A.** Sequence of stimuli and subject instructions. After being tilted to an initial position (p_1), the subject (seen from behind) made saccades to indicate the world cardinal directions (measurement of A-errors). Next, a world-fixed target was flashed for 1s onto the retinal periphery, after which the subject was rotated to a new position (p_2). Then the fixation point disappeared, cueing the subject to make a saccade towards the remembered target location. Finally, the A-errors were measured at the new position, after which the subject was rotated back to the upright position. **B.** Temporal order of events for body rotation, stimuli and saccadic responses during a trial.

the central fixation LED was on (9 seconds). During this period, he or she was rotated to a new orientation p_2 . When the central LED disappeared, this signaled the subject to make an eye movement to the memorized target location, and fixate there for a short moment. Two seconds later, with the subject still at the new orientation, the central

fixation LED was turned on again and another set of A-error measures was obtained. Finally, the subject was rotated back to the upright position and the room lights were switched on. Each trial lasted 33 seconds; between trials there was 27 seconds of rest.

STATIC PARADIGM

This paradigm was identical to the rotation paradigm, except that the subject was kept at the same orientation ($\rho_2 = \rho_1$) after target presentation. These trials served as memory controls, to test the subject's performance in the absence of an intervening rotation.

In both paradigms, initial and final orientations were chosen from the interval of -120° to $+120^\circ$, with steps of 30° (9 possible tilt positions). This resulted in a total number of 81 possible combinations, of which 75 were tested (68 rotation trials, 7 static trials). In the static paradigm, the two orientations -30° and $+30^\circ$ were not tested. In the dynamic paradigm, rotations from 120° to -30° , -120° to 30° , and vice versa were excluded. We never used tilt angles beyond 120° to avoid the complex behavior of the A-effect at very large tilt angles (beyond 135° , see Kaptein and Van Gisbergen 2004). For each trial, target location was chosen pseudo-randomly from 12 possible locations, so that all possible target positions were used at least six times. Trials with initial left-ear down and right-ear down rotations were interleaved. For illustration purposes (Figures 4.4, 4.6 and 4.8), we collected, in addition, some repeated measurements in one subject, using a specific selection of the stimulus set. All subjects were given a few practice runs to get used to the vestibular stimulation and the paradigm. The experiment was divided into 3 sessions of 25 trials, tested on different days. Each session lasted for about 40 minutes, including calibration and practice. Typically 2 out of the 75 trials per subject had to be excluded from analysis because the subject failed to maintain fixation or make saccades at the requested times.

After completion of all experimental sessions, 5 of 6 subjects were tested on their rotation perception in a slightly modified version of the rotation paradigm. Subjects performed 50 trials with pseudo-random initial and final body orientations, chosen from the 240 possible integer values in the interval of -120 to 120° , without assessing the A-errors or spatial memory performance. Instead, when the fixation light went off, subjects had to report verbally the amount of perceived intervening rotation, as minutes on a clock face ($+15 \text{ min} = +90^\circ$).

Analysis

Data analyses were performed off-line using Matlab software (the MathWorks). Using the fixation data of the calibration run, horizontal and vertical eye-coil signals were

calibrated with two neural networks, one for each position component. Each network consisted of two input units (representing the raw horizontal and vertical signal), three hidden units, and one output unit (representing the desired calibrated horizontal or vertical position signal). Raw eye-coil signals were calibrated by applying the resulting feedforward networks. Average calibration errors were typically $<0.5^\circ$. Off-line saccade detection was performed manually by the experimenter, on basis of the calibrated eye position signals.

In each trial, subjects made saccadic eye movements at both initial and final tilt position to indicate their perceived earth-centric cardinal directions. The average directional error of these oculomotor responses from the true earth-centered directions was computed following the method used by Van Beuzekom and Van Gisbergen (2000). In short, the direction of the saccadic endpoints of each arm of the resulting cross-like figure of saccades was determined. The difference of the mean of these directional settings from zero represents a distortion in the subjective earth-referenced frame, known as the Aubert-effect (A-effect) when in the same direction as the body tilt, and acknowledged as an E-effect when in a direction opposite to the body rotation. Since the A-effect is most systematic and substantial, as mentioned in the introduction, we will refer to both distortions as A-errors.

Subject performance in memorizing a target location, and retaining that location while being rotated during the memory period, was determined by the accuracy of the saccade made to that location at the end of the memory period. As such eye responses often contained several corrective saccades (Medendorp et al. 2002), we measured the directional error of the endpoint of the most eccentric saccade toward the remembered target. Clockwise errors were taken positive. We explored the relationship between the saccadic targeting response, the size of the A-error, and amount of intervening rotation across all trials in order to test which of the spatial memory models - illustrated in Figure 4.1 - would best fit our data.

Results

Subjective earth-centric frame

Our experiments were designed to test whether the location of a target, briefly presented before an intervening whole-body rotation, is stored in egocentric or allocentric coordinates. To make this distinction, we exploited the fact that subjects, when tilted sideways in darkness, make systematic errors when asked to indicate the direction of the gravity vector and the orthogonal horizon (see e.g. Mittelstaedt 1983; Van Beuzekom and Van Gisbergen 2000).

Figure 4.3 quantifies this measure (the A-errors) for all our subjects. In every trial, at each tested tilt position, subjects made four self-paced saccades to indicate the earth-centric cardinal directions. The misalignment of these earth-referenced saccades with the true Earth-centric directions specifies the subjective earth-reference frame. Figure 4.3A shows the gaze trajectories of these saccades for one subject, for all tilt angles




Figure 4.3. Quantification of the A-errors. **A.** Trajectories of saccades from one subject (JG) directed along the subject's percept of the Earth's cardinal axes, for the nine different tilt positions tested. For illustration purposes, data from 2-3 trials per rotation angle (out of 8 or 9 measured) are presented. U, D, L, R indicate the subjective upward, downward, leftward and rightward directions in space, respectively. Substantial A-errors were present for large tilt angles ($|\rho| > 60^\circ$). **B.** Average A-error (\pm SD) as a function of body tilt angle, for the same subject. **C.** A-error for all subjects. For small rotations, some subjects showed an E-effect, i.e. errors in a direction opposite to the body tilt.

separately. For (absolute) tilt angles up to 60° , the subjective frame corresponds closely to the true earth-frame, with distortion errors (A-errors) smaller than 5° . However, for tilt angles $> 60^\circ$, settings were far from flawless, with distortions of up to 50° in this subject. Statistical analysis revealed that these settings were independent of the preceding amount of rotation ($p < 0.01$, t-test), i.e., they were only related to the static body tilt position at which the measurements were taken. Figure 4.3B quantifies these data further by showing the mean A-error (± 1 SD) as a function of tilt angle. Similar spatial distortion profiles were observed in all our subjects (C), in correspondence with previous studies (Van Beuzekom and Van Gisbergen 2000; Kaptein and Van Gisbergen 2004). As we will show in the next sections, these data can be used to test the reference frame used in spatial memory.

Static paradigm

How well can stationary, but tilted subjects memorize locations of world-fixed targets? Using the static control condition, we quantified the subjects' performance in making saccades to target locations, remembered for 9 s, without being moved during the memory delay period. Figure 4.4A shows the results for one subject, tilted at 120° , performing 12 different spatial memory trials, by superimposing the saccade trajectories of all trials towards the four remembered targets, located on the cardinal axes in world-space. The responses were fairly accurate, even though the subject had memorized the target for 9 s. Figure 4.4B and C depict the subject's percept of the Earth-reference frame (A-errors) in the same trials, as indicated by saccades, respectively




Figure 4.4. Saccade trajectories of one subject (JG), tilted at 120° in the static paradigm. Results of 12 different trials. **A.** Spatial memory readout: saccades directed at memorized targets, presented on any of the four world-fixed cardinal axes (indicated as hours on a clock scale; e.g. 12 is an upward target location in space). Saccade directions are fairly accurate. **B.** A-errors before target presentation. U, D, as in Fig 4.3. **C.** A-errors after spatial memory readout was collected.

before and after the spatial memory task was performed. These measures showed a high degree of reproducibility: the saccadic cross-like trajectories, made in separate trials, but at the same tilt angle, are all closely aligned. Thus, the A-error remains constant within and across trials, in correspondence with Figure 4.3. It is noteworthy that the severe distortion in the subject's perception of external, world-centered space (shown in Figure 4.4B, C) bears no resemblance in the saccade directions to remembered visual targets presented in external space (Figure 4.4A).

Figure 4.5 shows the saccade accuracy averaged across subjects for each tilt angle tested in the static control condition. Saccades directed from tilted body position showed higher variability than saccades made from an upward, natural body position. Across all tilt angles the mean error ranged from -4.8° to $+6.6^\circ$, and was only significantly different from zero when the subjects were at $+120^\circ$ ($p=0.0068$, t-test). For the purpose of the presented study, we take from this that spatial memory is not systematically degraded by our 9 s memory delay,

Rotation paradigm

To what extent are subjects able to look at remembered locations of world-fixed targets, briefly presented before an intervening whole-body rotation? Figure 4.6A shows the performance of a subject in 12 trials in the intervening rotation condition. Here, the subject viewed the flashed targets at a 120° leftward tilted body position before he was rotated to a final body position of 120° rightward, and subsequently performed the saccades. As Figure 4.6A shows, the saccades to the remembered targets show dramatic directional errors. For example, a target initially presented upward in Earth space (the 12 o'clock direction), is localized, after the body rotation, by a saccade directed rightward in space. In other words, unlike the static condition (compare Figure 4.4A), a clear deterioration of performance occurs when remembered target locations must be indicated after an intervening body rotation.




Figure 4.5. Response error, averaged across subjects, as a function of tilt angle in the static paradigm. Positive errors denote errors made in the same direction as body orientation. Error bars denote SD.




Figure 4.6. Response saccades from one subject (JG) tested in the rotation paradigm. Performance in 12 trials in each of which the subject perceived the target at a 120° leftward tilt, then was rotated to a 120° rightward tilt, and subsequently made a saccade to the memorized target. **A.** Spatial memory readouts: saccades directed at memorized targets, presented on any of the four world-fixed cardinal axes (indicated as hours on a clock scale), after the intervening body rotation. Note the large directional errors (compare Fig 4.4A). **B.** A-errors at 120° leftward tilt (p_1), before target presentation. U, D, as in Fig 4.3. **C.** A-errors at 120° rightward tilt (p_2), obtained after the spatial memory readout was collected.

For completeness, Figure 4.6B illustrates the subject's perception of the earth-centric directions at the initial body orientation (-120°) for the same trials, while Figure 4.6C demonstrates the A-error saccades at the final body orientation (120°). As expected, the A-errors are quite different for these tilt angles, which conforms to the patterns observed in Figures 4.3 and 4.4.

Why do subjects make such severe errors in looking at world-fixed targets after an intervening body rotation? As outlined in the Introduction, if spatial memory in this task is coded in an allocentric, world-centered frame of reference, it will be affected by the perceived distortion of this frame when storing target location (Figure 4.1A). In that case, this will also shape the readout of this memory trace, coded in this frame, after the body rotation, based on the perceived distortion of the frame at the new body position (see Figure 4.1B). As such, we would expect the directional errors in the memory-saccade in a given trial be related to the amount of subjective distortion in the perceived earth-frame (A-error) at both the initial and final tilt angle in that trial. In other words, an allocentric coding scheme would predict a saccadic response error equal to the difference in A-error at the final and initial tilt angle ($A_{\text{net}} = A_2 - A_1$), as determined by the saccadic indications of the world's cardinal axes made at final and initial tilt angle, respectively. Figure 4.7A shows the data from one subject by plotting the directional error of the saccade versus the net A-settings for all trials tested (A_{net}).




Figure 4.7. Directional errors of saccades plotted against the size of the net A-error ($A2-A1$) or the amount of the intervening body rotation ($p2-p1$) in a dynamic paradigm. **A.** Response error plotted as a function of A_{net} ($A2-A1$) for one subject. The best-fit line had a high correlation ($r=0.96$) and a slope near one, consistent with the allocentric model. **B.** Response error plotted as a function of intervening body rotation ($\Delta p=p2-p1$) for the same subject. Regression line had a significant correlation ($r=0.92$), but a slope larger than zero, consistent with under-compensation for body rotation in an egocentric model. **C.** Regression lines of all subjects in an allocentric explanation scheme. All slopes are near one. **D.** Regression lines of all subjects in an egocentric explanation scheme. Slopes vary between 0.14 and 0.38, corresponding to an underestimation of Δp . Analyses on basis of both schemes reveal significant correlations, suggesting that readout errors of spatial memory can be explained by either of them.

Using linear regression, we quantified the relationship between the saccadic errors and the amount of subjective allocentric distortion (solid line). This analysis revealed a high correlation ($r=0.96$). As to the slope of the fitted line, the allocentric coding scheme

(Figure 4.1A) would predict a value of 1 (dashed line). For this subject, the slope had a value of 1.15 (SD = 0.04), which clearly seems to favor this model. Figure 4.7C shows the linear regression results from all subjects. In all subjects, we found high correlations ($0.73 < r < 0.96$). Across subjects, the slopes ranged between 0.81 and 1.15. On average, the slope was not significantly different from 1 ($p=0.74$, t-test), indicating almost perfect adherence to the allocentric model.

Could an egocentric model also explain these data? Recall our assumption that an egocentric model would predict the saccadic errors to be related to the amount of intervening body rotation (see Figure 4.1C, D). In Figure 4.7B we have plotted this relationship for the same data as presented in Figure 4.7A. A linear regression, which quantified this relationship, showed a high correlation ($r=0.92$). If there had been no updating, the slope of the fitted line would have been 1 (dashed line). Perfect updating would yield a slope of 0. For this subject we found a slope of 0.38 (SD = 0.02), which was significantly different both from 0 (t-test, $p<0.001$) and from 1 (t-test, $p<0.001$). Thus, according to this analysis, updating was not perfect – the change in body position seems underestimated in this subject. Similar results were found in all subjects (Figure 4.7D), with slopes, ranging from 0.14 to 0.38, that were significantly different from both 0 ($p<0.01$, t-test) and 1 ($p<0.001$, t-test). Furthermore, statistical analysis revealed that the observed correlations in these egocentric fits ($0.71 < r < 0.92$) were not significantly different from those observed for the allocentric model ($p=0.32$, t-test). In fact, this means that our analysis so far does not rule out the alternative coding scheme, i.e., an egocentric model. That is, one could still argue that an egocentric model with imperfect updating could describe the present results. How then can we distinguish between the putative models underlying the spatial memory computations in the present experiment?

Distinguishing between allocentric and egocentric coding

Based on our analyses so far, the directional error in the response saccades to the remembered target locations could be equally well predicted by a distorted allocentric coding mechanism as well as by an imperfect updater in an egocentric model. Why would both models perform about equally well? A confounding effect is that the amount of spatial distortion (A-errors) and tilt angle are so tightly related (see Figure 4.3). Therefore, a clear dissociation between our models is probably masked by the clear interaction between these two factors.

Can we remove these confounding effects and perform a more sensitive analysis to discriminate between the two models? In the following we will take advantage of the nonlinearity in the relationship between A-effect and body tilt angle (see Figure 4.3), to determine which model is best. In the analysis, we assume that an egocentric model

would predict similar errors in the memory-saccades for two trials with the same amount of intervening body rotation, irrespective of initial tilt position (ρ_1) and final tilt position (ρ_2). In contrast, the allocentric, world-centered coding scheme allows that the saccadic errors in trials with the same amount of intervening rotation may be different, depending on the A-error at initial and final tilt angle.

This idea is illustrated in Figure 4.8, in one subject, for two different testing conditions that have the same amount of intervening rotation ($\Delta\rho=120^\circ$) but different combinations of initial and final tilt position, i.e., -60° and 60° in the first condition, and -120° and 0° in the second condition. Fig. 8A and B display the subject's saccade trajectories towards targets flashed at the Earth-centric cardinal directions (indicated by numbers as hours on a clock face). As can be seen, the two conditions lead to different response saccades, with more pronounced errors in the latter condition (Figure 4.8B). Since the amount of intervening rotation ($\Delta\rho$) was the same in both cases, this result seems to argue against egocentric coding with updating of target location.

To further quantify this for all trials, we computed the difference in the saccadic response direction ($\Delta\text{Response Error}$) for all possible pairs of trials that had identical amounts of intervening body rotation, but different combinations of initial and final body tilt angle. Egocentric updating would predict this difference be zero ($\Delta\text{Response Error}=0$), whereas allocentric, world-centered coding would predict this difference be equal to the difference in the net A-settings for the two trials ($\Delta\text{Response Error} = \Delta A_{\text{net}}$, with $\Delta A_{\text{net}} = [A_2 - A_1]_{\text{trial2}} - [A_2 - A_1]_{\text{trial1}}$).

Using this pair-wise comparison, we analyzed all possible combinations of trials contained in our dataset. Figure 4.9 presents the results of this analysis for each subject. Each data point depicts the difference in saccadic error as function of the difference in net A-setting between two trials with identical amount of body rotation. In case of egocentric coding, the data should scatter around a line with slope 0 ($y=0$), whereas the data should fall along the diagonal ($y=x$) in case of allocentric coding (dashed line). For all subjects, we found significant correlations, $0.36 < r < 0.77$ (ANOVA, $p < 0.001$). Averaged across subjects, the slope was significantly different from zero (t-test, $p < 0.001$), but not significantly different from 1 (t-test, $p = 0.15$). Thus, this analysis shows that our data are most consistent with an allocentric coding of visuospatial memory, even though it does not entirely rule out an egocentric contribution.

Finally, one could argue that the conclusion of the latter analysis would be invalid if the amount of perceived intervening body rotation depends on the initial tilt position (ρ_1) and final tilt position (ρ_2). Even in an egocentric model, this would allow for various amounts of updating in situations in which there is equal amount of rotation ($\Delta\rho$) but different initial and final body orientations. Therefore, we performed a final experiment to control for this contingency (see Methods). Five subjects were tested




Figure 4.8. Spatial memory readouts for two testing conditions that had identical intervening body rotation (+120°), but different combinations of initial and final body positions, for one subject (JG). Traces show the saccade trajectories (12 trials) towards memorized targets that were presented, prior to rotation, on one of the four world-fixed cardinal axes. **A.** $\rho_1 = -60^\circ$; $\rho_2 = +60^\circ$. **B.** $\rho_1 = 0^\circ$; $\rho_2 = +120^\circ$. Saccades in A and B show large but different directional errors, despite the equal amount of rotation.

again, similarly as in the rotation paradigm, and verbal reports about the amount of their perceived rotation were obtained. Figure 4.10A shows these estimates of $\Delta\rho$ as function of actual $\Delta\rho$, for one subject. A linear fit with a slope of 0.89, shown by the solid line, captures this pattern very well ($r=0.99$). In all subjects tested, we found high correlations ($r>0.98$), and slopes in the range of 0.89 to 1.07, as demonstrated by Figure 4.10B. This means that whatever the computations or signals involved, our subjects can estimate the amount of intervening rotation in roll quite accurately, much better than could be expected on the performance in the spatial memory task (see Figure 4.7). To investigate whether the amount of estimated rotation had any systematic relationship with starting orientation, we fitted the following relationship to the data:

$$\text{Estimated } \Delta\rho = a \cdot \Delta\rho + b \cdot \rho_1 + c \cdot (\rho_1)^2 \quad (\text{Eq. 4.1})$$

Coefficient a specifies the linear dependence of the perceived $\Delta\rho$ on the actual amount of rotation, while parameters b and c represent any first- or second-order effect of ρ_1 on




Figure 4.9. Difference in response errors ($\Delta\text{Response Error}$) for pairs of trials that had identical egocentric but different allocentric predictions. Scatterplots for all six subjects display data of all possible comparisons. Plots show the difference in response error as a function of the difference in the net A-error ($\Delta A_{\text{net}} = [A_2 - A_1]_{\text{trial1}} - [A_2 - A_1]_{\text{trial2}}$). Egocentric scheme predicts a slope of zero; allocentric coding predicts a slope of one. Linear regressions (solid lines) favor the allocentric predictions.




