Conditional visuomotor behavior in the Parietal Reach Region and Dorsal Premotor Cortex

Dissertation

for the award of the degree

"Doctor rerum naturalium"

Division of Mathematics and Natural Sciences

of the Georg-August-Universität Göttingen

submitted by

Christian Klaes

from Wickede

Göttingen 2010

1. Reviewer: Dr. Alexander Gail

Sensorimotor Group, German Primate Center

2. Reviewer: Prof. Dr. Stefan Treue

Cognitive Neuroscience Laboratory, German Primate Center

3. Reviewer: Prof. Dr. Fred Wolf

Department of Nonlinear Dynamics, Max Planck Institute for Dynamics and Self-Organization

Date of the oral examination: 1.12.2010

I hereby declare that this thesis has been written independently and with no other sources and aids than quoted.

_____

(Christian Klaes)

Acknowledgements

Science is a collaborative effort, which can only succeed through intensive discussions, sharing of ideas and team work. At this point I want to thank all the people that have supported me throughout this project. First and foremost I want to thank Alexander Gail for the opportunity to work in his laboratory and his continuous support throughout the years. I also want to thank Stefan Treue and Fred Wolf for their helpful suggestions and discussions.

Several people have contributed to the studies presented in this thesis and made it possible to successfully accomplish the experiments. First, I want to thank Sina Plümer for her invaluable support in handling the animals. Also very important to mention are all the other people whose daily routine it is to care for the wellbeing of the monkeys and for the smooth workflow in the lab: Dirk, Klaus, Leonore, Ludwig and Janine

On the administrative side I want to thank Beatrix Glaser, Ralf Brockhausen and Kevin Windolph for taking care of all these forms and a (mostly ;) working network. Very important for a successful dissertation are all the other PhD students and post docs in the lab that provide helpful comments, moral support and entertainment: Shubo, Shenbing, Steffi, Bahareh, Lalitta, Anja, Valeska, Vlad, Elena, Ioulia, Suresh, Bettina and Phillip.

But all this would not have succeeded without a good social network outside the lab. I want to thank my parents and my brother for their support and also all my friends for making live enjoyable. Special thanks to the Göttingen gaming crew: Billi, Axel, Björn, Gerd and Julian. And also to my long term friends in the Ruhrgebiet: Jan, Tugi, Sandra & Sandra, Gerd and Heiko.

Last but not least I want to thank Bella.

Contents

1 Introduction

The physical world is made up from particles, which form the objects we can interact with. To do so we most often use our hands for reaching and grasping them. Making voluntary goal-directed arm movements is a basic ability of all primates. This act seems to be very simple - almost trivial - to us, but it involves a complex translation of sensory information into muscle activation patterns. It is not only interesting how we plan to reach an object but also how we combine and represent information about our environment and our internal states to form a decision to do so.

The following work consists of three scientific articles and one submitted manuscript, which addresses the question how reach planning is accomplished by the brain. The model system that was used to answer this question was the behaving macaque monkey. To investigate the complex system of reach planning extracellular activity was recorded from neurons in two distinct cortical areas, the parietal reach region (PRR) and the dorsal premotor cortex (PMd), while the animals performed a behavioral task. A computational model was developed in accordance with findings from the electrophysiological recordings, which was then used to interpret the recorded data and to understand the underlying mechanisms that lead to the findings.

1.1 Reach planning in the brain

Two major factors are important for planning a reach movement. The first one is the sensory input, which provides information about the surrounding environment and the possible objects that can be reached for or that have to be avoided. If you are sitting in front of your desk, for example, there a usually several objects you could reach for, like the phone, the computer mouse or a paper. The second factor is the current behavioral context, which provides the incentives and constraints to act in a specific way within this environment. If you want to make a call you reach for the phone, but depending on your schedule you might also reach for the mouse instead because you want to work on your latest project. The same visual input can therefore lead to very different decisions and actions, depending on the behavioral state and internal goals.

The planning and execution of reach movements are distributed tasks, to which many areas in the brain contribute. Several areas which have been identified to play a major role in this tasks form a frontoparietal network for reaching (Fig. 1). This frontoparietal network is thought to transform sensory input into an output of muscle activity patterns needed to move the hand to the target position (Burnod et al., 1999; Wise et al., 1997). The network receives information that influences where to reach mainly from two sources. On the one hand spatial sensory input from visual areas enters the frontoparietal network via the posterior parietal cortex (PPC) (Blatt et al., 1990; Colby et al., 1988; Marconi et al., 2001). On the other hand top-down modulating input about behavioral constraints and incentives is believed to come from prefrontal and premotor areas (Toni et al., 2001; Wallis and Miller, 2003; Wise and Murray, 2000). Other non-visual signals, as for example eye and arm positional information, also influence the network (Burnod et al., 1999). These signals are important for visuomotor transformations (see below).

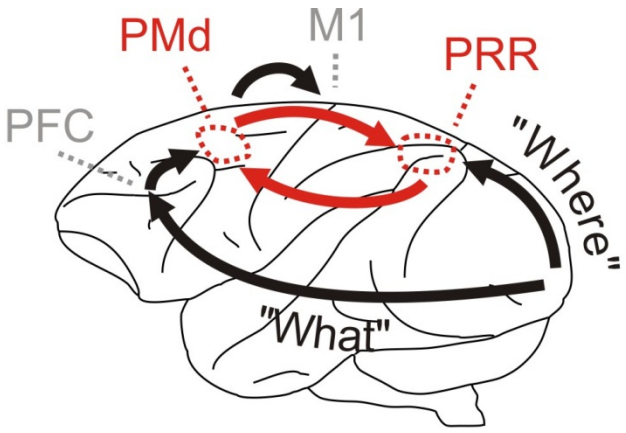

Figure 1: Schematic drawing of the frontoparietal network for reaching. Arrows indicate the presumed flow of information. The two areas in which electrophysiological recordings were taken are shown in red.

Parts of the frontoparietal network belong to the PPC, which is viewed as a sensory and motor association area, which transforms visual information into actions. It is considered to be a central node for the so called "where" pathway of visual information processing. The "where" pathway is supposed to be involved in the spatial relations of objects and their motion in visual scenes (Goodale and Milner, 1992; Ungerleider and Mishkin, 1982). It is seen complementary to the "what" pathway, which is supposed to deal with processing of orientations, shapes and color (Goodale and Milner, 1992; Gross, 1992; Logothetis et al., 1995).

The distinction between "where" and "what", which was defined originally from the sensory side, has been replaced by a more appropriate distinction of "action" and "object"

pathways (Goodale and Milner, 1992). There are some studies that even provide evidence that neurons in the PPC can also represent abstract task rules (Grol et al., 2006; Stoet and Snyder, 2004), which are usually associated with the frontal cortex (FC). Other parts of the frontoparietal network belong to the FC and prefrontal cortex (PFC), and presumably provide top-down influence. These structures are associated with executive control, learning and representation of abstract behavioral rules (Miller, 2000; Muhammad et al., 2006; Toni and Passingham, 1999; Wallis et al., 2001; White and Wise, 1999). The basis of our investigations were electrophysiological recordings in the PRR, which is part of the PPC, and the PMd, which is part of the FC. These two structures are important parts of the frontoparietal network for reaching and are described in more detail within the next two sections.

1.1.1 The parietal reach region (PRR)

PRR has been functionally defined as an area that shows reach related activity. It is especially known for its involvement in the planning of reach movements (Batista et al., 1999; Mountcastle et al., 1975; Snyder et al., 1997). Anatomically it includes parts of the medial intraparietal area (MIP), which belongs to Brodmann area 5, and area V6A (Snyder et al., 2000). It is located on the medial bank of the intraparietal sulcus (IPS) close to the posterior end of the superior parietal lobule (SPL). PRR receives input from visual extrastriate and somatosensory areas and projects to the PMd (Blatt et al., 1990; Johnson et al., 1996). Anatomically very close to PRR are the lateral intraparietal area (LIP) and the anterior intraparietal area (AIP), which seem to fulfill similar functions for saccadic eye movements and prehension (grasping), respectively. PRR also receives somatosensory input and it has been proposed to be involved in visuomotor transformations (see below).

1.1.2 The dorsal premotor cortex (PMd)

The area between the primary motor cortex (M1) and the PFC is known as the premotor cortex (PM). The region is further subdivided in a dorsal part (PMd) and a ventral part (PMv) (Kurata, 1991; Kurata, 1994; Preuss et al., 1996). It is known that PMd receives

most of its parietal input from the superior parietal lobe (SPL), including MIP (Kurata, 1991; Wise et al., 1997) and therefore also PRR. Most of these connections are reciprocal and therefore contribute to a sensorimotor loop. In addition to that PMd also receives input from PFC (Wise et al., 1997), which is supposed to mediate behaviorally acquired rules. PMd is typically associated with the representation and learning of abstract behavioral rules. For example it has been shown that the activity of PMd neurons changes during learning of a new stimulus-response association (Mitz et al., 1991). The caudal part of PMd projects to M1 and reciprocally connects the frontoparietal reach network to its output stage. Several studies have shown that PMd is involved in reach planning (Batista et al., 2007; Crammond and Kalaska, 1996; Pesaran et al., 2006; Weinrich and Wise, 1982) and, like in PRR, many neurons in PMd represent the upcoming reach direction (Crammond and Kalaska, 2000). A recent study suggested that not only definite goals are represented, but also potential goals as long as they remain valid possibilities (Cisek, 2006; Cisek and Kalaska, 2005).

1.2 Motor goal tuning

In the following studies we will often deal with properties of single neurons or populations of neurons. Therefore a short overview of some central terms is given here. Neurons are functionally characterized by measuring their firing rate in response to an externally controlled parameter. Neurons in sensory areas like V1, for example, have increased firing rates if a salient stimulus appears at a particular position on the retina. The so called 'receptive field' (RF) of a neuron is the region of space in which a particular stimulus elicits a change of firing rate of that neuron. Sometimes also the term response field is used, especially if the relationship depends on non-sensory parameters.

If the firing rate of the neuron is systematically changing when a parameter, like the orientation of a grating, is changed then typically it is said the neuron is tuned for this parameter. The functional dependency of the firing rate on the specific parameter then is expressed as the tuning function or also tuning profile of a neuron. In PRR and PMd the firing rates of neurons often depend on the direction of a planned movement. A particular direction which elicits the highest firing rate is called the preferred reach direction (PD) for that neuron. In the following neurons that have a systematic relationship between their

firing rate and the movement direction will be called motor goal tuned neurons. It is not trivial though to find out if neurons are truly motor goal tuned, as we will see in the next section.

1.3 The anti-reach paradigm

In a classical task to analyze reach related activity of neurons in the planning stage, a visual stimulus (cue) is briefly shown at some peripheral position on a screen. After a delay, the subject has to perform a reach toward that stimulus, starting from a common position (usually the center of the screen). We will call this a 'pro reach', because stimulus position and motor goal are equivalent (Fig. 2; top). In this case the neuronal response can be either attributed to the motor goal or the visual stimulus.

To disentangle this response, a so called 'anti reach' paradigm can been used (Crammond and Kalaska, 1994; Gail and Andersen, 2006; Zhang et al., 1997), which is based on a similar paradigm for fast ballistic eye movements (saccades) termed 'anti-saccade' task (Amador et al., 1998; Bell et al., 2000; Hallett, 1978). Similar to the 'pro reach' paradigm, a visual cue is shown in the periphery, but the reach that follows after the delay has to be directed towards the diametrically opposite direction of the cue (Fig. 2; bottom). If a specific neuron has the highest firing rate for a certain cue position, in pro- and anti-reach trials alike, its response can be attributed to the visual stimulus. If its firing rate is highest for a certain movement direction, it can be attributed to the motor goal.

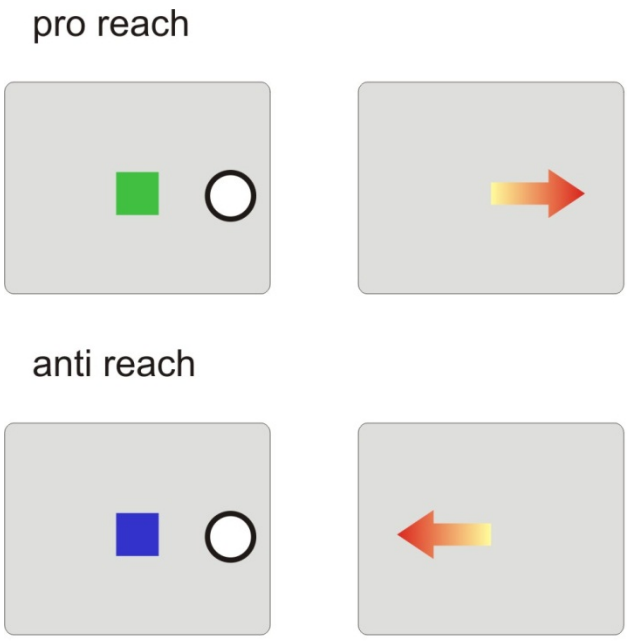

Figure 2: Schematic pro- and anti-reach task. A contextual cue indicates if it is a pro (top; green) or anti (bottom; blue) task. A spatial cue (white circle) indicates a position. After a delay a movement has to be made (right column), either towards the spatial cue (top; pro) or towards the diametrically opposite direction (bottom; anti).

1.4 Visuomotor transformations

One very important task that has to be solved by the frontoparietal reach network is to translate the position of objects from one coordinate system to another. For example, when a visible object is the goal of a reach movement, the movement goal has to be transformed from a retinal position to motor space, so that the reach is directed properly. Imagine again the situation of sitting in front of a desk. When you reach for the phone, for example, it does not matter if your gaze is directed towards it or towards the computer screen, to perform the task. This means the reach is planned in relation to the current arm position even though the object is projected on different parts of the retina. The necessary transformations that have to be performed depend on two factors. Firstly, on the effector that is to be used. In the example you could reach your target with the left or right arm. Secondly, on the sensory modality which is used to sense the environment, in the example it could be your eyes if the phone was in your visual field or your ears if you

want to reach for the ringing phone without looking at it. These so called 'reference frame transformations' may be accomplished in several stages, from an eye-centered coordinate frame to head-centered then to body-centered and finally to limb-centered. Planning of reach movements might be realized in any of these stages.

Several studies have provided evidence that reference frame transformations between eye and hand coordinates take place in PRR (Buneo et al., 2002; Chang et al., 2009; Cohen and Andersen, 2002). It is also assumed that many parietal areas use a common eye-centered reference frame, which suggests that the planning of movements might be accomplished in this reference frame (Andersen and Buneo, 2002). How the transformation between reference frames is computationally implemented is not fully understood, but one prominent mechanism that might mediate the transformation is gain modulation.

1.5 Gain modulation

To achieve a visuomotor transformation as described in the previous section a universal mechanism, called gain modulation, has been proposed (Zipser and Andersen, 1988). This mechanism has been shown to be widely used throughout the central nervous system (for a review see Salinas and Thier, 2000). Gain modulation can be described very generally as a nonlinear integration of information from different sources. Usually it refers to the change of the response function of a neuron, without altering its selectivity, by input from a different modality. This modulation can be multiplicative or otherwise nonlinear and changes the response amplitude in an ordered way according to a modulatory parameter. A neuron in PRR, for example, could have a preferred motor goal tuning for a reach movement to the left. The preferred direction will stay the same, independent of the

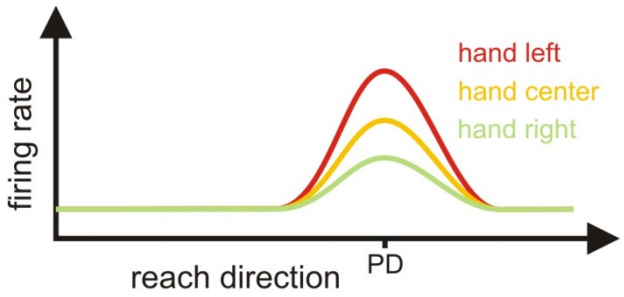

Figure 4: Schematic drawing showing the hypothetic tuning profile of a PRR neuron. The firing rate depends on the initial hand position, while the preferred reach direction (PD) remains unchanged, if the eye position stays constant.

initial position of the arm, but the tuning profile will be differently scaled, dependent on that initial position.

The first areas in which modulations of neural activity by eye position were described, were areas 7a and the lateral intraparietal area (LIP) (Andersen et al., 1990; Andersen and Mountcastle, 1983). LIP contains neurons that are active during the planning phase of a saccade. These neurons have response fields for different saccade directions in a retinocentric reference frame. A specific neuron, for example, could prefer saccades to the left side of the visual field and increases its firing rate if a saccade in that direction is planned. If the initial position of the gaze is now systematically changed by changing the fixation position, the firing rate of the neurons is modulated (either increased or decreased) in proportion to the change of eye position. Gain modulation was also found in many other cortical and subcortical structures (Batista et al., 1999; Boussaoud et al., 1998; Connor et al., 1996; Van Opstal et al., 1995). From a theoretical point of view it is interesting that neuronal network models could demonstrate that a gain modulation mechanism can realize reference frame transformations (Brozovic et al., 2007; Pouget and Sejnowski, 1997; Zipser and Andersen, 1988), which emphasizes its universal usage in the brain.

1.6 Decision making

How and when are motor goals represented in the premotor and posterior parietal cortex? In the previous sections, we discussed that motor goals are represented in these areas and that they are gain modulated. One open question is, if the motor goals are definite, meaning that they are the final decision, or if they represent an intermediate state of potential goals. At this point we have to go into more detail about the theory of decision making. Decision making is defined as the process of choosing appropriate actions from a set of possible alternatives. Examples include very basic animal behaviors, such as foraging, fight-or-flight decisions and more complicated human decisions, such as deciding how to dress or whom to vote. Recently the term "Neuroeconomics" has been coined to name a whole new field of research which combines economy, psychology and neuroscience in an effort to explain how and why we decide for one thing or another (Glimcher and Rustichini, 2004). In this work we will only discuss decisions of a more

limited scope for simplicity and scientific accessibility. The sensory input, on which many decisions are based, is often ambiguous and the outcome of the decision might be uncertain as well.

The decision to stop or go at a traffic light, for example, depends on the visual perception of a red or green signal, which might be disturbed by direct sunlight. The outcome of the decision can be uncertain as a consequence of the disturbed sensory information or because it depends on external factors. In the example, the outcome of our decision to stop or go also depends on the behavior of the other traffic participants and their decisions. This means many factors might influence our decision, including anticipating the behavior of others. To better analyze this complexity, the decision process is often split up into two parts. First, all possible options one can choose have to be evaluated. Second, among these options one is selected. To select a specific action it is necessary to assign values to all possible options and weighing them accordingly. For example you may value a 100 Euro note more than a sandwich under normal conditions, but when you are starving it might be the other way round. Sometimes the valuation might also be based on known or unknown probabilities, so that value of an option is only an estimate. After valuation a mechanism has to somehow select among the valuated options.

The processes of representation, valuation and selection have traditionally been thought to be computed separately and in a serial manner. In contrast to that several studies show that parietal and premotor areas represent several of these signals that are related to decision making at the same time (Cisek and Kalaska, 2005; Gold and Shadlen, 2000; Platt and Glimcher, 1999; Sugrue et al., 2004). This could mean that there is no distinct line that separates representation, valuation and selection. One recent study proposed the parallel processing of action specification (representation) and selection and presented a computational model of the decision making process for reach movements (Cisek, 2006). According to this model even while possible motor goals are still being specified, they already compete for selection in an integrated manner. It is also argued that a parallel system might be more plausible and efficient from an evolutionary point of view.

According to the Bayesian brain theory (Knill and Pouget, 2004), the brain computes and maintains probabilities, in the form of probability density functions, throughout the decision process, without the need to compute intermediate results. Only at the very last step, when the motor action is immediate, one of all possible outcomes is realized. In our real life analogy this would mean that as soon as you see the phone and the paper on the desk, within your brain a competition between the two actions of reaching for the phone and reaching for the paper starts. For this all behavioral incentives and constraints like the urgency of the phone call you want to make and your curiosity about the article in the paper would be taken into account and result in a probability density distribution that describes all these factors. In the end only one possibility survives and you eventually reach for the phone (Fig. 3).

Figure 3: Schematic drawing of the probability distribution of two valuated potential motor goals. In this example the right peak has the higher probability to be realized and the direction it indicates is more likely to become the direction of the reach.

An analogy can be drawn to another research field, which deals with systems that behave in a similar manner. Properties of particles on the quantum level can be described as probability density functions. A quantum system can be described as the superposition of all possible states that the system can take. The position of a particle in space, for example, is undetermined and can only be described in probabilistic terms. As long as no observer takes a measurement of one of the systems properties, all states exist in parallel (though not all have the same probability). Only when a measurement is taken, that is when the system is "asked to make a decision", one state becomes reality, while all the others vanish.

2 Original articles and manuscripts

This chapter contains the following articles and manuscripts:

1. Gail, A; Klaes, C; Westendorff, S (2009). Implementation of Spatial Transformation Rules for Goal-Directed Reaching via Gain Modulation in Monkey Parietal and Premotor Cortex. Journal of Neuroscience *29*, 9490-9499.
2. Westendorff, S; Klaes, C; Gail, A (2010). The Cortical Timeline for Deciding on Reach Motor Goals. Journal of Neuroscience *30*, 5426-5436.
3. Klaes, C; Westendorff, S; Chakrabarti, S; Gail, A (2011). Choosing goals, not rules: Deciding among rule-based action plans. Neuron *70*, 536-548
4. Klaes, C; Schneegans, S; Schöner, G; Gail, A. Sensorimotor learning biases choice behavior: A learning neural field model for decision making. Submitted

Author's contributions:

1. A.G., C.K. and S.W. designed the experiment, did the analysis and wrote the manuscript. C.K. and S.W. collected the data. All authors discussed the results and commented on the manuscript at all stages.
2. S.W., C.K. and A.G. designed the experiment. S.W. and C.K. collected the data. S.W. did the analysis and wrote the manuscript. A.G. edited the manuscript. All authors discussed the results and commented on the manuscript at all stages.
3. C.K., S.W., and A.G. designed the experiment. C.K., S.C. and S.W. collected the data. C.K. did the analysis and wrote the manuscript. A.G. edited the manuscript. All authors discussed the results and commented on the manuscript at all stages.
4. C.K., S.S. G.S. and A.G. designed the model. S.S. programmed the original model. C.K. did the analysis and wrote the manuscript. A.G. and S.S. edited the manuscript. C.K and S.S contributed equally to this work. All authors discussed the results and commented on the manuscript at all stages.

2.1 Implementation of Spatial Transformation Rules for Goal-Directed Reaching via Gain Modulation in Monkey Parietal and Premotor Cortex

Spatial transformations between different reference frames might be implemented via gain modulation, as described in the introduction (1.5). Are learned arbitrary transformation rules, like the anti-reach paradigm implemented in a similar manner? In this study we wanted to test if motor-goal representations in PRR and PMd during reach planning are gain modulated by the behavioral context. This could give us an indication on how a spatial cue is remapped onto an arbitrary motor goal. In a monkey experiment, we interspersed pro and anti reach trials to compare two different mapping rules, in which we could measure neuronal activity.

In PRR and in PMd we found that the behavioral context affects the tuning profile of individual neurons via gain modulation. On the one hand we could show that PRR has stronger directional selectivity during the planning of a reach toward a directly cued goal (pro reach) compared with an inferred target (anti reach). On the other hand PMd showed stronger overall activity during reaches toward inferred goals when compared with directly cued goals. These results support the idea that gain modulation is the computational mechanism that implements rule-based remappings of motor goals. The described different effects of the context in PRR and PMd fit well into the supposed information flow in the frontoparietal reach network. PRR is influenced more strongly by the sensory input which defines the possible reach goals. PMd seems to be influenced more by the top-down arbitrary rules, which presumably come from PFC.

Behavioral/Systems/Cognitive

Implementation of Spatial Transformation Rules for Goal-Directed Reaching via Gain Modulation in Monkey Parietal and Premotor Cortex

Alexander Gail, Christian Klaes, and Stephanie Westendorff
Bernstein Center for Computational Neuroscience, German Primate Center–Leibniz Institute for Primate Research, 37077 Göttingen, Germany

Planning goal-directed movements requires the combination of visuospatial with abstract contextual information. Our sensory environment constrains possible movements to a certain extent. However, contextual information guides proper choice of action in a given situation and allows flexible mapping of sensory instruction cues onto different motor actions. We used anti-reach tasks to test the hypothesis that spatial motor-goal representations in cortical sensorimotor areas are gain modulated by the behavioral context to achieve flexible remapping of spatial cue information onto arbitrary motor goals. We found that gain modulation of neuronal reach goal representations is commonly induced by the behavioral context in individual neurons of both, the parietal reach region (PRR) and the dorsal premotor cortex (PMd). In addition, PRR showed stronger directional selectivity during the planning of a reach toward a directly cued goal (pro-reach) compared with an inferred target (anti-reach). PMd, however, showed stronger overall activity during reaches toward inferred targets compared with directly cued targets. Based on our experimental evidence, we suggest that gain modulation is the computational mechanism underlying the integration of spatial and contextual information for flexible, rule-driven stimulus–response mapping, and thereby forms an important basis of goal-directed behavior. Complementary contextual effects in PRR versus PMd are consistent with the idea that posterior parietal cortex preferentially represents sensory-driven, "automatic" motor goals, whereas frontal sensorimotor areas are stronger engaged in the representation of rule-based, "inferred" motor goals.

Introduction

The same visual scene can lead to very different actions taken depending on the behavioral context. In a real boxing match, the face of your opponent is the direct goal of your arm movement. In a show fight, the face is not the goal of the movement but still guides it. Planning a reach toward a visual target object or inferring a reach goal from the position of an object according to a spatial transformation rule requires context-specific sensorimotor transformations (Gail and Andersen, 2006). Here, we test how the frontoparietal reach network flexibly creates context-specific motor-goal representations.

Spatially flexible cue–response mapping, as in a pro-reach/anti-reach task, requires integration of spatial sensory information with the context. Computational models suggest that this flexibility can be achieved with contextual gain modulation of spatially selective neurons (Salinas, 2004; Brozović et al., 2007), equivalent to gain field mechanisms suggested for multisensory integration (Andersen et al., 1985; Zipser and Andersen, 1988; Boussaoud et al., 1993, 1998; Brotchie et al., 1995; Galletti et al., 1995; Salinas and Abbott, 1996; Snyder et al., 1998; Salinas and Thier, 2000). If the idea of gain modulation for space–context integration holds true, then spatially selective neurons in primate cortical sensorimotor areas, like the parietal reach region (PRR) and the dorsal premotor cortex (PMd), should be upregulated and downregulated by the behavioral context.

Spatial sensory information presumably reaches the frontoparietal sensorimotor network via the posterior parietal cortex (Colby et al., 1988; Blatt et al., 1990; Marconi et al., 2001). Associative goal selection criteria or arbitrary transformation rules are believed to exert their influence on motor planning via prefrontal and premotor areas (Rushworth et al., 1997; Wise and Murray, 2000; Toni et al., 2001; Wallis and Miller, 2003). Both PMd and PRR have been shown to express sustained spatially selective activity during movement planning (Weinrich and Wise, 1982; Andersen et al., 1985; Boussaoud and Wise, 1993; Kalaska, 1996; Wise et al., 1997; Andersen and Buneo, 2002; Cisek, 2007), predominantly representing motor goals (di Pellegrino and Wise, 1993; Crammond and Kalaska, 1994; Gail and Andersen, 2006). We compare contextual modulations in PRR and PMd to test whether the two areas represent inferred versus directly cued motor goals differently. Posterior parietal cortex, for example, was hypothesized to mainly represent fast "automatic" motor goals (Desmurget et al., 1999; Pisella et al., 2000), as in pro-reaches.

We tested whether and how motor-goal representations in PRR and PMd are modulated by context-specific spatial transformation rules during reach planning. We used a memory-guided

Received March 5, 2009; revised June 13, 2009; accepted June 16, 2009.

This work was supported by the Federal Ministry for Education and Research (Germany) Grants 01GQ0433 and 01GQ0814. We thank Sina Plümer and Ludwig Ehrenreich for their extensive support in data collection, and Axel Lindner for helpful discussions on previous versions of this manuscript.

Correspondence should be addressed to Dr. Alexander Gail, Bernstein Center for Computational Neuroscience, German Primate Center–Leibniz Institute for Primate Research, Kellnerweg 4, 37077 Göttingen, Germany. E-mail: agail@gwdg.de.

DOI:10.1523/JNEUROSCI.1095-09.2009

 0270-6474/09/299490-10$15.00/0

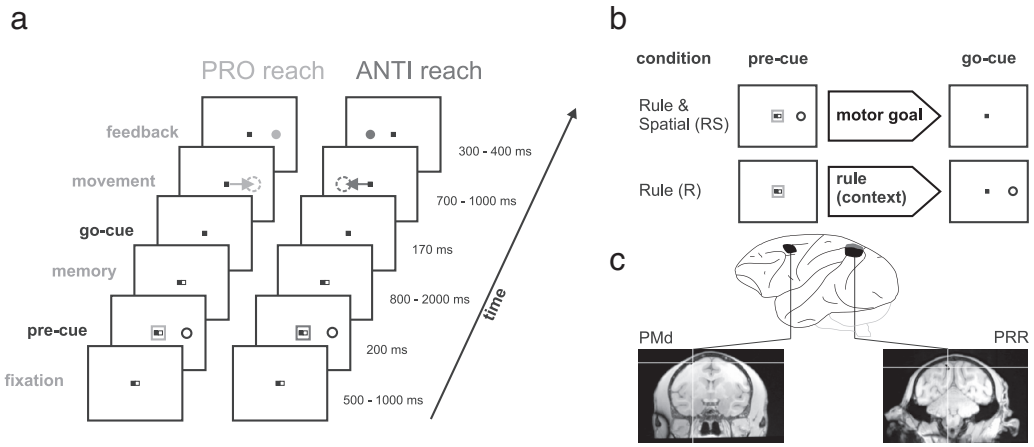

Figure 1. Memory-guided anti-reach task with precuing. ***a***, The left sequence shows an example of a pro-trial, and the right, an example of an anti-trial. Subjects had to maintain eye fixation (central dark spot) throughout the trial and hand fixation (central white spot) until the go instruction (disappearance of the white spot). The reach goal was defined by the combination of a context cue (colored central frame), indicating the pro/anti transformation rule, and a color-neutral spatial cue, which was presented at any of four different peripheral screen positions (0, 90, 180, 270°). Cues could be presented before (precue period) or after a variable memory period (go cue period), simultaneously or separately. In the example, the context and the spatial cue are both presented during the precue period. To be rewarded, the subject had to make a reach toward the previous spatial cue position (pro-reach) or to the opposite side (anti-reach). The arrows and dotted circles are for demonstration purpose only and were not shown to the subjects. ***b***, In this study, we compare conditions in which the subjects either had full information on the upcoming motor goal (transformation rule and spatial cue presented as precue; RS condition) or only the context information available (transformation rule presented as precue, spatial cue during go cue period; R condition). ***c***, Extracellular signals from multiple individual neurons were recorded simultaneously in PRR and PMd (regions of interest for monkey S) while the monkeys performed the task.

pro-reach/anti-reach task with partial precuing to investigate contextual modulations that were either independent or dependent of the spatial tuning. We tested separately the effect of context on either the response amplitude or the directional selectivity. Our results are consistent with the idea of space–context integration by a gain modulation mechanism. Additionally, we found complementary effects of context on motor-goal representations in PRR and PMd, which indicate different mutual roles of the two areas in context-specific visuomotor integration and motor-goal representation.

Materials and Methods

Tasks and control of behavioral parameters. Two male rhesus monkeys (*Macaca mulatta*; S and A) were trained to perform a memory-guided anti-reach task with partial precuing (Fig. 1*a*). Two visual cues, one spatial and one contextual, were presented to the subjects either simultaneously or at different points in time (19 inch ViewSonic LCD VX922; 5 ms off–on–off response time). The contextual cue (transformation rule) consisted of a green (pro-reach) or blue (anti-reach) frame around the eye and hand fixation points and instructed the subject to reach either toward (pro) or diametrically opposite to the spatial cue (anti). The peripheral spatial cue was located at one of four possible positions (0, 90, 180, and 270°) with an eccentricity of 9 cm [14.5° visual angle (VA)] relative to the fixation point. The motor goal was only defined uniquely, once both cues were known to the subject.

The monkey initiated a trial, by fixating a small red square in the center of the screen (eye fixation tolerance, 2.5–4.0° VA; 224 Hz CCD camera; ET-49B; Thomas Recording) and touching an adjacent white square of the same size (hand fixation tolerance: 4.0° VA, touch screen mounted directly in front of the video screen; IntelliTouch; ELO Systems). After a short period of 500–1000 ms (fixation period), the precue appeared for 200 ms. The precue could be either the contextual and the spatial cue at once [rule and spatial condition (RS)], only the spatial cue [spatial-only condition (S)], only the contextual cue [rule-only condition (R)], or nothing [null condition (N)], leading to four different memory conditions. After the precue, the monkey had to keep eye and hand fixation for 800–2000 ms (memory period). Depending on the memory condition, additional information was given to the monkey at the end of memory period. In the R, S, and N conditions, the missing information required to fully specify the reach goal (the spatial cue, the contextual cue, or both) was briefly flashed during the 170 ms go cue period. In the RS condition, no additional information was shown. Simultaneously, with the presentation of the second cue, the hand fixation square disappeared ("go" signal) and the monkey had to reach toward the instructed goal within a maximum of 700–1000 ms (movement period; 4.9° VA reach tolerance) and hold the goal position for 300–400 ms (feedback period). The monkey received visual feedback about the correct movement goal, which consisted of a filled circle of the same color as the contextual cue and was presented at the goal location during the feedback period. Eye fixation had to be kept throughout the trial; otherwise, the trial was aborted immediately. Liquid reward and acoustic feedback indicated correct (high pitch tone, reward) or incorrect (low pitch tone, no reward) behavior.

Pro-trials and anti-trials were randomly interleaved from trial to trial (~400 trials per data set) and the four memory conditions were randomly mixed in small blocks (10–20 trials of same condition per block). To test the contextual influence on motor planning, only the results of the RS and R conditions (transformation rule available during the planning phase of the movement) are meaningful and presented here (Fig. 1*b*).

Animal preparation and neural recordings. After training of the task, both monkeys were implanted with titanium head holders that were custom-fit to their heads based on computer-tomographical surface reconstruction of the skull (3di GmbH). After recovery and additional training, each monkey was implanted with two magnetic resonance imaging (MRI)-compatible recording chambers, also custom-fit to the skull. One chamber was placed above PRR [Horsley Clarke coordinates: 6 mm/−8.5 mm (monkey S/A) lateral; −10 mm/−9.5 mm anterior], the other chamber above PMd (13 mm/−13.5 mm lateral; 17 mm/19.8 mm anterior). Presurgical structural MRI was used to position the chambers (Fig. 1*c*). Postsurgical MRIs, showing the chamber relative to the brain, verified the correct chamber positions and allowed precise targeting of the desired anatomical structure. Sustained, direction-selective, neural responses during center-out reach planning was used as a physiological signature in both areas to confirm the imaging-based positioning. Both chambers were implanted contralaterally to the handedness of the monkey (A, left hemisphere; S, right hemisphere). All surgical and imaging procedures were conducted under general anesthesia.

For the extracellular recordings, up to four microelectrodes in each cortical area in a five-channel microdrive ("mini-matrix"; Thomas Recording) were used simultaneously. In most sessions, simultaneous recordings were conducted in both areas. The raw signals were preamplified (20×; Thomas Recording), bandpass filtered, and amplified (154

Hz to 8.8 kHz; 400–800×; Plexon), before on-line spike-sorting was conducted (Sort Client; Plexon). In addition to spike times, the spike waveforms were recorded and later subjected to off-line sorting for optimal isolation quality (Offline Sorter; Plexon).

Animal care and all experimental procedures were conducted in accordance with German law (Animal Welfare Act).

Selection of neural data. All recorded and sufficiently well isolated units, regardless of their tuning properties, were included in the neural data analyses unless explicitly stated otherwise. We analyzed the neural activity during the last 300 ms of the memory period (i.e., activity succeeding the precue with a time lag of at least 500 ms, and immediately preceding the go cue). This period was chosen to extract movement planning activity without confounding effects of (1) immediate visual input from the cue stimuli, or (2) transition phases from visual to motor-goal tuning (Gail and Andersen, 2006), or (3) visual and somatosensory input and motor-control signals related to movement initiation.

Analyses of neuronal directional selectivity. Directional selectivity was quantified with a directional tuning vector (DTV), which is defined as the vector average across all center-out cue directions $\vec{u}_i$ (unit vectors) weighted with the corresponding mean neural spike rates across trials with this cue direction (r_{ij}) and normalized to the total mean spike rate across all trials for the neuron *j* as follows:

$$\mathrm{DTV}_j = \sum_{i=1}^{4} r_{ij}\vec{u}_i / \sum_{i=1}^{4} r_{ij}.$$

The length of the resulting vector is between 0 and 1 and is a measure of the tuning strength. Its direction will be referred to as preferred direction (PD) of a neuron. The DTV was computed separately for pro-trials and anti-trials and was defined relative to the position of the spatial cue, not relative to the movement goal.

Significance of the directional selectivity was tested with a nonparametric one-way ANOVA (Kruskal–Wallis; Matlab; Mathworks) with the four different visual cue directions as groups and sample sizes defined by the number of identical trial repetitions. The ANOVA was calculated independently within each transformation rule (pro/anti) and precuing condition (total of $2 \times 4 = 8$ combinations).

The relative difference in PD between pro-trials and anti-trials in the RS condition indicates whether the tuning of a neuron in the given time window reflects the visual cue or the reach goal position. We quantified this difference with the tuning direction difference index (DD) as follows:

$$\mathrm{DD} = \mathrm{PD}_A - \mathrm{PD}_P.$$

PD_P and PD_A are the preferred directions of a neuron in pro-trials and anti-trials, respectively. A DD of ±0° indicates idealized visual tuning. A DD of ±180° means opposing PDs (relative to the cue position), which indicates idealized motor-goal tuning.

Analyses of contextual modulations. We tested three different effects of the context on the neuronal activity and quantified each of them with a contextual modulation index: gain modulation of spatially selective neurons (i.e., amplitude changes independent of changes in directional selectivity [gain modulation (GM)], changes in directional selectivity independent of changes in overall neural response strength [selectivity modulation (SM)], and direct modulation of neuronal activity levels independent of spatial tuning [direct modulation (DM)]).

Contextual gain modulation is characterized by amplitude changes of spatial tuning, independent of changes in directional selectivity. It is defined as follows:

$$\mathrm{GM} = \frac{r_P^{\mathrm{RS}} - r_A^{\mathrm{RS}}}{r_P^{\mathrm{RS}} + r_A^{\mathrm{RS}}},$$

where r_P^{RS} and r_A^{RS} are the mean spike rates for pro-trials and anti-trials in the RS condition. Since most neurons in PRR and PMd are spatially highly selective with a DD of around ±180° (see Results), large absolute GM values indicate strong gain modulation effects of the spatial motor-goal tuning by the context.

Contextual selectivity modulation is characterized by changes in directional selectivity independent of changes in response amplitude. To quantify contextual selectivity modulation, we computed the contrast of the normalized DTV length between pro-trials and anti-trials in the RS condition as follows:

$$\mathrm{SM} = \frac{|\mathrm{DTV}_P| - |\mathrm{DTV}_A|}{|\mathrm{DTV}_P| + |\mathrm{DTV}_A|},$$

where DTV_P and DTV_A are the normalized directional tuning vectors for pro-trials and anti-trials. A positive SM indicates that a neuron is more strongly tuned in pro-trials than in anti-trials, and vice versa for negative values. Values close to zero indicate that tuning selectivity is not modulated by the transformation rule.

Both GM and SM quantify modulations of spatial tuning by context. GM and SM are independent of each other in that each can change without affecting the other by varying parameters of hypothetical tuning functions (supplemental material S1, available at www.jneurosci.org). This does not mean that a single cell could not simultaneously have a high GM and SM. In fact, for neurons that are very strongly modulated by the context the activity in one of the two conditions (pro or anti) may be so weak that the tuning becomes arbitrarily shaped, which can lead to large GM and SM values and "random" DD values.

Direct context modulation can be quantified by differences in the memory activity between pro-reaches and anti-reaches in the R condition, since these differences reflect contextual influence independent of any spatial tuning. We defined the direct context modulation as the contrast of the average activities between pro-trials and anti-trials as follows:

$$\mathrm{DM} = \frac{r_P^R - r_A^R}{r_P^R + r_A^R}.$$

r_P^R and r_A^R are the mean spike rates in the R condition for pro-trials and anti-trials, respectively. A positive DM indicates neurons that are more active in pro-trials. A negative DM indicates stronger activity in anti-trials. Values close to zero indicate neurons indifferent to the transformation rule.

Note that all modulation indices are nonparametric and do not require fitting of any predefined tuning functions.

All contextual modulations were statistically analyzed on the level of individual neurons and, additionally, on the population level. We used bootstrap methods ($n = 100$ samples) to estimate the confidence limits of all indices (GM, SM, DM, and DD) for each individual neuron. The trial-by-trial spike rates were randomly permuted with repetitions, whereas the assignment of each trial to a certain direction, transformation rule, and memory condition was kept unchanged. GM, SM, and DM were considered significant if the 95% confidence limit did not overlap with zero. DD was considered to significantly deviate from motor-goal tuning if the circular confidence interval did not overlap with 180°.

On the population level, we estimated the to-be-expected distributions of all indices when assuming random data (shuffle test; $n = 100 *$ number of used neurons). We randomly permuted trial-by-trial spike rates across the two transformation rules, while keeping the cue direction and memory condition unchanged. This procedure eliminated any effect of the transformation rule. For the GM, SM, and DM, *t* tests were used to quantify deviations of the population means from zero. Additionally, a Bartlett test was used to test for deviations of the original distribution's variance from the normally distributed shuffled data. A positive Bartlett test indicates contextual modulations beyond what is expected because of random variations, even if the distribution of indices is centered at zero (i.e., even when there is no bias of the index toward either propreference or antipreference). To test whether a distribution of DD across the population of neurons deviates significantly from the predefined ±180° direction, we used circular statistics to compute the 95% confidence limits according to the following:

$$d = \arccos\left[\frac{\sqrt{\frac{2n(2R^2 - n\chi^2_{\alpha,1})}{4n - \chi^2_{\alpha,1}}}}{R}\right],$$

where $R = nr$. *n* is the number of samples, *r* is the mean resultant vector length of the angular direction data (here, DD indices), and χ is the

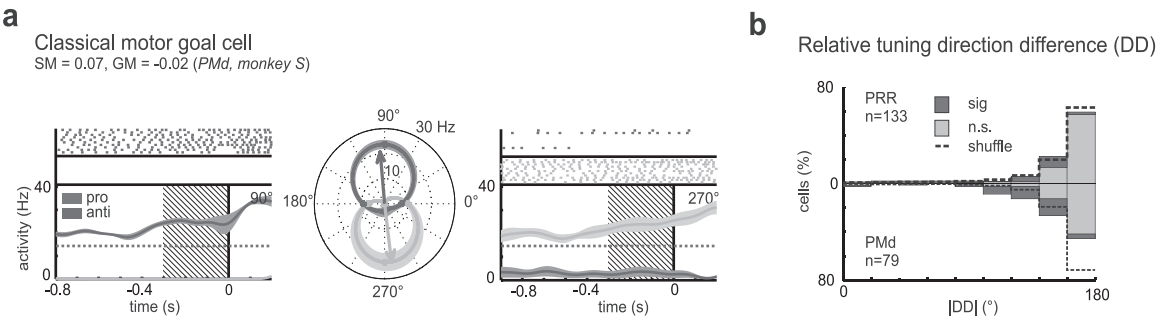

Figure 2. Motor-goal tuning in PRR and PMd during the late planning phase of reach movements. The example cell (***a***) shows classical motor-goal tuning, characterized by a PD depending on the motor goal and not the cue (DD not significantly different from 180°), as shown in the polar plot (middle). The response strength and selectivity were independent of the behavioral context (GM and SM not different from zero). Raster plots and spike density functions for pro (light gray) and anti (dark) trials are shown for the two most active directions (here, 90 and 270°) during the late memory period. Time 0 marks the onset of the go cue period. The mean firing rate during the fixation period is provided for comparison (dashed line). Spike density curves are smoothed with a Gaussian kernel ($\sigma = 50$ ms; dark lines, mean; light area, SE) for presentation purposes only. Also, polar tuning functions are interpolated for presentation purposes only, whereas all analyses are conducted nonparametrically on the original spike data. ***b***, Tuning DDs in PRR (top) and PMd (bottom) indicate that neurons in these areas are preferentially motor-goal tuned. Most individual neurons do not deviate significantly from 180° (light gray bars); a few do according to their bootstrap confidence limits (dark bars). The circular distributions of the DD for all neurons in each area do not significantly deviate from 180°. Note that because of symmetry the distribution of absolute values (|DD|) are plotted, whereas circular statistics are applied to the original DD values.

inverse of the χ^2 cumulative distribution function with a confidence limit of $1 - \alpha$ (Zar, 1999).

To quantify the correlation between the signs and strengths of different forms of contextual modulation in individual neurons, we computed Pearson's cross-correlation coefficients of the different indices across the population of neurons. To make the modulation indices scale-invariant, we normalized the modulation indices to the variances of their respective distributions as follows:

$$z_i = \frac{x_i}{\sigma_i}.$$

Here, i is the considered index (GM, SM, DM), x is the value of the index, and σ is the variance of the index distribution over the cell population. Note that this normalization is only for easier comparison of the modulation strengths in the scatter plots of Figure 7. The cross-correlation coefficient is independent of this linear scaling, which is applied within each index distribution. To test whether different types of modulation tend to occur in complementary neuronal subpopulations, regardless of the "direction" of modulation (i.e., propreference vs antipreference), we also compared the absolute values of the different modulation indices. We used Spearman's rank correlation to account for the non-normal distribution of absolute modulation indices.

Results

Both monkeys performed the task with high accuracy. The percentage of overall correct trials was 78% (monkey S) and 86% (monkey A) in the RS condition, and 75% (S) and 83% (A) in the R condition. Error trials were mainly attributable to early trial abortion (mostly breaks in ocular fixation) but not attributable to a confusion of the "pro" and "anti" rules. The percentage of correct pro-reach and anti-reach decisions in otherwise-correct trials was 98% (S) and 99% (A) in the RS condition, and 94% (S) and 99% (A) in the R condition. The performance of both monkeys showed no significant difference between pro-trials and anti-trials in the RS condition ($p > 0.05$, paired t test), and a minimal difference in the R condition (S: $m_{pro} = 95\%$, $m_{anti} = 94\%$, $p = 0.008$; A: $m_{pro} = 99\%$, $m_{anti} = 98\%$, $p = 0.003$, paired t test).

A total of 258 neurons from PRR (monkey S, 99; A, 159) and 193 from PMd (S, 75; A, 118) were recorded. Regardless of their spatial tuning properties, we tested all recorded neurons for direct context modulation (i.e., modulations of neuronal response amplitude by the context independent of directional selectivity) during the late memory period in the R condition (DM) (see Materials and Methods). For analyzing the effect of context on directional tuning properties (GM and SM) (see Materials and Methods and supplemental material S1, available at www.jneurosci.org, for details on modulation indices), we used a subset of neurons that had significant directional selectivity in at least one of the two spatial transformation conditions (pro/anti) during the late memory period of the RS condition. A total of 205 (79%) neurons in PRR [S, 71 (71%); A, 134 (84%)] and 132 (68%) neurons in PMd [S, 56 (75%); A, 76 (64%)] fulfilled this criterion. Data from both monkeys showed the same results and will be presented jointly (for comparison of the two monkeys, see supplemental material S2, available at www.jneurosci.org).