Figure 4.10. Verbal reports on the perceived amount of rotation in the rotation paradigm. **A.** Results as function of actual $\Delta\rho$, for one subject. A linear fit line had slope 0.89 and correlation $r=0.99$. **B.** Best-fit lines for all subjects tested. **C.** Coefficient values of the fit (\pm SD) that relates the perceived amount of rotation to the actual amount of intervening rotation $\Delta\rho$ and the initial body position (ρ_1), following $\Delta\rho = a \cdot \Delta\rho + b \cdot \rho_1 + c \cdot (\rho_1)^2$. *Coefficient significantly different from zero.

the estimation of $\Delta\rho$. The values of each of these coefficients are presented in Figure 4.10C, for all subjects separately. As can be seen, the estimated $\Delta\rho$ depends solely on the actual $\Delta\rho$, by values for a that were significantly different from zero (t-test, $p<0.001$). This was not the case for the b and c coefficients (t-test, $p>0.05$), indicating that the estimates of $\Delta\rho$ did not depend on the starting orientation. Thus, our control experiment shows that the perception of the amount of rotation is linearly related to the actual amount of rotation, and not dependent on the starting and finishing orientations of the body relative to gravity. This confirms the assumption of the egocentric model, and the conclusions above.

Discussion

We have designed a novel paradigm to test whether human subjects code spatial memories of space-fixed targets during whole-body rotations in an egocentric or in an allocentric frame of reference. Our test exploited the well-documented fact that subjects, when tilted sideways in the dark, make systematic errors in indicating the world-centered cardinal directions (Mittelstaedt 1983; Kaptein and Van Gisbergen 2004). We observed and quantified this distortion of the world-fixed reference frame in all our subjects, as illustrated by Fig. 3. We then investigated whether this distortion would be incorporated in the accuracy of saccades directed at memorized locations of visual targets briefly presented prior to a whole-body rotation in roll. Our results show clear evidence for this (Figure 4.7). The memory-guided saccades showed an error pattern that was qualitatively and quantitatively predicted by the combination of subjective distortion of the Earth-frame when storing the memory (at initial tilt angle) and probing the memory (at final tilt angle). This is suggestive for the use of an allocentric frame of reference to represent the location of the target in the present experiments, as we will further argue below.

In contrast with the observed systematic errors in spatial localization in our experiments, previous studies have shown that spatial memory copes well with horizontal and vertical body rotation (Mergner et al. 2001; Israel et al. 1999; Baker et al. 2003; Blouin et al. 1998), provided that the vestibular afferents veridically reflect head rotation (Blouin et al. 1998; Mergner et al. 2001). Is there a discrepancy between these results and our current findings? Perhaps, but it should be emphasized that the present study tested for large body rotations in roll, in which the body changes orientation relative to gravity, which may complicate sensorimotor processing. So far, few studies have tested object localization in space during torsional body movements. In fact, our previous study was the first in this domain (Medendorp et al. 2002), in which we found that subjects made almost no systematic errors when compensating for active head torsions within their anatomical limitations ($-45^\circ - 45^\circ$). One could explain this difference by the availability of additional signals, such as neck efference copies or neck proprioception, that would allow subjects to improve their performance. However, the fact that the current report demonstrates marked response errors in a similar paradigm is more likely due to testing within a wider motion range ($-120^\circ - 120^\circ$). More specifically, when we analyzed performance in the limited amount of trials comparable to those of Medendorp et al. (2002) (here, trials with rotations from $\pm 60^\circ$ and $\pm 30^\circ$ to the upright position) we virtually found no errors either. Moreover, it is known that A-errors also occur when subjects actively tilt their bodies such that the head adopts orientations beyond 90° (Van Beuzekom et al. 2001). It is known from

previous studies, and confirmed here (Figure 4.3), that the distortion of the subjective Earth-frame is small for head tilts $< 60^\circ$ (Mittelstaedt 1983; Van Beuzekom and Van Gisbergen 2000). Because performance is accurate for these head tilts, studies that examine spatial memory accuracy in this tilt range cannot readily distinguish between the underlying frames of reference.

Our evidence that, in the present task conditions, it is indeed an allocentric frame of reference in which a target location for a saccade is encoded is as follows.

First, saccadic response errors cannot be attributed to a memory degradation effect since they did not occur in the absence of intervening body rotation (Figure 4.4). The errors that occurred were virtually negligible and did not show any systematic relationship with tilt angle (Figure 4.5). From these stationary tilt results, it is interesting to note that saccadic performance is rather accurate, while yet the subjects' perception of external space is so severely distorted. Should these observations be seen as a one more demonstration that visual perception and action are dissociated and employ different frames of reference (Goodale and Milner 1992)? Not necessarily. It could also mean that the processing errors that occur to establish a world-centered memory representation of the visual stimulus are cancelled by the errors involved in transforming this allocentric representation into an eye-centered saccadic command at later time. At least, as we will further argue, during torsional whole-body rotations, the brain seems to store world-fixed object locations in a (perhaps perceptual) allocentric reference frame (Wexler 2003), rather than in an action-oriented egocentric frame of reference.

Second, the size of the errors in the direction of the saccades in the rotation paradigm is in good correspondence with the predictions of the allocentric coding scheme (Figure 4.7A,C). In contrast, the egocentric model would predict much smaller errors, taking the rather accurate rotation percepts in the control experiment in Figure 4.10 as measure for the quality for the rotation signal. It is interesting to note that in this control test, subjects perceived the amount of rotation so well, despite the constant velocity stimulation, which is a less effective stimulus for the semicircular canals.

Third, we observed clear differences in the saccadic responses for two trials requiring equal amounts of egocentric updating (see Figures 4.8 and 4.9) which makes it problematic to accept this coding scheme as interpretation of our results. In two such trials, the semicircular canals receive identical stimulation, while the otoliths are stimulated differently. Therefore, saccade performance in our task seems more related to a spatial reference frame, established by the otoliths, than to one constructed using the canal signals. One could still argue in favor of egocentric updating assuming the rotational updating signal is detected by the otoliths in a tilt-dependent fashion. Following this interpretation, the otolith signals are gain-modulated with a factor

smaller than one for tilt-positions further away from upright. This would allow for various amounts of updating in situations in which there is equal amount of rotation ($\Delta\rho$) but different initial and final body orientations. However, it has been shown that subjects can estimate their body orientation in space rather well, at each tilt angle (Kaptein and Van Gisbergen 2004). Moreover, as our control, rotation estimation experiment showed (see Figure 4.10*B,C*), there was no significant effect of initial tilt position on perceived $\Delta\rho$, making the idea of a variable tilt-dependent gain factor rather unlikely.

Fourth, we can rule out a possible under-compensation for ocular-counter roll (OCR) as an explanation for our results. Previous studies have reported that the eyes counter-rotate by less than 10° within their orbits when the head rotates (Bockisch and Haslwanter 2001; Klier and Crawford 1998). This value is virtually negligible in relation to the size of errors that we have observed here. Moreover, the fact that we observed no or only small response errors in the static condition, where subjects were tilted but underwent no intervening body motion (Figure 4.5), argues against the OCR as an important factor in the explanation of our results.

Taken together, our results are most consistent with an allocentric, world-centered coding of a spatial memory during whole-body rotations in roll. Previous work has strongly suggested that the brain is capable of constructing an inertial, world-centered representation of head velocity and position during tilt rotations (Merfeld et al. 1993a, 1993b, 1999; Angelaki and Hess 1994; Hess and Angelaki 1997; Pettorossi et al. 1998). Indeed, earlier reports have shown that subjects have a nearly veridical percept of their self-orientation in space (Mast and Jarchow 1996; Mittelstaedt 1983; Kaptein and Van Gisbergen 2004). Since vestibular afferent signals are coded in head-centered coordinates, these signals must be centrally transformed relative to an inertial, gravitocentric reference frame (Hess and Angelaki 1999). This also means that the distorted Earth reference frame reflects a property of a central computation for a gravitocentric, inertial representation rather than an inaccuracy in the underlying source signals (Van Beuzekom and Van Gisbergen 2000; Eggert 1998). The striking relationship between the saccadic targeting responses and the A-errors, even in individual trials, suggests that the same internal representation, perhaps anchored to the direction of gravity, may underlie these observations (Figure 4.7). We reiterate that the response measures here were saccades. It would be useful to perform experiments that exploit this paradigm in other motor systems

Our conclusions agree with results from other behavioral studies that claimed that the brain incorporates allocentric information when directing saccades to visual stimuli (Honda 1999; Dassonville et al. 1995). Using pointing, Carrozzo et al. (2002) have shown that when allocentric cues are given about the relation of targets, these cues are

used to code target location. In this study, subjects who were instructed that all possible targets would fall along an (imaginary) line showed errors in the endpoints of pointing movements to these remembered locations that were aligned with this line. Similarly, when a visual target is presented within the context of a large frame whose center is located left or right of the observer's midline, the perceived location of the target is biased towards the opposite direction, an allocentric phenomenon called the induced Roelofs effect (Bridgeman et al. 1997; Dassonville et al. 2004).

However, our findings seem at odds with previous psychophysical studies suggesting egocentric, often eye-centered, updating in other task conditions, including saccadic targeting and pointing (e.g. Henriques et al. 1998; Baker et al. 2003; Medendorp and Crawford 2002). Baker et al. (2003) directly compared saccadic precision after horizontal whole-body rotations, smooth-pursuit eye movements, and saccadic eye movements to memorized targets that remained either fixed in the world or fixed to gaze. Based on the assumption of noise propagation at various processing stages in the brain, they rejected explicit world- or head-centered representations as an explanation of their results. An important difference between our study and the one by Baker et al. (2003) is the change of the body relative to gravity. It is conceivable that the brain relies on this allocentric cue, if readily available, and employs the invariant direction of gravity as a reference for storing target locations during torsional body motion. This would also be compatible with the recent results of Klier et al. (2005), who found reduced performance in a spatial updating task that lacks useful gravitational cues. In this respect, more work is needed to elucidate how the present results on allocentric, world-centered coding generalize to other movement situations. It could be argued that in conditions lacking such allocentric cues, e.g., during horizontal body rotations, the brain resorts to sole egocentric coding (and updating) of remembered target locations (Baker et al. 2003; Medendorp et al. 2003b).

To reconcile these various findings, the suggestion can be made that the reference frame used to encode a spatial memory is not fixed, but may depend on the sensory context and the task at hand (Battaglia-Mayer et al. 2003; Hayhoe et al. 2003). Alternatively, these findings could imply that the brain can concurrently define information in multiple frames of reference, co-existing at the same time (Bridgeman et al. 1997; Carrozzo et al. 2002; Snyder et al. 1998). Egocentric representations by themselves may not always represent the most efficient means to code information. In most favorable circumstances, the brain may interchange information between allocentric maps and egocentric representations to optimize motor behavior (Crawford et al. 2004). In this respect, one should note that an allocentric representation alone cannot drive the motor response. Ultimately, allocentric information must be transformed backwards into an egocentric representation, to limb-, eye- or head-related

coordinates, for motor planning, requiring non-linear operations to deal with properties of 3-D rotations (Klier and Crawford 1998; Crawford et al. 2004; Medendorp et al. 2002).

Finally, it remains a matter of speculation how and where the allocentric representation that we have probed here, using saccadic eye movements, is encoded in the brain. Hippocampal regions are known to construct allocentric memory representations. However, they are only implicated in the active control of long-term spatial memory, for delays longer than 20 s, while the present study probed a short-term memory representation (Pierrot-Deseilligny et al. 2002). Moreover, it is not very likely that hippocampal regions are involved in the memory of single target locations for saccades (Muri et al. 1994). A more likely place to look for is an area where information about body position in space is integrated with information coded in the coordinates of the retinal frame. A major multisensory center for this integration seems to be the parietal cortex. There is currently evidence for separate body and world-referenced coding of stimulus locations in parietal cortex, based on implicit representations constructed by gain-modulation of visual signals (Snyder et al. 1998). A possible role could also be attributed to the parieto-insular vestibular cortex (PIVC), which has been implicated in the perception of verticality and self-motion (Brandt and Dieterich 1999). Neurons in the PIVC receive inputs from both the semicircular canals and the otoliths, as well as visual and neck proprioceptive inputs (Grusser et al. 1990; Brandt and Dieterich 1999). In addition, patients with lesions in the PIVC have been shown a distorted perception of verticality (Brandt et al. 1994; Yelnik et al. 2002). Thus the PIVC may have the signals necessary to implement the spatial memory representation that we have revealed.

To conclude, the present study clearly showed that humans make errors in directing saccades to remembered target locations, presented prior to a whole-body rotation in roll. The errors could be linked to an internal mechanism that keeps target locations in an allocentric reference frame, rather than an action-oriented egocentric frame of reference. It remains a challenge to understand how and where the central computations underlying this finding are implemented by the brain.

Chapter 5

Repetition suppression dissociates spatial frames of reference in human saccade planning

This chapter has been submitted for publication

To understand how the brain processes and transforms spatial information for movements, the notion of a reference frame is indispensable (Soechting and Flanders 1992). Using this concept, monkey electrophysiological evidence has shown that movement-related neurons employ a variety of reference frames, anchored to eyes, head, other body-parts, or world (Andersen and Buneo 2002; Martinez-Trujillo et al. 2004; Olson 2003). However, it is unclear to what extent this information, which is extracted from post-synaptic action potentials of relatively few task-related pyramidal neurons, can be related to the computations of larger neuronal populations (Logothetis 2008) and to other species, including humans.

Data on spatial reference frames of large neuronal assemblies in the human brain are still scarce. A few recent fMRI studies addressed this issue using topographic mapping procedures. Examining how topographic maps of target locations change as a function of eye position allows to distinguish between retinal (eye-centered) or non-retinal (head/body/space centered) reference frames (Medendorp et al. 2003a; Sereno and Huang 2006; Gardner et al. 2008). As a result, Medendorp et al. (2003a) demonstrated the existence of an eye-centered saccade-and-reach area in parietal cortex.

However, neurons may not always be topographically arranged along the dimensions of the reference frame they employ. A brain area could encode information in a particular reference frame even if the respective neurons do not show an orderly spatial organization according to the value of that particular parameter. This is likely the case for regions involved in movement control, where multidimensional motor constraints must be organized into a two-dimensional map (Graziano and Aflalo 2007).

Repetition suppression (RS) offers a potential solution to investigate the reference frames used in the neural control of movement without relying on the special case of an orderly topographic arrangement of the relevant neurons. RS is based on the observation that repeated processing of a given stimulus feature leads to a reduction of neural activity in neurons tuned to that particular feature (Desimone 1996). By varying the property of the stimulus across different dimensions, the features processed in a given brain region can be uncovered. While many fMRI studies have successfully used this technique in studies of perceptual representation (see Grill-Spector et al. 2006, for review) and action observation (Hamilton and Grafton 2006, 2008; Dinstein et al. 2008), to date this method has not been applied to examine sensorimotor control.

Applying RS methods to study the motor system is not trivial. Most types of movements are intrinsically variable (Churchland et al. 2006), which would easily violate the RS assumption of repeated processing. Here, we circumvent this problem by studying saccades, which are highly stereotyped movements that can be well replicated in subsequent trials.

In this study, we used RS methods to investigate the reference frames for saccade shifts in the human brain. We found a reduction of the BOLD signal in the main cortical centers for saccades (IPS, FEF, and SEF) when the target location was repeated in eye-centered coordinates, but not during a repetition in a non-retinal frame. We conclude that motor commands from these centers, of which only some have a topographic distribution of spatially-tuned neurons (IPS and FEF), encode saccade shifts in eye-centered coordinates.

Methods

Subjects and ethical approval

Eighteen healthy subjects with normal or corrected-to-normal vision participated in the study (8 female, 10 male, aged 20-37 years). Three subjects were left-handed; one subject was aware of the exact purpose of the experiment. All gave written informed consent in accordance with the guidelines of the local ethics committee (CMO Committee on Research involving Human Subjects, region Arnhem-Nijmegen, The Netherlands). Subjects practiced the task 1-2 days in advance in a mock setup outside the scanner to ensure that the task and paradigm were correctly understood. In addition, a few practice trials were performed inside the scanner just prior to the experiment.

Experimental setup

Subjects were lying supine in the scanner, with their heads tilted 30° with respect to the scanner bed by means of a wooden support board that was attached to the bed. This enabled the subjects to view all stimuli directly without mirrors, making the task as natural as possible. Their head was fitted inside a phased-array receiver head coil. The head and neck were stabilized within the head coil using foam blocks and wedges. A foam block was also placed underneath the knees, and in some subjects the elbows and neck were further supported by cushions to make them feel more comfortable.

A stimulus device consisting of seven horizontally placed yellow-colored light-emitting diodes (LEDs), was attached to an arch of about 40 cm height that was placed over the subject's hip, at a viewing distance of 34 cm. The central LED was aligned with the subject's body midline; three peripheral LEDs were located on either side, at eccentricities of 4.5, 9 and 18° from the central LED. This setup allowed subjects to view all stimuli with a comfortable, slightly downward gaze direction relative to the head.

Stimulus LEDs were controlled using Presentation software (Neurobehavioral Systems, San Francisco, CA, USA). Position of the left eye was recorded using a long-range infrared video-based eyetracker (SMI, Teltow, Germany) at a frequency of 50 Hz.

MR settings

Anatomical and functional images were obtained on a Siemens 3 Tesla MRI scanner (Siemens Trio, Erlangen, Germany). Functional images consisted of 32 axial slices acquired by a gradient-echo planar imaging sequence using an eight-channel phased-array receiver head coil (slice thickness 3.0 mm, gap = 17%, in-plane pixel size 3.5 x 3.5 mm, TR = 2000 ms, TE = 35 ms, FOV = 224 mm, flip angle = 80°). In total, 1140 functional images were obtained in one run, lasting 35 minutes. Hereafter, high-resolution anatomical images were acquired using a T1-weighted MP-RAGE sequence (192 sagittal slices, voxel size 1.0 x 1.0 x 1.0 mm, TR = 2300 ms, TE = 2.02 ms, FOC = 256 mm, flip angle = 8°).

Experimental paradigm

The experiment took place in complete darkness; only the stimulus LEDs were visible. Subjects performed a memory-guided saccade task, using a rapid event-related repetition suppression (RS) design (Figure 5.1A, upper panel). A trial started with a subject fixating an illuminated stimulus LED (Fixation Point, F). Then, after a period of 3 s, one of the other stimulus LEDs flashed for 200 ms, which served as the target stimulus (S) for the pending saccade. This was followed by a 3.8 s memory delay during which the subject maintained fixation on F. Subsequently, F was extinguished, which was the go-cue for the subject to make the saccade to S, as accurately as possible. Then, 1 s later, the next trial started, with F at a different location than S in the previous trial. Each trial lasted eight seconds. The total experiment consisted of 36 blocks of 4 trials each.

In each trial, both F and S could be presented at one of five possible locations, at -18°, -9°, 0°, +9° or +18° from the center. Combinations of F and S were chosen pseudo-randomly; we did not test trials in which S=F since this implied no saccadic response. In the majority of trials (85 %), the angular separation between F and S was 9° to exploit the fact that 9° saccades may drive higher BOLD responses than 18° saccades, based on the overrepresentation of the central visual field in several visual and oculomotor regions (Ben Hamed et al. 2001).

Repetition suppression effects were elicited by systematically manipulating target location over successive trials in a 2x2 design, with conditions retinal and non-retinal coordinates (labeled as R and N, respectively), and levels novel and repeated (labeled as n and r, respectively). E.g., as illustrated in Figure 5.1A, the retinal location




Figure 5.1 A. Experimental paradigm. Upper panel. A typical novel trial t started with the illumination of a fixation LED (F). After 3 s, a saccadic target LED (S) was flashed for 200 ms in the visual periphery, while subjects kept fixation at F. After a memory delay period of 3.8 s, F was extinguished, which cued the subject to make a saccade to S. 1 s later the next trial started. Lower panels. In a subsequent repetition trial $t+1$, S could be presented at either the same retinal location as in the previous trial, while the location was novel in non-retinal (head-centered) coordinates (left), or at a novel retinotopic position, but at the same non-retinal location (right). Alternatively, the targets location could be either novel or repeated in both coordinate frames (not shown). Both fixation and target stimulus LEDs were yellow-colored and had the same luminance (difference in LED colors in the figure is for clarification purposes only). **B.** Eye traces of one subject over the time course of 20 trials with F at 0° and S either at -9° (black traces) or 9° (red traces). The subject keeps fixation throughout the trial, also during target stimulus presentation. After the go cue, response saccades are consistently made toward the location of the remembered target.

of a target presented in trial t , could be repeated in the next trial $t+1$, while the non-retinal location was novel (lower left panel; retinal repeated, non-retinal novel; RrNn). Alternatively, the retinal location of the target in trial $t+1$ could be novel compared to the preceding trial t , while the non-retinal location was repeated (RnNr, lower right panel). Finally there were two types of trials (not shown) in which the location of the target was either repeated or novel in both coordinate frames (RrNr and RnNn, respectively).

The first trial of each block was not included in the RS analysis in order to avoid carry-over effects from the previous block. The remaining 108 trials consisted of 36 RnNn trials, and 24 trials of each of the other three types of trials (RrNn, RnNr, RrNr). A target's retinal or non-retinal location was never repeated more than once in a row in order to get the strongest RS effects and avoid adaptation fatigue (Van Turenout et al. 2003). Target directions were balanced across the visual and craniotopic hemifields; average amplitudes were the same across the four conditions.

After each block of four trials, subjects performed a so-called washout task to allow the BOLD signal to return to baseline level after several RS trials, alleviating possible longer lasting RS effects. The start of this washout task was indicated by three brief subsequent flashes of two targets (first $-4.5^\circ/+4.5^\circ$, then $-9^\circ/+9^\circ$, finally $-18^\circ/18^\circ$), followed by the onset of the central LED for a jittered duration (1.4-12.6 s). Subjects were instructed to fixate this LED and track it as it subsequently jumped to different locations after each 250 ms, eight times in total. These locations were balanced across directions and were evenly distributed across the 7 LEDs on the stimulus device. The washout task ended by a period of central fixation (1.4-14.0 s) followed by again the same three short flashes, but now in opposite order. Each washout period lasted 15.2 – 32.0 s (mean 23.1 s). After each 6 blocks and their associated washouts, subjects had a rest period of 30 s, during which there was no visual stimulation and they could freely move their eyes. The total experiment lasted 60 minutes, including practice and anatomical scanning.

Behavioral analysis

Eye movement data (horizontal component) were processed separately per block of four trials and calibrated in degrees based on the fixation data of the following washout period. This generally yielded calibration accuracies better than 1.5° . Figure 5.1B show the eye traces of a typical subject from central fixation to a remembered target location at either 9° (red) or -9° (black), in relation to the temporal order of events (see Figure 5.1A). As shown, this subject maintained fixation during the presentation of the target cue, and made eye movements with latencies of about 200 ms in the correct directions after the fixation target was turned off. Due to technical problems, eye-movement data

of one subject were lost for the last 12 blocks of trials. We used the eye recordings to identify error trials, which were defined as trials in which subjects did not keep fixation when required, or made saccadic responses that were anticipatory or into the wrong direction. Eye traces were also used to determine reaction times. On average, 9 ± 4 (SD) trials per subject were discarded based on these criteria. For the remaining trials, average fixation accuracy was 1.8° (SD = 1.4°) across subjects. Accuracy of saccades to the remembered targets, in degrees of visual angle, was 3.0° (SD = 1.2°) across conditions. This confirmed that the saccades were driven by the memory of the actual targets and were not simply guided stereotypically to the left or right.

Preprocessing of fMRI data

fMRI data were analyzed using BrainVoyager QX (Brain Innovation, Maastricht, The Netherlands). Subsequent analyses were performed using Matlab (The Mathworks). The first five volumes of each subject's data set were discarded to allow for T1 equilibration. Functional data were first corrected for slice scan time acquisition and motion. Subsequently, the data were temporally filtered using a high-pass filter with a cutoff frequency of 1/268 s. The functional images were co-registered with the anatomical scan and transformed into Talairach coordinate space using the nine-parameter landmark method (Talairach and Tournoux 1988). Finally, the images were smoothed with an isotropic Gaussian kernel of 8-mm full-width-at-half-maximum.

Statistical inference and regions of interest

We used a standard general linear model (GLM) in a first analysis of the data. For each subject we defined 16 regressors. The first modeled the 2-s fixation periods at the beginning of each trial, starting at one second after the presentation of the fixation LED, as well as the fixation periods during the washout periods. A second regressor characterized the periods of 0.2 s during which the target stimulus was presented. Four other regressor functions characterized the subsequent working memory interval according to the 2 x 2 design of conditions Retinal and Non-retinal locations with levels Novel and Repeated. These regressors (RnNn, RrNn, RnNr, and the RrNr) covered the 3.8 s delay period starting with target offset until fixation point offset (go cue). Saccade periods were modeled by a seventh regressor. This included the first second after the go cue, the first second after presentation of the fixation LED of the next trial and the saccade periods during the washout periods.

In addition to these seven regressors, we used nine regressors of non-interest. One modeled the delay period of the first trial of each block, another regressor was composed to capture the delay periods of error trials. We also designed one regressor to model the periods of rest and the intervals in which the cues for the start and end of

the washout period were presented. All regressors were defined as boxcar-functions over the time interval they described and were convolved with a hemodynamic response function (modeled using a two-gamma model function with response undershoot ratio of 6, time to response peak of 5 s and time to undershoot peak of 15 s). The last six regressor functions represented the head motion, based on the six parameters provided by BrainVoyager's motion-correction algorithm.

Individual subject GLMs were corrected for serial correlations in the time courses. Random effects group analyses were performed to test effects across subjects, using the false discovery rate (FDR) controlling procedure to correct for multiple comparisons, at the $q(\text{FDR}) < 0.05$ significance level (Genovese et al. 2002). Using a random-effects group analysis, we first determined the regions that show significant activity during the working memory period for the saccade relative to baseline, thus by computing the contrast $(RnNn + RrNn + RnNr + RrNr)/4 > \text{fixation}$. From the activation maps, we selected three bilateral regions of interest (ROI), known to be important regions in saccade generation: FEF, SEF and a region in the intraparietal sulcus (IPS). Each ROI was defined as all the contiguous voxels that exceeded a threshold of $q(\text{FDR}) < 0.05$ within a cubic cluster of $8 \times 8 \times 8$ mm, centered at the points of peak activation.

Linear deconvolution

In a second analysis, we used finite impulse response deconvolution to extract the activation profiles in the ROIs for each of the four RS conditions (RnNn, RrNn, RnNr, and RrNr). In this approach, the BOLD data were first resampled into 0.5 s time intervals. Then, for each condition, a set of 31 impulse responses (one impulse per 0.5-s volume) was aligned to the start of each trial in the group. Together, the 31 impulse regressors for a given condition modeled the activation time course for trials in this condition with one point per second over 15 s. Thus, each group of trials yielded 31 columns to a subject's GLM design matrix, with ones at the appropriate locations, to model the 31 impulse functions for that trial group (Serences 2004; Brown et al. 2006). Fitting this design matrix to the resampled data automatically deconvolves the time series of each RS condition (Brown et al. 2006), without making any assumption about the shape of the activation profile. Next, for each RS condition and each ROI, a mean signal and standard deviation were computed across subjects. Differences between conditions capture the RS effects in either reference frame. That is, retinal RS follows from $(RnNn + RnNr) - (RrNn + RrNr)$ and non-retinal RS is computed as $(RnNn + RrNn) - (RnNr + RrNr)$. Statistical significance was tested using paired t-tests and repeated-measures ANOVAs at the $p < 0.05$ confidence level.

Results

Behavioral performance

Table 5.1 shows performance (defined as correct fixation and saccade direction) and saccade latencies for each target condition (RnNn, RrNn, RnNr, and RrNr, see Methods). Across subjects performance was >93% correct, in all conditions. A 2x2 repeated-measures ANOVA with repeated versus nonrepeated trials and retinal versus non-retinal target locations as factors revealed no significant main ($F(1,17) < 3.98$, $p > 0.062$) or interaction effect ($F(1,17) = 1.30$, $p = 0.27$). The mean latencies of the saccadic responses were 217 ± 69 ms (mean + SD) across the four conditions. The differences among the four conditions were not statistically significant ($F(1,17) < 0.86$, $p > 0.36$). Together, the behavioral results indicate that possible differences in corresponding fMRI activations cannot be related to different levels of task performance.

fMRI activation data

ACTIVATION MAPS DURING DELAY PERIOD

Using a random-effects group GLM analysis across all 18 subjects, we first identified the cortical areas involved in saccade preparation by comparing activity during the memory period (from target presentation to go cue), with normal fixation periods (see Methods). Figure 5.2A and B show two anatomical views of these results, in neurological convention, thresholded at $q(\text{FDR}) < 0.05$. In Figure 5.2C and D, this activation map is rendered onto an inflated representation of the left hemisphere of one of the subjects. Consistent with previous results, a bilateral network of eye-movement related cortical areas was activated (Schluppeck et al. 2005; Curtis and D'Esposito 2006; Brown et al. 2004; Connolly et al. 2002). This included a region along the

Table 5.1. Percentage correct performance and mean reaction times for each condition

Target location condition	Performance (%)	RT (ms)
Novel retinal, novel non-retinal	94.5 ± 5.3	215 ± 69
Repeated retinal, novel non-retinal	93.8 ± 7.3	220 ± 74
Novel retinal, repeated non-retinal	94.7 ± 6.2	215 ± 83
Repeated retinal, repeated non-retinal	96.3 ± 3.5	220 ± 64




Figure 5.2. Brain activation during the delay period, averaged across all 18 subjects ($P < 0.05$, FDR-corrected), presented in 2 anatomical views in neurological convention (A, B), and on an inflated representation of the left hemisphere of one of the subjects (C, D). A parietofrontal network is activated, including areas on the banks of the intraparietal sulcus (IPS), the frontal eye field (FEF), supplementary eye fields (SEF) and the ventral premotor area (PMv). Furthermore, bilateral activity was found in the frontal insular cortex (Ins), as well as a small region within the ventral part of the postcentral gyrus (C) and midbrain (D).

intraparietal sulcus (IPS), which might be the human analog of monkey area LIP (Medendorp et al. 2003a, Connolly et al. 2007; Sereno et al. 2001). In the frontal cortex, we found significant voxels at the junction of the precentral sulcus and the superior frontal sulcus, probably corresponding to the frontal eye fields (FEF; Paus et al. 1996; Brown et al. 2004). More medially, significant voxels were found along the interhemispheric fissure, extending onto the dorsal cortical surface, which can be classified as the supplementary eye fields (SEF; Picard and Strick 2001; Grosbras et al. 1999; Brown et al. 2004). More laterally in the frontal cortex, significant responses were found in voxels covering the precentral sulcus, corresponding to the ventral premotor area (PMv; Picard and Strick 2001; Beurze et al. 2007). Finally, significant

activations were observed bilaterally in the frontal insular cortex (Ins; Grefkes et al. 2004; Beurze et al. 2007).

Table 5.2 lists the mean Talairach coordinates (in mm) of the peak voxel within each region, together with the corresponding t-values across subjects. From these regions, we subjected the bilateral IPS, FEF, and SEF to a careful investigation of the RS effects.

REFERENCE FRAME-DEPENDENT REPETITION SUPPRESSION

Can repetition suppression reveal which frames of reference are used to code the representation in these oculomotor regions? Given our hypotheses, we may predict that, when the retinal location of a target is repeated in subsequent trials, voxels will show an attenuation of their BOLD-activation when the underlying neuronal populations code target location in a retinal reference frame, but not if they code in a non-retinal reference frame. Conversely, regions that code the non-retinal (e.g. craniotopic) location of a target will only show BOLD adaptation when the non-retinal location of the target is repeated. Of course, it is also possible that a region would be best characterized by a mixture of these two frames.

Figure 5.3A shows the reconstructed BOLD response of the left and right IPS region over a time course of 11 s, averaged across subjects (see Methods). Repeated trials (gray) had the same target location as the previous trial (black) in retinal coordinates. Time $t=0$ s denotes the onset of the target stimulus; $t=4$ s the go-cue for

Table 5.2. Brain regions activated during delay period (delay>fixation)

Anatomical region	Functional label	Side	x	y	z	t-value
Frontal insular cortex	Insula	L	-31	21	11	4.48
		R	28	20	10	4.37
Intraparietal sulcus	IPS	L	-24	-57	48	3.37
		R	9	-61	47	3.15
Superior frontal sulcus	FEF	L	-25	-11	52	7.12
		R	20	-9	48	7.61
Medial frontal cortex	SEF	L	-6	-4	58	5.12
		R	2	-3	56	5.08
Precentral sulcus	PMv	L	-56	-3	37	4.82
		R	56	9	27	5.63




Figure 5.3. Group results. **A, B.** Reconstruction of the hemodynamic responses in the IPS averaged across all subjects, for novel (black traces) and repeated trials (gray traces) in retinal (**A**) and non-retinal (**B**) coordinates. **C, D.** Average difference between repeated and novel trials, together with 95% confidence intervals. LH, left hemisphere; RH, right hemisphere. Gray areas indicate the periods over which the differences between the novel and repeated trials were taken.

the saccade. As shown, in both novel and repeated trials, after the brief presentation of the target stimulus ($t=0$ s), cortical activation during the first delay period shows first a phasic response (time interval 0 to 4 s), followed by a tonic response (time 4 - 6 s). Then, at time 7 - 10 s, there is again a strong increase in cortical activation, caused by the execution of planned saccade and the subsequent saccade to fixate a new fixation point (see Methods). The activity, in particular the early phasic and tonic activity is suppressed in repeated trials compared to novel trials, in both hemispheres, which would be consistent with the prediction of the retinal model. Figure 5.3C illustrates this more clearly, by showing the mean difference (\pm 95% confidence intervals) between the activation patterns during novel and repeated trials (average repetition suppression in retinal coordinates). Across the entire period, BOLD activation during repeated trials is significantly lower than during novel trials (paired t-test, $p<0.001$), with the suppression effects most pronounced during the tonic memory phase.

To investigate whether the retinal representation in the IPS is intermingled with a non-retinal representation, we compared novel and repeated trials with the same target location in non-retinal coordinates. As shown in Figure 5.3B, activation patterns during novel and repeated trials are quite similar. Their difference is plotted in Figure 5.3D, together with the 95% confidence intervals (gray area). Across the entire time

course, the difference in activation does not significantly deviate from zero ($p>0.07$). This suggests the presence of a retinal spatial representation in the IPS, while there is no evidence for a non-retinal representation.

The results of the IPS are exemplary for those in the FEF and SEF. Therefore, to analyze the findings quantitatively for each ROI, we computed in each subject the average difference between the novel and repeated signals at three phases of the trial, indicated by the vertical gray boxes in Figure 5.3A. The resulting value is a measure for the amount of repetition suppression (RS value). We computed these RS values (corrected for the fMRI hemodynamic lag) for the stimulus-related activity (1-3s), the delay period (4-6s), and the execution phase (7-10 s). For each ROI, the amount of RS was determined across hemispheres, in both reference frames.

Figure 5.4 plots the average results of this analysis across the entire group of subjects. As shown, brain activations are significantly suppressed when a target location is repeated in retinal coordinates (black bars), for all ROIs and trial phases (repeated measures ANOVA; $F(1,17)>6.4$, $p<0.05$ in all cases). Retinal suppression was strongest during the delay phase. This confirms the observations in Figure 5.3 and illustrates the role of these regions in saccade planning. In contrast, we found no significant




Figure 5.4. Repetition suppression effects in the IPS, FEF, and SEF, at various trial phases in relation to a retinal (black bars) and non-retinal (white bars) reference frame. Data combined across hemispheres. Error bars: SE. * $P<0.05$; ** $P<0.01$, *** $P<0.001$.

suppression effects when a target location is repeated in non-retinal coordinates (red bars), during any of the three trial phases (repeated measures ANOVA; $F(1,17) < 2.6$, $p > 0.12$). Thus, the results in Figure 5.4 provide evidence for the existence of, at least, a sustained eye-centered representation in the selected saccade regions, whereas higher-order (non-retinal) representations are absent.

LATERALIZATION

To what extent are the RS findings of a retinal coding of target location consistent with the topographic organization of these areas, as revealed by lateralized cortical activity? Because we varied eye position, our paradigm allows us to distinguish between lateralized activity in retinal and non-retinal coordinates. If the spatially-selective retinal neurons are topographically organized in the selected ROIs, we would expect that targets in the contralateral visual field will generate a higher BOLD response than targets presented in the ipsilateral hemifield. Alternatively, it is possible that the retinal RS effects are not embedded in a neural map with an orderly spatial organization. Because only retinal RS effects were seen, we anticipate that none of the regions will demonstrate head-centered laterality.

To test the presence of lateralized activity in our data, we performed two GLM analyses, each using two regressors to describe target location (left or right in retinal or non-retinal coordinates) during the delay period (see also Methods). We compared the resulting beta-weights of these regressors in both GLMs, separately for each ROI. Figure 5.5A presents the differences between the activity elicited by contralateral and ipsilateral targets. For the IPS and FEF, a strong retinal lateralization was found, which was significant across hemifields (repeated measures ANOVA; $F(1,17) > 28.4$, $p < 0.001$ in both regions). In the SEF, however, there was no significant lateralized activity ($F(1,17) = 0.45$, $p = 0.51$). In combination with our RS results, this suggests that, although retinal RS effects are present in the SEF, there is no topographic distribution of these spatially selective neurons in this area.

For completeness, when targets were sorted according to their non-retinal (head-centric) location, there was no significant difference between contralateral and ipsilateral activity in any of the regions (Figure 5.5B; repeated measures ANOVA; $F(1,17) < 0.74$; $p > 0.4$ in all regions). This compares well to the RS results, which do not favor the non-retinal reference frame either. All together, our results show that repetition suppression can be used as a tool to distinguish between reference frames in frontoparietal areas involved in spatial memory processing for saccades, even when those regions lack a clear topographic organization.




Figure 5.5. Lateralized activity in IPS, FEF and SEF during the delay period, averaged across subjects. **A.** Difference in % BOLD signal, across hemispheres, between contralateral and ipsilateral target locations in retinal coordinates. A contralateral bias exists in the IPS and FEF ($p < 0.001$), but not in the SEF ($p = 0.51$). **B.** Lateralized activity when targets are expressed in terms of their non-retinal location. No directional preference for head-centric targets is observed in any of the regions. Error bars: SE.

Discussion

Identifying the computational architecture of the human brain has been a major aim in neuroscience research over the last decades. One of the key questions concerns the internal organization of the various brain regions involved in sensorimotor processing, i.e., how and why different regions provide different solutions to the underdetermined problem of mapping multidimensional motor constraints into a two-dimensional neuronal matrix (Graziano and Aflalo 2007; Kohonen 2001).

Using repetition suppression (RS) effects, we addressed a particular instance of this general issue by studying the spatial reference frames employed by three human oculomotor areas (IPS, FEF, and SEF) in the context of a delayed-saccade task (Pierrot-Desilligny et al. 2004). To do so, subjects performed trials of delayed-saccades that were repeated with the remembered target at the same location in either retinal or non-retinal coordinates. Within all regions, significant suppression effects were only observed in relation to repetition of the target location in retinal coordinates, not when the target reappeared in non-retinal coordinates (Figures 5.3-5.4). We found the time course of retinal suppression to show the strongest attenuation effects during the delay period, reflecting the important role of these regions in preparing the saccade.

We also investigated the lateralization of activity in the hemispheres when targets were presented ipsi- or contralateral in either retinal (eye-centered) or non-retinal (head/ body/space centered) coordinates. This revealed a representation of contralateral

target location in the IPS and FEF, defined in reference to the eye, which is consistent with the retinal repetition suppression effects (Figure 5.5). These findings confirm previous fMRI results on the topographic representation of saccadic movements in IPS and FEF (Sereno et al. 2001; Schluppeck et al. 2005; Kastner et al. 2007; Hagler and Sereno 2006; Medendorp et al. 2006; Curtis and D'Esposito 2006; Curtis and Connolly 2008), providing a novel empirical validation of using RS methods for studying the motor system.

Our data provides no evidence for a topographic organization in the SEF, in either retinal or non-retinal coordinates (see Figure 5.5), which is consistent with recent fMRI findings by Kastner et al. (2007). Nevertheless, just as LIP and FEF, the human SEF appears to encode saccadic movements in an eye-centered frame of reference (see Figure 5.4). These findings illustrate that, whereas these three visuomotor areas process eye-centered saccadic information, their topographic layouts suggest different use of this information. Under the assumption that the structural organization of the cerebral cortex follows the principle of maximizing smoothness of neurally encoded features (Graziano and Aflalo 2007; Durbin and Mitchison 1990), we infer that spatial features constitute a relevant dimension for IPS and FEF computations and not for the SEF, in line with a role of the latter region in guiding eye movements according to arbitrary sets of visual elements (Olson 2003).

In support of our interpretations, the lack of non-retinal suppression effects indicates that the observed retinal suppression effects cannot be due to general motor habituation or fatigue, but mark the identity of the underlying neural organization. It has been proposed that RS may be the result of a 'sharpening' of cortical representations (Wiggs and Martin 1998; Desimone 1996; Grill-and Malach 2001). A repeating stimulus can be coded more efficiently by employing fewer active neurons (Desimone 1996; Friston 2005). From a Bayesian perspective (Ma et al. 2006; Vaziri et al. 2006), this can be understood in terms of a target location of the last trial serving as a prior probability distribution for the next trial. When this prior is integrated with the new sensory evidence, the network may settle to a tighter distribution in neural space at the second repetition.