Contextual gain modulations of motor-goal tuning

When the reach goal was known to the monkeys during the memory period (RS condition), most of the spatially selective neurons in PRR (82.0%) and PMd (72.2%) were tuned for the motor goal, as was expected from previous results (Crammond and Kalaska, 1994; Gail and Andersen, 2006). Some individual neurons in the current study showed "classical" motor-goal tuning, which was independent of the spatial transformation rule that led to this motor goal (i.e., independent of the context). In the example neuron (Fig. 2*a*), the preferred directions for procondition and anticondition were opposite to each other when measured relative to the position of the cue, which means they were aligned with the motor goal (tuning direction difference DD not different from 180°; $p > 0.05$; bootstrap test) (see Materials and Methods), and their response strength as well as the directional selectivity did not differ between the pro and the anti condition (GM and SM not different from 0; $p > 0.05$; bootstrap test) (see Materials and Methods and below). The population of neurons in both areas on average represented the direction of the motor goal and not the cue position. In other words, the distribution of DD across the population of neurons did not deviate from 180° neither in PRR (circular mean, 179 ± 8° SEM) (see Materials and Methods) nor in PMd (mean, 184 ± 10°) (Fig. 2*b*). To be included in this analysis, the cells had to be significantly tuned in both spatial transformation conditions (pro/anti). This criterion was met by 133 (52%) of all neurons in PRR [S, 50 (51%); A, 83 (52%)] and 79 (41%) of all neurons in PMd [S, 35 (47%); A, 44 (37%)].

To test for contextual effects on spatial tuning properties, we distinguished between contextual GM and contextual SM (see

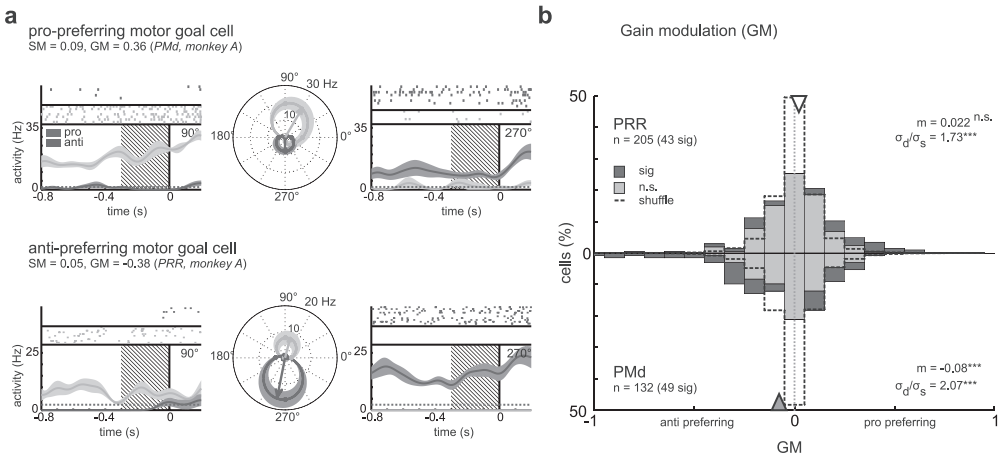

Figure 3. Contextual gain modulation of motor-goal tuning in PRR and PMd. ***a***, Examples of neurons with a high absolute GM indicating stronger activity for either pro-reaches (top) or anti-reaches (bottom), with invariant motor-goal tuning preference (DD not significantly different from 180°) and selectivity (SM not significantly different from zero; conventions are as in Fig. 2*a*). The anti-preferring neuron (bottom) was also one that fulfilled the strict criteria for an ideal contextual gain modulation of motor-goal tuning: the neuron was significantly tuned in pro and anti trials, had a significant GM, the DD was not significantly different from 180°, and the SM not significantly different from zero. (Note that the tuning vectors shown in the polar plot are not normalized, resulting in unequal length for pro and anti trials despite equally strong tuning; the SM is computed with normalized tuning vectors.) ***b***, Distribution of GM across the neuronal populations (conventions are as in Fig. 2*b*). In PMd (bottom), there was a significant bias toward negative values (mean, −0.08; gray triangle; *t* test) indicating on average stronger antipreference in PMd, but not in PRR (top; white triangle). The GM distributions in both areas were broader than for the shuffled data (dashed line; $\sigma_d > \sigma_s$; Bartlett's test) indicating significant modulation effects in both areas. Significance levels were as follows: $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Materials and Methods and supplemental material S1, available at www.jneurosci.org). Figure 3*a* shows two example neurons that were gain modulated by the behavioral context. The first neuron (top) was highly active in pro-trials, but responded only weakly in anti-trials. The preferred direction, representing the motor goal, remained unchanged (DD not significantly different from 180°), as did the directional selectivity (SM not different from zero). The second neuron (bottom row) showed the opposite preference and responded strongly in anti-trials, but only weakly in pro-trials. Again the spatial tuning encoded the motor goal in both rule conditions in a similar manner. Twenty-one percent (43 of 205) of neurons in PRR and 37% (49 of 132) in PMd had a significant GM when tested in individual neurons (Fig. 3*b*). There was a significant difference between the numbers of significant pro-preferring and anti-preferring neurons in PRR (31 pro, 12 anti; $p = 3.7 \times 10^{-3}$; χ^2 test) but not in PMd (28:21; $p > 0.05$) (for the number ratios, see Fig. 6*c*). On the population level, in both areas the GM values were more broadly distributed (i.e., on average, had higher absolute values) than the GM of the shuffled data (PRR, $p < 10^{-10}$; PMd, $p < 10^{-10}$; Bartlett's test), indicating contextual gain modulation (Fig. 3*b*) (see also Fig. 6*a*). In PMd, the GM distribution additionally had a significant bias toward anti-preferring neurons ($m = -0.08$; $p = 8 \times 10^{-4}$; *t* test) (i.e., neurons in PMd were on average more active during the planning of anti-reaches). In PRR, the GM distribution was centered on zero ($m = 0.021$; $p > 0.05$; *t* test) (see also Fig. 6*b*).

Contextual modulations of directional selectivity

In addition to a modulatory effect on the gain of the motor-goal tuning, context also modulated the directional selectivity of many neurons in our experiment (Fig. 4). In the first example (Fig. 4*a*, top), the tuning of the neuron, which was highly selective for pro-reaches to the left, became bimodal in the case of anti-reaches (i.e., was active for leftward and rightward anti-reaches). The second example (Fig. 4*a*, bottom) shows another neuron that was spatially tuned in pro-trials but not tuned in anti-trials. Twenty-six percent (54 of 205) of neurons in PRR and 25% (33 of 132) in PMd had a significant SM (Fig. 4*b*). There was a significant difference between the numbers of significant pro-preferring and anti-preferring neurons in PRR (45 pro, 9 anti; $p = 3.2 \times 10^{-7}$; χ^2 test) but not in PMd (18:15; $p > 0.05$) (see also Fig. 6*c*). Correspondingly, the population distribution of SM (Fig. 4*b*) in PRR shows a shift toward positive values ($m = 0.11$; $p = 2.0 \times 10^{-7}$; *t* test), indicating a bias toward greater directional selectivity in pro-trials than in anti-trials (see also Fig. 6*b*). In PMd, there was no shift of the population mean toward pro-preferring or anti-preferring selectivity ($m = 0.02$; $p > 0.05$; *t* test). SMs were more widely distributed in the original data than in the shuffled data in both areas (Bartlett's test: PRR, $p = 2.4 \times 10^{-4}$; PMd, $p = 4.7 \times 10^{-6}$), indicating significantly stronger modulation of tuning selectivity than expected by chance (see also Fig. 6*a*). The SM values in PMd had a trend to be bimodally distributed, indicating simultaneous presence of neurons with rather strong modulation of directional selectivity preferring either pro-reaches or anti-reaches.

Direct context modulation

Modulation of motor-goal tuning in PMd and PRR requires that the contextual information about the currently valid transformation rule is accessible to these cortical regions. How strongly is this transformation rule represented in PRR and PMd neurons at a time when only the rule (pro/anti) and not the complete motor goal is known to the monkey (i.e., independent of any spatial representations)? The direct context modulation index (DM) compares the level of neuronal activity during the memory period in the R condition, when contextual but not spatial information is available. Sixteen percent (42 of 258) of neurons in PRR and 25% (49 of 193) in PMd had a significant DM (Fig. 5). There was no significant difference between the numbers of significant pro-preferring and anti-preferring neurons in either PRR (24 pro, 18 anti; $p > 0.05$; χ^2 test) or PMd (23:26; $p > 0.05$) (Fig. 6*c*). The balance between propreference and antipreference was also reflected in the fact that the distributions of DM values were centered at zero in PRR (mean $m = 0.011$; $p > 0.05$; *t* test) and PMd ($m = 0.013$; $p > 0.05$) (Fig. 6*b*). The variances of the DM distributions were significantly larger than the variances of the

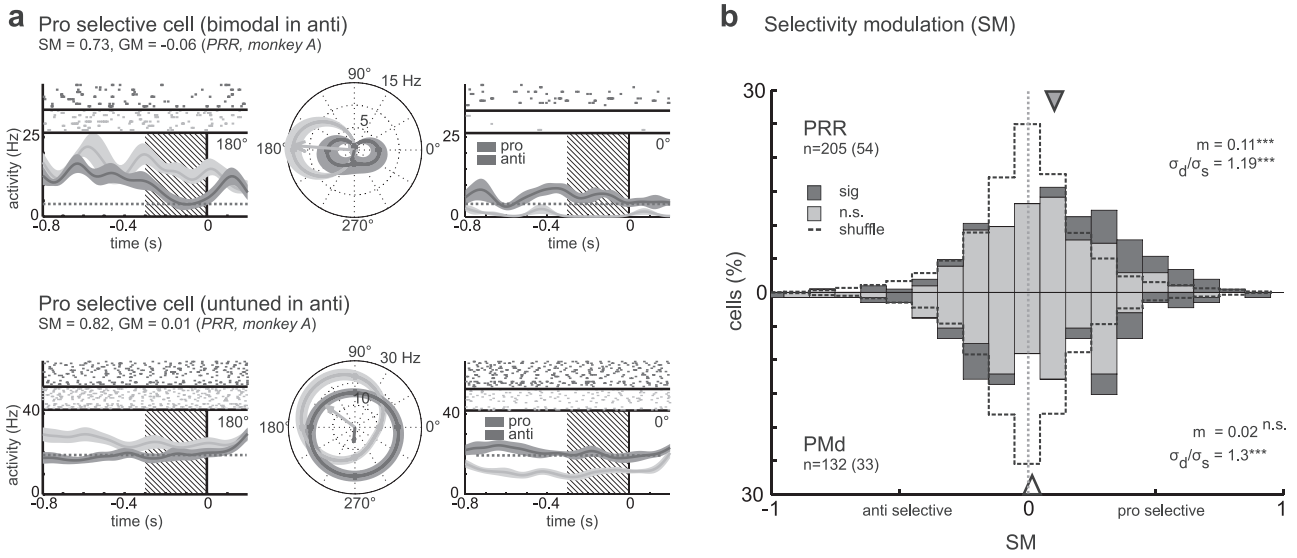

Figure 4. Contextual selectivity modulation of motor-goal tuning in PRR and PMd. ***a***, Examples of cells with a high absolute SM, indicating stronger directional selectivity for either pro-reaches or anti-reaches, but having the same response strength for both conditions on average across all reach directions (GM not significantly different from zero; conventions are as in Fig. 2*a*). The modulation resulted in different forms of tuning [e.g., bimodal (top) or nonselective (bottom) tuning functions during anti-reach planning]. ***b***, Distribution of SM across the neuronal populations (conventions are as in Fig. 2*b*). There was a significant bias for positive (pro-preferring) values in PRR (top), but not in PMd (bottom), indicating on average stronger directional selectivity during pro-reach planning in PRR, like the examples shown in ***a***. The SM distributions in both areas were broader than those of the shuffled data (dashed line; $\sigma_d > \sigma_s$; Bartlett's test) indicating significant modulation effects in both areas (conventions are as in Fig. 3).

shuffled DM data in both areas (PRR, $p = 9.1 \times 10^{-4}$; PMd, $p = 1.1 \times 10^{-7}$; Bartlett's test) (Fig. 6*a*).

Complementary contextual modulations in PRR and PMd

Figure 6 summarizes the contextual modulation effects on the gain of tuned neural responses (GM) and their directional selectivity (SM), together with the direct neural responses to the context information itself (DM). When compared with the shuffle prediction, contextual gain modulation showed the strongest effect, both in PRR and PMd (Fig. 6*a*). Yet directional selectivity modulation and direct contextual modulation indices were also significantly larger than expected by chance in both areas. The bias of the contextual modulation effects in terms of a propreference/antipreference were complementary in the two areas (Fig. 6*b*): Contextual gain modulations in PMd were biased toward antipreference [i.e., tuned responses were on average 15% (GM = −0.08) higher during anti-reaches compared with pro-reaches, whereas in PRR gain modulations on average were balanced]. In contrast, directional selectivity modulations in PRR were biased toward propreference; i.e., spatial tuning was on average stronger in pro-reaches compared with anti-reaches by 20% (SM = 0.11), whereas in PMd, selectivity modulations on average were balanced. In PRR, the fraction of individual neurons being biased toward higher response amplitudes or stronger directional selectivity was larger in pro-reaches compared with anti-reaches, whereas in PMd the number of individually modulated neurons was balanced (Fig. 6*c*).

Different types of modulation in different neuronal populations?

The contextual gain and selectivity modulation indices (GM and SM) per se are independent of each other (supplemental material S1, available at www.jneurosci.org). But one could imagine that, in individual neurons, the different modulation effects are correlated. This would be the case if two types of modulation are the phenomenological consequences of one and the same underlying computational mechanism, or if the underlying processes causing the different types of modulation interfere with each other. There was no correlation in either PRR or in PMd between GM and DM ($r_{PRR} = 0.1$, $p_{PRR} > 0.05$; $r_{PMd} = 0.03$, $p_{PMd} > 0.05$), SM and DM ($r_{PRR} = 0.13$, $p_{PRR} > 0.05$; $r_{PMd} = 0.07$, $p_{PMd} > 0.05$), or SM and GM ($r_{PRR} = 0.06$, $p_{PRR} > 0.05$; $r_{PMd} = 0.03$, $p_{PMd} > 0.05$) (Fig. 7*a–c*).

Two reasons could account for a lack of correlation between GM, SM, and DM. Either there was no interaction between the underlying processes that caused the gain and the selectivity modulations, but they still occurred within the same neurons; or these processes affected complementary neuronal subpopulations. If different neuronal subpopulations were affected by contextual gain and selectivity modulations, then high absolute values in GM and DM should be paired with low absolute values of the SM, and vice versa. In PRR, absolute values of SM were negatively correlated with absolute values of both GM ($r = -0.3$; $p = 0.005$; Spearman's rank correlation) and DM ($r = -0.32$; $p = 0.003$) (Fig. 7*d–f*). In PMd, only SM and DM were negatively correlated ($r = -0.36$; $p = 0.006$).

High absolute GM values do not necessarily mean that the neurons have to be gain modulated in an "ideal" way (i.e., in a multiplicative manner, as depicted in the example of Fig. 3*a*). Neurons with substantial GM might also have a high absolute SM value, which means, in principle, they might not have the typical unimodal directional tuning. In fact, a high absolute GM indicates a weak response of the neuron in either pro-reaches or anti-reaches. This can lead to weak or undefined directional tuning in the nonresponsive condition, and hence to high SM values. The lack of correlation between GM and SM, together with the anticorrelation of absolute GM and SM values in PRR, indicates that such interdependencies of GM and SM were not very common. To identify those cells that fulfilled all criteria for an ideal contextual gain modulation, we used an additional strict set of constraints that is compatible with the idea of gain modulation proper. We found that 48% (44 of 92) of cells that had a signifi-

cant GM were (1) tuned in both context conditions, but did not have a significant SM and did not show a change in preferred direction, or (2) were only tuned in the corresponding context condition (i.e., in the pro condition if they had a positive GM, or in the anti conditions if they had a negative GM) (PRR, 26 of 43, 60%; PMd, 18 of 49, 37%).

Discussion

We found strong modulation effects of behavioral context on the predominant motor-goal tuning in PRR and PMd. We propose gain modulation as a mechanism to achieve flexible goal-directed visuomotor remapping in a context-specific manner. Second, we found different types of modulation with complementary biases in PRR and PMd. PRR showed stronger directional selectivity during the planning of reaches toward directly cued (pro) compared with inferred (anti) goals. PMd, however, showed stronger overall activity during reaches toward inferred compared with directly cued goals.

Contextual modulations in PRR and PMd

Behavioral context affected neural activity in PRR and PMd in two major ways, either by directly driving neurons (direct context modulation, DM) or by modulating spatial motor-goal representations (contextual gain/selectivity modulations, GM/SM). Both types of context modulation can be predicted from the hidden-layer properties of a neural network model (Brozović et al., 2007). Gain-field modulation evolved in this model as a consequence of learning context-specific spatial cue–response mapping, equivalent to our anti-reach task. Gain modulation of spatially selective neurons is the key principle underlying the spatial transformation mechanism in this model, similar to previous models of multisensory integration for spatial reference frame transformations (Zipser and Andersen, 1988; Salinas and Abbott, 1996). Gain modulation during multisensory integration for eye or hand movements had previously been found in the posterior parietal cortex (Andersen et al., 1985; Brotchie et al., 1995; Galletti et al., 1995; Snyder et al., 1998; Batista et al., 1999; Nakamura et al., 1999; Buneo et al., 2002) and frontal areas (Boussaoud et al., 1993, 1998; Mushiake et al., 1997; Cisek and Kalaska, 2002). In contrast, here we have shown gain modulation effects in PRR and PMd for remapping visuospatial information onto reach motor goals according to abstract cognitive transformation rules.

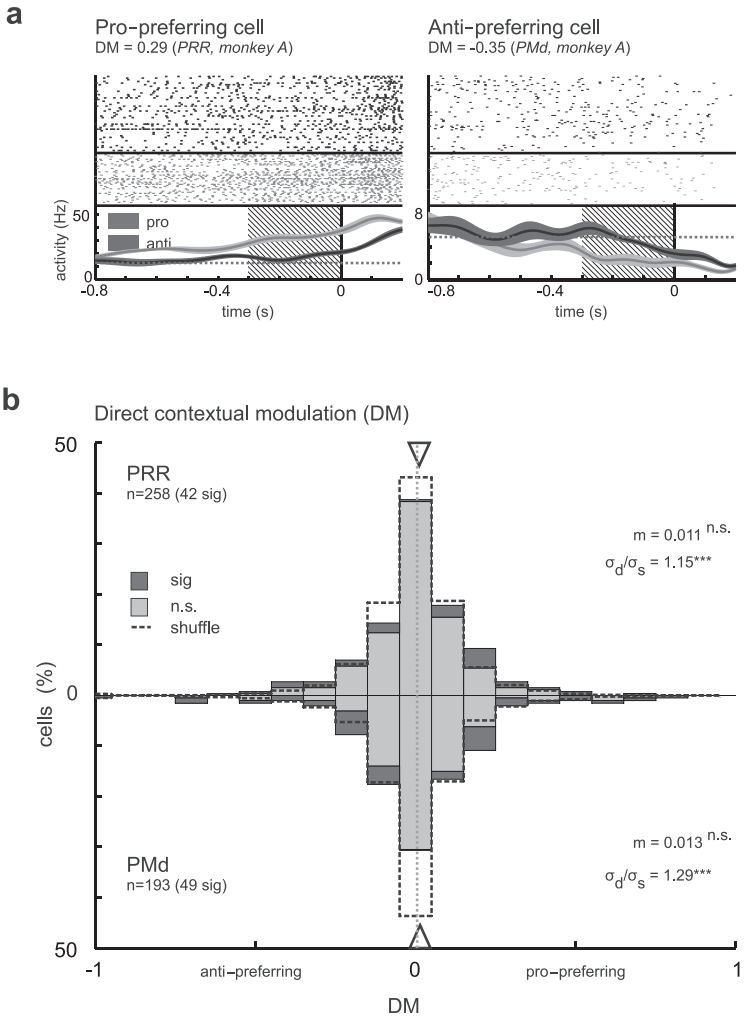

Figure 5. Direct context modulation in PRR and PMd independent of spatial tuning. ***a***, Examples show a pro-preferring cell (left), which was continuously more active in the memory period of pro-trials than during anti-trials (R condition). In the anti-preferring cell (right), the contextual modulation appeared later (i.e., only ~500 ms before the go cue). ***b***, Distribution of DM across the neuron populations. In PRR (top) as in PMd (bottom), the distributions were unbiased (i.e., centered at zero). The distribution of DM in both areas was broader than that of the shuffled data (dashed line; $\sigma_d > \sigma_s$; Bartlett's test) indicating contextual modulation effects (all conventions are as in Fig. 3).

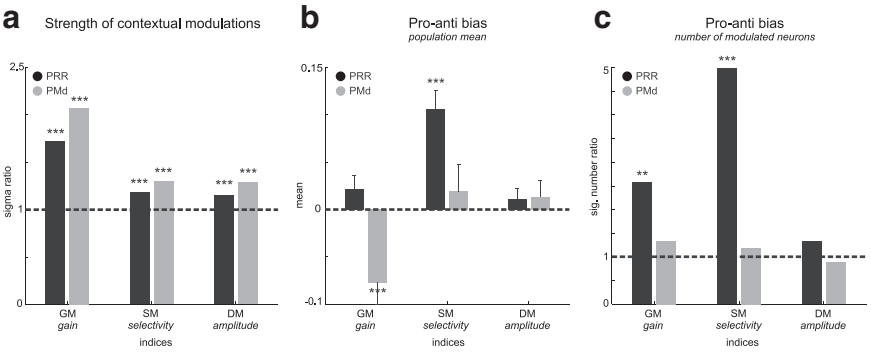

Figure 6. Summary of the strength and bias of contextual modulations in PRR (black) and PMd (gray). ***a***, Strength of contextual modulations. Both areas showed contextual modulations of all three types (GM, SM, and DM; Bartlett's test; significance levels are as in Fig. 3), with GM being the strongest. The σ ratio indicates the width of the distribution of modulation indices relative to the width of the distribution of the shuffled data (see Materials and Methods) (Figs. 3–5). ***b***, Pro/anti bias in the population mean. PRR and PMd showed complementary pro/anti biases in the mean contextual gain and selectivity modulation (GM and SM; mean ± SEM; *t* test). ***c***, Pro/anti bias in the number of modulated neurons. The relative number of neurons with a propreference versus an antipreference was higher in PRR for GM and SM, but balanced otherwise (χ^2 test).

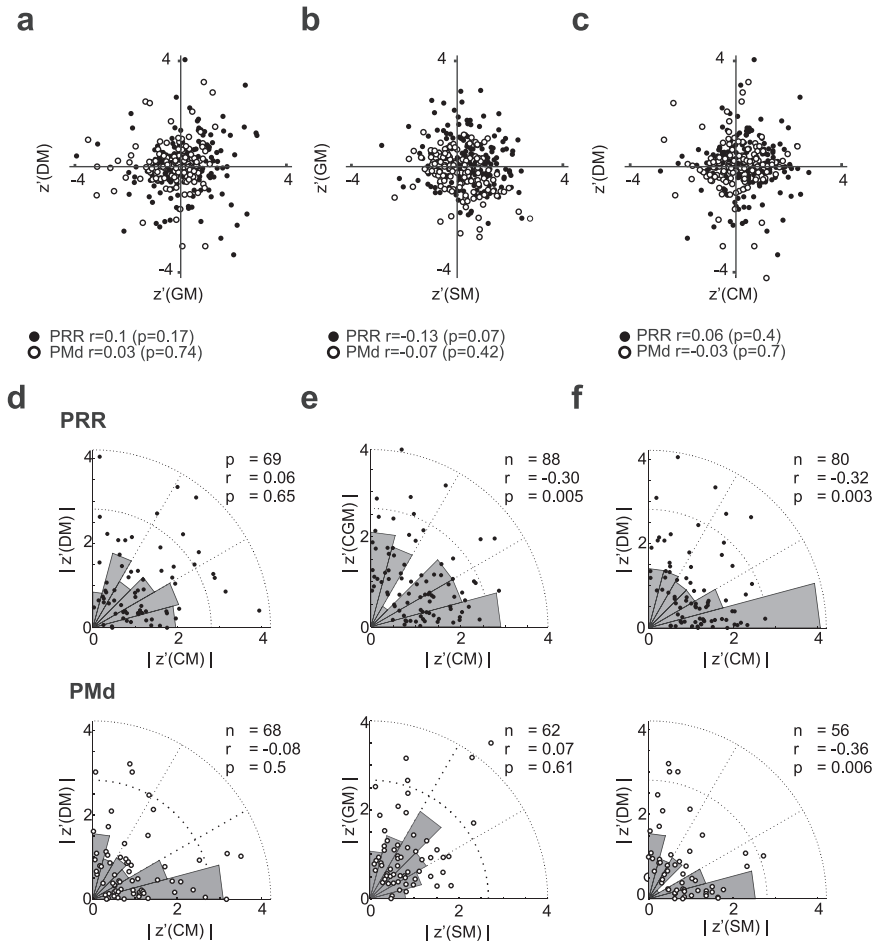

Figure 7. Interdependence of contextual modulations in PRR and PMd. ***a–c***, Pairwise cross-correlations between the three modulation indices. The Pearson correlation of the three modulation indices (*r* and *p* values given at the bottom of the diagrams) across neurons shows that none of the indices in either PRR or PMd were significantly correlated. ***d–f***, Pairwise cross-correlation of the absolute values of the modulation indices for PRR (top) and PMd (bottom) separately. In PRR, SM–DM (***f***; top) and SM–GM (***e***; top) were negatively correlated [i.e., strong GM and DM modulations implied weak SM modulations, and vice versa (Spearman's rank correlation)]. In PMd, only SM and DM showed negative correlations (***f***). The underlying rose plots (gray circular histograms) illustrate the frequencies of relative modulation strengths.

Gain modulation of neurons with motor-goal tuning was partly accompanied by modulation of the directional selectivity. This is consistent with the idea that a contextual gain modulation network builds the computational basis for the remapping process. Gain modulation then does not imply multiplicative modulations in the strict mathematical sense, but rather a summation of bell-shaped tuning functions (directional tuning) with a gain function (contextual tuning), which is passed through a recurrent network with nonlinear transfer properties (Salinas and Abbott, 1996; Brozović et al., 2007, 2008). Such a computational architecture results in gain modulations, which may not only affect the amplitude but also the shape of the tuning functions. This can be observed to varying extend for neurons within the same sensorimotor layer (Brozović et al., 2007).

Direct context modulation, as observed during the R condition, is also very plausible in the light of a gain modulation network. It is predicted for the sensorimotor ("hidden") layer of the network model, when only contextual information but no spatial information is available (Brozović et al., 2007), and all inputs are a priori additive. The impact of the context input in directly driving the neurons might be relatively weak, and only unfold to its full modulatory potential when the neuron is additionally driven by strong spatial input (Abbott and Chance, 2005). Because of their latency (Fig. 5, right example), we do not attribute the direct context modulations to the sensory properties of the context cue, but rather to the behavioral relevance of the cue (Boussaoud and Wise, 1993; Snyder et al., 1997; Toth and Assad, 2002; Stoet and Snyder, 2004; Nakayama et al., 2008).

Neurons in PMd had previously been reported to be modulated by context, since task rules had an effect on spatial selectivity (di Pellegrino and Wise, 1993; Crammond and Kalaska, 1994). Spatially selective responses were shown to be dependent on whether the motor goal was at the same position as the visual cue or at a default location (di Pellegrino and Wise, 1993). In contrast to our data, the strong condition dependency in the study of di Pellegrino and Wise is attributable to the fact that, in the default condition, the visual cue did not contain any relevant spatial information. Hence the modulation did not reflect an effect of contextual-specific spatial remapping, but rather an effect of motor planning based on spatial remapping versus nonspatial stimulus–response associations (Crammond and Kalaska, 2000). In contrast, in our experiment, the visual cue location was always relevant for the reach direction and the context modulations are the consequence of two different spatial transformation rules. Crammond and Kalaska reported a substantial fraction of cells in PMd with a significant interaction effect of a remapping rule on the reach-related response during the late phase of an instructed delay period (Crammond and Kalaska, 1994). Their finding could reflect a similar gain mechanism as observed here. However, they could not analyze specific effects of context on gain or selectivity of directional tuning because of their experimental protocol. Also, both previous studies could not compare the effects in PMd with PRR, in which gain modulation effects on directional tuning by spatial transformation rules, to our knowledge, have not been shown before.

Differences between PRR and PMd

Spatial motor-goal tuning in PRR and PMd, in part, was differently affected by contextual modulations. First, PMd neurons were on average stronger active by 15% during planning of anti-reaches, but only if spatial information was already available, i.e., in the RS but not in the R condition, negative bias of GM, but not DM. We interpret this as an indication for the need of overruling a "default" movement plan (pro-reach) induced by the visual cue (Schlag-Rey et al., 1997; Everling et al., 1999), which might otherwise dominate. The stronger activity in PMd during anti-reaches in this view reflects the learned counterbalance of the network, which is necessary to compensate for the imbalance between pro-reaches and anti-reaches in terms of cue–response congruency. The counterbalance is only necessary in conditions

in which a spatial cue was already presented, to suppress a default pro-movement (RS condition), but not when only the rule is known and no specific default plan yet exists (R condition). Enhanced activation because of spatial cue–response incongruence was found previously in the supplementary eye fields during the planning of antisaccades (Schlag-Rey et al., 1997; Amador et al., 2004). Stronger responses during planning of prosaccades compared with antisaccades, however, were found in the posterior parietal area LIP (lateral intraparietal area) (Gottlieb and Goldberg, 1999). We found a higher fraction of individual neurons in PRR with a pro-preferring bias in response strength, but no bias in the average population response.

Second, neurons in PRR were more directionally selective during planning of pro-reaches (positive bias of SM corresponding to 20% stronger selectivity). This might reflect conflicting input to PRR specifically during planning of anti-reaches. Opposing input could result from bottom-up visuospatial information about the (remembered) cue location or default reach goal, on the one hand, and top-down projections of spatial motor-goal information from motor-tuned output stages, on the other hand (Brozović et al., 2007). Despite predominant motor-goal tuning on the population level (Gail and Andersen, 2006), the combination of such opposing inputs could explain the strongly reduced directional selectivity, which was found in many PRR neurons mainly in the anti-reach condition.

Together, the complementary effects of context on motor-goal representations indicate that posterior parietal cortex might more strongly represent stimulus-driven default or automatic movement plans (Desmurget et al., 1999; Pisella et al., 2000), whereas premotor cortex counterbalances this with predominant representations of rule-guided "inferred" movement plans.

Different contextual modulations in different neuronal populations

In PRR, those neurons that show strong gain modulation (high absolute GM) tended to have weak directional selectivity modulation (low absolute SM), and vice versa. This suggests that gain and selectivity modulations are the result of two separate underlying mechanisms, which affect different neuronal subpopulations. Gain and direct modulations, on the other hand, did not show such mutual exclusiveness. These findings are consistent with our predictions based on the previous model (Brozović et al., 2007). If the reduced directional selectivity in anti-reaches is indeed a consequence of conflicting independent bottom-up visuospatial and top-down motor-goal input to PRR, then we would not expect the GM/DM versus SM to necessarily be large in the same neurons. In PMd, we could not identify a systematic relationship or correlation pattern between the different types of contextual modulations.

Gain modulation as universal mechanism for flexible remapping

The contextual gain modulation of motor-goal representations in PMd and PRR strongly suggests gain-field mechanisms for flexibly mapping spatial sensory information onto spatial motor-goal representations according to arbitrary transformation rules. Based on our data, we propose that gain modulation is used by the brain not only in sensorimotor areas to achieve reference frame transformations driven by multisensory input (Andersen et al., 1985; Zipser and Andersen, 1988; Boussaoud et al., 1993, 1998; Brotchie et al., 1995; Galletti et al., 1995; Buneo et al., 2002), but also in the frontoparietal reach network to achieve contextually modulated, goal-directed visuomotor remapping, as previously suggested theoretically (Salinas, 2004; Brozović et al., 2007). Our current data support this idea by providing first experimental evidence for context-specific gain modulations of spatial motor-goal tuning in PRR and PMd, which could denote the key underlying principle of flexible goal-directed behavior.

References

Abbott LF, Chance FS (2005) Drivers and modulators from push-pull and balanced synaptic input. Prog Brain Res 149:147–155.

Amador N, Schlag-Rey M, Schlag J (2004) Primate antisaccade. II. Supplementary eye field neuronal activity predicts correct performance. J Neurophysiol 91:1672–1689.

Andersen RA, Buneo CA (2002) Intentional maps in posterior parietal cortex. Annu Rev Neurosci 25:189–220.

Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons. Science 230:456–458.

Batista AP, Buneo CA, Snyder LH, Andersen RA (1999) Reach plans in eye-centered coordinates. Science 285:257–260.

Blatt GJ, Andersen RA, Stoner GR (1990) Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. J Comp Neurol 299:421–445.

Boussaoud D, Wise SP (1993) Primate frontal-cortex—effects of stimulus and movement. Exp Brain Res 95:28–40.

Boussaoud D, Barth TM, Wise SP (1993) Effects of gaze on apparent visual responses of frontal-cortex neurons. Exp Brain Res 93:423–434.

Boussaoud D, Jouffrais C, Bremmer F (1998) Eye position effects on the neuronal activity of dorsal premotor cortex in the macaque monkey. J Neurophysiol 80:1132–1150.

Brotchie PR, Andersen RA, Snyder LH, Goodman SJ (1995) Head position signals used by parietal neurons to encode locations of visual stimuli. Nature 375:232–235.

Brozović M, Gail A, Andersen RA (2007) Gain mechanisms for contextually guided visuomotor transformations. J Neurosci 27:10588–10596.

Brozović M, Abbott LF, Andersen RA (2008) Mechanism of gain modulation at single neuron and network levels. J Comput Neurosci 25:158–168.

Buneo CA, Jarvis MR, Batista AP, Andersen RA (2002) Direct visuomotor transformations for reaching. Nature 416:632–636.

Cisek P (2007) Cortical mechanisms of action selection: the affordance competition hypothesis. Philos Trans R Soc Lond B Biol Sci 362:1585–1599.

Cisek P, Kalaska JF (2002) Modest gaze-related discharge modulation in monkey dorsal premotor cortex during a reaching task performed with free fixation. J Neurophysiol 88:1064–1072.

Colby CL, Gattass R, Olson CR, Gross CG (1988) Topographical organization of cortical afferents to extrastriate visual area PO in the macaque—a dual tracer study. J Comp Neurol 269:392–413.

Crammond DJ, Kalaska JF (1994) Modulation of preparatory neuronal activity in dorsal premotor cortex due to stimulus-response compatibility. J Neurophysiol 71:1281–1284.

Crammond DJ, Kalaska JF (2000) Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. J Neurophysiol 84:986–1005.

Desmurget M, Epstein CM, Turner RS, Prablanc C, Alexander GE, Grafton ST (1999) Role of the posterior parietal cortex in updating reaching movements to a visual target. Nat Neurosci 2:563–567.

di Pellegrino G, Wise SP (1993) Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. J Neurosci 13:1227–1243.

Everling S, Dorris MC, Klein RM, Munoz DP (1999) Role of primate superior colliculus in preparation and execution of anti-saccades and pro-saccades. J Neurosci 19:2740–2754.

Gail A, Andersen RA (2006) Neural dynamics in monkey parietal reach region reflect context-specific sensorimotor transformations. J Neurosci 26:9376–9384.

Galletti C, Battaglini PP, Fattori P (1995) Eye position influence on the parieto-occipital area PO (V6) of the macaque monkey. Eur J Neurosci 7:2486–2501.

Gottlieb J, Goldberg ME (1999) Activity of neurons in the lateral intraparietal area of the monkey during an antisaccade task. Nat Neurosci 2:906–912.

Kalaska JF (1996) Parietal cortex area 5 and visuomotor behavior. Can J Physiol Pharmacol 74:483–498.

Marconi B, Genovesio A, Battaglia-Mayer A, Ferraina S, Squatrito S, Molinari M, Lacquaniti F, Caminiti R (2001) Eye-hand coordination during reaching. I. Anatomical relationships between parietal and frontal cortex. Cereb Cortex 11:513–527.

Mushiake H, Tanatsugu Y, Tanji J (1997) Neuronal activity in the ventral part of premotor cortex during target-reach movement is modulated by direction of gaze. J Neurophysiol 78:567–571.

Nakamura K, Chung HH, Graziano MS, Gross CG (1999) Dynamic representation of eye position in the parieto-occipital sulcus. J Neurophysiol 81:2374–2385.

Nakayama Y, Yamagata T, Tanji J, Hoshi E (2008) Transformation of a virtual action plan into a motor plan in the premotor cortex. J Neurosci 28:10287–10297.

Pisella L, Gréa H, Tilikete C, Vighetto A, Desmurget M, Rode G, Boisson D, Rossetti Y (2000) An "automatic pilot" for the hand in human posterior parietal cortex: toward reinterpreting optic ataxia. Nat Neurosci 3:729–736.

Rushworth MF, Nixon PD, Passingham RE (1997) Parietal cortex and movement. 1. Movement selection and reaching. Exp Brain Res 117:292–310.

Salinas E (2004) Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. J Neurosci 24:1113–1118.

Salinas E, Abbott LF (1996) A model of multiplicative neural responses in parietal cortex. Proc Natl Acad Sci U S A 93:11956–11961.

Salinas E, Thier P (2000) Gain modulation: a major computational principle of the central nervous system. Neuron 27:15–21.

Schlag-Rey M, Amador N, Sanchez H, Schlag J (1997) Antisaccade performance predicted by neuronal activity in the supplementary eye field. Nature 390:398–401.

Snyder LH, Batista AP, Andersen RA (1997) Coding of intention in the posterior parietal cortex. Nature 386:167–170.

Snyder LH, Grieve KL, Brotchie P, Andersen RA (1998) Separate body- and world-referenced representations of visual space in parietal cortex. Nature 394:887–891.

Stoet G, Snyder LH (2004) Single neurons in posterior parietal cortex of monkeys encode cognitive set. Neuron 42:1003–1012.

Toni I, Rushworth MF, Passingham RE (2001) Neural correlates of visuomotor associations—spatial rules compared with arbitrary rules. Exp Brain Res 141:359–369.

Toth LJ, Assad JA (2002) Dynamic coding of behaviourally relevant stimuli in parietal cortex. Nature 415:165–168.

Wallis JD, Miller EK (2003) From rule to response: neuronal processes in the premotor and prefrontal cortex. J Neurophysiol 90:1790–1806.

Weinrich M, Wise SP (1982) The premotor cortex of the monkey. J Neurosci 2:1329–1345.

Wise SP, Murray EA (2000) Arbitrary associations between antecedents and actions. Trends Neurosci 23:271–276.

Wise SP, Boussaoud D, Johnson PB, Caminiti R (1997) Premotor and parietal cortex: corticocortical connectivity and combinatorial computations. Annu Rev Neurosci 20:25–42.

Zar JH (1999) Biostatistical analysis. Upper Saddle River, NJ: Prentice Hall.

Zipser D, Andersen RA (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. Nature 331:679–684.

Supplemental Material

S1: Relation between GM and SM

We use the GM and the SM to measure two different types of contextual modulation of directionally tuned cells: decreased (or increased) overall activity for one of the two transformation rules, on the one hand, and reduced (or enhanced) spatial selectivity, on the other hand. To be good indicators for these modulations, the indices should be selective and independent. Figure S1 shows how GM and SM independently change as a function of either tuning amplitude *A* or width κ of idealized von Mises tuning curves. The von Mises distribution is given by:

$$f(x \mid \mu, \kappa) = \frac{e^{\kappa \cos(x-\mu)}}{2\pi\, I_0(\kappa)}$$

where I_0 is the modified Bessel function of order zero. Figure S1 shows that for an idealized motor-goal cell with identical tuning width in both task conditions, a change of the amplitude in one task condition only results in a change of the GM without influencing the SM. Note that non-normalized tuning vectors are plotted. Normalized tuning vectors in the left half of the graph would have equal lengths for pro and anti. Note also that in case the response completely vanishes for one transformation rule, the SM becomes arbitrarily high even without systematic change in κ (far left). With additional random fluctuations, as typical for empirical neural data, this effect of high SM values would have to be expected for GM values smaller but close to one. If, on the other hand, we only modify the tuning width without changing the overall activity (in a normalized von Mises distribution the area under the curve does not change if κ is changed), then only the SM changes and the GM remains unchanged. Since empirical neuronal tuning

functions not necessarily exactly fit von Mises functions in all cases, we decided to use our non-parametric approach. Our modulation indices (GM, SM) reflect contextual modulations of average amplitude and selectivity independently and without the requirement of a certain exact tuning shape.

S2: Results in individual monkeys

Population data for the individual monkeys (Fig. S2) show that contextual modulations in both animals are very similar. Due to the similarity we pooled the data for the main manuscript. In both monkeys DCMi distributions were unbiased in PRR (S: m = 0.006, p = 0.69; A: m = 0.014, p = 0.34; t-test) and PMd (S: m = 0.018, p = 0.44; A: m = 0.010, p = 0.67), but broader than the distribution of the shuffled data, except for PRR in monkey S (PRR S: p = 0.17, A: p = 0.0022; PMd S: $p = 6.3 \times 10^{-4}$, A: $p = 3.38 \times 10^{-4}$; Bartlett test). CGM distributions were unbiased in PRR (S: m = 0.0026, p = 0.88; A: m = 0.035, p = 0.064; t-test) and negatively (anti) biased in PMd (S: m = -0.095, p = 0.013; A: m = -0.064, p = 0.025) for both animals. In both areas CGM distributions are broader than expected by chance in both monkeys (PRR S: $p = 3.1 \times 10^{-4}$, A: $p < 10^{-4}$; PMd S: $p < 10^{-10}$, A: $p < 10^{-10}$; Bartlett test). The CSM distributions consistently showed a positive bias in PRR (S: m = 0.14, $p = 7.56 \times 10^{-5}$; A: m = 0.086, $p = 4.81 \times 10^{-4}$) and no bias in PMd (S: m = -0.028, p = 0.53; A: m = 0.054, p = 0.13). Again CSM distributions in both areas are broader than expected by chance (PRR S: p = 0.0011, A: p = 0.032; PMd S: p = 0.0018, A: p = 0.0014; Bartlett test).

S3: Gain-modulation of motor-goal versus visual tuning

The idea that visual cue locations are flexibly mapped onto arbitrary motor-goals by gain-modulation mechanisms implies that at the intermediate level of the visuomotor transformation one should expect gain-modulated visual tuning, rather than gain-modulated motor-goal tuning. After all, motor-goal tuning is considered the outcome of the transformation process, not the intermediate representation. Why did we then see gain-modulation of motor-goal tuning in PRR and PMd? The key to understanding this seeming contradiction is the differentiation between a pure feed-forward network, for achieving the desired transformation, compared to a network with top-down feedback. Brozovic et al. (Brozovic et al., 2007) showed that with both network architectures one can perform context-specific visuomotor transformations via gain-modulation. But the feed-forward network developed gain-modulated visual tuning in the 'hidden' sensorimotor layer, while the top-down network produced the same output with gain-modulated motor-goal tuning in the sensorimotor layer. The latter architecture was considered more physiological in the sense that PRR and PMd had been shown to be tuned for the motor-goal during instructed-delay periods (Crammond and Kalaska, 1994;Gail and Andersen, 2006) and connectivity between the posterior parietal cortex and premotor areas is known to be reciprocal.

One can assume a certain delay before the top-down motor-goal projections exert their influence on the spatial tuning in the sensorimotor part of the network. Hence, the predominant motor-goal tuning during the instructed-delay is preceded by a short period of 'visual' tuning (or default motor-goal tuning for the pro-target) during spatial cue

presentation, both in the model (Brozovic et al., 2007), and in corresponding physiological data (Gail and Andersen, 2006;Zhang et al., 1997;Zhang and Barash, 2000;Zhang and Barash, 2004). Once the transformation is achieved and the output stage of the sensorimotor network shows motor-goal tuning, this motor-goal tuning will be fed back to the intermediate sensorimotor stages and replace the predominant visual tuning quickly with predominant motor-goal tuning, as has been found. The observed gain-modulation of the motor-goal tuning in this view is a reflectance of the general context-induced gain-modulation at this processing stage, which has its major functional relevance during the transformation process itself, i.e. when visual tuning dynamically changes to motor-goal tuning, but is still visible in the resulting motor-goal tuning itself. Importantly, if the motor-goal tuning constitutes the outcome of a transformation process based on gain-modulation of visual cue representations then we should expect the short-term visual tuning during cue presentations to be gain-modulated. This was the case in the model (Brozovic et al., 2007). It is technically very difficult to test this prediction in a physiological experiment like our current study. This is because of the following potential confounds: Since the emergence of a motor goal is a dynamic process which is influenced by the sensory input, it is difficult to find a time window that only accounts for visual tuning – gain-modulated or not – and does not also include those dynamical processes which ultimately lead to a motor goal representation. For example, a neurons early response might have smaller amplitude in anti-reaches because the instruction cue was in its response field (e.g., right) but the motor-goal is on the opposite side, i.e. outside the response field. Hence, we have to expect down-regulation of this neurons' response as soon as the transformation process is initiated (Supplementary Fig. 3). The

exact timing of this is hard to identify (Supplementary Fig. 3 top and bottom), which makes it impossible to judge if any early difference in response amplitude is an expression of a gain-modulated 'visual' response or the consequence of the transformation process itself. Albeit desirable, we could not achieve a readily interpretable analysis of this early trial period equivalently to the more stationary late memory period.

Reference List

1. Brozovic,M., Gail,A., and Andersen,R.A. (2007). Gain Mechanisms for Contextually Guided Visuomotor Transformations. Journal of Neuroscience *27*, 10588-10596.

2. Crammond,D.J. and Kalaska,J.F. (1994). Modulation of preparatory neuronal activity in dorsal premotor cortex due to stimulus-response compatibility. Journal of Neurophysiology *71*, 1281-1284.

3. Gail,A. and Andersen,R.A. (2006). Neural Dynamics in Monkey Parietal Reach Region Reflect Context-Specific Sensorimotor Transformations. Journal of Neuroscience *26*, 9376-9384.

4. Zhang,J., Riehle,A., Requin,J., and Kornblum,S. (1997). Dynamics of single neuron activity in monkey primary motor cortex related to sensorimotor transformation. J Neurosci *17*, 2227-2246.

5. Zhang,M. and Barash,S. (2000). Neuronal switching of sensorimotor transformations for antisaccades. Nature *408*, 971-975.

6. Zhang,M. and Barash,S. (2004). Persistent LIP Activity in Memory Antisaccades: Working Memory For a Sensorimotor Transformation. Journal of Neurophysiology *91*, 1424-1441.