It is further important to emphasize that the eye-centered suppression effects in the IPS, FEF, and SEF are more consistent with a role in saccade planning than execution. First, suppression was more pronounced in the delay period (Figure 5.4), which is the period of saccade preparation. Also, for eye movement execution, eye-centered representations must be further transformed, as a function of eye position, by downstream mechanisms into head-centered commands for the ocular muscles (Crawford and Guitton 1997). As Figures 5.3-5.4 show, we did not find head-centered suppression effects in these regions. A third reason is that two physically identical eye

movements require also the same patterns of muscle innervations, which simply would not allow for any suppression of activity.

When comparing our results to monkey neurophysiological findings, we should keep in mind that BOLD-imaging and single unit recording are different techniques that show different things. The current notion is that fMRI informs about local information processing, whereas unit recordings report about the output stage of those computations (Logothetis 2008; Bartels et al. 2008). Despite these reservations, the present findings are for the most part quite consistent with previous mapping and single-unit experiments in monkeys (Koyama et al. 2004). Single-unit studies report evidence for an eye-centered topographic organization of saccade targets in the lateral intraparietal sulcus (Blatt et al. 1990; Ben Hamed et al. 2001) and the FEF (Bruce and Goldberg 1985; Robinson and Fuchs 1969; Schall 1991). Although many earlier human studies have reported topographic maps in the IPS and FEF (see above), the underlying reference frame has been much less studied. The present study, examining the spatial organization across different eye positions, provides solid evidence for an eye-centered topographic organization of both regions.

Debate exists about a topographic organization of saccade goals in monkey SEF (Schlag and Schlag-Rey 1987; Tehovnik and Lee 1993; Russo and Bruce 2000). Furthermore, various single-unit studies have provided evidence that SEF neurons can encode target locations in a continuum from eye-, to head-, to body- and object-centered reference frames (Martinez-Trujillo et al. 2004; Olson 2003; Schlag and Schlag-Rey 1987), perhaps to represent all possible contingencies for different task-related motor functions (Martinez-Trujillo et al. 2004). In contrast, our study has revealed a retinal code only in the human SEF, and a clear absence of topographic structure. A possible explanation for the apparent discrepancy is that the head-fixed saccade conditions here have constrained us probing representations other than those referenced to the eyes.

In conclusion, the present study exploited fMRI-RS to unveil the frames of reference employed by frontal and parietal areas during saccade planning. While our findings advance the understanding of how the human brain processes spatial information for saccades, they also support the feasibility and validity of using RS methodology in the sensorimotor domain.

Chapter 6

Summary and discussion

How the brain represents the outside world so that we can successfully interact with objects in it has been a major research question in neuroscience over the last decades. Seemingly effortlessly, we can keep track of objects in our surroundings even though the visual input to our retinas changes as much as 3-5 times per second, whenever we make a saccade. Even when objects have disappeared from sight due to movements of our eyes, head or body, we are still able to make correct eye or reaching movements towards them, assisted by our spatial memory system. In recent years, much experimental evidence has been reported that various cortical and sub-cortical brain structures store these locations in retinotopic maps, i.e., as positions relative to the current viewing direction. A necessary consequence of this coding regime is that these neural maps must be updated for every subsequent eye movement, in order to preserve the spatial representation of the objects in the world. Although it has been demonstrated that this retinal updating indeed takes place, most of this evidence was obtained using simple horizontal eye rotation manipulations. It still is to be revealed if these findings on spatial constancy can be generalized to more complex motion conditions, such as movements of the entire body like during walking or riding a car, or during disconjugate rather than conjugate eye movements. The experiments described in this thesis have extended upon our knowledge of spatial memory computations by investigating the mechanisms of spatial constancy for motor control during more intricate conditions than simple eye rotations. In addition, novel ways of detecting the characteristic signatures of spatial representations within the human brain were explored. The most important contributions that this thesis makes to the field are the following:

- 1) During active whole-body motion (lateral translations), humans use a retinal (eye-centered) representation to guide reaching movements, which is updated according to the predicted nonlinear patterns.
- 2) The distance of reach targets is coded and updated in retinal coordinates, relative to the depth of fixation, in a similar fashion as for target direction.
- 3) Gravity can serve as an anchor in spatial coding. When the body's orientation changes relative to the Earth vertical, the direction of gravity is used as a reference to code target locations for saccades in the fronto-parallel plane.
- 4) Repetition suppression (RS) in fMRI is a feasible and valid technique to investigate the nature of spatial representations in the human brain. The main cortical centers for saccades show repetition suppression effects in eye-centered coordinates, not in head- or body-centered coordinates.

In the following section, we will provide a detailed summary of each of these findings.

Spatial constancy control during lateral body motion

Chapter 2 describes a study in which we investigated how target locations are coded during whole body translations, a situation which puts major challenges on the computations that preserve spatial constancy. When our bodies, and thus our eyes, translate, visual objects behind and in front of the eyes' binocular fixation point shift in opposite directions on the retina due to motion parallax. It is not known if the brain uses retinal coordinates to compute parallax in the translational updating of remembered space or if it uses non-retinal (e.g. head-centered) coordinates to maintain spatial constancy across translational motion. We tested this by having subjects ($n=12$) view targets, flashed in darkness in front of or behind fixation, then translate their body sideways, and subsequently reach to the memorized target. Reach responses showed parallax-sensitive updating errors: errors increased with depth from fixation and reversed in lateral direction for targets presented at opposite depths from fixation. In a series of control experiments we ruled out possible biasing factors such as the presence of a fixation light during the translation, the eyes accompanying the hand to the target, and the presence of visual feedback about hand position. Quantitative geometrical analysis confirmed that updating errors were better described by using retinal than non-retinal coordinates. From these results we concluded that spatial updating for translational motion operates in retinal coordinates. Neural network simulations were presented that suggest that the brain relies on ego-velocity signals and stereoscopic depth and direction information in spatial updating during self-motion.

Spatial representations in depth constancy

Keeping track of spatial direction is often sufficient to guide motor responses such as eye or pointing movements. In some other situations, however, it is also required to have correct distance information available, for example, in reaching and grasping. How this depth constancy is achieved was examined in a study described in Chapter 3. We tested between two coding mechanisms that the brain may use to retain distance information about a target for a reaching movement across vergence eye movements. If the brain was to encode a retinal disparity representation (retinal model), i.e., target depth relative to the plane of fixation, each vergence eye movement would require an active update of this representation to preserve depth constancy. Alternatively, if the brain was to store an egocentric distance representation of the target by integrating retinal disparity and vergence signals at the moment of target presentation, this representation should remain stable across subsequent vergence shifts (non-retinal

model). We tested between these schemes by measuring errors of human reaching movements ($n=14$ subjects) to remembered targets, briefly presented before a vergence eye movement. For comparison, we also tested their directional accuracy across version eye movements. With intervening vergence shifts, the memory-guided reaches showed an error pattern that was based on the new eye position and on the depth of the remembered target relative to that position. This suggests that target depth is recomputed after the gaze shift, supporting the retinal model. Our results also confirm earlier literature showing retinal updating of target direction. Furthermore, regression analyses revealed updating gains close to one for both target depth and direction, suggesting that the errors arise after the updating stage during the subsequent reference frame transformations that are involved in reaching.

Gravitational effects on spatial constancy

One argument of why the brain may rely on retinal (or, for this matter, egocentric) representations to code the location of targets for future action is that when deprived from visual information, relevant extrapersonal (allocentric) spatial cues are usually absent. Under some conditions, however, there is a cue that may become available in spatial constancy control – i.e., gravity. When we change our head's orientation relative to gravity by means of tilting, our vestibular sensors can sense how the world moves relative to ourselves, by reconstructing the direction of gravity. Chapter 4 reports about the role of gravitational signals in the coding and updating of remembered visual space. We used a memory-saccade task to test whether the location of a target, briefly presented before a whole-body rotation in roll, is stored in egocentric (e.g. eye-, head- or body-centered) or in allocentric, gravity-based coordinates. To make this distinction, we exploited the fact that subjects, when tilted sideways in darkness, make systematic errors when indicating the direction of gravity (an allocentric task), even though they have a veridical percept of their self-orientation in space. We hypothesized that, if spatial memory is coded allocentrically, these distortions would affect the coding of remembered targets and their readout after a body rotation. Alternatively, if coding would be egocentric, updating for body rotation becomes essential and errors in performance should be related to the amount of intervening rotation. Subjects ($n=6$) were tested making saccades to remembered world-fixed targets after passive body tilts. Initial and final tilt angle ranged between -120° CCW to 120° CW. The results showed that subjects made large systematic directional errors in their saccades (up to 90°). These errors did not occur in the absence of intervening body rotation, ruling out a memory degradation effect. Regression analysis showed that the errors were closely

related to the amount of subjective allocentric distortion at both the initial and final tilt angle, rather than to the amount of intervening rotation. We conclude that the brain uses an allocentric reference frame, possibly gravity-based, to code visuospatial memories during whole-body tilts. This supports the notion that the brain can define information in multiple frames of reference, depending on sensory inputs and task demands.

Unveiling spatial representations using neuroimaging techniques

Finally, in a functional magnetic resonance imaging (fMRI) experiment, described in Chapter 5, we identified the frame of reference used by various cortical regions involved in coding spatial representations for saccades. Monkey neurophysiological evidence suggests that the cortical centers involved in saccade control employ a variety of reference frames to represent saccade targets. The goal of this study was to determine the reference frames for saccade shifts in the human brain. Using an fMRI repetition suppression paradigm, we distinguished between retinal (eye-centered) and non-retinal (e.g., head-centered) coding frames in three key regions: the intraparietal sulcus (IPS), frontal eye fields (FEF) and supplementary eye fields (SEF). Subjects ($n=18$) made delayed-saccades to one of five possible peripheral targets, separated at intervals of 9° visual angle. Target locations were chosen pseudo-randomly, based on a 2×2 factorial design with factors retinal and non-retinal coordinates and with levels novel and repeated. Thus, in subsequent trials, the location of the target was either repeated in retinal coordinates, repeated in non-retinal coordinates, novel in retinal coordinates or novel in non-retinal coordinates. In all three regions, analysis of the BOLD dynamics revealed an attenuation of the fMRI signal in trials repeating the location of the target in retinal coordinates, which was strongest during the preparatory phase of the saccade. We found no significant suppression effects during repeated trials in non-retinal coordinates. Further analyses showed an orderly, topographic representation of retinal target location in the IPS and FEF, but not in the SEF. All together, these results provide evidence that motor commands from these centers, irrespective of their topographic nature, encode saccade shifts in eye-centered coordinates.

Concluding remarks

The experiments described in this thesis provide behavioral evidence that the brain employs dynamic (updated) eye-centered spatial representations to keep track of the

3D location of targets for action, even when geometry puts challenging requirements on the updating system, such as caused by translational body movements or eye movements in depth. Interestingly, when body orientation is changed relative to the Earth vertical, gravity has appeared to be another important reference in spatial coding. Together, these observations paint a picture that the brain can code spatial information in multiple frames of reference, depending on the sensory systems that are involved and the task conditions that are employed.

Why would the eye-centered reference frame be so dominant in spatial memory operations in cue-poor conditions? We can consider several arguments for this. First, the visual system is the main sensor for spatial information. This appears clearly from the fact that we are much more impaired in spatial orientation when vision is lost, e.g. at night, than when other modalities like hearing provide no input. In the brain, there is a large amount of regions involved in visual processing, many more than for any other sensory system. This bottom-up retinotopic signature of most sensory spatial information may bias the representational organization in other cortical areas during brain development towards an eye-centered coding, because of computational and energetic efficiency. Second, in everyday eye-hand coordination, we generally make combine eye-hand movements towards a spatial target, making it favorable to use a common coding to guide all effector systems. In such combined eye-hand movements, the eyes usually land on the target before the hand does, which increases the spatial resolution of target location information, because of the relative neural overrepresentation of the central part of vision. This may provide a very sensitive feedback control mechanism to guide the movements of the hand and other effectors.

The studies presented in this thesis show that our brain is able to properly perform spatial updating computations under both active (self-initiated) and passive complex movement conditions. This ability is not trivial, since it requires an integration of multiple sources of information about self-motion that have different spatial and signal characteristics (see Figure 1.6), involving all kinds of non-linear computations. For example, the computation of the retinotopic consequences of body translation on basis of the amount of the linear acceleration and rotational velocity of the head, as detected by the vestibular system, requires one or several reference frame transformations. Aside from the experimental results that the brain takes movements into account to keep internal representations up-to-date, our neural network simulations described in Chapter 2 also explore how the neural systems may effectuate the required associations for this.

In everyday situations, more cues are available to indicate a target's position than those in the impoverished conditions used in the experiments described in this thesis. In daily life, for example, we can see the desk on which we have put our cup of coffee

and we are often aware and familiar with the spatial layout of the room we are in. These cues might be of assistance to infer current spatial locations of objects that have disappeared from sight. Although it lies beyond the scope of this thesis to give a judgment about the extent to which the brain makes use of such spatial cues, it is reasonable to assume that when multiple spatial cues are available these are also used in spatial coding. They might be combined with internal spatial constancy signals, probably in an optimal, Bayesian fashion. This is illustrated by the whole-body rotation experiment discussed in Chapter 4, when body posture was changed relative to gravity. In this case, gravity appears a relevant spatial cue as opposed to when we are in a normal upright position, and influences the internal coding of a spatial memory. However, in cue-poor environments such as in the majority of our experimental test conditions, all this additional spatial information is absent. In these conditions, the brain's basic representational system in visuomotor behavior can be revealed, which seems to be a retinal frame of reference. Also the fMRI results, described in Chapter 5, are in support of this notion.

In conclusion, the results presented in this thesis shed more light on how the human brain implements spatial constancy in motor control. Central in our observations is the dominance of retinal coordinates, although we show that spatial representations can exist in other coordinate systems as well. While retinal coordinates may be dominant in spatial constancy, further reference frame transformations are required to execute a movement to a remembered target. The neural basis of these processes may be found in parietal and/or frontal regions, but it should also be clear that more work is needed in assigning the computational functions to specific physiological substrates, and identifying the extraretinal signals involved.

Chapter 7

Nederlandstalige inleiding en samenvatting

Misschien wel het meest karakteristieke kenmerk van leven is dat organismen in staat zijn om zich te handhaven in hun milieu door adequaat te reageren op zintuiglijke informatie. In het dierenrijk is de ontwikkeling van een neuronaal systeem dat sensorische informatie verwerkt, leert en bewegingen aanstuurt een groot evolutionair voordeel gebleken in de overlevingsstrijd. Dit heeft organismen in staat gesteld om te vluchten voor gevaar, actief naar voedsel te zoeken en nieuwe niches en habitats in te nemen op een snellere manier dan ooit tevoren in de evolutie.

Hoe complexer diersoorten werden, hoe uitgebreider en gespecialiseerder hun zenuwstelsel werd (zie Randell e.a. 1997). Zo hebben sommige simpele ongewervelden zoals echinodermata geen centraal zenuwstelsel, maar slechts een ring van onderling verbonden neuronen om sensorische informatie door te geven. Bij gewervelde dieren – waaronder zoogdieren – heeft zich daarentegen een gespecialiseerd neuronaal netwerk ontwikkeld dat bestaat uit een centraal en perifeer zenuwstelsel. Hierin hebben alle onderdelen specifieke functies om het lichaam aan te sturen. Zo zijn het ruggenmerg en de hersenstam betrokken bij het controleren van geautomatiseerde interne vegetatieve processen zoals het laten kloppen van het hart, de ademhaling en reflexen. Het prosencefalon van de mens (waaronder de grote hersenen) is daarentegen gespecialiseerd in zogenoemde hogere functies zoals waarneming, het uitvoeren van handelingen, leren, spraak en cognitie (Kandel e.a. 2000). De specialisatie van het aansturen van bewegingen is een belangrijke eigenschap waarmee primaten zich van andere dieren onderscheiden. Dit heeft geleid tot zeer gespecialiseerde en onderscheidende vaardigheden zoals reiken, grijpen, gebruik maken van hulpmiddelen zoals gereedschap en focussen met de ogen.

Sensorisch-motorische omzetting

Ondanks het schijnbare automatisme en gemak waarmee we doelgerichte handelingen uitvoeren zoals reiken en grijpen, moeten de hersenen complexe rekenkundige problemen oplossen om deze bewegingen goed uit te kunnen voeren. Stelt u zich bijvoorbeeld eens voor dat u een espressokopje wilt pakken terwijl u de krant leest zoals geïllustreerd staat in Figuur 7.1. Om de juiste grijpbeweging uit te kunnen voeren moet de visuele informatie die uw ogen geven over de positie van het kopje omgezet worden in specifieke spiersamentrekkingen die nodig zijn om uw hand naar de espresso te laten bewegen. Dit proces wordt ook wel een *sensorimotor transformatie* genoemd, ofwel sensorisch-motorische omzetting. Hoe dit uitgevoerd wordt is een belangrijke vraag binnen de neurowetenschappen. Om te begrijpen hoe de hersenen deze



Figuur 7.1. Sensorimotor transformatie in een alledaagse situatie: het pakken van een kopje espresso tijdens het lezen van de krant. De zintuiglijke informatie over de positie van het kopje (in het rechter deel van het blikveld) moet omgezet worden naar het juiste commando om de arm mee aan te sturen ('beweeg naar links'). Hiervoor is het nodig dat informatie van het ene type coördinaten omgezet wordt naar het andere – een referentiekader transformatie.

omzettingen uitvoeren maken neurowetenschappers vaak gebruik van concepten als coördinaatsystemen en referentiekaders, zoals oog-, hoofd-, lichaams-, of wereldgecentreerde coördinaatsystemen (Soechting en Flanders 1992). Wiskundig gezien is een referentiekader niets anders dan een aantal assen die loodrecht op elkaar staan en in een punt kruisen, de oorsprong. Met behulp van deze assen kan de ruimtelijke locatie van objecten beschreven worden door middel van een aantal coördinaten. Dit concept wordt op vele terreinen toegepast, ook buiten de neurowetenschappen. Zo zullen Nederlandse GPS-ontvangers (zoals in autonavigatiesystemen) u vertellen waar u zich bevindt volgens de Nederlandse Rijksdriehoeksmeting (RD-grid), in Cartesische coördinaten (meters) ten opzichte van de Onze Lieve Vrouwentoren in Amersfoort. Internationaal is het WGS84-grid echter de standaard, en deze vertelt u waar u bent

ten opzichte van de meridiaan van Greenwich, in graden oosterlengte en noorderbreedte. Dus één locatie kan door middel van verschillende coördinaten beschreven worden, afhankelijk van het referentiekader dat gebruikt wordt om deze positie mee te definiëren (in dit geval, RD of WGS84).


Het concept van referentiekaders kan ook gebruikt worden om ruimtelijke representaties (weergaves) in de hersenen mee te beschrijven. Alle visuele informatie die op onze netvliesen terechtkomt kan worden uitgedrukt in polaire coördinaten (graden) ten opzichte van de fovea, het midden van het netvlies – oftewel in een ooggecentreerd referentiekader. Wanneer we een beweging maken moeten we onze armen en handen echter verplaatsen ten opzichte van ons lichaam, dus moeten de bijbehorende bewegingscommando's uitgedrukt worden in een lichaamsgecentreerd referentiekader. In ons voorbeeld bevindt de espresso zich in het rechter deel van het blikveld (dus de sensorische coördinaten zijn 'aan de rechterkant'), maar staat links ten opzichte van de rechterhand (dus het bewegingscommando moet zijn: 'beweeg naar links'). Op deze manier kunnen we sensorimotor transformaties begrijpen, door te beschrijven hoe de hersenen de positie van het kopje omzetten van de coördinaten van het oog naar de lichaamscoördinaten die bij het aansturen van de hand horen.