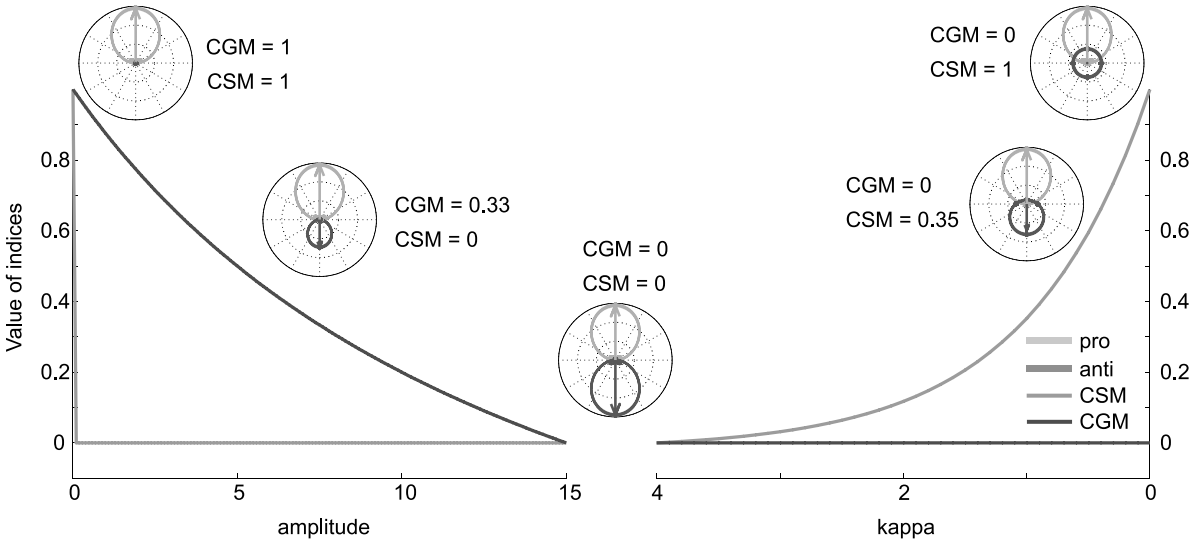

supplementary figure 1

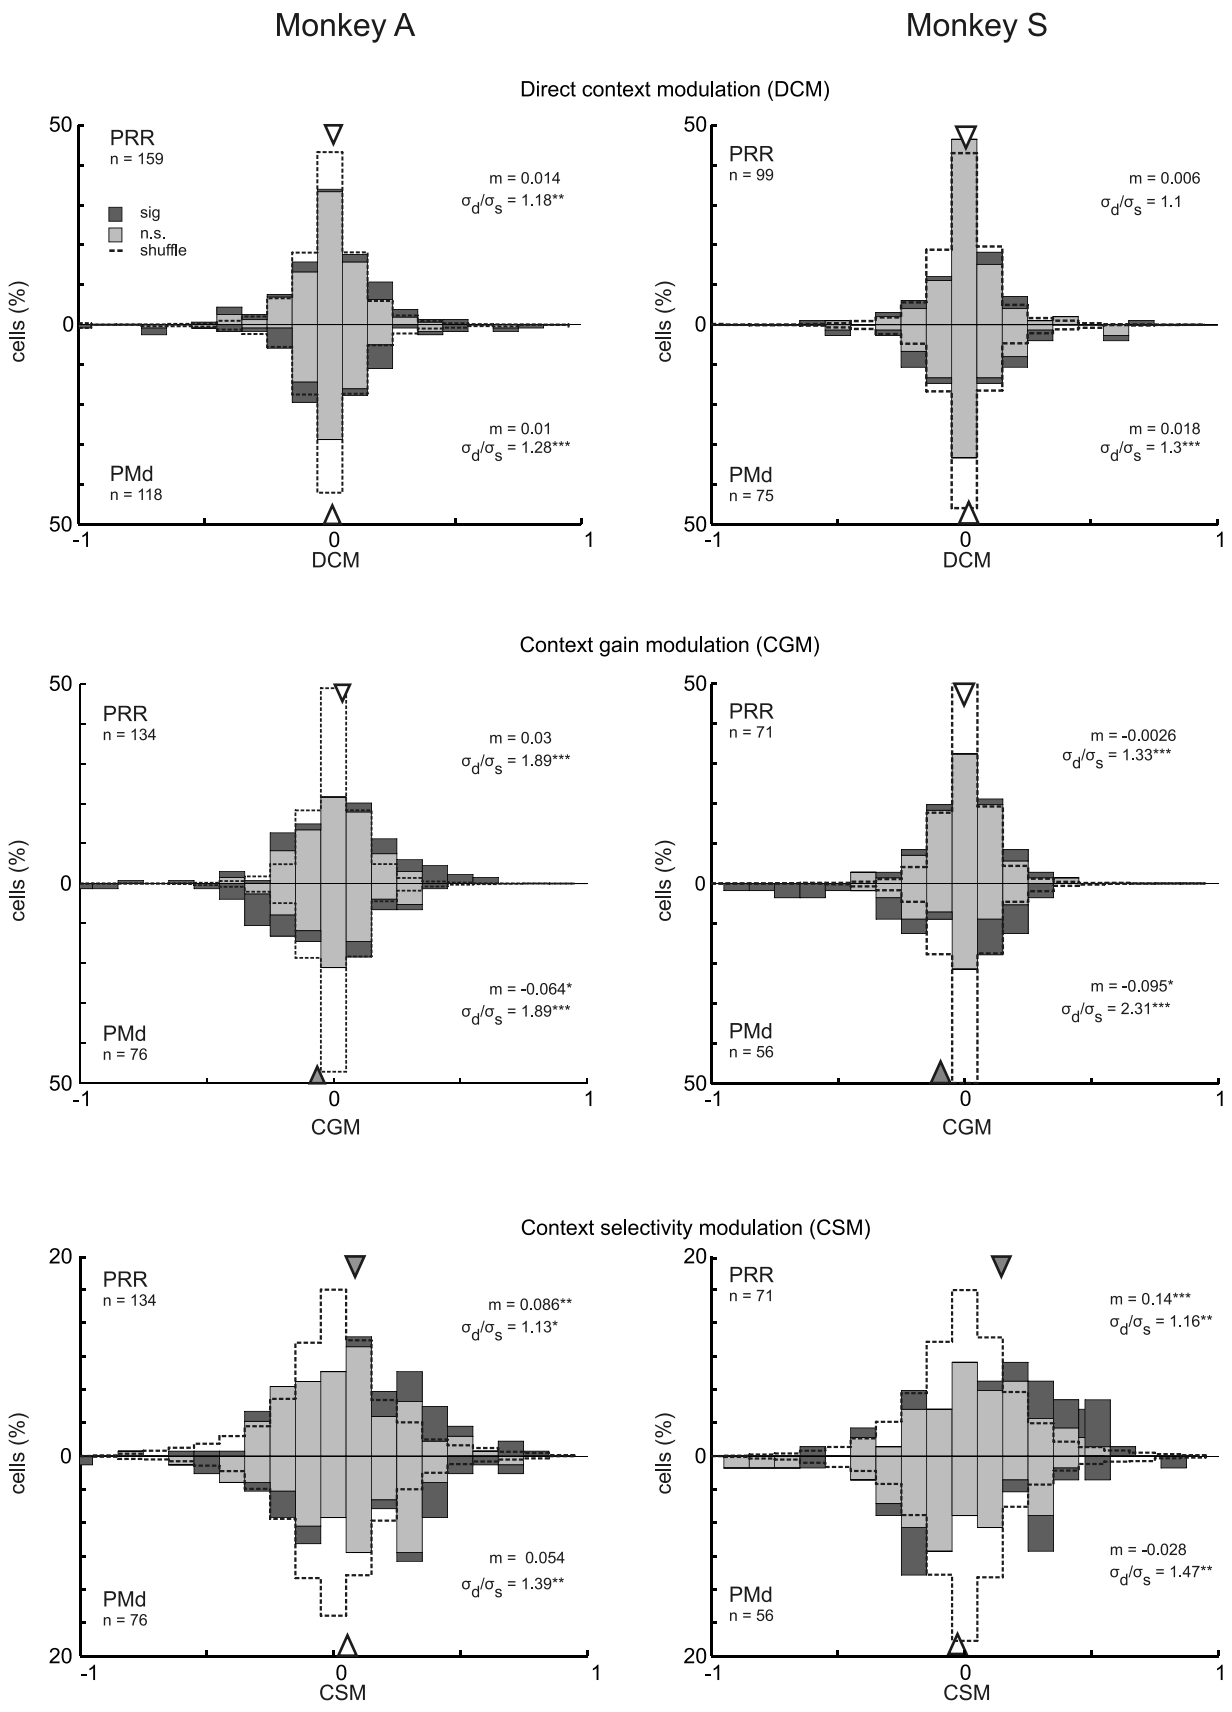

supplementary figure 2

visual cell example 1
(*PMd, monkey A*)

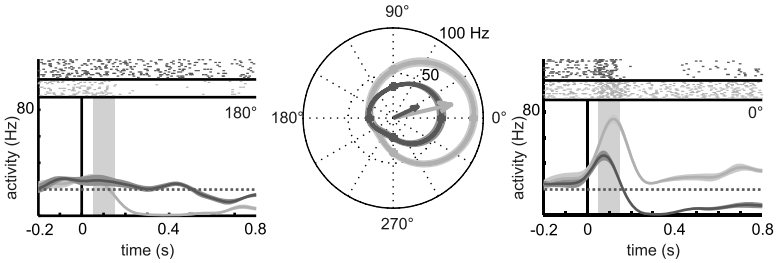

visual cell example 2
(*PRR, monkey S*)

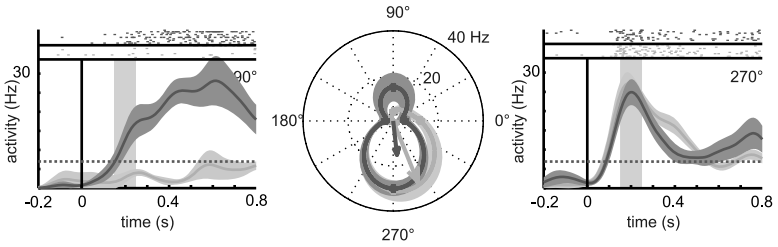

supplementary figure 3

2.2 The Cortical Timeline for Deciding on Reach Motor Goals

The previous study dealt with the mechanism by which transformation rules are implemented and showed differences between PRR and PMd. In this study we investigated the dynamics of decisions between motor goals in the two areas. Since both areas are reciprocally connected we wanted to know where the integration of spatial and contextual information takes place first. To accomplish this, we used a modified pro- and anti-reach task in which we separated the rule and the spatial information in time. In that way we could compare cue-related and motor-related activity and also motor planning activity and activity related to motor initiation.

We could show that the selectivity for the goal of an upcoming movement appeared earlier in PMd than PRR. This timing difference was only observed, when the task required a spatial remapping. That means if the spatial representation of a cue or a preliminary motor plan had to be transformed into a spatially incongruent motor goal representation, like in an anti-reach task. If the spatial representation was congruent (pro-reach task), no transformation had to take place and accordingly no frontoparietal latency differences were found. From this we concluded that the latency differences are the result of a dynamic reorganization process in PRR. Again we find a difference between PRR and PMd, which depends on the arbitrary remapping rule.

Behavioral/Systems/Cognitive

The Cortical Timeline for Deciding on Reach Motor Goals

Stephanie Westendorff, Christian Klaes, and Alexander Gail
Bernstein Center for Computational Neuroscience, German Primate Center–Leibniz Institute for Primate Research, 37077 Göttingen, Germany

Flexible sensorimotor planning is the basis for goal-directed behavior. We investigated the integration of visuospatial information with context-specific transformation rules during reach planning. We were especially interested in the relative timing of motor-goal decisions in monkey dorsal premotor cortex (PMd) and parietal reach region (PRR). We used a rule-based mapping task with different cueing conditions to compare task-dependent motor-goal latencies. The task allowed us a separation of cue-related from motor-related activity, and a separation of activity related to motor planning from activity related to motor initiation or execution. The results show that selectivity for the visuospatial goal of a pending movement occurred earlier in PMd than PRR whenever the task required spatial remapping. Such remapping was needed if the spatial representation of a cue or of a default motor plan had to be transformed into a spatially incongruent representation of the motor goal. In contrast, we did not find frontoparietal latency differences if the spatial representation of the cue or the default plan was spatially congruent with the motor goal. The fact that frontoparietal latency differences occurred only in conditions with spatial remapping was independent of the subjects' partial a priori knowledge about the pending goal. Importantly, frontoparietal latency differences existed for motor-goal representations during movement planning, without immediate motor execution. We interpret our findings as being in support of the hypothesis that latency differences reflect a dynamic reorganization of network activity in PRR, and suggest that the reorganization is contingent on frontoparietal projections from PMd.

Introduction

Goal-directed behavior is guided by spatial sensory input in combination with context-specific rules. This way, the same sensory input can lead to a wide variety of motor outputs. The reciprocally connected posterior parietal cortex (PPC) and dorsal premotor cortex (PMd) are believed to mediate such space–context integration (Mitz et al., 1991; Kalaska, 1996; Wise et al., 1996; 1997; Kalaska et al., 1997; Passingham et al., 2000; Eliassen et al., 2003; Wallis and Miller, 2003; Stoet and Snyder, 2004; Buneo and Andersen, 2006; Cisek, 2007; Pesaran et al., 2008). For example, neurons in the parietal reach region (PRR) (Gail and Andersen, 2006; Gail et al., 2009) and PMd (Boussaoud and Wise, 1993; di Pellegrino and Wise, 1993; Crammond and Kalaska, 1994, 2000) show sustained motor-goal tuning when the reach goal has to be inferred from a spatial cue by applying a context-specific transformation rule.

How is such space–context integration accomplished, and what is the functional role of reciprocal frontoparietal projections? Network models of space–context integration have suggested frontoparietal projections to help sustain motor-goal memory (Cisek, 2006) or to mediate context-specific visuomotor transformations via gain modulation (Brozović et al., 2007; Gail et al., 2009). Also, frontoparietal neural latency differences (LDs) were shown for motor-related activity around the time of movement initiation (Kalaska et al., 1983; Johnson et al., 1996), which nourished the ideas that frontoparietal LDs could reflect corollary discharge signals (Mountcastle et al., 1975; Bioulac and Lamarre, 1979; Kalaska et al., 1983; Seal and Commenges, 1985; Johnson et al., 1996) and parietal cortex might be involved in internal forward modeling for current state estimation during motor control (Mulliken et al., 2008; Shadmehr and Krakauer, 2008). Finally, for motor planning activity before movement initiation frontoparietal LDs and transient bidirectional frontoparietal interactions are pronounced specifically during decision-making tasks, but not for tasks with direct cueing (Pesaran et al., 2008). This suggests that the functional interdependencies between frontal and parietal areas are task dependent (Kalaska and Crammond, 1992; Cisek, 2006; 2007).

We tested how compatible frontoparietal LDs are with the above ideas. We asked specifically whether frontoparietal LDs are the consequence of a corollary discharge signal, and in which way they depend on cognitive decision processes. Different from previous studies, we explicitly isolated the latencies of motor planning activity from cue- or movement-related activity by using a pro-reach/anti-reach task (Crammond and Kalaska, 1994; Zhang et al., 1997; Gail and Andersen, 2006) in combination with partial precueing (Lecas et al., 1986; Riehle and Requin, 1989; Riehle, 1991; Hoshi and Tanji, 2000; Nakayama et al., 2008).

The pattern of frontoparietal LDs that we found in our study suggests that slower motor-goal representations in PRR than PMd reflect a need for the buildup or reorganization of spatial representations in PRR during spatial remapping tasks, independent of corollary discharge or decision signals. We will discuss how far the results could indicate the existence of prospective internal forward model predictions triggered by motor planning rather than by motor execution.

Received Sept. 17, 2009; revised Jan. 22, 2010; accepted Jan. 29, 2010.

This work was supported by German Federal Ministry for Education and Research (Bundesministerium für Bildung und Forschung) Grants 01GQ0433 and 01GQ0814. We thank Sina Plümer and Ludwig Ehrenreich for their help in data collection.

Correspondence should be addressed to Dr. Alexander Gail at the above address. E-mail: agail@gwdg.de.

DOI:10.1523/JNEUROSCI.4628-09.2010

 0270-6474/10/305426-11$15.00/0

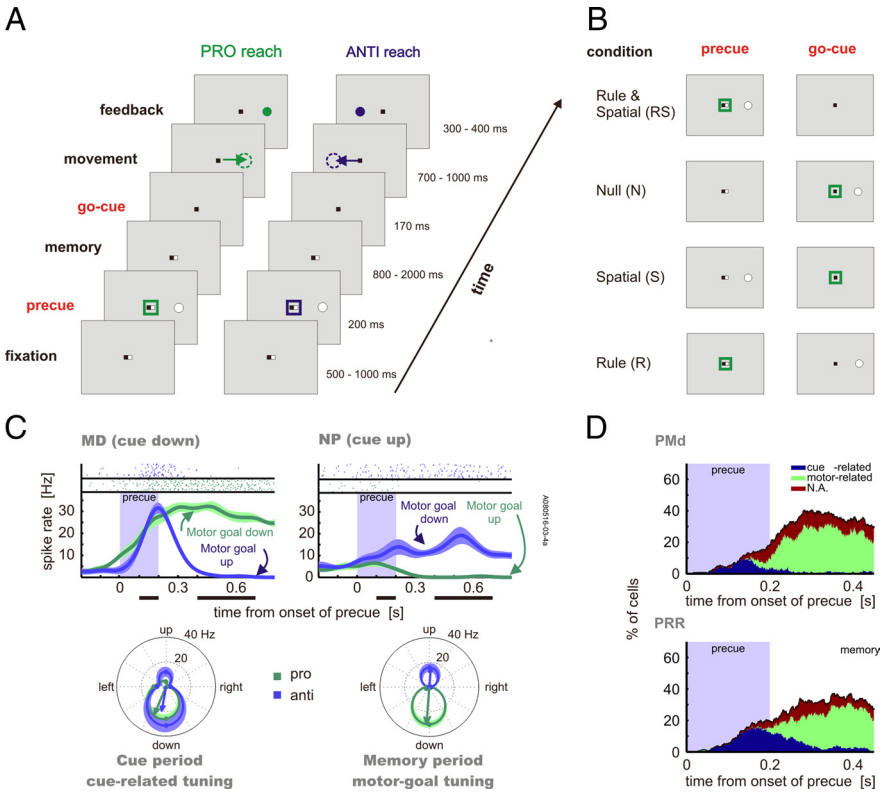

Figure 1. Anti-reach task with different cueing conditions and visuomotor tuning in PMd and PRR. ***A***, Memory-guided anti-reach task. The left sequence shows an example of a pro-trial, the right an example of an anti-trial in the RS condition. Black and white central squares represent the eye- and hand-fixation stimuli. The central colored frame represents the contextual cue instructing the transformation rule pro or anti, and the peripheral white circle represents the spatial cue. The target position is indicated by the dotted circle. Arrows and dotted circles are for demonstration purpose only and were not shown to the subjects. Note that a go-cue period of 170 ms is only relevant for other cueing conditions, as in ***B***. ***B***, In four cueing conditions, the contextual and spatial cue could be presented simultaneously or separately either before or after a variable memory period: transformation rule and spatial cue presented as precue (RS condition), no precue (N; rule and spatial cue as go-cue), only rule as precue (R; spatial cue as go-cue), or only spatial cue as precue (S; rule as go-cue). ***C***, PSTHs and raster plots of a PRR example neuron. This example motor-goal neuron shows a dynamic switching for cue- to motor-related tuning in response to the precue (shaded area) in the RS condition. Neuronal activities are shown for spatial cues presented at 270° (left panel; MD) and 90° (right panel; NP). Polar plots show the directional tuning of the neuron between 100 and 200 ms (cue period), and 400 and 700 ms (memory period) relative to the cue position separately for pro- and anti-trials. ***D***, Visuomotor dynamics in the population of PMd and PRR neurons. The percentage of spatially tuned neurons of all recorded neurons is plotted for PMd (top) and PRR (bottom) for the first 450 ms after onset of the precue in the RS condition. Blue bars reflect the percentage of neurons with cue-related tuning, green bars reflect motor-related tuning, and red bars reflect spatially undefined tuning (see Materials and Methods).

Materials and Methods

Tasks and control of behavior. With our task, we wanted to test how the dynamics of motor-goal decisions and their cortical representations depend on whether the subjects had to perform a memory or reaction-time (RT) task. Further, we aimed to test how partial advance information (precueing) about the to-be-performed reach influences the dynamics of motor-goal decisions.

To separate neural motor-goal representations from other neural signals, we used an anti-reach task design (Gail and Andersen, 2006). The anti-reach task required the subjects to map a spatial instruction cue onto one of two reach goals, either at the location of the cue (pro) or opposite to it (anti). The spatial transformation rule was instructed with a colored context cue (Fig. 1*A*) (see below for details). The motor goal could only be defined uniquely by the combination of both cues. The eight reach conditions (two context conditions × four cue directions) were randomly interleaved.

Different cueing conditions were realized by presenting the cues either before (precue period) or after (go-cue period) an instructed delay. The precue consisted of the contextual and the spatial cue together [rule and spatial condition (RS)], only the spatial cue [spatial-only condition (S)], only the contextual cue [rule-only condition (R)], or nothing [null condition (N)]. The four cueing conditions were randomly mixed in small blocks of 10–20 trials (Fig. 1*B*).

Stimuli were presented on an LCD screen (19 inch ViewSonic LCD VX922; onset latencies corrected). The contextual cue consisted of a colored frame around the small, central red eye and white hand fixation points, and instructed the subject to reach toward (green = pro) or diametrically opposite (blue = anti) of the spatial cue. The peripheral and color-neutral spatial cue was located left, right, above, or below the central fixation points at 9 cm eccentricity [14.5° visual angle (VA)].

The exact timeline of the trials was as follows (Fig. 1*A*): Monkeys initiated the trials by acquiring central eye fixation (tolerance: 2.5–4.0° VA; 224 Hz CCD camera, ET-49B, Thomas Recording) and hand fixation at a touch screen (tolerance: 4.0° VA; IntelliTouch screen, ELO Systems). A variable-length fixation period (500–1000 ms) was followed by the brief precue period (200 ms). The monkeys had to keep eye and hand fixation for the memory period (800–2000 ms), before the hand fixation stimulus disappeared (go-signal). In non-RS trials, a simultaneous go-cue was flashed (170 ms) and completed the missing information to uniquely specify the reach goal. Monkeys had to reach toward the instructed goal within a maximum of 700–1000 ms (movement period; reach tolerance: 5.7° VA) and hold the goal position for 300–400 ms (feedback period). Visual feedback indicated the correct movement goal (filled circle of the same color as the contextual cue at the goal location) during the last 300 ms of the feedback period, i.e., only once the monkey had finished the reach. Eye fixation had to be kept throughout the trial. Liquid reward and acoustic feedback indicated correct (high pitch tone, reward) or incorrect (low pitch tone, no reward) behavior.

Animal preparation and neural recordings. Two male rhesus monkeys (*Macaca mulatta*; monkey S and monkey A) were trained to perform the anti-reach task with partial precueing. The recording procedure was as in previous experiments (Gail et al., 2009). In brief, monkeys were implanted with a titanium head holder and two magnetic resonance imaging (MRI)-compatible recording chambers, custom-fit to the monkeys' heads (3di). Chamber positioning above PRR [Horsley Clarke coordinates: 6 mm/8.5 mm (monkey S/monkey A) contralateral; 10 mm/9.5 mm posterior], and PMd (13 mm/13.5 mm contralateral; 17 mm/19.8 mm anterior) was guided by presurgical structural MRI and confirmed by postsurgical MRI. Sustained direction-selective neural responses during center-out reach planning (memory period) were used as a physiological signature in both areas to confirm the region of interest. Both chambers were implanted contralaterally to the handedness of the monkey (monkey A, left hemisphere; monkey S, right hemisphere). All surgical and imaging procedures were conducted under general anesthesia.

We used two five-channel microdrives ("mini-matrix"; Thomas Recording) for extracellular recordings, mostly simultaneously in each chamber. The raw signals of the electrodes were preamplified (20×; Thomas Recording), bandpass filtered, and amplified (154 Hz to 8.8 kHz; 400–800×; Plexon) before online spike sorting was conducted (Sort Client; Plexon). Spike times and spike waveforms were recorded and later subjected to additional offline sorting (Offline Sorter; Plexon).

Animal care and all experimental procedures were conducted in accordance with German laws governing animal care.

Time-resolved directional and rule selectivity in individual neurons. Time-resolved analysis of directional and rule selectivity was conducted to reveal the cortical dynamics of motor-goal decisions. The analysis time window was between 200 ms before onset of the relevant cue and 450 ms after cue onset. In the RS condition, the time around the precue was analyzed. In the N, R, and S conditions, the time around the go-cue was analyzed. All recorded and sufficiently well isolated units were included in the analysis, regardless of their tuning properties, unless explicitly stated otherwise.

For time-continuous analyses, we used spike densities, which were computed by convolving each spike with a causal kernel, which was defined as follows:

$$R(t) = \frac{\tau_g + \tau_d}{{\tau_d}^2} * (1 - e^{\frac{-t}{\tau_g}}) * e^{\frac{-t}{\tau_d}},$$

where $R(t)$ is the spike density at time point t. The rise time constant τ_g was set to 2 ms, and the decay time constant τ_d was set to 20 ms (Thompson et al., 1996; Monosov et al., 2008). Average spike densities across trials with identical conditions (9–15 repetitions per condition) were sampled at 1 ms resolution and aligned to the onset of either the precue or the go-cue. To reduce variability in the time-resolved tuning analyses of individual neurons (see below), we additionally used a sliding window of 20 ms length, within which the spike densities were averaged and which was shifted in steps of 1 ms.

Time-resolved rule selectivity (context tuning) was defined by a significant difference (rank sum test) between the average spike densities in pro- and anti-trials in the memory period of the R condition. The analysis was conducted in each time bin.

Time-resolved directional selectivity was quantified with a directional tuning vector (DTV). It is defined as the vector average across all center-out cue directions $\vec{u}_i$ (unit vectors) weighted with the corresponding mean spike rates r_{ij} of neuron j as follows:

$$\mathrm{DTV}_j(t) = \sum_{i=1}^{4} r_{ij}(t)\vec{u}_i.$$

The direction α of the DTV can take any value between 0° and 360°. The DTV was defined relative to the location of the spatial cue, i.e., 0° corresponds to a rightward cue. The significance of directional tuning was tested with a nonparametric one-way ANOVA (Kruskal–Wallis, $\alpha = 0.05$, unless stated otherwise) plus an additional criterion. The criterion required that the length of the DTV had to be >0.2, after the DTV was normalized to the spike density at the direction with the maximal response in that time bin. The fixed criterion of 0.2 was arbitrary, but varying it between 0.15 and 0.25 did not change our conclusions. With this criterion, we excluded neurons with a bimodal tuning close to symmetric. Such symmetric bimodal tuning could represent two potential motor goals instead of one selected motor goal and would be characterized by a significant ANOVA but a DTV length close to zero (see S-4, available at www.jneurosci.org as supplemental material). The DTV was computed in every time bin, separately for pro- and anti-trials, and separately within each cueing condition.

Cue-related versus motor-related tuning. The anti-reach task design allowed separating cue-related from motor-related tuning by comparing the directional selectivity in pro- and anti-reaches (Gail and Andersen, 2006). A neuron's instantaneous (bin-to-bin) directional tuning was categorized according to the absolute difference in DTV direction between pro- and anti-trials in each time bin: cue-related tuning ($|\Delta\alpha| < 45°$), motor-related tuning ($|\Delta\alpha| > 135°$), or spatially undefined tuning (else). This definition of instantaneous tuning depended on a direct bin-wise comparison of spatial tuning between pro- and anti-trials in each cueing condition.

As known from previous studies (Crammond and Kalaska, 1994; Gail and Andersen, 2006), spatiotemporal tuning profiles of individual neurons in PMd and PRR can change over time from cue related to motor related (Fig. 1*C*), whereas other neurons may show motor-related tuning from the earliest onset of their response. The activity profile of the example neuron in Figure 1*C* shows the changing tuning properties of the neuron between the time windows 100–200 ms (cue period) and 400–700 ms (memory period). Figure 1*D* shows the number of cue- and motor-related neurons as a fraction of all recorded neurons over the time of the trial in PMd and PRR.

We characterized the neurons according to their spatiotemporal response profiles. We considered neurons to have a significant cue-related response when the instantaneous cue-related tuning in the RS condition persisted in at least 90% of 30 consecutive time bins within 200 ms after the onset of the precue (an equivalent definition of cue-related response could be achieved for responses following the go-cue in the N condition, which did not change our results). We defined neurons as motor-goal neurons when they showed significant movement planning activity. Motor planning activity was defined as motor-related tuning for at least 90% of 50 consecutive time bins within 450 ms after precue onset in the RS condition. The different number of time bins for a cue-related response and motor planning activity was set because the cue response is transient, whereas the motor-related response is more sustained. We defined neurons as perimovement neurons that did not show motor planning activity but that at some point in time after the go-cue were motor tuned.

Finally, we defined a time-invariant, motor-related preferred direction (PD) for each neuron. The PD was computed by averaging the tuning vectors in pro-reaches over all time bins of the analysis time window in which the neuron showed significant instantaneous motor-related tuning. The PD was computed separately within each cueing condition.

Determination of motor-goal latencies separately in pro- and anti-trials. If motor-goal latencies in a cortical area were different in pro- and anti-trials, then the analysis of instantaneous motor-related tuning would reveal the time course of the motor-goal tuning in the slower of two conditions (pro or anti), since instantaneous tuning is based on an bin-to-bin comparison between pro- and anti-trials.

To analyze motor-goal latencies separately in pro- and anti-trials, we needed to extract the time courses of cue- versus motor-related tuning in pro- and anti-trials separately. We accomplished this with two measures. First, we restricted our data set to neurons that did not show cue-related responses. Cue-related responses are indistinguishable from motor-related responses in pro-reaches, since cue and motor goals are spatially identical in this case. Hence, cue-related responses would confound measurements of motor-goal latencies in pro-trials. Second, we compared the instantaneous tuning direction of a neuron in the currently considered context (pro or anti) with the time-invariant PD of this neuron (see definition in the previous section) to decide on the time-resolved cue versus motor relatedness within this context.

Determination of motor-goal latencies in reaction-time conditions. Motor-related tuning could result either from motor planning activity or from perimovement activity (corollary discharge or sensory feedback about the movement). In the RT conditions (N, R, S), the isolation of motor-goal latencies (planning activity) from motor feedback (perimovement activity) was not guaranteed without further precautions. (Only the motor-related tuning in the memory period of the RS condition could be directly considered motor-goal tuning, since movements in these trials were initiated much later.) To isolate motor-goal latencies from latencies of perimovement activity in RT conditions, we separately conducted our latency analyses within either a population of exclusive motor-goal neurons, i.e., neurons that showed significant motor planning activity, or within the complementary population of perimovement neurons.

Population analyses of motor-related tuning. We characterized the motor-related tuning in neural populations in two ways. First, we computed average peristimulus time histograms (PSTHs) across neurons of each cortical area and within each different cueing condition. We did this to illustrate general neural response properties as a function of time. Second, we computed the time-resolved fraction of motor-tuned neurons (recruitment curves) in each cortical area and cueing condition for statistical analyses. We excluded neurons with cue-related responses to show the motor-related time courses.

PSTHs across neurons were computed as the average response in the direction of each neuron's maximal response [PSTH for maximum direction (MD)] and, separately, in the opposite direction [nonpreferred direction (NP)]. The maximum direction for each neuron was defined as

the direction with the maximal average response across time and trials in pro-trials within the time window of 200–350 ms after cue onset (precue in RS; go-cue in N, R, and S conditions). A time window shorter than the analysis time window was chosen because in PMd the activity of the neurons starts to decrease around 300 ms after the go-cue (see Fig. 2*C*,*D*). Note that the MD and NP directions are restricted to the discrete cue directions left, right, up, and down, since only for these directions PSTHs can be computed. The response of each neuron was normalized to its maximal response within the 200–350 ms time window and within each task condition before averaging across neurons. Average PSTHs across neurons were calculated separately for pro- and anti-trials in MD and NP for each cueing condition. Neurons that did not show significant directional selectivity in any time bin during the analysis time window in either pro- or anti-trials were excluded from this analysis. Because of the necessary normalization, random fluctuations in such nontuned neurons would be inappropriately amplified, and would contribute to the average PSTH with similar weight as strongly tuned neurons.

Recruitment curves denote the fraction of neurons that were motor tuned in a specific time bin relative to the total number of neurons that were motor tuned in any time bin of the analysis time window (see Fig. 2*B*). Recruitment curves in our experiment have two advantages over PSTH comparisons. First, normalization and averaging of neural responses were not needed for computing recruitment curves. Averaging across neurons, as done for population PSTHs, typically cannot account for the different response profiles of the individual neurons. For example, many neurons have very low baseline firing and show brisk response onsets after cue presentation, whereas others have high baseline levels and might exhibit their directional selectivity by small excitatory responses to one cue direction and strong inhibitory responses to other directions. Also, normalization tends to boost weak, noisy responses. Second, the recruitment curves are based on tuning functions that interpolate between all four measured directions. The comparison between MD and NP in PSTHs is restricted to two of the four measured directions. As a consequence, PSTHs are less sensitive for neurons with oblique PDs and broad tuning, which is not the case for recruitment curves. For these reasons, we based our statistical analyses on the recruitment curves and use PSTHs only for illustration purposes.

Quantification of neural LDs. We determined motor-goal latencies by a threshold criterion for the neuronal recruitment curves (for controls and alternatives, see S1–S3, available at www.jneurosci.org as supplemental material). Latency of motor tuning within each condition was defined as the average time relative to cue onset at which the fraction of motor-tuned neurons exceeded each of seven threshold levels for at least 100 ms. Seven equidistant threshold levels were defined relative to the following baseline: $Th(p) = \text{baseline} + (p \times (1 - \text{baseline}))$, where "baseline" refers to the average fraction of tuned neurons in the 200 ms before cue onset. The p value was set to 0.05, 0.1, 0.15, 0.2, 0.25, 0.3, and 0.35 to characterize the early tuning onset of the fastest ~35% of neurons. The average across seven thresholds was built to become less sensitive to minor random fluctuations. Note that absolute latencies depend on the choice of thresholds, and hence are of limited conclusiveness. We statistically analyzed only latency differences between conditions or neural populations.

Pairwise neural LDs between task conditions or cortical areas were computed and tested for being different from zero. Averaging across absolute latencies measured with multiple threshold levels would induce artificial variance and compromise statistical power. Therefore, we computed the LDs separately for each threshold level and only then computed the average LD across thresholds. Permutation tests ($N = 10{,}000$ samples) were used to estimate the probability that the original sample indicates a significant difference from zero. In each permutation, each neuron's tuning data were randomly assigned to one of the two task conditions or cortical areas in consideration, such that the numbers of neurons in both populations were the same as in the original sample. The resulting p value represents the percentage of random permutations leading to an LD larger or equal to the original sample. A difference was considered significant if the p value was <5%.

Neural response latencies and behavioral reaction times. Behavioral RTs were defined as the median time from onset of the go-cue until the release of the touch screen within each recording session and task condition. Average RTs were computed across recording sessions. RT differences between pro- and anti-trials were tested with a paired t test across sessions and within each task condition.

To relate motor-goal latencies of individual neurons to behavioral RTs, we analyzed motor-goal latencies after the go-cue in all RT conditions (N, R, and S) on a trial-by-trial basis. We used a threshold criterion to define latencies in those trials where the neurons showed a good-enough response to the pending motor goal: post-go-cue, single-trial activity needed to exceed baseline activity (= mean spike density of the last 200 ms of the N condition memory period) by at least three times its SD for at least 100 ms within the analysis time window. To cope with the sparse spike events in low-firing neurons, we estimated spike densities in single trials with a Gaussian kernel of width $\sigma = 50$ ms, as follows:

$$R(t) = \frac{1}{\sqrt{2\pi\sigma^2}} * e^{\frac{-t^2}{2\sigma^2}}.$$

The Gaussian kernel resulted in stronger smoothing than the EPSP kernel used for the trial-average data above. In the RS condition, the onset of motor-goal tuning (after the precue) is not related to the RT (after the go-cue); therefore, single-trial latencies in the RS condition were not taken into account. Nontuned neurons and neurons with a cue-related response (see above) were excluded from the analysis to avoid confounds.

We removed RT and single-trial neural latency outliers for the trial-to-trial regression analysis between RTs and neural latencies. A value was classified as an outlier if its distance from the median was >1.5 times the interquartile range (distance between the 75th and 25th percentiles). Since only a small fraction of trials allowed determination of neural response latencies, data from all RT conditions were merged, and linear regression curves were calculated for all neurons in which single trial latencies could be estimated in at least 10 trials.

Results

Extracellular single-neuron spiking activities were recorded from 258 neurons in PRR (monkey S, 99; monkey A, 159) and 192 neurons in PMd (monkey S, 74; monkey A, 118) while the monkeys performed the anti-reach task with partial precueing (Fig. 1). The average performance of monkey S was 77%/77% (pro/anti) in the RS condition, 80%/79% in the N condition, 75%/75% in the R condition, and 79%/79% in the S condition. The average performance of monkey A was 86%/86% in the RS condition, 86%/85% in the N condition, 85%/83% in the R condition, and 85%/84% in the S condition for pro- and anti-reaches, respectively. Errors were mainly caused by ocular fixation breaks and not by incorrect target choices. The choice of the monkey for a reach target was correct in 97% for monkey S, and in 99% for monkey A.

Motor-goal latencies after simultaneous cueing of transformation rule and spatial information

Is motor planning activity represented earlier in either PMd or PRR? If there is an LD, do the LDs depend on a motor command or motor execution, or do they also exist for planning activity proper? To answer these questions, we first analyzed motor-goal tuning latencies in PMd and PRR, when the monkeys were instructed simultaneously with both the transformation rule (pro/anti) and the spatial cue. We compared motor-goal latencies in response to the simultaneous cues in instructed-delay trials (precue in the RS condition) and in reaction-time trials (go-cue in the N condition). If motor-goal latency differences exist in both conditions—especially if they exist after the precue at the beginning of an instructed delay—then they must be independent of motor execution and represent dynamics of motor planning. If latency differences only existed after a go-cue immediately before or during motor execution, then they should be related to corollary discharge or motor-related sensory feedback. The anti-reach task allowed us to separate cue- from motor-related

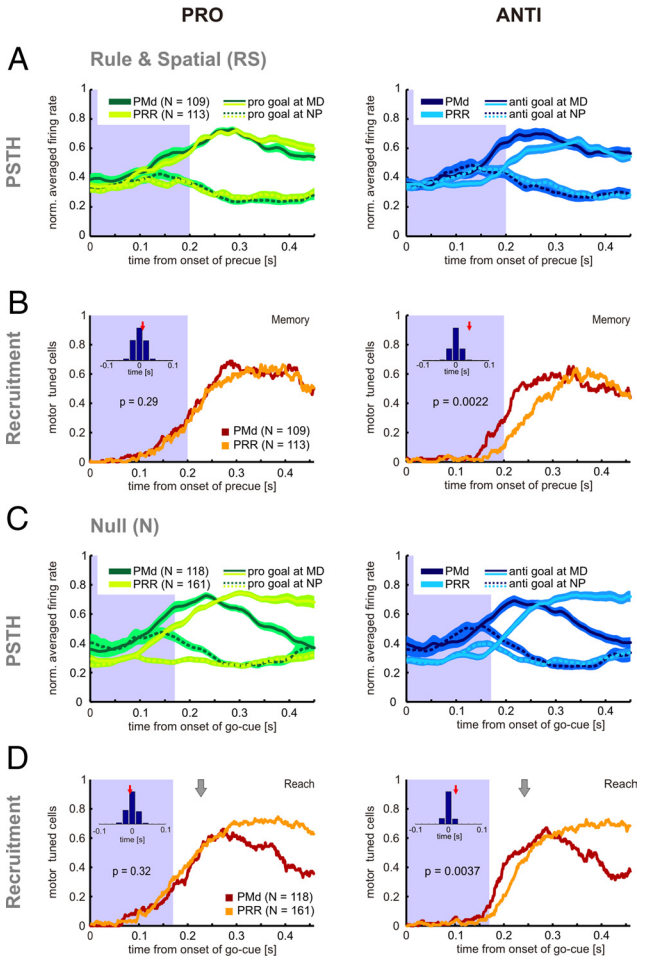

Figure 2. Dynamics of motor-related tuning in PMd and PRR in memory and reaction-time conditions. ***A***, ***B***, Data for pro-trials (left column, green) and anti-trials (right column, blue) aligned to the onset of the precue in the RS condition (memory). Shaded areas indicate the time of cue visibility. ***C***, ***D***, Equivalent data aligned to the onset of the go-cue in the N condition (reaction time). Gray arrows in ***D*** indicate mean RT across 176 recording sessions. ***A***, ***C***, PSTHs: Mean ± SE (shaded) of PMd (dark colors) and PRR (light colors) normalized population activity in MD (solid) and NP (dotted). ***B***, ***D***, Recruitment curves: fraction of neurons tuned in PMd (red) and PRR (yellow) in each time bin relative to the number of neurons tuned for the pro- or anti-goal in any time bin of the analysis time window (numbers in legend). Latencies in each condition were determined via a threshold criterion. Frontoparietal LDs between PMd and PRR were tested with a randomization test (see Materials and Methods). The histogram insets show the distribution of LDs with randomized data compared with the LD of the original sample (red arrow). The *p* values indicate the significance of the LDs. PMd showed earlier motor-related tuning than PRR in anti-reaches but not in pro-reaches, both in the RS and N condition.

responses (Fig. 1*C*,*D*) and to extract motor-goal latencies separately in each context (pro and anti). For analyses of motor-goal latencies separately within each context we had to restrict the neurons to those that did not show cue-related responses (see Materials and Methods). In PMd, 175 of 192 neurons (91%) fulfilled this criterion. In PRR, this was the case for 215 of 258 neurons (83%, RS condition). Because of this necessary constraint, our conclusions—strictly speaking—are limited to neurons without a cue-related response. However, since the fraction of neurons with a cue-related response is small, and only a subpopulation of these would qualify for our analyses because of additional selection criteria, we do not expect that these few neurons would change the results in any significant way.

Figure 2 shows the population PSTHs and fraction of significantly motor-tuned neurons (recruitment curves) in PMd and PRR for pro- and anti-reaches in the RS condition (Fig. 2*A*,*B*)

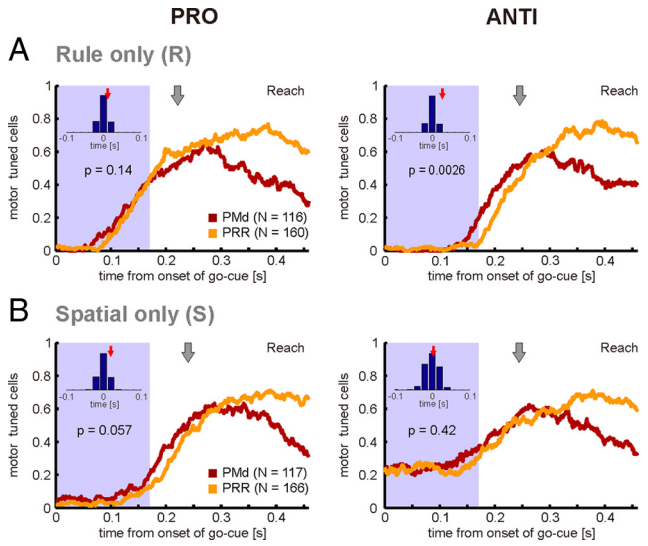

Figure 3. Dynamics of motor-related tuning in conditions with partial precueing. All the conventions for the recruitment curves are as in Figure 2*D*. ***A***, Recruitment curves for the R condition. ***B***, Recruitment curves for the S condition. Conditions that required spatial remapping (R-anti, S-pro) showed earlier motor-related tuning in PMd than PRR, whereas conditions without remapping (R-pro, S-anti) showed no LDs. See Results and Figure 5 for details on why S-pro, and not S-anti, required spatial remapping from the monkeys. Note that the LD in S-pro was significant when selecting neurons according to a stricter criterion (supplemental Fig. S1, available at www.jneurosci.org as supplemental material).

and N condition (Fig. 2*C*,*D*). In the pro-trials of the RS condition, the population PSTHs in PMd and PRR were practically identical (Fig. 2*A*, left). In the anti-trials, on the other hand, the average PRR response to the MD had a higher latency than the PMd response (Fig. 2*A*, right). Correspondingly, there was no significant LD between PMd and PRR for pro-reaches ($p > 0.05$, randomization test), but there was for anti-reaches ($LD_{PRR—PMd} = 41$ ms, $p = 0.0022$) in the memory period of the RS condition (Fig. 2*B*). In the movement period of the N condition, the results look similar (for a more in-depth discussion of potential differences, see supplemental Fig. S4, available at www.jneurosci.org as supplemental material). There was no LD between PMd and PRR for pro-reaches, whereas in anti-reaches the motor tuning emerged earlier in PMd than in PRR ($LD_{PRR—PMd} = 22$ ms, $p = 0.0037$). The fact, that the LD is present at the beginning of motor-goal memory trials (RS condition), at a time when the corresponding motor command will be issued at earliest ~1–1.5 s later, suggests that the LD is related to the decision process about the motor goal or movement planning, rather than to movement execution.

Spatial remapping induces differences in motor-goal latencies between PMd and PRR

In the N as well as the RS condition, spatial and rule information were presented together, and in both conditions motor-goal tuning was earlier in PMd than PRR, but only for anti-reaches. What factors determine the time course of motor-goal decisions in the frontoparietal network? Is it possible to influence the LDs between PMd and PRR by changing the temporal sequence in which spatial and rule information are provided? We tested the influence of presenting rule information before spatial information (R condition) and of presenting spatial information before rule information (S condition). The results are shown in Figure 3. In the R condition (Fig. 3*A*), there was no frontoparietal LD for pro-goals ($p > 0.05$), but for anti-goals motor tuning occurred earlier

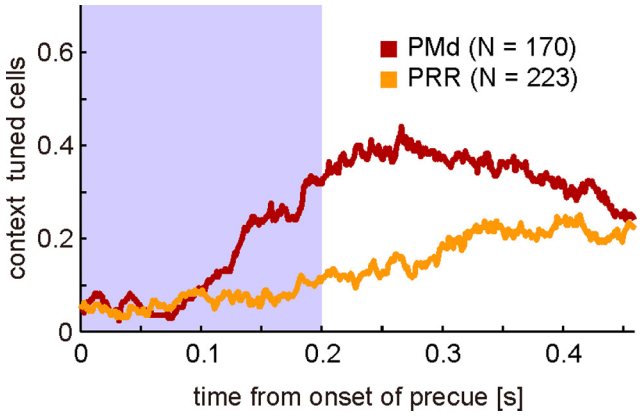

Figure 4. Dynamics of contextual tuning after precueing of the transformation rule. Recruitment curves show that the fraction of neurons with significantly different spike densities between pro- and anti-trials was higher in PMd (red) than PRR (yellow). This was the case especially early in the memory period after the R condition precue. Other conventions are the same as in Figure 2*D*.

in PMd than PRR ($LD_{PMd—PRR} = 27$ ms, $p = 0.0026$). In contrast, in the S condition (Fig. 3*B*) there was a trend for earlier motor tuning in PMd than PRR in pro-reaches. The LD was close to significant when we selected the neurons according to the standard criterion for spatial tuning, as defined in Materials and Methods ($LD_{PRR—PMd} = 19$ ms, $p = 0.057$) (Fig. 3*B*). It was significant ($LD_{PRR—PMd} = 25$ ms, $p = 0.026$) when neurons were selected according to a stricter tuning criterion, which we did as a control (supplemental Fig. S1, available at www.jneurosci.org as supplemental material). There was no significant LD in anti-reaches in the S condition ($LD_{PRR—PMd} = 3$ ms, $p > 0.05$) for either selection criterion (Fig. 3*B*; supplemental Fig. S1, available at www.jneurosci.org as supplemental material). Note that the higher baseline level in the S anti-trials can be explained by the spatial tuning during the S condition memory period, as will be discussed below.

The results of the RS, N, and R conditions, which showed LDs only for anti-goals but not for pro-goals, suggest that the fronto-parietal LDs could reflect the process of spatial remapping required in anti- but not pro-reaches. The LDs that we found in the S condition during pro- but not anti-reaches seem to be at odds with this remapping hypothesis. In the following, we will demonstrate that the seemingly counterintuitive results of the S condition are consistent with the idea that spatial remapping is the relevant factor for inducing motor-goal LDs between PRR and PMd. For this, we analyzed the neural encoding and putative cognitive strategy of the monkeys in the memory periods of the R and S conditions, first with respect to the context encoding, and second—and more importantly—with respect to the spatial encoding.

Figure 4 shows the comparative dynamics of context tuning in PMd and PRR. The context was represented earlier and more prevalently in PMd than PRR. Toward the end of the R condition memory period, the fraction of rule-tuned cells was only slightly higher in PMd than PRR, consistent with earlier reports (Gail et al., 2009). This finding does not imply any specific effect on the relative time courses of motor-goal encoding in PMd versus PRR after the spatial go-cue at the end of the R condition memory period.