Ruimtelijke constantie voor het aansturen van bewegingen

Om te kunnen overleven is het essentieel dat je kunt vertrouwen op een juiste weergave van locaties van objecten in de wereld om ons heen. Een roofdier heeft dit nodig om zijn prooi te kunnen vangen, en wij mensen bijvoorbeeld om gereedschap te kunnen hanteren. Het proces dat de positie van objecten om ons heen volgt wordt ook wel ruimtelijke constantie genoemd (Von Helmholtz 1867; Von Holst en Mittelstaedt 1950) en is belangrijk bij het maken van doelgerichte bewegingen. Ook al lijkt het misschien makkelijk om voor ruimtelijke constantie te zorgen zolang objecten zichtbaar zijn, het is minder triviaal (maar nog steeds essentieel) dat dit ook gebeurt als deze visuele informatie verdwenen is. We voeren namelijk niet altijd meteen een beweging uit naar objecten die we zien; soms moet een beweging naar een zojuist waargenomen positie dus nog even worden uitgesteld. Bijvoorbeeld als we in de situatie van Figuur 7.1 de espresso wel gezien hebben maar er nog niet naar reiken omdat deze nog te heet is en onze ogen eerst op de krant richten. In dat geval moet de positie van het kopje opgeslagen worden in ons geheugen, in het zogenoemde ruimtelijk geheugen (Baddeley 1996; Goldman-Rakic 1996; Pierrot-Deseilligny et al. 2002). Hiervan kan later gebruik van gemaakt worden om een handbeweging naar het kopje mee aan te sturen. Als de hersenen deze informatie nu in ooggecentreerde coördinaten zouden opslaan –

als ‘rechts’ – dan moeten ze er rekening mee houden dat deze waarden al niet meer kloppen op het moment dat we ergens anders heen kijken. Om toch voor ruimtelijke constantie te zorgen is het dus nodig dat doellocaties op een manier opgeslagen worden dat ze ofwel niet beïnvloed worden door oogbewegingen, ofwel intern aangepast (herberekend) worden door te compenseren voor oogbewegingen (Crawford e.a. 2004).

De laatste tientallen jaren heeft een groot aantal studies laten zien dat we behoorlijk goed in staat zijn om ruimtelijke constantie te bewaren. Hallett en Lightstone waren in 1976 de eersten die dit systematisch onderzochten, door middel van de nu klassieke zogenoemde dubbel-stap saccadetaak (zie Figuur 7.2). Hierbij moeten proefpersonen twee snelle oogbewegingen (saccades) na elkaar maken (een ‘dubbele [oog]stap’). In dit specifieke experiment zagen proefpersonen heel kort na elkaar twee lichtflitsen in de periferie van hun blikveld (de doelen, aangegeven met D1 en D2 in A), waarna ze gevraagd werden om achtereenvolgens naar beide doellocaties te kijken. De onderzoekers redeneerden dat op het moment dat de proefpersonen de eerste oogbeweging (naar D1) hadden gemaakt, dat de oorspronkelijke richting van het tweede doel ($\vec{D2}_i$) hierdoor niet meer klopte met de richting waarin de tweede



Figuur 7.2. De dubbel-stap saccadetaak. **A.** Terwijl de proefpersoon naar een visueel fixatiepunt (F) kijkt, worden na elkaar twee doelen (D1 en D2) kort geflitst in de visuele periferie van het blikveld. De opdracht aan de proefpersoon is om een dubbele (‘dubbel-stap’) oogbeweging (saccade) te maken: eerst naar D1, en vanuit daar naar D2. **B.** Na de oogbeweging naar D1 moet de proefpersoon rekening houden met de grootte en richting van deze eerste saccade om de tweede oogbeweging goed uit te kunnen voeren. Uitgedrukt in vectoren betekent dit dat de retinale positie van D1 ($\vec{D1}_i$) afgetrokken moet worden van die van D2 ($\vec{D2}_i$) om de juiste nieuwe bewegingsrichting en -afstand te berekenen: $\vec{D2}_f = \vec{D2}_i - \vec{D1}_i$. De proefpersoon zal een foutieve oogbeweging maken, naar D2', als de retinale positie van D2 niet wordt herberekend.

oogbeweging gemaakt zou moeten worden (zie Figuur 7.2B). Dus na de eerste oogbeweging moeten de proefpersonen de richting van de tweede oogbeweging herberekenen om het juiste doel te bereiken, door middel van de berekening $\overrightarrow{D2_i} - \overrightarrow{D1_i}$. In het experiment van Hallett en Lightstone bleken proefpersonen inderdaad in staat om de dubbel-stap taak correct uit te voeren, wat als bewijs werd gezien dat het inderdaad mogelijk is om ruimtelijke constantie te behouden tijdens oogbewegingen. Deze bevindingen zijn vervolgens bevestigd in een aantal neurofysiologische experimenten met apen door Mays en Sparks (1980, 1983). Sindsdien hebben vele onderzoekers deze resultaten gerepliceerd. Meer recentelijk is in een groot aantal andere experimenten voortgeborduurd op deze resultaten waarin ruimtelijke constantie onderzocht is aan de hand van andere taken. Hiermee heeft men nu laten zien dat proefpersonen ook in staat zijn om herinnerde doelen te onthouden tijdens het maken van langzame volgbewegingen met de ogen (McKenzie en Lisberger 1986; Gellman en Fletcher 1992; Blohm e.a. 2003, 2006; Baker e.a. 2003; Raffi e.a. 2007) en oogbewegingen in diepte (Krommenhoek en Van Gisbergen 1994). Ook bewegingen van het hoofd (Guitton 1992; Medendorp e.a. 2002, 2003b; Vliegen e.a. 2005) en het gehele lichaam (Baker e.a. 2003; Li e.a. 2005; Li en Angelaki 2005; Klier e.a. 2008) bleken geen groot verstorend effect te hebben op ruimtelijke constantie.

Referentiekaders voor ruimtelijke constantie bij bewegingsaansturing

Computationeel gezien zijn er veel manieren om ruimtelijke constantie te behouden om bewegingen aan te kunnen sturen, afhankelijk van welk onderliggend referentiekader wordt gebruikt. Één mogelijkheid is dat ruimtelijke constantie wordt bewaard door middel van een lichaamsgecentreerd referentiekader. Ruimtelijke representaties die lichaamsgecentreerde coördinaten gebruiken zijn namelijk redelijk stabiel omdat ze niet beïnvloed worden door oog- of hoofdbewegingen, maar alleen door een verplaatsing van het lichaam (Flanders e.a. 1992; zie voor een overzicht Battaglia-Mayer e.a. 2003). Van de andere kant zouden de hersenen ook een ooggecentreerd referentiekader kunnen gebruiken om doellocaties in weer te geven. Het gebruik van zo'n ooggecentreerde representatie (ook wel aangeduid met de termen retinaal [naar het netvlies – retina], retinotopisch, kijkrichtinggecentreerd of kijkrichtingafhankelijk) heeft als consequentie dat deze elke keer wanneer er een oogbeweging gemaakt wordt aangepast moet worden. Als de persoon in Figuur 7.1 naar haar horloge kijkt, verplaatst de espresso zich van de rechter- naar de linkerkant van het blikveld. In de hersenen zou een ooggecentreerde representatie van de positie van het espressokopje dan ook aangepast moeten worden van 'rechts' naar 'links'.

Een andere mogelijkheid nog is dat het referentiepunt buiten het lichaam ligt. Doellocaties kunnen bijvoorbeeld opgeslagen worden ten opzichte van andere objecten (bijvoorbeeld ten opzichte van de tafel in Figuur 7.1), of ten opzichte van de zwaartekracht. Deze zogenoemde allocentrische of wereldvaste representaties hoeven niet aangepast te worden na oog-, hoofd- of lichaamsbewegingen, dus is er altijd sprake van ruimtelijke constantie. Natuurlijk moet deze allocentrische informatie wel omgezet worden naar lichaamsgecentreerde signalen, waarmee de spieren aangestuurd moeten worden.

Welk referentiekader het gunstigst is voor dagelijkse handelingen als grijpen en kijken naar dingen is niet meteen duidelijk en hoeft ook niet hetzelfde te zijn voor deze verschillende soorten bewegingen. Bij een gebrek aan ruimtelijke informatie over de wereld om ons heen lijkt een lichaamsgecentreerde weergave bijvoorbeeld optimaal voor reiken en wijzen, vanwege zijn stabiliteit. Maar er zijn ook argumenten voor het gebruik van een ooggecentreerde representatie om alle typen beweging mee aan te sturen. Allereerst is ons visuele systeem het belangrijkste zintuig waarmee ruimtelijke informatie verkregen kan worden, en veel hersengebieden zijn dan ook betrokken bij het verwerken van visuele informatie. Hierdoor zou het computationeel en energietechnisch gunstig kunnen zijn om zoveel mogelijk een retinaal referentiekader te gebruiken. Een tweede reden heeft te maken met het verschil in spatiële resolutie tussen deze coördinaatsystemen. De resolutie van de retinale informatie zou kunnen verslechteren als deze omgezet wordt naar lichaamsgecentreerde coördinaten. Een ander voordeel van een ooggecentreerd referentiekader zou kunnen zijn dat dit het makkelijker maakt om de aansturing van verschillende onderdelen van het bewegingsapparaat (bijvoorbeeld de ogen en de hand) op elkaar af te stemmen (Andersen e.a. 1997; Cohen en Andersen 2002). Ten slotte kan het gunstig zijn om één algemene geheugenkaart te hebben, in ooggecentreerde coördinaten, waarin alle potentiële doelen opgeslagen zijn. Hierbij hoeft alleen het doel waar een beweging naartoe wordt gepland omgezet te worden naar lichaamscoördinaten ('omzetting op verzoek', Henriques e.a. 1998).

Samenvatting van dit proefschrift

De laatste jaren is er veel bewijs geleverd dat verschillende hersengebieden doellocaties voor bewegingen op lijken te slaan in een ooggecentreerd referentiekader, oftewel als posities ten opzichte van onze kijkrichting. Dit heeft zoals gezegd als consequentie dat om voor ruimtelijke constantie te zorgen de ooggecentreerde ('retinale') coördinaten van deze objecten aangepast moeten worden bij elke oogbeweging. Verschillende

onderzoeken hebben laten zien dat deze retinale herberekening inderdaad plaatsvindt. Deze experimenten staan voor een deel hierboven beschreven en in Hoofdstuk 1, en zullen hier niet verder en detail besproken worden. In de meeste van deze experimenten bestonden de lichaamsbewegingen echter slechts uit veranderingen van kijkrichting van de ogen. Het was tot nu toe nog onbekend of deze resultaten wel gegeneraliseerd mochten worden naar ingewikkeldere situaties. Voorbeelden hiervan zijn bewegingen van het hele lichaam, zoals tijdens lopen of autorijden, of oogbewegingen in diepte (van veraf naar dichtbij of andersom). In dit laatste geval bewegen onze ogen niet allebei dezelfde kant op (zoals wel het geval is bij verandering van *kijkrichting*) maar juist in tegengestelde richting. De experimenten die beschreven staan in dit proefschrift hebben onderzocht hoe onze hersenen voor ruimtelijke constantie zorgen tijdens complexere bewegingen dan alleen een verandering van kijkrichting. Ze hebben nieuwe kennis opgeleverd over welke berekeningen hieraan ten grondslag liggen binnen ons ruimtelijk geheugen. Daarnaast hebben we ook nieuwe methodes getest waarmee verschillende typen ruimtelijke representaties geïdentificeerd kunnen worden op basis van specifieke hersenactiviteit die bij mensen gemeten kan worden.


De belangrijkste bijdragen die dit proefschrift aan het onderzoeksveld levert zijn de volgende:

- 1) Tijdens zelf-geïnitieerde bewegingen van het lichaam (zijwaartse stapbeweging) wordt een retinale (ooggecentreerde) representatie gebruikt voor het aansturen van reikbewegingen; deze representatie wordt correct aangepast op basis van de onderliggende niet-lineaire geometrie.
- 2) De afstand (diepte) van doelen voor reikbewegingen wordt opgeslagen in retinale dispariteitscoördinaten (als afstand ten opzichte van de fixatiediepte), en herberekend bij oogbewegingen in diepte, vergelijkbaar met wat het geval is voor doelrichting.
- 3) Zwaartekracht kan als referentie gebruikt worden in ruimtelijke representatie. Als onze lichaamsoriëntatie verandert ten opzichte van wat rechtop is in de wereld, slaan we de richting van doelen in het fronto-parallelle (verticale) vlak op ten opzichte van de richting van de zwaartekracht.
- 4) Repetitie suppressie in fMRI is een bruikbare techniek om ruimtelijke representaties mee te identificeren in de menselijke hersenen. De belangrijkste gebieden in de grote hersenen die bij oogbewegingen betrokken zijn, laten repetitie suppressie zien in ooggecentreerde coördinaten, maar niet in hoofd- of lichaamsgecentreerde coördinaten.

Hieronder zullen we elk van deze resultaten gedetailleerder bespreken.

Ruimtelijke constantie tijdens zijwaartse lichaamsbewegingen

In veel situaties verplaatsen onze ogen zich door de ruimte, zoals tijdens bewegingen van ons hoofd of hele lichaam. Onder zulke omstandigheden moeten onze hersenen met complexe geometrie rekening houden bij het aanpassen van ooggecentreerde representaties. Tijdens zijwaartse bewegingen (translaties) bewegen objecten die op een verschillende afstand liggen ten opzichte van onze kijkafstand (fixatiediepte) zich namelijk met verschillende snelheden over ons netvlies. Dit wordt bewegingsparallax genoemd. Ook een ooggecentreerd ruimtelijk geheugen moet met deze parallaxgeometrie rekening houden als herinnerde doellocaties aangepast worden voor een zijwaartse beweging. Dit staat geïllustreerd in Figuur 7.3. De figuren aan de linkerkant laten zien hoe twee herinnerde visuele doellocaties D1 en D2 (A) herberekend worden na een oogrotatie naar links (B). In dit geval moeten beide locaties eenvoudigweg met dezelfde hoeveelheid rotatie aangepast worden om ruimtelijke constantie te behouden (C). Een oogtranslatie daarentegen heeft als gevolg dat de twee doellocaties met verschillende hoeveelheden aangepast moeten worden (de rechterfiguren). Dit houdt in dat onze hersenen de geometrie die bij bewegingsparallax hoort moeten simuleren. Als onze hersenen echter niet-retinale (bijvoorbeeld



Figuur 7.3. Geometrische consequenties van oogrotaties (linkerfiguren) en –translaties (rechterfiguren) voor herberekening van ruimtelijke representaties. De coördinaten van twee doelen D1 en D2 met dezelfde retinale positie (A) moeten tijdens een oogrotatie evenveel aangepast worden (B, linkerfiguur), namelijk met de inverse van de hoeveelheid oogrotatie (C, linkerfiguur). Bij een oogtranslatie daarentegen (B, rechterfiguur) moeten de doelen met verschillende hoeveelheden aangepast worden, afhankelijk van de doelafstand (C, rechterfiguur). Aangepaste figuur uit Medendorp e.a. (2008).

hoofdgecentreerde) coördinaten zouden gebruiken tijdens translaties, is het computationeel veel eenvoudiger om ruimtelijke constantie te waarborgen, omdat parallaxsimulatie niet nodig is.

Hoofdstuk 2 beschrijft een studie waarin onderzocht werd in welk type referentiekader doellocaties gerepresenteerd worden tijdens dit soort zijwaartse bewegingen van het lichaam. Hierbij werden in het donker aan 12 proefpersonen heel kort visuele doelen getoond die ofwel voor ofwel achter het fixatiepunt van de ogen lagen. Vervolgens maakten de vrijwilligers een stap zijwaarts, waarna ze de locatie waarop ze het doel hadden gezien zo nauwkeurig mogelijk moesten aanwijzen. Het bleek dat er systematische fouten zaten in deze reikbewegingen die gerelateerd waren aan de bewegingsparallax: de reikfouten namen toe met de afstand tot het fixatiepunt. Daarnaast waren ze tegengesteld in richting voor de doelen aan weerszijden van het fixatiepunt. In een aantal controle-experimenten hebben we vervolgens de mogelijke invloed van een aantal externe factoren op deze resultaten kunnen uitsluiten. Deze factoren waren de aanwezigheid van een fixatielicht tijdens de translatie, het maken van een gecombineerde oog-hand beweging en de aanwezigheid van visuele terugkoppeling van de handpositie. Kwantitatieve geometrische analyses bevestigden uiteindelijk dat de fouten beter verklaard konden worden door middel van retinale dan door niet-retinale coördinaten. Hieruit concludeerden we dat een retinaal (ooggecentreerd) referentiekader ook gebruikt wordt om voor ruimtelijke constantie te zorgen tijdens zijwaartse stapbewegingen. Daarnaast lieten simulaties met neurale netwerken zien dat onze hersenen voor deze berekeningen vermoedelijk gebruik maken van snelheidssignalen en informatie van stereoscopische kijkafstand en -richting.


Ruimtelijke representaties voor diepte constantie

Voor bewegingen zoals wijzen en het maken van saccades is het vaak voldoende om bij te houden in welke richting objecten zich bevinden. In andere situaties hebben we echter ook afstands-informatie nodig, bijvoorbeeld tijdens reiken of grijpen. Het is algemeen geaccepteerd dat de hersenen visuele richting- en afstands-informatie apart van elkaar verwerken (Cumming en DeAngelis 2001; DeAngelis 2000; Flanders e.a. 1992; Vindras e.a. 2005), dus het mag aangenomen worden dat er ook gescheiden processen ten grondslag liggen aan het bewaren van ruimtelijke constantie voor beide dimensies. Tot nu toe was nog niet veel bekend over hoe voor ruimtelijke constantie van doelaafstand gezorgd wordt. Vooral wanneer we weinig zintuiglijke informatie krijgen en als we geen zijwaartse bewegingen maken is *retinale dispariteit* de belangrijkste bron van afstands-informatie (Howard en Rogers 1995; Julesz 1971; Wei

vergelijking hebben we ook onderzocht hoe nauwkeurig de richtingscomponent van hun reikbeweging was. We zagen wanneer er een tussenliggende oogbeweging was gemaakt, dat de reikbewegingen een foutenpatroon hadden dat afhing van de nieuwe oogpositie en van de afstand van het herinnerde doel ten opzichte van deze positie. Dit wijst erop dat doelafstand herberekend wordt na oogbewegingen, conform het retinale model. Daarnaast volgde uit regressieanalyses dat de aanpassing van de relatieve doelpositie bijna helemaal correct was, voor zowel de richting als de diepte van het doel. Dit suggereert dat de waargenomen reikfouten pas na dit herberekeningsproces optreden, tijdens de daaropvolgende referentiekader transformaties die nodig zijn om de reikbeweging uit te kunnen voeren.

De rol van zwaartekracht op ruimtelijke constantie

Één reden waarom de hersenen gebruik zouden kunnen maken van retinale (of andere lichaamsgecentreerde) representaties om doellocaties voor toekomstige handelingen in op te slaan, is dat als we geen visuele informatie hebben, zoals 's nachts, ruimtelijke informatie over de wereld buiten onszelf (zogenoemde allocentrische informatie) over het algemeen afwezig is. In sommige gevallen is er echter ook dan nog allocentrische informatie beschikbaar die gebruikt kan worden voor ruimtelijke constantie, te weten de zwaartekracht. Onze evenwichtsorganen voelen namelijk hoe de wereld beweegt ten opzichte van ons als we ons hoofd kantelen, doordat ze de richting van de zwaartekracht reconstrueren. In Hoofdstuk 4 hebben we onderzocht of deze zwaartekrachtssignalen van invloed zijn op ons ruimtelijk geheugen. Proefpersonen zagen kort een visueel doel waarna hun lichaam gekanteld werd over een grote hoek (begin- en eindoriëntaties lagen tussen -120° en $+120^\circ$). Vervolgens moesten ze een saccade maken naar de herinnerde doellocatie. Als het doel in oog- of lichaamsgecentreerde coördinaten opgeslagen zou zijn, dan zouden deze herberekend moeten worden na de lichaamsrotatie. In het geval van een zwaartekrachts-afhankelijke oftewel allocentrische (wereldvaste) doelrepresentatie zou deze aanpassing niet nodig zijn. Om onderscheid te kunnen maken tussen deze twee mogelijke representaties, maakten we gebruik van het feit dat proefpersonen grote systematische fouten maken als ze de richting van de zwaartekracht moeten aangeven wanneer ze gekanteld zijn in het donker. Deze fout wordt A-fout genoemd, naar zijn ontdekker (Aubert 1861), en is aanzienlijk bij grote kantelhoeken (zie Figuur 7.5; Van Beuzekom en Van Gisbergen 2000; Kaptein en van Gisbergen 2004). We verwachtten dat deze verstoringen het opslaan en uitlezen van herinnerde doelen zouden beïnvloeden in het geval van een allocentrisch ruimtelijk geheugen, maar niet bij een egocentrische codering.




Figuur 7.5. Systematische fouten die proefpersonen maken wanneer ze moeten bepalen wat verticaal is in de wereld, als functie van de kantelhoek. Bij grote kantelhoeken worden grote fouten gemaakt, wat het A-effect genoemd wordt. Aangepaste figuur uit Van Beuzekom en Van Gisbergen (2000).