The situation is different in the S condition. The S condition leaves some ambiguity as to what the animals memorize during the memory period of these trials: (1) a memory of the spatial precue; (2) the two potential motor goals (pro/anti), which are

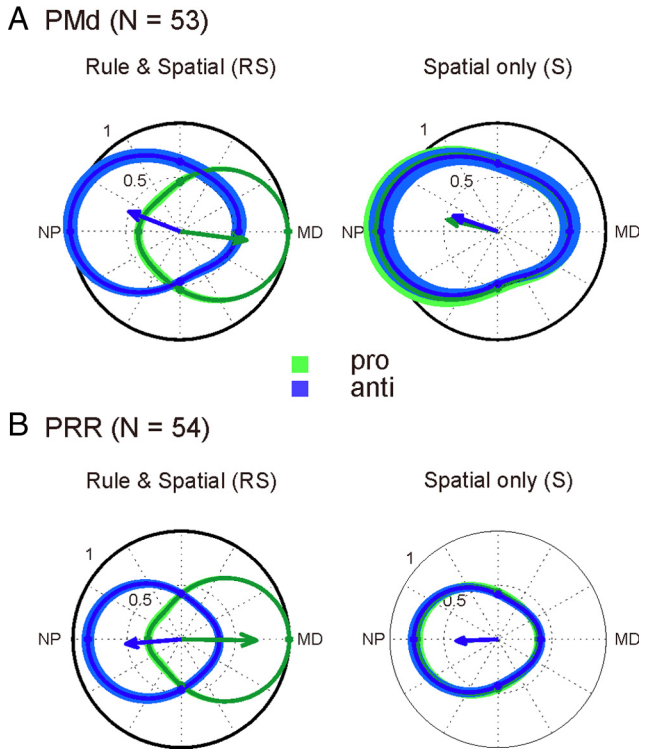

Figure 5. Predominant spatial tuning preference for the anti-goal in the memory period of the S condition. ***A***, ***B***, The normalized population tuning in pro-trials (green) and anti-trials (blue) is shown relative to the preferred spatial cue position in PMd (***A***) and PRR (***B***). Tuning curves of each neuron were aligned such that their MD in pro-reaches corresponded to 0° before averaging across neurons. The population tuning in the RS memory periods (left) confirms the typical expected motor-goal tuning in PMd and PRR. The population tuning in the S conditions (right) indicates preference for the anti-goal. Note that the memory period of the S condition precedes the go-cue. Since the transformation rule (pro or anti) could not be known by the monkeys, the tuning curves for pro- and anti-trials are practically identical. Only motor-goal neurons were used for this analysis (see Materials and Methods).

associated with each spatial cue as long as the context is not known; and (3) a preliminary default movement plan to either the pro- or anti-goal. If the monkeys memorized the spatial precue or by default planned a pro-reach during the memory period of the S condition, then a spatial remapping would have been required after the contextual go-cue during anti- but not pro-reaches, as was the case in the other cueing conditions. If, on the other hand, the monkeys by default planned an anti-reach during the memory period of the S condition, then a spatial remapping would have been required after the contextual go-cue during pro- but not anti-reaches. To test whether such a need for spatial remapping during S condition pro-reaches can explain the observed LDs in this condition, we analyzed the population tuning during the memory period of the S condition before the contextual go-cue was shown. Only the population of motor-goal neurons (see Materials and Methods) was used for this analysis, because for other neurons we could not expect a sustained response during the memory period. In PMd, two of the motor-goal neurons dropped out of the analysis because of a lack of activity at the end of the S condition memory period (spike density at MD, <1 Hz).

Importantly, and against intuition, in both monkeys neurons in PMd and PRR in this experiment almost exclusively encoded the anti-goal during the memory period of the S condition. Figure 5 shows the population tuning of motor-goal neurons for PMd and PRR in the last 200 ms of the memory period in the RS

and the S conditions (the baseline period for the recruitment curves). The time window was chosen to analyze the activity in the S condition immediately before the contextual information was given to the monkey. The tuning curves of all motor-goal neurons were aligned and normalized with respect to their maximal response (plotted at 0°), which was calculated as the average spike density in the 200 ms across pro-trials in the RS condition. The motor-goal tuning in the RS condition was taken as a reference to be able to differentiate in the S condition between tuning to the pro- and to the anti-goal. Neurons in PMd and PRR in the RS condition (Fig. 5*A*,*B*, left column) on average show the typical opposing tuning curves between pro- and anti-trials, reflecting their motor-goal encoding (Fig. 1*C*). The right columns of Figure 5, *A* and *B*, show that the same neurons in the memory period of the S condition are spatially tuned toward the anti-reach direction. We interpret this tuning as a preliminary default movement plan of both monkeys to the anti-goal. It suggests that after the contextual go-cue in the S condition a remapping is indeed required in pro- but not in anti-reaches, consistent with the neural LD data above.

Figure 6, *A* and *B*, summarizes the latency results for PMd and PRR in all conditions. There was a significant difference between PMd and PRR motor-related latencies in only those conditions where a spatial remapping was required (RS-anti, N-anti, R-anti, S-pro). There was no LD in conditions without remapping (RS-pro, N-pro, R-pro, S-anti).

Motor planning versus motor feedback in reaction-time trials

Sensorimotor areas, especially in the posterior parietal cortex, are activated by movement-induced somatosensory or visual feedback (Sakata et al., 1973; Mountcastle et al., 1975; Kalaska et al., 1983; Colby and Duhamel, 1991), not only by sensory cues or motor planning activity. The motor-goal LDs between PMd and PRR in the three RT conditions (N, R, and S) could be confounded by such motor feedback, including corollary discharge signals, if, for example, PRR was more strongly driven by movement-induced feedback signals, whereas PMd was more strongly driven by motor planning signals. To test whether during movement initiation in RT conditions the LDs between PMd and PRR were related to motor planning rather than motor feedback, we split our data set into motor-goal neurons (N_PMd/PRR = 65/75) and a complementary group of perimovement neurons (N_PMd/PRR = 127/183). Motor-goal neurons were characterized by motor-related tuning during the memory period of the RS condition and perimovement neurons by motor-related tuning in the movement period, but not during the memory period (see Materials and Methods). Figure 6*C* shows the LDs between PRR and PMd averaged across all RT conditions without spatial remapping (N-pro, R-pro, S-anti) and with spatial remapping (N-anti, R-anti, S-pro). The average LDs were computed separately for the population of all neurons, the population of motor-goal neurons, and the population of perimovement neurons. In the conditions without remapping, there were no significant LDs for either group. In the conditions with remapping, there were significant LDs between PMd and PRR for the population of all neurons (23 ms, $p = 0.017$) and the motor-goal neurons (21 ms, $p = 0.03$). For the perimovement neurons, there was only a nonsignificant trend (21 ms, $p = 0.084$), despite the fact that the number of eligible neurons in this category was higher than that for the motor-goal neurons. The absolute latencies for the motor-goal tuning were shorter in the group of motor-goal neurons than the perimovement neurons (average across all three RT conditions and both contexts: $\Delta = 30$ ms; $p = 0.00049$) (Fig. 6*D*). This difference in absolute latency confirmed that the splitting of the two groups of neurons according to their tuning properties in the RS condition memory period was meaningful with respect to the differentiation between motor-goal and motor-feedback latencies.

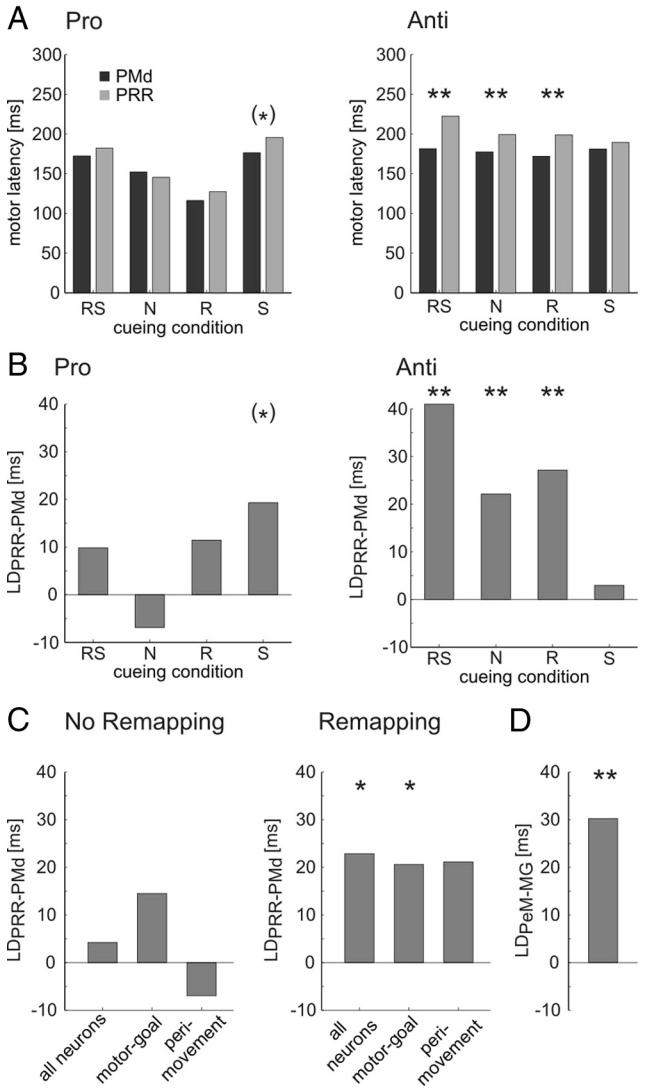

Figure 6. Summary of frontoparietal latencies and split analysis for different neuron populations. ***A***, ***B***, Summary of average absolute latencies (***A***) and LDs (***B***) between PRR and PMd for motor-related tuning for pro-trials (left) and anti-trials (right) for all cueing conditions (RS, N, R, S). ***C***, Comparison of remapping conditions with other conditions. LDs between PRR and PMd (PRR − PMd) averaged across all RT conditions with no required remapping (N-pro, R-pro, S-anti) and with required remapping (N-anti, R-anti, S-pro) for either all neurons or separately for motor-goal (MG) and perimovement (PeM) neurons. Asterisks indicate the level of significance (*$p < 0.05$; **$p < 0.01$, randomization test). See Results and supplemental material (available at www.jneurosci.org as supplemental material) for the significance of the S-pro condition. Only remapping conditions induce significant frontoparietal LDs. These differences are visible mainly in motor-goal neurons, less so in perimovement neurons. ***D***, Average tuning latencies are smaller in motor-goal compared with perimovement neurons. The bar shows LDs (PeM − MG) across all RT conditions and both cortical areas.

Effects of remapping on motor-goal latencies within PMd and PRR

We also compared motor-goal latencies between pro- and anti-reaches within each cortical area, rather than latencies between cortical areas (Fig. 7). Note that these are the same data as in the above analyses, just rearranged for different statistical comparisons. Pro-goal tuning emerged earlier than anti-goal tuning in

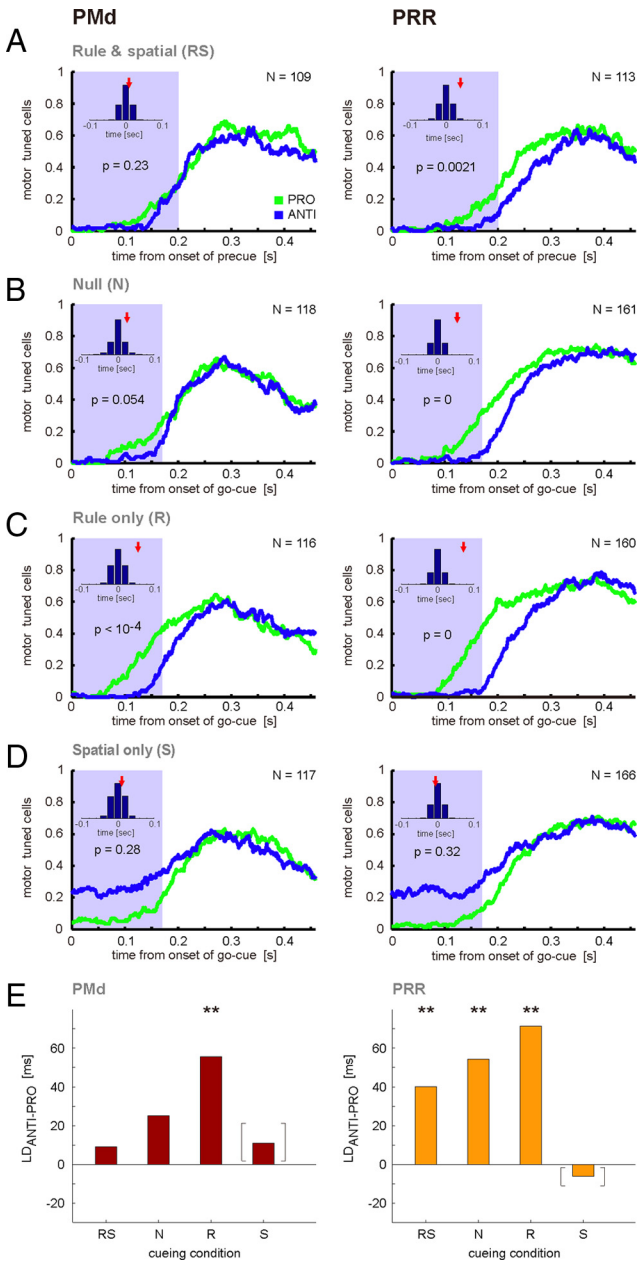

Figure 7. Comparison of motor-tuning latencies between pro- and anti-trials. ***A–D***, Recruitment curves for all four cueing conditions in pro-trials (green) and anti-trials (blue), separately for PMd (left) and PRR (right). Other conventions are as in Figures 2 and 3. Note that the data are the same as in Figures 2 and 3, but are grouped differently to emphasize and quantify the differences between pro- and anti-trials rather than PMd and PRR. ***E***, Summary of LDs between pro- and anti-trials (anti − pro). Brackets for the S condition indicate that the conclusiveness of LDs is questionable because of the different baseline levels in this condition.

PRR whenever the motor-goal decision was immediately preceded by a spatial cue: pro–anti-LDs were significant after the precue in the RS condition ($LD_{A\_P} = 40$ ms, $p = 0.0021$), after the go-cue in the N condition ($LD_{A\_P} = 54$ ms, $p = 0$), and after the go-cue in the R condition ($LD_{A\_P} = 71$ ms, $p = 0$). There was no pro–anti-LD in response to the go-cue in the S condition ($p = 0.32$). In PMd, we found faster representation of pro-goals compared with anti-goals only in response to the go-cue in the R condition ($LD_{P\_A} = 56$ ms, $p < 10^{-4}$). In PMd, none of the other conditions showed a significant difference between pro- and anti-goals (RS: $p = 0.27$; N: $p = 0.054$; S: $p = 0.42$). Note that for the S condition in either area the LD analysis between pro- and anti-trials is confounded by the default anti-tuning in the preceding memory period, which makes interpretation in this case difficult.

Motor-goal latencies and behavioral reaction times

If motor-goal latencies in PMd and PRR reflect decisions on arm movements, a correlation between motor-goal latencies and RTs should be expected. Figure 8*A* shows the median RTs in pro-reaches (light gray) and anti-reaches (dark gray) for the different cueing conditions averaged across all recording sessions ($N = 176$). The RTs for pro-reaches were 240 ± 2 ms (RS), 227 ± 2 ms (N), 221 ± 2 ms (R), and 240 ± 2 ms (S). The RTs in anti-reaches were 246 ± 2 (RS), 242 ± 2 ms (N), 245 ± 2 ms (R), and 244 ± 2 ms (S). The RT differences between pro- and anti-reaches were significant in each cueing condition (RS: 6 ± 1 ms, $p < 10^{-6}$; N: 15 ± 1 ms, $p < 10^{-6}$; R: 24 ± 1 ms, $p < 10^{-6}$; S: 4 ± 1 ms; $p = 4.6\ 10^{-5}$; data not shown). The RT differences between pro- and anti-reaches in the RS condition cannot be explained by motor-goal latencies, since motor-goal tuning in these conditions was already present during the memory period. To compensate for RT differences between pro- and anti-trials for the nonspecific effects of SR congruency observed in the RS condition, we also computed RT differences that were corrected for this nonspecific effect by subtracting the RT difference in the RS condition from the original RT differences (Fig. 8*B*). The corrected RT differences were 9 ± 1 ms (N condition, $p < 10^{-6}$), 18 ± 1 ms (R, $p < 10^{-6}$), and −2 ± 1 ms (S, $p = 0.015$).

We tested whether the neuronal motor-goal latencies better correlated with the time of the cue presentation or with the behavioral response onset (reaction time). The top left panel of Figure 8*C* shows as an example the slopes of the linear regression between single-trial RTs and neuronal onset latencies (relative to cue presentation; see Materials and Methods) for the motor-goal neurons in PRR. In a complementary manner, the slopes of the regression between (negative) single-trial RTs and neuronal lead times (onset of neuronal response relative to movement onset) were also analyzed (Fig. 8*C*, top right). Note that these two ways of quantifying the relationship between neural latencies and behavioral responses are not mutually independent, but help to illustrate the results. For the regression analysis, we only included trials of the RT conditions (N, R, and S) from neurons that did not show cue responses (see Materials and Methods). A slope close to 1 in the comparison between RTs and neuronal onset latencies (top left) indicates that the onset of activity in PRR motor-goal neurons was correlated to the movement onset. A slope close to 1 in the comparison between RTs and neuronal lead times (top right) would have indicated that the onset of activity was related to the cue onset. Overlap of the regression lines with the upper gray-shaded area indicates neural latencies larger than the RT, which only occurred in a small fraction of the motor-goal neurons. The pattern of slopes showed better correlation of the neural latencies with the movement onset compared with the cue onset for PRR motor-goal neurons, and also for perimovement neurons (Fig. 8*C*, bottom right). For motor-goal neurons in PRR, the median slope for relatedness to RT was 0.76, and for relatedness to cue onset 0.14 ($p = 0.00026$, rank sum test). For perimovement neurons in PRR, the median slope for RT was 0.62, and for cue onset 0.20 ($p < 10^{-5}$, rank sum test). In PMd, we could not find a significant bias in either neural population. For motor-goal neurons in PMd, the median slope for RT was 0.60, and for cue onset 0.25 ($p > 0.05$). For perimovement neurons, the median slope for RT was 0.47 and for cue onset 0.39 ($p >$

0.05). The distribution of slopes for both neural populations was unimodal and did not suggest that the moderate correlation with either cue onset or RT is the result of an averaging effect across an inhomogeneous distribution of neurons (data not shown). In summary, in PRR neural motor-goal latencies were best correlated with the time of the behavioral response; in PMd, there was no bias, which would indicate stronger correlation of the neuronal latencies with either the cue onset or the reaction time.

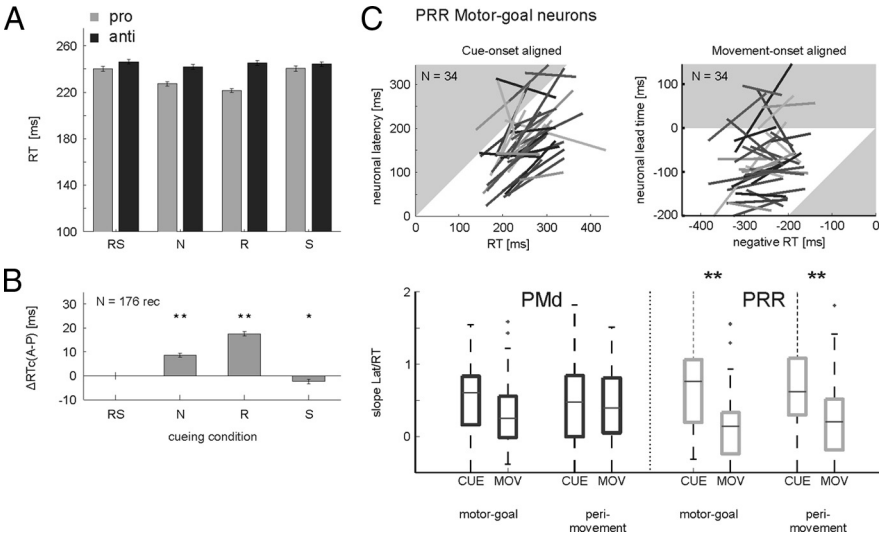

Figure 8. Correlation between neuronal latencies and behavioral reaction times. ***A***, Average RTs in pro- (light) and anti-trials (dark) in the four different cueing conditions. ***B***, RT differences between anti- and pro-reaches corrected for the nonspecific difference in the RS condition (see Results). ***C***, Regression analysis. Regression lines between neuronal latencies and RTs for each motor-goal neuron in PRR when aligned to the go-cue onset (top left). The equivalent regression data are shown (top right) between the "lead time" of neural tuning (= time of tuning onset before movement onset) and negative RTs (= time of cue onset before movement onset). Shaded areas indicate neuronal tuning latencies larger than RT (upper shaded areas in upper left and upper right) or neuronal tuning before cue onset (lower right shaded area in upper right). The slopes close to the unity line (upper left) and close to zero (upper right) indicate correlations of the neural latency with the movement onset, rather than with the cue onset. The boxplots in the lower panel shows the statistical analysis of the slope distributions for cue-onset alignment (CUE) and movement alignment (MOV) in both areas separately for motor-goal and perimovement neurons. In both groups of neurons, neural latencies in PRR but not PMd were correlated with behavioral response latencies.

Statistical controls

Latency measures depend on methodological procedures. For example, any statistical significance level (here, $\alpha = 5\%$) for motor-related tuning (significant spatial tuning in pro- and anti-reaches with opposite tuning vectors) marks an arbitrary threshold criterion. To avoid confounds caused by potential differences of the effect size in PMd and PRR, we performed our analysis using different threshold criteria for spatial tuning (supplemental Fig. S1, available at www.jneurosci.org as supplemental material). Also, we balanced the number of neurons between PMd and PRR by random subsampling (supplemental Fig. S2, available at www.jneurosci.org as supplemental material). Finally, we used alternative latency measures based on cumulative sums, as were used in previous studies from other groups (supplemental Fig. S3, available at www.jneurosci.org as supplemental material). None of the control measurements or alternative methods led to conclusions different from those presented in the main text.

Discussion

We compared motor-goal latencies in PMd and PRR during a nonstandard visuomotor association task. Neurons in PMd were tuned earlier for the pending motor goal than neurons in PRR, whenever the task required spatial visuomotor remapping. Frontoparietal LDs existed for motor goals during movement planning, independent of movement execution. We interpret our findings in support of the hypothesis that spatial reorganization of network activity in PRR causes frontoparietal LDs. We suggest that the initiation of the remapping process in PRR is contingent on frontoparietal projections from PMd.

Shorter motor-goal latencies in PMd than PRR in remapping conditions

We expected the relative timing of motor-goal decision processes between reach planning areas to change with precueing, as speculated previously (Kalaska and Crammond, 1992; Cisek, 2006; 2007). Against our expectation, PMd–PRR latency differences depended on the required visuospatial mapping instead. Variation of the precue did not change this basic dependency. Our results showed earlier motor-goal tuning in PMd than PRR whenever the reach goal had to be spatially remapped from the position of the preceding incongruent spatial cue (anti-trials in RS, N, and R conditions) or the incongruent preliminary movement plan (pro-trials in S condition).

Different reasons could account for the dependency of frontoparietal LDs on spatial remapping. It was previously suggested that frontoparietal interactions and corresponding LDs are the consequences of a decision-making process (Pesaran et al., 2008). In line with this view, one could emphasize the inherent difference in cognitive strategy required to perform either a pro-reach or an anti-reach task. To be consistent with our data, we would have to assume that all conditions that include a remapping require an active decision, whereas conditions without remapping do not. The function of such a decision could be the active updating of the preliminary motor goal, which was induced by the spatial cue or preliminary movement planning.