Onze resultaten lieten zien dat proefpersonen na een tussenliggende rotatie grote systematische fouten maakten in de richtingscomponent van hun saccades (tot wel 90°). Deze fouten waren afwezig wanneer er geen rotatie was geweest, wat aantoont dat deze niet simpelweg te wijten waren aan een verslechtering van het geheugen. Regressie-analyses lieten vervolgens zien dat de fouten overeenkwamen met een combinatie van de hoeveelheid subjectieve allocentrische verstoring op de begin- en eindoriëntatie van het lichaam. Hieruit hebben we geconcludeerd dat de hersenen tijdens lichaamskantelingen een allocentrisch ruimtelijk geheugen gebruiken dat waarschijnlijk gebaseerd is op de richting van de zwaartekracht. Dit ondersteunt de gedachte dat de hersenen ruimtelijke informatie in verschillende referentiekaders kunnen opslaan, afhankelijk van over welke zintuiglijke informatie ze beschikken en wat de taakspecifieke eisen zijn.

Identificatie van ruimtelijke representaties met behulp van hersen-beeldvormingstechnieken

Tenslotte hebben we in een experiment dat beschreven staat in Hoofdstuk 5 met functioneel magnetisch resonantie beeldvorming (fMRI – functional Magnetic Resonance Imaging) onderzocht welke referentiekaders gebruikt worden door drie hersengebieden die sterk betrokken zijn bij het maken van saccades: de intrapariëtale sulcus (IPS), frontale oogvelden (FEF) en supplementaire oogvelden (SEF). Eerder neurofysiologisch onderzoek heeft namelijk uitgewezen dat vergelijkbare gebieden bij apen verschillende referentiekaders gebruiken om doelen voor saccades te representeren. Veel fMRI-onderzoek naar referentiekaders maakt gebruik van het feit dat hersengebieden vaak een gelateraliseerde of topografische interne organisatie hebben. Dit houdt in dat neuronen die naast elkaar liggen, bijvoorbeeld in de primaire visuele schors (V1), delen van de buitenwereld representeren die ook naast elkaar liggen, zoals geïllustreerd staat in de linkerhelft van Figuur 7.6A. In een typisch ooggecentreerd, topografisch georganiseerd visueel hersengebied zullen stimuli in ons



Figuur 7.6. Twee mogelijke manieren van corticale organisatie van neuronen die specifieke ruimtelijke locaties (hun responsvelden) representeren. Bij een ordentelijke, topografische organisatie (zoals in de primaire visuele hersenschors) zijn neuronen netjes gerangschikt, op basis van hun responsvelden (**A**, linkerfiguur). Zo zal een stimulus op +30° de neuronen in de grijze rechthoek van de linker hersenhelft activeren, en een stimulus op -30° de cellen in de ellips in de andere hersenhelft. Met fMRI kan dit verschil zichtbaar worden gemaakt (**B**; links); voor doelen rechts zal in de linkerhersenelft hogere activiteit gemeten worden dan aan de rechterkant, en vice versa. Maar als neuronen willekeurig verdeeld zijn over beide hersenhelften (rechterfiguren) zullen er geen activiteitsverschillen meer zijn tussen beide hersenhelften. Een stimulus op bijvoorbeeld +30° zal dan namelijk ongeveer evenveel neuronen activeren in de linker- en rechterhersenelft, waardoor er geen verschil in signaal gemeten zal worden met fMRI (**B**; rechts). In dat geval kan het referentiekader van dit gebied dan ook niet hiermee geïdentificeerd worden. Een alternatieve methode hiervoor is om te kijken naar repetitie suppressie (**C**). Dit houdt in dat neuronale activiteit zal afnemen als een stimuluslocatie wordt herhaald in het referentiekader dat deze neuronen gebruiken, onafhankelijk van hoe de cellen verdeeld zijn over de hersenhelften.

linker blikveld (het cirkeltje) dan ook voornamelijk een reactie opleveren in de rechter hersenhelft, en andersom. Met fMRI is dit zichtbaar als een groter activiteitssignaal aan de rechterkant van de hersenen (Figuur 7.6B, linkerkant). Dit wordt een gelateraliseerde activiteit genoemd. Maar hersengebieden hebben lang niet altijd een topografische interne organisatie. In dat geval is het niet mogelijk om op bovenstaande manier het referentiekader ervan te identificeren. Bij een willekeurige ruimtelijke verdeling van de neuronen (Figuur 7.6A, rechts), zal er geen gelateraliseerde activiteit meer zijn (Figuur 7.6B, rechts) en kan het onderliggende referentiekader niet meer bepaald worden.

In Hoofdstuk 5 hebben we een alternatieve methode gebruikt waarbij geen topografische hersenorganisatie vereist is om een onderscheid te kunnen maken tussen retinale (ooggecentreerde) en niet-retinale (bijvoorbeeld hoofdcintrische) ruimtelijke representaties: repetitie suppressie (RS). Dit is het fenomeen dat een hersenrespons in sterkte afneemt bij herhaling van dezelfde stimulus. RS was nog niet eerder gebruikt om referentiekaders te identificeren of om bewegingsplanning te onderzoeken. Onze hypothese was dat de activiteit van neuronen die locaties van doelen representeren in één bepaald referentiekader afneemt als een doellocatie wordt herhaald in dat referentiekader, maar niet als het wordt herhaald in een ander coördinaatsstelsel. Deze afname zou onafhankelijk moeten zijn van het wel of niet aanwezig zijn van een topografische organisatie (Figuur 7.6C). In ons experiment maakten 18 proefpersonen saccades naar één van vijf mogelijke herinnerde doelen in de visuele periferie met een onderlinge afstand van 9° visuele hoek. Deze doellocaties waren in opeenvolgende tests ofwel nieuw in retinale dan wel niet-retinale coördinaten, of werden herhaald in retinale dan wel niet-retinale coördinaten. In alle drie de hersengebieden nam de hersenactiviteit af wanneer het doel herhaald werd in retinale coördinaten. Dit was het sterkst tijdens de voorbereidingsfase die voorafgaat aan het maken van de saccade. We vonden geen significante suppressie-effecten wanneer het doel herhaald werd in niet-retinale coördinaten. Daarnaast liet een andere analyse zien dat er wel een regelmatige, topografische representatie van retinale doellocatie was in de IPS en FEF, maar niet in de SEF. Alles bij elkaar wijzen deze resultaten er op dat motorische commando's vanuit deze hersencentra oogbewegingen in ooggecentreerde coördinaten representeren onafhankelijk van de topografische eigenschappen van deze gebieden.

Conclusies

De experimenten die in dit proefschrift staan beschreven bieden bewijs voor de hypothese dat onze hersenen vaak dynamische (up-to-date gehouden) ooggecentreerde

ruimtelijke representaties gebruiken om de driedimensionale locaties van doelen te volgen. Dit is zelfs het geval als de benodigde herberekeningen geometrisch gezien erg complex zijn, zoals bijvoorbeeld bij zijwaartse lichaamsverplaatsingen (Hoofdstuk 2) of oogbewegingen in diepte (Hoofdstuk 3). Ook de fMRI-resultaten beschreven in Hoofdstuk 5 zijn in overeenstemming met een ooggecentreerd ruimtelijk geheugen. Daarnaast zagen we echter dat ook zwaartekracht een belangrijke referentie kan zijn voor ons ruimtelijk geheugen, namelijk wanneer ons lichaam gekanteld wordt (Hoofdstuk 4). Alles bij elkaar laten deze resultaten zien dat onze hersenen ruimtelijke informatie in verschillende referentiekaders kunnen opslaan, afhankelijk van de betrokken zintuigen en taakcondities.

Waarom zou bij minimale zintuiglijke informatie een ooggecentreerd referentiekader het ruimtelijk geheugen domineren? We kunnen hiervoor verschillende argumenten aanvoeren. Allereerst zijn de ogen onze belangrijkste sensorische organen om ruimtelijke informatie te krijgen. Zo hebben we veel meer problemen met ruimtelijke oriëntatie wanneer we niets kunnen zien – bijvoorbeeld 's nachts – dan wanneer andere zintuigen, zoals ons gehoor, ons geen signalen geven. In onze hersenen is dan ook een groot aantal gebieden betrokken bij het verwerken van visuele signalen, veel meer dan van andere zintuiglijke informatie. Dat het merendeel aan sensorische informatie een retinotopisch karakter heeft zou er toe kunnen leiden dat ook andere, niet-visuele hersengebieden ooggecentreerde referentiekaders hebben ontwikkeld, vanwege computationele en energetische efficiëntie. Ten tweede maken we in de dagelijkse praktijk meestal gecombineerde oog- en handbewegingen naar objecten. Daarom zou het gunstig zijn om hetzelfde type representatie te gebruiken om alle onderdelen van het bewegingsapparaat mee aan te sturen. Tijdens oog-hand bewegingen zijn de ogen gewoonlijk eerder bij het doel dan de hand. Hierdoor wordt de spatiële resolutie van de informatie over doellocatie vergroot, omdat het centrale deel van ons visuele veld relatief sterk gerepresenteerd is in onze hersenen. Dit zou voor een nauwkeurige terugkoppeling kunnen zorgen bij het aansturen van de hand of andere ledematen.

De experimenten die in dit proefschrift beschreven staan laten zien dat onze hersenen in staat zijn om ruimtelijke representaties aan te passen tijdens zowel actieve (zelf-geïnitieerde) als passieve complexe bewegingen. Met neurale netwerksimulaties (Hoofdstuk 2) hebben we vervolgens laten zien dat snelheidssignalen hierbij cruciaal zijn. Deze vaardigheid is niet vanzelfsprekend omdat voor de onderliggende herberekeningen een niet-lineaire integratie van meerdere bronnen van bewegingsinformatie vereist is die ieder verschillende ruimtelijke en signaalkarakteristieken hebben (zie Figuur 1.6). Voor de berekening bijvoorbeeld van de ooggecentreerde consequenties van een zijwaartse lichaamsverplaatsing zijn één of

meerdere referentiekadertransformaties nodig. De hoeveelheid lineaire versnelling en draaisnelheid van het hoofd wordt namelijk in hoofdgecentreerde coördinaten door het evenwichtsorgaan waargenomen. Om het effect hiervan op een retinale representatie te berekenen, zullen deze eerst omgezet moeten worden naar ooggecentreerde coördinaten.

In ons dagelijks leven hebben we meer informatie over de positie van objecten dan in de minimalistische omstandigheden die gebruikt zijn in de experimenten in dit proefschrift. Normaal gesproken kunnen we bijvoorbeeld de tafel zien waarop we ons kopje koffie neergezet hebben en zijn we ook bekend met de ruimtelijke indeling van de kamer waarin we ons bevinden. Deze informatie kan nuttig zijn om de locaties te kunnen herleiden van objecten die zich niet meer in ons blikveld bevinden. Alhoewel we in dit proefschrift niet onderzocht hebben in hoeverre de hersenen gebruik maken van dit soort ruimtelijke informatie, is het aannemelijk dat wanneer meerdere bronnen van ruimtelijke informatie aanwezig zijn, deze ook gebruikt worden voor het intern representeren van doellocaties. Deze externe informatie zou op een optimaal gewogen manier gecombineerd kunnen worden met interne signalen omtrent ruimtelijke constantie. Dit wordt bijvoorbeeld geïllustreerd in Hoofdstuk 4, waarin we lichaamsoriëntatie veranderden ten opzichte van de zwaartekracht. Hierdoor werd deze een relevante bron van ruimtelijke informatie – in tegenstelling tot alledaagse situaties waarin we gewoon rechtop staan of zitten – en beïnvloedde het ruimtelijk geheugen. In omstandigheden waarin de hoeveelheid ruimtelijke informatie echter minimaal is, zoals in de meeste van onze experimentele condities, is dit soort aanvullende informatie afwezig. Onder dergelijke omstandigheden kan blootgelegd worden welk referentiekader de hersenen als basis gebruiken om visueel-motorische informatie te verwerken. Dit referentiekader lijkt een ooggecentreerde representatie te zijn.

Concluderend werpen de resultaten die in dit proefschrift gepresenteerd zijn een nieuw licht op hoe de menselijke hersenen voor ruimtelijke constantie zorgen die nodig is voor het aansturen van bewegingen. Centraal in onze observaties staat de dominantie van ooggecentreerde coördinaten, alhoewel we ook hebben laten zien dat onder sommige omstandigheden ook andere ruimtelijke representaties gebruikt worden. Maar ook al lijken retinale coördinaten algemeen te zijn voor het behouden van ruimtelijke constantie, verdere referentiekadertransformaties zijn altijd nodig om bewegingen naar een herinnerd doel uit te voeren. De neurale basis van deze processen kunnen wellicht gevonden worden in pariëtale en/of frontale hersengebieden, maar het moge duidelijk zijn dat meer onderzoek nodig is om deze computationele functies aan specifieke neurofysiologische substraten toe te wijzen, en om de hierbij betrokken niet-retinale signalen te identificeren.

References

- Andersen RA, and Buneo CA. Intentional maps in posterior parietal cortex. *Annu Rev Neurosci* 25: 189–220, 2002.
- Andersen RA, Essick GK, and Siegel RM. Encoding of spatial location by posterior parietal neurons. *Science* 230: 456–458, 1985.
- Andersen RA, Snyder LH, Bradley DC, and Xing J. Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annu Rev Neurosci* 20:303–330, 1997.
- Angelaki DE, and Hess BJ. Inertial representation of angular motion in the vestibular system of rhesus monkeys. I. Vestibuloocular reflex. *J Neurophysiol* 71: 1222–1249, 1994.
- Aubert H. Eine scheinbare bedeutende Drehung von Objekten bei Neigung des Kopfes nach rechts oder links. *Virchows Arch* 20: 381–393, 1861.
- Avillac M, Denève S, Olivier E, Pouget A, and Duhamel JR. Reference frames for representing visual and tactile locations in parietal cortex. *Nat Neurosci* 8: 941–949, 2005.
- Baddeley A, and Della Sala S. Working memory and executive control. *Philos Trans R Soc Lond B Biol Sci* 351: 1397–1403, 1996.
- Baker JT, Harper TM, and Snyder LH. Spatial memory following shifts of gaze. I. Saccades to memorized world-fixed and gaze-fixed targets. *J Neurophysiol* 89: 2564–2576, 2003.
- Bartels A, Logothetis NK, and Moutoussis K. fMRI and its interpretations: an illustration on directional selectivity in area V5/MT. *Trends Neurosci* 31: 444–453, 2008.
- Batista AP, Buneo CA, Snyder LH, and Andersen RA. Reach plans in eye-centered coordinates. *Science* 285: 257–260, 1999.
- Battaglia—Mayer A, Caminiti R, Lacquaniti F, and Zago M. Multiple levels of representation of reaching in the parieto-frontal network. *Cereb Cortex* 13: 1009–1022, 2003.
- Baud—Bovy G and Viviani P. Pointing to kinesthetic targets in space. *J Neurosci* 18: 1528–1545, 1998.
- Ben Hamed S, Duhamel JR, Bremmer F, and Graf W. Representation of the visual field in the lateral intraparietal area of macaque monkeys: a quantitative receptive field analysis. *Exp Brain Res* 140: 27–44, 2001.
- Best PJ, White AM, and Minai A. Spatial processing in the brain: the activity of hippocampal place cells. *Annu Rev Neurosci* 24: 459–486, 2001.
- Beurze SM, de Lange FP, Toni I, and Medendorp WP. Integration of target and effector information in the human brain during reach planning. *J Neurophysiol* 97: 88–99, 2007.
- Beurze SM, Van Pelt S, and Medendorp WP. Behavioral reference frames for planning human reaching movements. *J Neurophysiol* 96: 352–362, 2006.
- Bhattacharyya R, Musallam S, and Andersen, RA. Reaching in depth: neural activity in posterior parietal cortex reflects distance to target and vergence angle. *Soc Neurosci Abstr*, 289.18, 2005.

- Blatt GJ, Andersen RA, and Stoner GR. Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. *J Comp Neurol* 299: 421–445, 1990.
- Blohm G, Missal M, and Lefèvre P. Interaction between smooth anticipation and saccades during ocular orientation in darkness. *J Neurophysiol* 89: 1423–1433, 2003.
- Blohm G, Optican LM, and Lefèvre P. A model that integrates eye velocity commands to keep track of smooth eye displacements. *J Comput Neurosci* 21: 51–70, 2006.
- Blouin J, Labrousse L, Simoneau M, Vercher JL, and Gauthier GM. Updating visual space during passive and voluntary head-in-space movements. *Exp Brain Res* 122: 93–101, 1998.
- Bock O. Contribution of retinal versus extraretinal signals towards visual localization in goal-directed movements. *Exp Brain Res* 64: 467–482, 1986.
- Bockisch CJ, and Haslwanter T. Three-dimensional eye position during static roll and pitch in humans. *Vision Res* 41: 2127–2137, 2001.
- Brandt T, and Dieterich M. The vestibular cortex. Its locations, functions, and disorders. *Ann NY Acad Sci* 871: 293–312, 1999.
- Brandt T, Dieterich M, and Danek A. Vestibular cortex lesions affect the perception of verticality. *Ann Neurol* 35 (4): 403–412, 1994.
- Bridgeman B, Peery S, and Anand S. Interaction of cognitive and sensorimotor maps of visual space. *Percept Psychophys* 59: 456–469, 1997.
- Brown MR, DeSouza JF, Goltz HC, Ford K, Menon RS, Goodale MA, and Everling S. Comparison of memory- and visually guided saccades using event-related fMRI. *J Neurophysiol* 91: 873–889, 2004.
- Brown MR, Goltz HC, Vilis T, Ford KA, and Everling S. Inhibition and generation of saccades: rapid event-related fMRI of prosaccades, antisaccades, and nogo trials. *Neuroimage* 33: 644–659, 2006.
- Bruce CJ, and Goldberg ME. Primate frontal eye fields. I. Single neurons discharging before saccades. *J Neurophysiol* 53: 603–635, 1985.
- Burgess N, Maguire EA, and O'Keefe J. The human hippocampus and spatial and episodic memory. *Neuron* 35: 625–641, 2002.
- Carrozzo M, Stratta F, McIntyre J, and Lacquaniti F. Cognitive allocentric representations of visual space shape pointing errors. *Exp Brain Res* 147: 426–437, 2002.
- Cassanello CR, and Ferrera VP. Visual remapping by vector subtraction: analysis of multiplicative gain field models. *Neural Comput* 19: 2353–2386, 2007.
- Churchland MM, Afshar A, and Shenoy KV. A central source of movement variability. *Neuron* 52: 1085–1096, 2006.
- Cohen YE, and Andersen RA. A common reference frame for movement plans in the posterior parietal cortex. *Nat Rev Neurosci* 3: 553–562, 2002.
- Colby CL, and Goldberg ME. Space and attention in parietal cortex. *Annu Rev Neurosci* 22: 319–349, 1999.

- Connolly JD, Goodale MA, Cant JS, and Munoz DP. Effector specific fields for motor preparation in the human frontal cortex. *NeuroImage* 34: 1209–1219, 2007.
- Connolly JD, Goodale MA, Menon RS, and Munoz DP. Human fMRI evidence for the neural correlates of preparatory set. *Nat Neurosci* 5: 1345–1352, 2002.
- Crawford JD, and Guitton D. Visual–motor transformations required for accurate and kinematically correct saccades. *J Neurophysiol* 78: 1447–1467, 1997.
- Crawford JD, Medendorp WP, and Marotta JJ. Spatial transformations for eye–hand coordination. *J Neurophysiol* 92: 10–19, 2004.
- Cuijpers RH, Kappers AM, and Koenderink JJ. Visual perception of collinearity. *Percept Psychophys* 64: 392–404, 2002.
- Culham JC, and Valyear KF. Human parietal cortex in action. *Curr Opin Neurobiol* 16: 205–212, 2006.
- Cumming BG, and DeAngelis GC. The physiology of stereopsis. *Annu Rev Neurosci* 24: 203–238, 2001.
- Curtis CE, and Connolly JD. Saccade preparation signals in the human frontal and parietal cortices. *J Neurophysiol* 99: 133–45, 2008.
- Curtis CE, and D'Esposito M. Selection and maintenance of saccade goals in the human frontal eye fields. *J Neurophysiol* 95:3923–3927, 2006.
- Dassonville P, Bridgeman B, Kaur Bala J, Thiem P, and Sampanes A. The induced Roelofs effect: two visual systems or the shift of a single reference frame? *Vision Res* 44: 603–611, 2004.
- Dassonville P, Schlag J, and Schlag-Rey M. The use of egocentric and exocentric location cues in saccadic programming. *Vision Res* 35: 2191–2199, 1995.
- DeAngelis GC. Seeing in three dimensions: the neurophysiology of stereopsis. *TICS* 4: 80–90, 2000.
- Desimone R. Neural mechanisms for visual memory and their role in attention. *Proc Natl Acad Sci USA* 93: 13494–13499, 1996.
- Dinstein I, Hasson U, Rubin N, and Heeger DJ. Brain areas selective for both observed and executed movements. *J Neurophysiol* 98: 1415–1427, 2007.
- Dobbins AC, Fiser J, and Allman JM. Distance modulation of neural activity in the visual cortex. *Science* 281: 552–555, 1998.
- Duhamel JR, Colby CL, and Goldberg ME. The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* 255: 90–92, 1992.
- Durbin R, and Mitchison G. A dimension reduction framework for understanding cortical maps. *Nature* 343: 644–647, 1990.
- Eggert T. *Der Einfluss orientierter Texturen auf die subjektive Vertikale und seine systemtheoretische Analyse* (PhD thesis). Munich, Germany: Munich Technical University, 1998.
- Ferraina S, Paré M, and Wurtz RH. Comparison of cortico-cortical and cortico-collicular signals for the generation of saccadic eye movements. *J Neurophysiol* 87: 845–858, 2002.

- Ferraina S, Paré M, and Wurtz RH. Disparity sensitivity of frontal eye field neurons. *J Neurophysiol* 83: 625–629, 2000.
- Flanders M, Tillery SI, and Soechting JF. Early stages in a sensorimotor transformation. *Behav Brain Sci* 15: 309–362, 1992.
- Friston K. A theory of cortical responses. *Philos Trans R Soc Lond B Biol Sci* 360: 815–836, 2005.
- Fukushima K, Yamanobe T, Shinmei Y, Fukushima J, Kurkin S, and Peterson BW. Coding of smooth eye movements in three-dimensional space by frontal cortex. *Nature* 419: 157–162, 2002.
- Gardner JL, Merriam EP, Movshon JA, and Heeger DJ. Maps of visual space in human occipital cortex are retinotopic, not spatiotopic *J Neurosci* 28: 3988–3999, 2008
- Gellman RS, and Fletcher WA. Eye position signals in human saccadic processing. *Exp Brain Res* 89: 425–434, 1992.
- Genovese CR, Lazar NA, and Nichols T. Thresholding of statistical maps in functional neuroimaging using the false discovery rate. *Neuroimage* 15: 870–878, 2002.
- Genovesio A, and Ferraina S. Integration of retinal disparity and fixation–distance related signals toward an egocentric coding of distance in the posterior parietal cortex of primates. *J Neurophysiol* 91: 2670–2684, 2004.
- Genovesio A, Brunamonti E, Giusti MA, and Ferraina S. Postsaccadic activities in the posterior parietal cortex of primates are influenced by both eye movement vectors and eye position. *J Neurosci* 27: 3268–3273, 2007.
- Gibson JJ, Olum P, and Rosenblatt F. Parallax and perspective during aircraft landings. *Am J Psychol* 68: 372–385, 1955.
- Glasauer S, Amorim MA, Vitte E, and Berthoz A. Goal-directed linear locomotion in normal and labyrinthine-defective subjects. *Exp Brain Res* 98: 323–335, 1994.
- Gnadt JW, and Andersen RA. Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res* 70: 216–220, 1988.
- Gnadt JW, and Beyer J. Eye movements in depth: what does the monkey's parietal cortex tell the superior colliculus? *Neuroreport* 9: 233–238, 1998.
- Gnadt JW, and Mays LE. Neurons in monkey parietal area LIP are tuned for eye-movement parameters in three-dimensional space. *J Neurophysiol* 73: 280–297, 1995.
- Goldberg ME, and Bruce CJ. Primate frontal eye fields. III. Maintenance of a spatially accurate saccade signal. *J Neurophysiol* 64: 489–508, 1990.
- Goldman-Rakic PS. The prefrontal landscape: implications of functional architecture for understanding human mentation and the central executive. *Philos Trans R Soc Lond B Biol Sci* 351: 1445–1453, 1996.
- Gonzalez F, Rivadulla C, Perez R, and Cadarso C. Depth perception in random dot stereograms is not affected by changes in either vergence or accommodation. *Optom Vis Sci* 75: 743–747, 1998.

- Goodale MA, and Milner AD. Separate visual pathways for perception and action. *Trends Neurosci* 15: 20–25, 1992.
- Graziano MS, and Aflalo TN. Mapping behavioral repertoire onto the cortex. *Neuron* 56: 239–251, 2007.
- Green AM, and Angelaki DE. Coordinate transformations and sensory integration in the detection of spatial orientation and self-motion: from models to experiments. *Prog Brain Res* 165:155–180, 2007.
- Grefkes C, Ritzl A, Zilles K, and Fink GR. Human medial intraparietal cortex subserves visuomotor coordinate transformation. *Neuroimage* 23: 1494–1506, 2004.
- Grefkes C, Weiss PH, Zilles K, and Fink GR. Crossmodal processing of object features in human anterior intraparietal cortex: an fMRI study implies equivalencies between humans and monkeys. *Neuron* 35: 173–184, 2002.
- Grill-Spector K, and Malach R. fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta Psychologica* 107:293–321, 2001.
- Grill-Spector K, Henson R, and Martin A. Repetition and the brain: neural models of stimulus-specific effects. *Trends Cogn Sci* 10:14–23, 2006.
- Grosbras MH, Lobel E, Van de Moortele PF, Lebihan D, and Berthoz A. An anatomical landmark for the supplementary eye fields in human revealed with functional magnetic resonance imaging. *Cereb Cortex* 9: 705–711, 1999.
- Grusser OJ, Pause M, and Schreier U. Vestibular neurones in the parieto-insular cortex of monkeys (*Macaca fascicularis*): visual and neck receptor responses. *J Physiol* 430: 559–583, 1990.
- Guitton D. Control of eye-head coordination during orienting gaze shifts. *Trends Neurosci* 15: 174–179, 1992.
- Hagler DJ Jr, and Sereno MI. Spatial maps in frontal and prefrontal cortex. *Neuroimage* 29: 567–577, 2006.
- Hallett PE, and Lightstone AD. Saccadic eye movements towards stimuli triggered by prior saccades. *Vision Res* 16: 99–106, 1976.
- Hamilton A, and Grafton S. Action outcomes are represented in human inferior fronto-parietal cortex. *Cereb Cortex* 18:1160–1168, 2008.
- Hamilton A, and Grafton S. Goal representation in human anterior intraparietal sulcus. *J Neurosci* 26:1133–1137, 2006.
- Hayhoe MM, Shrivastava A, Mruczek R, and Pelz JB. Visual memory and motor planning in a natural task. *J Vis* 3: 49–63, 2003.
- Heide W, Blankenburg M, Zimmermann E, and Kömpf D. Cortical control of double-step saccades: implications for spatial orientation. *Ann Neurol* 38: 739–748, 1995.
- Heiser LM, and Colby CL. Spatial updating in area LIP is independent of saccade direction. *J Neurophysiol* 95: 2751–2767, 2006.

- Henriques DY, Medendorp WP, Gielen CC, and Crawford JD. Geometric computations underlying eye-hand coordination: orientations of the two eyes and the head. *Exp Brain Res* 152: 70–78, 2003.
- Henriques DY, Klier EM, Smith MA, Lowy D, and Crawford JD. Gaze-centered remapping of remembered visual space in an open-loop pointing task. *J Neurosci* 18:1583–1594, 1998.
- Henson RN, and Rugg MD. Effects of stimulus repetition on latency of the BOLD impulse response. *NeuroImage* 13: 683, 2001.
- Herter TM, and Guitton D. Human head-free gaze saccades to targets flashed before gaze-pursuit are spatially accurate. *J Neurophysiol* 80: 2785–2789, 1998.
- Hess BJ, and Angelaki DE. Inertial processing of vestibule-ocular signals. *Ann NY Acad Sci* 871: 148–161, 1999.
- Hess BJ, and Angelaki DE. Inertial vestibular coding of motion: concepts and evidence. *Curr Op Neurobiol* 7: 860–866, 1997.
- Honda H. Modification of saccade-contingent visual mislocalization by the presence of a visual frame of reference. *Vision Res* 39: 51–57, 1999.
- Howard IP, and Rogers BJ. *Binocular Vision and Stereopsis*. New York: Oxford Claridon, 1995.
- Howard IP. *Human visual orientation*. New York: Wiley, 1982.
- Israel I, and Berthoz A. Contribution of the otoliths to the calculation of linear displacement. *J Neurophysiol* 62: 247–263, 1989.
- Israel I, Chapis N, Glasauer S, Charade O, and Berthoz A. Estimation of passive horizontal linear whole-body displacement in humans. *J Neurophysiol* 70: 1270–1273, 1993.
- Israel I, Ventre-Dominey J, and Denise P. Vestibular information contributes to update retinotopic maps. *Neuroreport* 10: 3479–3483, 1999.
- Jack AI, Patel GH, Astafiev SV, Snyder AZ, Akbudak E, Shulman GL, and Corbetta M. Changing human visual field organization from early visual to extra-occipital cortex. *PLoS ONE* 2:e452, 2007.
- Jackson SR, and Husain M. Visuomotor functions of the posterior parietal cortex. *Neuropsychologia* 44: 2589–2593, 2006.
- Julesz B. *Foundations of Cyclopean Perception*. Chicago: Chicago University Press, 1971.
- Kaiser M, and Lappe M. Perisaccadic mislocalization orthogonal to saccade direction. *Neuron* 41: 293–300, 2004.
- Kandel ER, Schwartz JH, and Jessell TM (eds). *Principles of Neural Science*. McGraw-Hill New York, 2000.
- Kaptein RG, and Van Gisbergen JAM. Interpretation of a discontinuity in the sense of verticality at large body tilt. *J Neurophysiol* 91: 2205–2214, 2004.
- Karn KS, Moller P, and Hayhoe MM. Reference frames in saccadic targeting. *Exp Brain Res* 115: 267–282, 1997.

- Kastner S, DeSimone K, Konen CS, Szczepanski SM, Weiner KS, and Schneider KA. Topographic maps in human frontal cortex revealed in memory-guided saccade and spatial working-memory tasks. *J Neurophysiol* 97: 3494–3507, 2007.
- Keith GP, and Crawford JD. Saccade-related remapping of target representations between topographic maps: a neural network study. *J Comput Neurosci* 24: 157–178, 2007.
- Keith GP, Smith MA, and Crawford JD. Functional organization within a neural network trained to update target representations across 3-D saccades. *J Comput Neurosci* 22: 191–209, 2007.
- Khan AZ, Pisella L, Rossetti Y, Vighetto A, and Crawford JD. Impairment of gaze-centered updating of reach targets in bilateral parietal-occipital damaged patients. *Cereb Cortex* 15: 1547–1560, 2005.
- Klier EM, and Crawford JD. Human oculomotor system accounts for 3-D eye orientation in the visual-motor transformation for saccades. *J Neurophysiol* 80: 2274–2294, 1998.
- Klier EM, Angelaki DE, and Hess BJ. Human visuospatial updating after noncommutative rotations. *J Neurophysiol* 98: 537–544, 2007.
- Klier EM, Angelaki DE, and Hess BJ. Roles of gravitational cues and efference copy signals in the rotational updating of memory saccades. *J Neurophysiol* 94: 468–478, 2005.
- Klier EM, Hess BJ, and Angelaki DE. Human visuospatial updating after passive translations in three-dimensional space. *J Neurophysiol* 99: 1799–1809, 2008.
- Kohonen T. *Self-Organizing Maps*. Berlin: Springer, 2001.
- Koyama M, Hasegawa I, Osada T, Adachi Y, Nakahara K, and Miyashita Y. Functional magnetic resonance imaging of macaque monkeys performing visually guided saccade tasks: comparison of cortical eye fields with humans. *Neuron* 41:795–807, 2004.
- Krommenhoek KP, and Van Gisbergen JA. Evidence for nonretinal feedback in combined version-vergence eye movements. *Exp Brain Res* 102: 95–109, 1994.
- Kudoh N. Dissociation between visual perception of allocentric distance and visually directed walking of its extent. *Perception* 34: 1399–1416, 2005.
- Laurens J, and Droulez J. Bayesian processing of vestibular information. *Biol Cybern* 96: 389–404, 2007.
- Li N, and Angelaki DE. Updating visual space during motion in depth. *Neuron* 48: 149–158, 2005.
- Li N, Wei M, and Angelaki DE. Primate memory saccade amplitude after intervened motion depends on target distance. *J Neurophysiol* 94: 722–733, 2005.
- Logothetis NK. What we can do and what we cannot do with fMRI. *Nature* 453: 869–878, 2008.
- Ma WJ, Beck JM, Latham PE, and Pouget A. Bayesian inference with probabilistic population codes. *Nat Neurosci* 9: 1432–1438, 2006.
- Martinez-Trujillo JC, Medendorp WP, Wang H, and Crawford JD. Frames of reference for eye-head gaze commands in primate supplementary eye fields. *Neuron* 44: 1057–1066, 2004.

- Mast F, and Jarchow T. Perceived body position and the visual horizontal. *Brain Res Bull* 40: 393–398, 1996.
- Mays LE, and Sparks DL. Saccades are spatially, not retinocentrically, coded. *Science* 208: 1163–1165, 1980.
- McIntyre J, Stratta F, and Lacquaniti F. A viewer-centered reference frame for pointing to memorized targets in three-dimensional space. *J Neurophysiol* 78: 1601–1618, 1997.
- McKenzie A, and Lisberger SG. Properties of signals that determine the amplitude and direction of saccadic eye movements in monkeys. *J Neurophysiol* 56: 196–207, 1986.
- Medendorp WP, and Crawford JD. Visuospatial updating of reaching targets in near and far space. *Neuroreport* 13: 633–636, 2002.
- Medendorp WP, Goltz HC, Crawford JD, and Vilis T. Integration of target and effector information in human posterior parietal cortex for the planning of action. *J Neurophysiol* 93: 954–962, 2005.
- Medendorp WP, Goltz HC, Vilis T, and Crawford JD. Gaze-centered updating of visual space in human parietal cortex. *J Neurosci* 23: 6209–6214, 2003a.
- Medendorp WP, Goltz HC, and Vilis T. Directional selectivity of BOLD activity in human posterior parietal cortex for memory-guided double-step saccades. *J Neurophysiol* 95: 1645–1655, 2006.
- Medendorp WP, Smith MA, Tweed DB, and Crawford JD. Rotational remapping in human spatial memory during eye and head motion. *J Neurosci* 22: 196RC, 2002.
- Medendorp WP, Tweed DB, and Crawford JD. Motion parallax is computed in the updating of human spatial memory. *J Neurosci* 23: 8135–8142, 2003b.
- Medendorp WP, Van Asselt S, and Gielen CC. Pointing to remembered visual targets after active one-step self-displacements within reaching space. *Exp Brain Res* 125: 50–61, 1999.
- Merfeld DM, Ypung LR, Oman CM, and Shelhamer MJ. A multi-dimensional model of the effect of gravity on the spatial orientation of the monkey. *J Vestib Res* 3: 141–161, 1993b.
- Merfeld DM, Young LR, Paige GD, and Tomko DL. Three dimensional eye movements of squirrel monkeys following postrotatory tilt. *J Vestib Res* 3: 123–139, 1993a.
- Merfeld DM, Zupan L, and Peterka RJ. Humans use internal models to estimate gravity and linear acceleration. *Nature* 398: 615–618, 1999.
- Mergner T, Nasios G, Maurer C, and Becker W. Visual object localisation in space. *Exp Brain Res* 141: 33–52, 2001.
- Merriam EP, Genovese CR, and Colby CL. Remapping in human visual cortex. *J Neurophysiol* 97: 1738–1755, 2007.
- Merriam EP, Genovese CR, and Colby CL. Spatial updating in human parietal cortex. *Neuron* 39: 361–373, 2003.
- Mittelstaedt H. A new solution to the problem of the subjective vertical. *Naturwissenschaften* 70: 272–281, 1983.

- Moschovakis AK, Dalezios Y, Petit J, and Grantyn AA. New mechanism that accounts for position sensitivity of saccades evoked in response to stimulation of superior colliculus. *J Neurophysiol* 80: 3373–3379, 1998.
- Muri RM, Rivaud S, Timsit S, Cornu P, and Pierrot–Deseilligny C. The role of the right medial temporal lobe in the control of memory–guided saccades. *Exp Brain Res* 101: 165–168, 1994.
- Naccache L, and Dehaene S. The priming method: imaging unconscious repetition priming reveals an abstract representation of number in the parietal lobes. *Cereb Cortex* 11:966–974, 2001.
- Nakamura K, and Colby CL. Updating of the visual representation in monkey striate and extrastriate cortex during saccades. *Proc Natl Acad Sci* 99: 4026–4031, 2002.
- Olson CR. Brain representation of object–centered space in monkeys and humans. *Annu Rev Neurosci* 26: 331–354, 2003.
- Palmer SE. *Vision science: Photons to Phenomenology*. Cambridge MA: MIT Press, 1999.
- Paus T. Location and function of the human frontal eye–field: a selective review. *Neuropsychologia* 34: 475–483, 1996.
- Pettorossi VE, Bambagioni D, Bronstein AM, and Gresty MA. Assessment of the perception of verticality and horizontality with self–paced saccades. *Exp Brain Res* 121: 46–50, 1998.
- Philbeck JW, and Loomis JM. Comparison of two indicators of perceived egocentric distance under full–cue and reduced–cue conditions. *J Exp Psychol Hum Percept Perform* 23: 72–85, 1997.
- Picard N, and Strick PL. Imaging the premotor areas. *Curr Opin Neurobiol* 11: 663–672, 2001.
- Pierrot–Deseilligny C, Milea D, and Müri RM. Eye movement control by the cerebral cortex. *Curr Opin Neurol* 17: 17–25, 2004.
- Pierrot–Deseilligny DC, Muri RM, Rivaud PS, Gaymard B, and Ploner CJ. Cortical control of spatial memory in humans: The visuoculomotor model. *Ann Neurol* 52: 10–19, 2002.
- Poggio GF, and Poggio T. The analysis of stereopsis. *Annu Rev Neurosci* 7: 379–412, 1984.
- Poljac E, and Van den Berg AV. Localization of the plane of regard in space. *Exp Brain Res* 163: 457–467, 2003.
- Pouget A, Ducom JC, Torri J, and Bavelier D. Multisensory spatial representations in eye–centered coordinates for reaching. *Cognition* 83: 1–11, 2002.
- Pouget A, and Snyder LH. Computational approaches to sensorimotor transformations. *Nat Neurosci* 3 Suppl: 1192–1198, 2000.
- Press WH, Flannery BP, Teukolsky SA, and Vetterling WT. *Numerical recipes in C (2nd ed.)*. Cambridge, MA: Cambridge, 1992.
- Quaia C, Optican LM, and Goldberg ME. The maintenance of spatial accuracy by the perisaccadic remapping of visual receptive fields. *Neural Netw* 11: 1229–1240, 1998.
- Raffi M, Squatrito S, and Maioli MG. Gaze and smooth pursuit signals interact in parietal area 7m of the behaving monkey. *Exp Brain Res* 182: 35–46, 2007.

- Randall D, Burggren W, and French K (eds). *Eckert Animal Physiology: Mechanisms and Adaptations*. WH Freeman and Company, 1997.
- Robinson DA, and Fuchs AF. Eye movements evoked by stimulation of frontal eye fields. *J Neurophysiol* 32: 637–648, 1969.
- Rogers B, and Graham M. Motion parallax as an independent cue for depth perception. *Perception* 8: 125–134, 1979.
- Rosenbluth D, and Allman JM. The effect of gaze angle and fixation distance on the responses of neurons in V1, V2, and V4. *Neuron* 33: 143–149, 2002.
- Russo GS, and Bruce CJ. Supplementary eye field: representation of saccades and relationship between neural response fields and elicited eye movements. *J Neurophysiol* 84: 2605–2621, 2000.
- Sakata H, Taira M, Kusunoki M, Murata A, and Tanaka Y. The TINS Lecture. The parietal association cortex in depth perception and visual control of hand action. *Trends Neurosci* 20: 350–357, 1997.
- Schall JD. Neuronal-activity related to visually guided saccadic eye-movements in the supplementary motor area of rhesus-monkeys. *J Neurophysiol* 66: 530–558, 1991.
- Schlag J, and Schlag-Rey M. Evidence for a supplementary eye field. *J Neurophysiol* 57: 179–200, 1987.
- Schlag J, Schlag-Rey M, and Dassonville P. Saccades can be aimed at the spatial location of targets flashed during pursuit. *J Neurophysiol* 64: 575–581, 1990.
- Schluppeck D, Glimcher P, and Heeger J. Topographic organization for delayed saccades in human posterior parietal cortex. *J Neurophysiol* 94: 1372–1384, 2005.
- Serences JT. A comparison of methods for characterizing the event-related BOLD timeseries in rapid fMRI. *Neuroimage* 21: 1690–1700, 2004.
- Sereno MI, and Huang RS. A human parietal face area contains aligned head-centered visual and tactile maps. *Nat Neurosci* 9: 1337–1343, 2006.
- Sereno MI, Pitzalis S, and Martinez A. Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science* 294: 1350–1354, 2001.
- Shadmehr R and Wise SP. *The computational neurobiology of reaching and pointing: a foundation for motor learning*. Cambridge, MA: MIT Press, 2005.
- Skavenski AA, and Steinman RM. Control of eye position in the dark. *Vision Res* 10: 193–203, 1970.
- Smith MA, and Crawford JD. Implications of ocular kinematics for the internal updating of visual space. *J Neurophysiol* 86: 2112–2117, 2001.
- Snyder LH, Batista AP, and Andersen RA. Intention-related activity in the posterior parietal cortex: a review. *Vision Res* 40: 1433–1441, 2000.
- Snyder LH, Grieve KL, Brotchie P, and Anderson RA. Separate body- and world-referenced representations of visual space in parietal cortex. *Nature* 394: 887–891, 1998.
- Snyder LH. Coordinate transformations for eye and arm movements in the brain. *Curr Opin Neurobiol* 10: 747–754, 2000.

- Soechting JF, and Flanders M. Moving in three-dimensional space: frames of reference, vectors, and coordinate systems. *Annu Rev Neurosci* 15: 167–191, 1992.
- Sommer MA, and Wurtz RH. A pathway in primate brain for internal monitoring of movements. *Science* 296: 1480–1482, 2002.
- Sommer MA, and Wurtz RH. Brain circuits for the internal monitoring of movements. *Annu Rev Neurosci* 31: 317–338, 2008.
- Sparks D, and Nelson J. Sensory and motor maps in the mammalian superior colliculus. *TINS* 10: 312–317, 1987.
- Sparks DL, and Mays LE. Spatial localization of saccade targets. I. Compensation for stimulation-induced perturbations in eye position. *J Neurophysiol* 49: 45–63, 1983.
- Sperry RW. Neural basis of the spontaneous optokinetic response produced by visual inversion. *J Comp Physiol Psychol* 43: 482–489, 1950.
- Talairach J, and Tournoux PA. *A co-planar stereotaxic atlas of the human brain*. New York: Thieme Medical Publishers, 1988.
- Tehovnik EJ, and Lee K. The dorsomedial frontal cortex of the rhesus monkey: topographic representation of saccades evoked by electrical stimulation. *Exp Brain Res* 96: 430–442, 1993.
- Tweed D, and Vilis T. Implications of rotational kinematics for the oculomotor system in three dimensions. *J Neurophysiol* 58: 832–849, 1987.
- Umeno MM, and Goldberg ME. Spatial processing in the monkey frontal eye field. I. Predictive visual responses. *J Neurophysiol* 78: 1373–1383, 1997.
- Van Beuzekom AD, and Van Gisbergen JAM. Properties of the internal representation of gravity inferred from spatial-direction and body-tilt estimates. *J Neurophysiol* 84: 11–27, 2000.
- Van Beuzekom AD, and Van Gisbergen JAM. The subjective vertical and the sense of self orientation during active body tilt. *Vision Res* 41: 3229–3242, 2001.
- Van Turenout M, Bielamowicz L, and Martin A. Modulation of neural activity during object naming: effects of time and practice. *Cereb Cortex* 13: 381–391, 2003.
- Vaziri S, Diedrichsen J, and Shadmehr R. Why does the brain predict sensory consequences of oculomotor commands? Optimal integration of the predicted and the actual sensory feedback. *J Neurosci* 26: 4188–4197, 2006.
- Vindras P, Desmurget M, and Viviani P. Error parsing in visuomotor pointing reveals independent processing of amplitude and direction. *J of Neurophysiol*, 94: 1212–1224, 2005.
- Vliegen J, Van Grootel TJ, and Van Opstal AJ. Gaze orienting in dynamic visual double steps. *J Neurophysiol* 94: 4300–4313, 2005.
- Von Helmholtz H. *Handbuch der Physiologischen Optik*. Leipzig: Voss, 1867.
- Von Holst E, and Mittelstaedt H. The reafferent principle: reciprocal effects between central nervous system and periphery. *Naturwissenschaften* 37, 464–476, 1950.

- Walker MF, Fitzgibbon J, and Goldberg ME. Neurons of the monkey superior colliculus predict the visual result of impending saccadic eye movements. *J Neurophysiol* 73: 1988–2003, 1995.
- Wei M, DeAngelis GC, and Angelaki DE. Do visual cues contribute to the neural estimate of viewing distance used by the oculomotor system? *J Neurosci* 23: 8340–8350, 2003.
- Wexler M. Anticipating the three-dimensional consequences of eye movements. *PROC NATL ACAD SCI USA* 102: 1246–1251, 2005.
- Wexler M. Voluntary head movement and allocentric perception of space. *Psychol Sci* 14: 340–346, 2003.
- White RL III, and Snyder LH. A neural network model of flexible spatial updating. *J Neurophysiol* 91: 1608–1619, 2004.
- Wiggs CL, and Martin A. Properties and mechanisms of perceptual priming. *Curr Opin Neurobiol* 8: 227–233, 1998.
- Yelnik AP, Lebreton FO, Bonan IV, Colle FM, Meurin FA, Guichard JP, and Vicaud E. Perception of verticality after recent cerebral hemispheric stroke. *Stroke* 33: 2247–2254, 2002.
- Zhang T, Heuer HW, and Britten KH. Parietal area VIP neuronal responses to heading stimuli are encoded in head-centered coordinates. *Neuron* 42: 993–1001, 2004.

Publication list

Journal publications

- Van Pelt S, Toni I, Diedrichsen J, and Medendorp WP. Repetition suppression dissociates spatial frames of reference in human saccade planning. *Submitted for publication*.
- Van Pelt S, and Medendorp WP. Updating target distance across eye movements in depth. *Journal of Neurophysiology* 99: 2281–2290, 2008.
- Medendorp WP, Beurze SM, Van Pelt S, and Van Der Werf J. Behavioral and cortical mechanisms for spatial coding and action planning. *Cortex* 44: 587–597, 2008.
- Van Pelt S, and Medendorp WP. Gaze-centered updating of remembered visual space during active whole-body translations. *Journal of Neurophysiology* 97: 1209–1220, 2007.
- Beurze SM, Van Pelt S, and Medendorp WP. Behavioral reference frames for planning human reaching movements. *Journal of Neurophysiology* 96: 352–362, 2006.
- Van Pelt S, Van Gisbergen JAM, and Medendorp WP. Visuospatial memory computations during whole-body rotations in roll. *Journal of Neurophysiology* 94: 1432–1442, 2005.
- Medendorp WP, Van Gisbergen JAM, Van Pelt S, and Gielen CCAM. Context compensation in the vestibuloocular reflex during active head rotations. *Journal of Neurophysiology* 84: 2904–2917, 2000.

Conference proceedings and talks

- Van Pelt S, Diedrichsen J, Toni I, and Medendorp WP (November 2008). Repetition suppression distinguishes between retinal and nonretinal spatial reference frames in human cortex [oral presentation]. *Society for Neuroscience Annual Meeting*, Washington, US.
- Van Pelt S, and Medendorp WP (October 2007). Retinal updating of target location across vergence eye movements [poster]. *ESF-EMBO Symposium – Three dimensional sensory and motor space: perceptual consequences of motor action*, Sant Feliu de Guixols, Spain.
- Van Pelt S, and Medendorp WP (September 2007). Spatial representations for action control. *Midlands & Wales Motor Meeting*, Bangor University, UK.

- Medendorp W.P, Van Pelt S (August 2007). Representation of target distance across eye movements in depth [poster]. *European Conference on Eye Movements (ECEM)*, Potsdam, Germany.
- Van Pelt S, and Medendorp WP (October 2006). Representation of target distance across eye movements in depth [poster]. *Society for Neuroscience Annual Meeting*, Atlanta, US.
- Van Pelt S, and Medendorp WP (May 2006). Reach targets are remembered in a gaze-centered reference frame during whole-body motion [oral presentation]. *Bi-annual Workshop on Motor Control RU-KUL*, Nijmegen, The Netherlands.
- Van Pelt S, and Medendorp WP (November 2005). Updating of reaching space for translational motion [poster]. *Society for Neuroscience Annual Meeting*, Washington, US.
- Van Pelt S, Van Gisbergen JAM and Medendorp WP (October 2004). Spatial memory computations during passive whole-body rotations in roll [poster]. *Society for Neuroscience Annual Meeting*, San Diego, US.
- Van Pelt S, Van Gisbergen JAM, and Medendorp WP (May 2004). Spatial memory computations during passive whole-body tilts [oral presentation]. *Bi-annual Workshop on Motor Control KUN-KUL*, Leuven, Belgium.

Dankwoord

Het cliché is waar, een proefschrift schrijf je niet alleen. Maar met hoeveel mensen dan wel? Ondanks dat het onmogelijk is een volledig beeld te geven van wie allemaal bijgedragen hebben aan de totstandkoming van dit boekje, zal ik op deze plek toch een poging wagen.

Pieter, aan jou ben ik zonder twijfel de meeste dank verschuldigd. Ik ben allereerst erg blij dat je me de kans gaf bij je te komen werken nadat ik een tijdje uit dit onderzoeksveld was geweest. Maar ook – of beter, met name – tijdens het onderzoekstraject was je van cruciaal belang. Je hebt een feilloos gevoel voor wat de belangrijke onderzoeksvragen van het moment zijn en hebt een dito visie op wat de juiste wijze is om deze op een degelijke wijze te benaderen. Met heldere, hypothesegestuurde experimenten van een hoge standaard en interpretaties zonder wol. Maar minstens net zo belangrijk was en is je betrokkenheid bij het onderzoek van je promovendi; ik kan me niet herinneren dat ik ooit *niet* bij je kon binnenlopen met een vraag, resultatenfiguren of andere zaken – je had nagenoeg altijd tijd. Het gezamenlijk actief brainstormen op papier of op het whiteboard bij aanvang van elk nieuw experiment en het bekijken van de eerste resultaten van de daaruit volgende pilots waren altijd spannende tijden. Ook waardeer ik het belang dat je hecht aan het bewaken van een realistische tijdsplanning, zonder op kwaliteit te willen inleveren; je intensieve ondersteuning van het schrijfsproces was daarbij erg behulp- en leerzaam. Afgezien hiervan was het ook gewoon erg prettig samenwerken omdat het op persoonlijk vlak goed klikte. Bovendien was het natuurlijk ook erg gezellig in je steeds verder uitdijende onderzoeksgroep, met Rens, Sabine, Maaïke en Jurrian, zeker tijdens onze bijna jaarlijkse SfN-uitjes naar de VS.

Harold, ik dank je voor je rol als promotor; ik heb bewondering voor hoe jouw ambitie, inspiratie en daadkracht de AIM-groep en het NICI een nieuw tijdperk binnen hebben geloodst als onderdeel van het Donders Institute. Jan en Ivan, bedankt voor jullie samenwerking bij de experimenten die ik op jullie afdelingen uit mocht voeren. Jan wil ik, samen met Raymond, ook bedanken voor de scherpe en deskundige commentaren bij het kritisch doorlezen van verschillende van mijn artikelen.

Goed onderzoek staat of valt met de kwaliteit van de technische infrastructuur. Wat dat betreft is het NICI goed toebedeeld. Norbert, Jos, Lee, Willem, Eric en Pascal, jullie overtroffen elke keer weer mijn verwachtingen als er een nieuwe experimentele opstelling gemaakt moest worden, vaak slechts op basis van een papieren schets of een prototype van karton dat ik vaak tot jullie afgrijzen als ‘functionele’ opstelling in pilotexperimenten gebruikte. Chris wil ik erg bedanken voor het schrijven van de specialistische software om al deze high-tech aan te sturen en voor het gedogen van de enigszins anarchistische organisatie in het lab. En zonder de professionele assistentie van Paul waren mijn fMRI-experimenten nooit zo vlot verlopen. Dankzij Gerard en

André kon ik altijd over de benodigde hard- en software beschikken, en Yvonne, Beppie en Thea wisten wel raad met alle regelvragen.

Ook zonder de hulp van stagiaires was dit boekje niet tot stand gekomen. Joke, Rob, Jesse en Margriet, het was erg leuk en leerzaam om jullie te mogen begeleiden. Jullie kritische vragen hielden me alert, en jullie vaak aanwezige enthousiasme en motivatie waren aanstekelijk. Het is goed om te zien dat jullie allemaal jullie eigen weg gevonden hebben, al dan niet binnen een academische setting. Dank ben ik ook verschuldigd aan de vele proefpersonen die aan mijn onderzoek mee wilden doen, veelal collega-promovendi van het NICI, de afdeling Biofysica of het FCDC.

Stan Gielen en Thom Oostendorp wil ik bedanken voor de mogelijkheden die ze me boden om voor hen onderwijs te verzorgen binnen de opleidingen cognitive neuroscience en geneeskunde.

Ook al is je onderzoeksonderwerp nog zo leuk, zonder leuke collega's is het werk toch zwaar. Ik heb erg genoten van de gezellige, informele sfeer binnen het NICI, aan de koffietafel (al dan niet gedomineerd door verbale zwaargewichten van het eerste NICI-uur), tijdens de vele vrijdagmiddagborrels, volleybalduels in de pauze en het zeilweekend. Als bioloog tussen de cognitieve wetenschappers moest ik in het begin wel wennen aan het andere jargon en theoretische concepten, en de 'tradities', maar ik ben me er erg thuis gaan voelen, op het zingen van Sinterklaas- en Kerstliedjes na wellicht. NICI-leden, ik wil jullie allemaal heel erg bedanken voor deze geweldige tijd waardoor het werk zelden als 'werk' voelde. Enkelen van jullie wil ik hier toch even met naam noemen. Sabine, ondanks onze verschillende karakters hebben we toch altijd een fijne manier van samenwerken gehad, misschien ook mede doordat we volgens mij op een vergelijkbaar nuchtere manier naar het fenomeen 'aioschap' kijken. Daarnaast heb ik erg veel gehad aan hoe je me wegwijs hebt gemaakt in de wondere wereld van fMRI en BrainVoyager. Janneke, door jou was het prettig vertoeven op kamer B.01.02, onder andere door onze levendige, persoonlijke en eigenwijze discussies over het wetenschappelijke bedrijf en onze mogelijkheden daarbinnen en –buiten. Ik vind het bijzonder knap hoe je de laatste fase van je promotieonderzoek in sneltreinvaart hebt doorlopen gezien de ups en downs tijdens sommige periodes van je project; ik vind het erg leuk dat we nu zo kort na elkaar promoveren. Daarnaast heb ik erg genoten als reisgenoot op één van je vele treinreisavonturen, naar ons congres aan de Costa Brava. Jasminka, bedankt voor je hulp bij het opzetten van mijn fMRI-experiment en het assisteren bij het scannen; ook vond ik het altijd erg leuk om met je van gedachten te wisselen over zowel de frustraties als de zegeningen van het promotieproces en het grote Daarna. Ellen, je bent één van de weinige jonge onderzoekers die er bij mijn start bij het NICI al waren; het aantal dingen dat ik in je waardeer is in de loop der jaren alleen maar toegenomen, zoals hoe je 10 verschillende onderzoeksbollen tegelijk in de lucht

lijkt te houden en je productieve, inspirerende en zelfverzekerde manier van werken, maar ook je verfrissende blik op carrière maken, onze discussies over quarterlifecrises, en natuurlijk je oude keuken. Fijn dat je mijn paranimf wilt zijn!

Jasper, bedankt dat je de kافت hebt willen ontwerpen; het is een eer om een werk van jou op de voorkant te hebben staan.

Verder ben ik in de gelukkige omstandigheid dat ik genoeg vrienden om me heen heb die me er voor hebben behoed dat ik me zou verliezen in het onderzoek – voor het geval daar überhaupt al gevaar voor was – door me te laten zien wat werkelijk belangrijk is in deze wereld. Lieve vrienden, ik waardeer jullie vriendschap zeer. Van jullie wil ik Jan in het bijzonder bedanken. Jan, met niemand van mijn vrienden heb ik zo veel, zo gedetailleerd, zo kritisch en met zoveel plezier over wetenschap in het algemeen en ons beider onderzoek in het bijzonder gediscussieerd als met jou, naast alle andere zaken des levens natuurlijk die langskwamen in de Mark, Maxim of Science Café. Jouw gezonde doses scepsis en humor houden me scherp in het leven. Ik ben blij dat je er vandaag als mijn paranimf bij kunt zijn.

Pap en mam, bedankt voor het creëren van een omgeving waar ik me altijd thuis en gewaardeerd voelde en mijn interesse voor alles wat met wetenschap te maken heeft kon floreren – jullie hebben me altijd gestimuleerd om te doen wat ik het leukst en interessantst vond.

Tenslotte, Inge, met woorden kan ik niet omschrijven hoeveel je voor me betekent. Bedankt voor je ondersteuning, begrip, reflectie, onbaatzuchtigheid, en, vooral, liefde. Jij en Hanna zijn belangrijker in mijn leven dan welke wetenschappelijk vraag dan ook.

Curriculum vitae

Stan van Pelt werd op donderdag 31 augustus 1978 geboren in Zevenbergschen Hoek. Na een verhuizing naar de andere kant van Noord-Brabant doorliep hij zijn middelbare schoolperiode aan het St. Willibrord Gymnasium in Deurne. Tijdens zijn opleiding Biologie (1996-2002) aan de toen nog Katholieke Universiteit Nijmegen (KUN – de huidige Radboud Universiteit, RU) volgde hij een eerste onderzoeksstage bij de afdeling Biofysica & Medische Fysica, waar hij de karakteristieken van de menselijke vestibulo-oculaire reflex onderzocht, onze interne *image stabilizer*. Een tweede stage liep hij bij de afdeling Aquatische Oecologie & Milieubiologie, en het Marine Science Institute van de University of the Philippines in Manilla. Hier onderzocht hij de relatie tussen de geografische verspreiding en genetische diversiteit van *Enhalus acoroides* populaties, een tropische zeegrassoort. Na deze studie cum laude afgerond te hebben was hij nog een jaar werkzaam als junior onderzoeker binnen een samenwerkingsproject van Rijkswaterstaat en KUN naar de mogelijkheden voor herintroductie van zeegrassen in de Waddenzee, vooraleer hij koos voor een promotieonderzoek binnen de neurowetenschappen. Dit laatste werk voerde hij uit aan het Nijmeegs Instituut voor Cognitie en Informatie (NICI) binnen de RU, sinds 1 september 2008 als Centre for Cognition onderdeel van het Donders Institute for Brain, Cognition and Behaviour. De resultaten van dit onderzoek staan beschreven in dit proefschrift. Sinds juli 2008 is Stan als postdoc-onderzoeker verbonden aan het Centre for Cognitive Neuroimaging binnen ditzelfde Donders Institute, waar hij onderzoek doet naar de genetische basis van hoogfrequente hersengolven.

Series Donders Institute for Brain, Cognition and Behaviour

1. Van Aalderen-Smeets, SI (2007). *Neural dynamics of visual selection*. Maastricht University, Maastricht, The Netherlands.
2. Schoffelen, JM (2007). *Neuronal communication through coherence in the human motor system*. Radboud University Nijmegen, The Netherlands.
3. De Lange, FP (2008). *Neural mechanisms of motor imagery*. Radboud University Nijmegen, The Netherlands.
4. Grol, MJ (2008). *Parieto-frontal circuitry in visuomotor control*. University Utrecht, The Netherlands.
5. Bauer, M (2008). *Functional roles of rhythmic neuronal activity in the human visual and somatosensory system*. Radboud University Nijmegen, The Netherlands.
6. Mazaheri, A (2008). *The influence of ongoing oscillatory brain activity on evoked responses and behaviour*. Radboud University Nijmegen, The Netherlands.
7. Hooijmans, CR (2008). *Impact of nutritional lipids and vascular factors in Alzheimer's Disease*. Radboud University Nijmegen, The Netherlands.
8. Gaszner, B (2008). *Plastic responses to stress by the rodent Edinger-Westphal nucleus*. Radboud University Nijmegen, The Netherlands.
9. Willems, RM (2009). *Neural reflections of meaning in gesture, language, and actions*. Radboud University Nijmegen, The Netherlands.
10. Van Pelt, S (2009). *Dynamic neural representations of human visuomotor space*. Radboud University Nijmegen, The Netherlands.