Alternatively, PRR motor-goal representations could depend on frontoparietal projections from PMd in any case, while becoming visible as interareal LDs only under specific circumstances. LDs of several tens of milliseconds, as seen here and also in previous studies (Riehle, 1991; Kalaska and Crammond, 1992; Johnson et al., 1996; Pesaran et al., 2008), seem rather long to be attributable to transmission delays between the monosynaptically connected areas PMd and PRR (Pandya and Kuypers, 1969; Jones and Powell, 1970; Kurata, 1991; Johnson et al., 1996; Caminiti et al., 1999; Marconi et al., 2001; Tanné-Gariépy et al., 2002; Battaglia-Mayer et al., 2003). Instead, we prefer to attribute the observed large frontoparietal LDs to the dynamic reorganization of network activity, which is required in visuospatially organized areas like PRR in the case of spatial remapping. According to this view, the dynamic reorganization is only started in PRR, once the motor goal has been selected in PMd. The dependency of PRR on PMd output would only become detectable as a frontoparietal LD in cases of spatially incongruent mapping, when an isolated spatial motor-goal representation is not directly induced by the spatial cue, but has to build up first. We predict that tasks with

nonspatial symbolic cueing could also induce frontoparietal LDs, similar to tasks with a spatially balanced distribution of cues (visual search task, Pesaran et al., 2008). What is common to all three tasks is the need for the buildup of a spatial goal representation in PRR that is not induced directly by a spatial cue. Spatially congruent mapping conditions, on the other hand, do not require dynamic remapping or buildup. The seamless and immediate transition from cue- to motor-goal-related tuning, even if its "initiation" in PRR might depend on PMd output, would render frontoparietal LDs invisible in this case.

The latter interpretation fits our and previous data, and explains long frontoparietal motor-goal LDs independently of interareal transmission delays. It is parsimonious since it requires neither an assumption that PMd-to-PRR projections become active only during decision tasks nor that remapping trials involve a decision process, whereas other trials do not. Note that we cannot exclude the alternative, but similar, possibility that both areas, PMd and PRR, receive a motor-goal decision signal from an independent third area, and that PRR in the case of remapping needs longer than PMd to converge onto a new dynamic state to represent the spatial motor goal.

Latency differences between PMd and PRR for motor planning activity

Previous physiological studies suggested that frontoparietal LDs between reach-related areas in premotor and parietal cortices exist for perimovement activity immediately preceding or accompanying motor execution (Bioulac and Lamarre, 1979; Georgopoulos et al., 1982; Kalaska et al., 1983; Riehle and Requin, 1989; Riehle, 1991; Johnson et al., 1996), in line with a previously stated corollary discharge hypothesis (Mountcastle et al., 1975; Bioulac and Lamarre, 1979; Kalaska et al., 1983; Seal and Commenges, 1985; Johnson et al., 1996). A previous study (Pesaran et al., 2008) showed that latency differences between PMd and PRR also exist for motor-goal representations during movement planning, i.e., independent of movement execution, whereas another study did not show such differences (Johnson et al., 1996). The fact that frontoparietal LDs depend on the type of task rather than the movement execution can explain these different results. We found shorter latencies in PMd than PRR for motor-goal activity in response to a preparatory stimulus (RS condition) and in response to a go-instruction (N, R, and S conditions). Even in the RT conditions, we could attribute the frontoparietal LDs to differences in the latencies of motor-goal signals, rather than motor feedback signals. Our findings contradict a corollary discharge hypothesis in its strict sense (see alternative suggestion below), in which PMd–PRR latency differences would be the consequence of a frontoparietal efference copy signal, i.e., would be dependent on explicit motor commands (Kalaska et al., 1983; Johnson et al., 1996).

Motor-goal latencies in relation to behavioral reaction times

To discuss the neural dynamics of sensorimotor integration in the context of behavioral performance, it is important to distinguish between general neural response onset latencies (Kalaska et al., 1983; Riehle, 1991; Kalaska and Crammond, 1992; Johnson et al., 1996; Pesaran et al., 2008) and latencies of neural activities that are more specifically related to the cue, the motor-goal decision, or motor execution.

First, with the anti-reach task design we separated cue- from motor-related tuning (Crammond and Kalaska, 1994; Zhang et al., 1997; Gail and Andersen, 2006). Neurons in PMd (Crammond and Kalaska, 1994) and PRR (Gail and Andersen, 2006) can dynamically switch from cue- to motor-related encoding, whereas others are motor related as soon as they become active. Population analysis of neural onset latencies including all neurons would lead to confounded motor latencies. Especially, frontoparietal LDs could be obscured because of the simultaneity of the fastest responding neurons (Johnson et al., 1996), which respond to the cue. We avoided this confound by analyzing neurons without cue-related responses.

Second, using different partial precueing conditions (Lecas et al., 1986; Riehle and Requin, 1989; Riehle, 1991; Hoshi and Tanji, 2000; Nakayama et al., 2008), we additionally separated motor-goal latencies related to motor planning (Riehle, 1991; Johnson et al., 1996) from latencies of perimovement activity related to motor-execution (Riehle and Requin, 1989; Crammond and Kalaska, 2000). Perimovement activities could comprise movement-induced corollary discharge, somatosensory, or visual signals, especially in parietal areas (Sakata et al., 1973; Mountcastle et al., 1975; Kalaska et al., 1983; Colby and Duhamel, 1991; Nakayama et al., 2008). Average response latencies from mixed motor-goal and perimovement activity (Kalaska et al., 1983; Riehle and Requin, 1989; Johnson et al., 1996; Pesaran et al., 2008) would overestimate motor-goal latencies in areas with strong motor-feedback signals. We avoided this confound by differentiating motor-goal from perimovement neurons.

The trial-by-trial motor-goal latencies in PRR were better correlated with manual RTs than those of PMd. This could mean that manual responses are triggered only once the motor-goal decision (which becomes visible in PMd before PRR) leads to the required (and slower) evolution of the proper motor-goal representations in PRR.

Frontoparietal projections for prospective forward model estimates?

Parietal cortex function is important for proper online correction of movements (Desmurget et al., 1999) and is suspected to use forward model predictions for this purpose (Wolpert et al., 1998; Desmurget et al., 1999; Mulliken et al., 2008; Shadmehr and Krakauer, 2008). We suggest that such a forward estimate is generated not only during movements, but prospectively, before a motor command is issued, as an integral part of the planning process. In visuomotor tasks, movement planning then would equal the process of analyzing the difference between the current sensory state and the prospective sensory state, which is predicted to result from the currently imagined movement in the future, once it is executed (Salinas, 2004; Shadmehr and Wise, 2005). This view is consistent with our data and previous studies, which showed activation of PPC during motor imagery (Decety et al., 1994; Stephan et al., 1995; Sirigu et al., 1996; Gerardin et al., 2000; Naito et al., 2002), and the fact that parietal electrical stimulation can trigger subjective movement intentions without actual movements being performed (Desmurget et al., 2009).

References

Battaglia-Mayer A, Caminiti R, Lacquaniti F, Zago M (2003) Multiple levels of representation of reaching in the parieto-frontal network. Cereb Cortex 13:1009–1022.

Bioulac B, Lamarre Y (1979) Activity of postcentral cortical neurons of the monkey during conditioned movements of a deafferented limb. Brain Res 172:427–437.

Boussaoud D, Wise SP (1993) Primate frontal-cortex—effects of stimulus and movement. Exp Brain Res 95:28–40.

Brozović M, Gail A, Andersen RA (2007) Gain mechanisms for contextually guided visuomotor transformations. J Neurosci 27:10588–10596.

Buneo CA, Andersen RA (2006) The posterior parietal cortex: sensorimotor interface for the planning and online control of visually guided movements. Neuropsychologia 44:2594–2606.

Caminiti R, Genovesio A, Marconi B, Mayer AB, Onorati P, Ferraina S,

Mitsuda T, Giannetti S, Squatrito S, Maioli MG, Molinari M (1999) Early coding of reaching: frontal and parietal association connections of parieto-occipital cortex. Eur J Neurosci 11:3339–3345.
Cisek P (2006) Integrated neural processes for defining potential actions and deciding between them: a computational model. J Neurosci 26:9761–9770.
Cisek P (2007) Cortical mechanisms of action selection: the affordance competition hypothesis. Philos Trans R Soc Lond B Biol Sci 362:1585–1599.
Colby CL, Duhamel JR (1991) Heterogeneity of extrastriate visual areas and multiple parietal areas in the macaque monkey. Neuropsychologia 29:517–537.
Crammond DJ, Kalaska JF (1994) Modulation of preparatory neuronal activity in dorsal premotor cortex due to stimulus-response compatibility. J Neurophysiol 71:1281–1284.
Crammond DJ, Kalaska JF (2000) Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. J Neurophysiol 84:986–1005.
Decety J, Perani D, Jeannerod M, Bettinardi V, Tadary B, Woods R, Mazziotta JC, Fazio F (1994) Mapping motor representations with positron emission tomography. Nature 371:600–602.
Desmurget M, Epstein CM, Turner RS, Prablanc C, Alexander GE, Grafton ST (1999) Role of the posterior parietal cortex in updating reaching movements to a visual target. Nat Neurosci 2:563–567.
Desmurget M, Reilly KT, Richard N, Szathmari A, Mottolese C, Sirigu A (2009) Movement intention after parietal cortex stimulation in humans. Science 324:811–813.
di Pellegrino G, Wise SP (1993) Visuospatial versus visuomotor activity in the premotor and prefrontal cortex of a primate. J Neurosci 13:1227–1243.
Eliassen JC, Souza T, Sanes JN (2003) Experience-dependent activation patterns in human brain during visual-motor associative learning. J Neurosci 23:10540–10547.
Gail A, Andersen RA (2006) Neural dynamics in monkey parietal reach region reflect context-specific sensorimotor transformations. J Neurosci 26:9376–9384.
Gail A, Klaes C, Westendorff S (2009) Implementation of spatial transformation rules for goal-directed reaching via gain modulation in monkey parietal and premotor cortex. J Neurosci 29:9490–9499.
Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. J Neurosci 2:1527–1537.
Gerardin E, Sirigu A, Lehéricy S, Poline JB, Gaymard B, Marsault C, Agid Y, Le Bihan D (2000) Partially overlapping neural networks for real and imagined hand movements. Cereb Cortex 10:1093–1104.
Hoshi E, Tanji J (2000) Integration of target and body-part information in the premotor cortex when planning action. Nature 408:466–470.
Johnson PB, Ferraina S, Bianchi L, Caminiti R (1996) Cortical networks for visual reaching: physiological and anatomical organization of frontal and parietal lobe arm regions. Cereb Cortex 6:102–119.
Jones EG, Powell TP (1970) An anatomical study of converging sensory pathways within the cerebral cortex of the monkey. Brain 93:793–820.
Kalaska JF (1996) Parietal cortex area 5 and visuomotor behavior. Can J Physiol Pharmacol 74:483–498.
Kalaska JF, Crammond DJ (1992) Cerebral cortical mechanisms of reaching movements. Science 255:1517–1523.
Kalaska JF, Caminiti R, Georgopoulos AP (1983) Cortical mechanisms related to the direction of two-dimensional arm movements—relations in parietal area 5 and comparison with motor cortex. Exp Brain Res 51:247–260.
Kalaska JF, Scott SH, Cisek P, Sergio LE (1997) Cortical control of reaching movements. Curr Opin Neurobiol 7:849–859.
Kurata K (1991) Corticocortical inputs to the dorsal and ventral aspects of the premotor cortex of macaque monkeys. Neurosci Res 12:263–280.
Lecas JC, Requin J, Anger C, Vitton N (1986) Changes in neuronal activity of the monkey precentral cortex during preparation for movement. J Neurophysiol 56:1680–1702.
Marconi B, Genovesio A, Battaglia-Mayer A, Ferraina S, Squatrito S, Molinari M, Lacquaniti F, Caminiti R (2001) Eye-hand coordination during reaching. I. Anatomical relationships between parietal and frontal cortex. Cereb Cortex 11:513–527.
Mitz AR, Godschalk M, Wise SP (1991) Learning-dependent neuronal activity in the premotor cortex: activity during the acquisition of conditional motor associations. J Neurosci 11:1855–1872.
Monosov IE, Trageser JC, Thompson KG (2008) Measurements of simultaneously recorded spiking activity and local field potentials suggest that spatial selection emerges in the frontal eye field. Neuron 57:614–625.
Mountcastle VB, Lynch JC, Georgopoulos A, Sakata H, Acuna C (1975) Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. J Neurophysiol 38:871–908.
Mulliken GH, Musallam S, Andersen RA (2008) Decoding trajectories from posterior parietal cortex ensembles. J Neurosci 28:12913–12926.
Naito E, Kochiyama T, Kitada R, Nakamura S, Matsumura M, Yonekura Y, Sadato N (2002) Internally simulated movement sensations during motor imagery activate cortical motor areas and the cerebellum. J Neurosci 22:3683–3691.
Nakayama Y, Yamagata T, Tanji J, Hoshi E (2008) Transformation of a virtual action plan into a motor plan in the premotor cortex. J Neurosci 28:10287–10297.
Pandya DN, Kuypers HG (1969) Cortico-cortical connections in the rhesus monkey. Brain Res 13:13–36.
Passingham RE, Toni I, Rushworth MF (2000) Specialisation within the prefrontal cortex: the ventral prefrontal cortex and associative learning. Exp Brain Res 133:103–113.
Pesaran B, Nelson MJ, Andersen RA (2008) Free choice activates a decision circuit between frontal and parietal cortex. Nature 453:406–409.
Riehle A (1991) Visually induced signal-locked neuronal activity changes in precentral motor areas of the monkey: hierarchical progression of signal processing. Brain Res 540:131–137.
Riehle A, Requin J (1989) Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. J Neurophysiol 61:534–549.
Sakata H, Takaoka Y, Kawarasaki A, Shibutani H (1973) Somatosensory properties of neurons in the superior parietal cortex (area 5) of the rhesus monkey. Brain Res 64:85–102.
Salinas E (2004) Fast remapping of sensory stimuli onto motor actions on the basis of contextual modulation. J Neurosci 24:1113–1118.
Seal J, Commenges D (1985) A quantitative analysis of stimulus- and movement-related responses in the posterior parietal cortex of the monkey. Exp Brain Res 58:144–153.
Shadmehr R, Krakauer JW (2008) A computational neuroanatomy for motor control. Exp Brain Res 185:359–381.
Shadmehr R, Wise SP (2005) The computational neurobiology of reaching and pointing: a foundation for motor learning. Cambridge, MA: MIT.
Sirigu A, Duhamel JR, Cohen L, Pillon B, Dubois B, Agid Y (1996) The mental representation of hand movements after parietal cortex damage. Science 273:1564–1568.
Stephan KM, Fink GR, Passingham RE, Silbersweig D, Ceballos-Baumann AO, Frith CD, Frackowiak RS (1995) Functional anatomy of the mental representation of upper extremity movements in healthy subjects. J Neurophysiol 73:373–386.
Stoet G, Snyder LH (2004) Single neurons in posterior parietal cortex of monkeys encode cognitive set. Neuron 42:1003–1012.
Tanné-Gariépy J, Rouiller EM, Boussaoud D (2002) Parietal inputs to dorsal versus ventral premotor areas in the macaque monkey: evidence for largely segregated visuomotor pathways. Exp Brain Res 145:91–103.
Thompson KG, Hanes DP, Bichot NP, Schall JD (1996) Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. J Neurophysiol 76:4040–4055.
Wallis JD, Miller EK (2003) from rule to response: neuronal processes in the premotor and prefrontal cortex. J Neurophysiol 90:1790–1806.
Wise SP, di Pellegrino G, Boussaoud D (1996) The premotor cortex and nonstandard sensorimotor mapping. Can J Physiol Pharmacol 74:469–482.
Wise SP, Boussaoud D, Johnson PB, Caminiti R (1997) Premotor and parietal cortex: corticocortical connectivity and combinatorial computations. Annu Rev Neurosci 20:25–42.
Wolpert DM, Goodbody SJ, Husain M (1998) Maintaining internal representations: the role of the human superior parietal lobe. Nat Neurosci 1:529–533.
Zhang J, Riehle A, Requin J, Kornblum S (1997) Dynamics of single neuron activity in monkey primary motor cortex related to sensorimotor transformation. J Neurosci 17:2227–2246.

Supplemental Material

S1 – Control of latency results for effect size

Inhomogeneity of average spike rate or effect size, e.g. differences in the neural response selectivity, could bias latency measures between different cortical areas. First, since our recruitment curves focused on only the significant neurons, different firing rates in both area could lead to a higher number of significantly tuned neurons in the area with higher firing rates due to statistical power. We analyzed firing rates in PMd and PRR on average across those neurons, which contributed to the analyses in the main text, i.e. neurons with motor-related tuning and no cue-response. For each neuron the average firing rate was determined in the same time window as was used for computing the maximum direction (MD), which was 200 to 350 ms after onset of the cue (precue in RS, go-cue in N, R, S). Firing rates in PMd during anti-trials were not significantly different from those in PRR in either cueing condition. Second, the effect size in each area was quantified with the maximum p-value for directional selectivity within the analysis time window averaged across all neurons (The negative logarithm of the p value to the basis of ten was used for statistical comparisons; p values smaller than 10^{-6} were set to 10^{-6}). There was either no significant difference in effect size between PMd and PRR (RS-pro, R-anti, S-pro) or the effect size was significantly higher ($p<0.05$, t-test) in PRR compared to PMd (RS-anti, N-pro, N-anti, R-pro and S-anti). Only in RS-anti the effect size in PMd was higher than in PRR. If at all, higher effect size in PRR should lead to opposite latency differences between both areas in most cueing conditions than the ones we found. Taken together, neither differences in average firing rate nor in average tuning strength could explain the observed LDs between PMd and PRR.

In the standard analysis neurons were included if they showed motor-related tuning at some point in time during the analysis time window. The significance criterion of being tuned ($p<0.05$ in Kruskal-Wallis test) implies an artificial threshold. As additional control we ran our analysis with different thresholds for the significance of an individual neuron's tuning ($\alpha = 0.01$, 0.1, or 0.2). In the S-pro condition the effect of LD(PRR-PMd), which only showed a non-significant trend ($p = 0.057$) in the standard analysis with tuning threshold $\alpha = 0.05$ (Fig. 3B and 6A), was significant ($p = 0.026$) when neurons were selected with a stricter significance criterion for tuning ($\alpha = 0.01$). In all other conditions the significances did not change (Fig. S-1). These findings corroborate our conclusions of the main manuscript, especially they provide evidence for earlier motor-goal representations in PMd than PRR in the S-pro condition.

S 2 – Control of latency results for sampling bias

Similar to effect size (Supplemental Material 1) differences in the size of the neuron sample between the two areas could affect latency measures based on recruitment curves. We randomly subsampled the same number of neurons in PMd and PRR (Pesaran et al. 2008), to balance the number of neurons analyzed in each area (Fig. S-2). The distributions of LDs in the different task conditions obtained from repeated random sampling of neurons were not different from the original values ($p>0.4$, confidence interval), which means the results were independent of the specific neuron samples. In summary, section 1 and 2 of Supplementary Material suggest that population size, effect size, or average spike rate differences did not confound our latency measures.

S 3 – Alternative latency measures

As an alternative latency measurement we compared the cumulative sums over all single-unit motor-related tuning onset latencies (Fig. S-3). Within each cueing condition and for each neuron the onset latency was defined as the first time point at which the tuning became motor-related and continued to do so for at least 90% of 50 consecutive time bins (equivalent to the definition of motor planning activity in the memory period of the RS condition; see Material and Methods). The time window of analysis for each cueing condition was identical to the standard analysis in the main manuscript, which was -200 ms to + 450 ms relative to the onset of the pre-cue (RS) or go-cue (N, R, S). The number of neurons which fulfilled the criterion for motor-tuning is provided in the legend of Fig. S-3. Median onset latencies for motor-tuning were lower in PMd than PRR ($p<0.01$, Wilcoxon signed rank test) in all four cueing conditions which required spatial remapping (RS-anti, N-anti, R-anti, S-pro) These results confirm the results in the main text.

S 4 – Temporary bimodal tuning properties

Population PSTHs around the time of the go-cue in the N condition (Fig. 2C) suggest that for a brief period of time, between approx. 100-250 ms after cue onset, neural activity in PMd did not only increase for reaches to the maximum direction (MD), but also for reaches to the non-preferred direction (NP). As a consequence, the separation between MD and NP PSTH curves occurred later in PMd than PRR during pro-

reaches, while the recruitment curves and statistical tests indicated synchronous onset of motor-goal tuning in PMd and PRR. Here we explain this seeming contradiction by showing that in PMd in response to the go-cue in the N condition a small fraction of neurons temporarily showed bimodal tuning, i.e. they briefly reflect both potential motor goals, the pro- and the anti-goal.

To test the significance of directional tuning in our standard analysis in the main text we used a non-parametric 1-way ANOVA with the additional criterion that the length of the normalized DTV had to be larger than 0.2. With this criterion we excluded neurons with a bipolar tuning close to symmetric (ANOVA significant, $|DTV| < 0.2$) from the recruitment curves and statistical analysis, but not from the population PSTHs. Such symmetric bipolar tuning could represent two potential motor goals instead of one selected motor goal and would be characterized by a significant ANOVA while the DTV has a length close to zero (see Material and Methods). The test for bimodality was conducted separately in pro- and anti-trials for all neurons which did not have a cue-related response. Figures S-4 A and B show the time-resolved number of bimodal neurons in PMd and PRR for the RS and N conditions. Especially in the N condition in PMd there was a clear increase in the number of bimodal neurons between 100-250 ms after cue onset, most pronounced for pro-reaches. Figure S4 C shows the tuning between 100-250 ms for PMd neurons in the N-pro condition (normalized within each neuron and averaged across neurons), which showed bimodal tuning for at least 10 time bins in this time window. The fact, that the average tuning shows opposing local maxima at 0 and 180 deg is consistent with the interpretation that the tuning reflects the encoding of two potential motor goals.

We speculate that bimodal tuning was stronger in the N condition than in the RS condition because of the time pressure for the monkey in the N condition. It takes longer to process the rule information than the spatial information (compare RTs for R and S conditions in pro-reaches, Fig. 6A). Evaluation of the spatial cue before the rule is known narrows down the choices to two out of four remaining spatial targets, the pro- and the anti-goal. Once the rule information was processed, one of the two putative activity peaks had to be suppressed (resulting in the motor-goal tuning of which we analyzed the latencies with our recruitment curves). This strategy might have been faster for the monkey than building up a new activity peak after the delayed evaluation of the transformation rule. In the RS condition after the pre-cue, there was no time pressure for the monkey and therefore no need to build up potential motor-goal tuning.

Supplemental Figures

Figure S-1: Control for the effect of tuning significance thresholds on LD measures. Frontoparietal LDs are shown for pro- (top row) and anti-trials (bottom row) in the different cueing conditions. A-C: LD results when different significance thresholds were applied to the spatial tuning of each neuron in pro- and anti-reaches. A: LDs if significance of spatial tuning in single neurons and time windows was defined by $p<0.2$ (Kruskal-Wallis). B and C: Same as A but significance defined with $p < 0.1$ and 0.01, respectively.

Figure S-2: Control for the effect of sample sizes in PMd and PRR on LD measures. Mean and standard deviation (1000 randomizations) of frontoparietal LDs for pro- (left panel) and anti-trials (right panel) in the different cueing conditions. In each randomization run a random sub-sample of the same number of neurons in PRR and PMd were taken. Asterisks indicate the level of significance (*: $p<0.05$, **: $p<0.01$). The dots indicate the original value with unequal number of neurons in PMd and PRR. The original LDs were not different from the LDs derived from balanced sample sizes ($p >0.4$).

Figure S-3: Alternative neural latency measure. The curves show cumulative sums over the onset latencies of motor-related tuning. PMd (red) and PRR (yellow) data are shown separately for pro- (left) and anti-trials (right) in the different cueing conditions (A-D). The motor-tuning onset latencies are derived for each single neuron within each condition. Additional to the cumulative sums the median onset latency and the p-value (ranksum-test) for the comparison between PMd and PRR onset latency

distribution are provided. Also, the numbers of neurons in each area are provided, for which an onset latency could be computed.

Figure S-4: Analysis of temporary bimodal tuning after cue presentation. A: Number of putatively bimodal neurons in PMd (red) and PRR (yellow) and separate for pro- (left) and anti-trials (right) in the RS-condition after onset of the precue. B: Same as in A, but in the N condition after onset of the go-cue. C: Average normalized tuning of putatively bimodal PMd neurons in pro-trials in the N condition (the condition with the strongest indication for bimodal tuning). The tuning curve is bi-lobed, indicating bimodal tuning, even though not fully symmetrically. The time window, which was used for computing the directional tuning, is indicated by the black bar in B. Data are aligned and normalized to the average maximal response in pro-trials in the time window of analysis.

Figure S-1

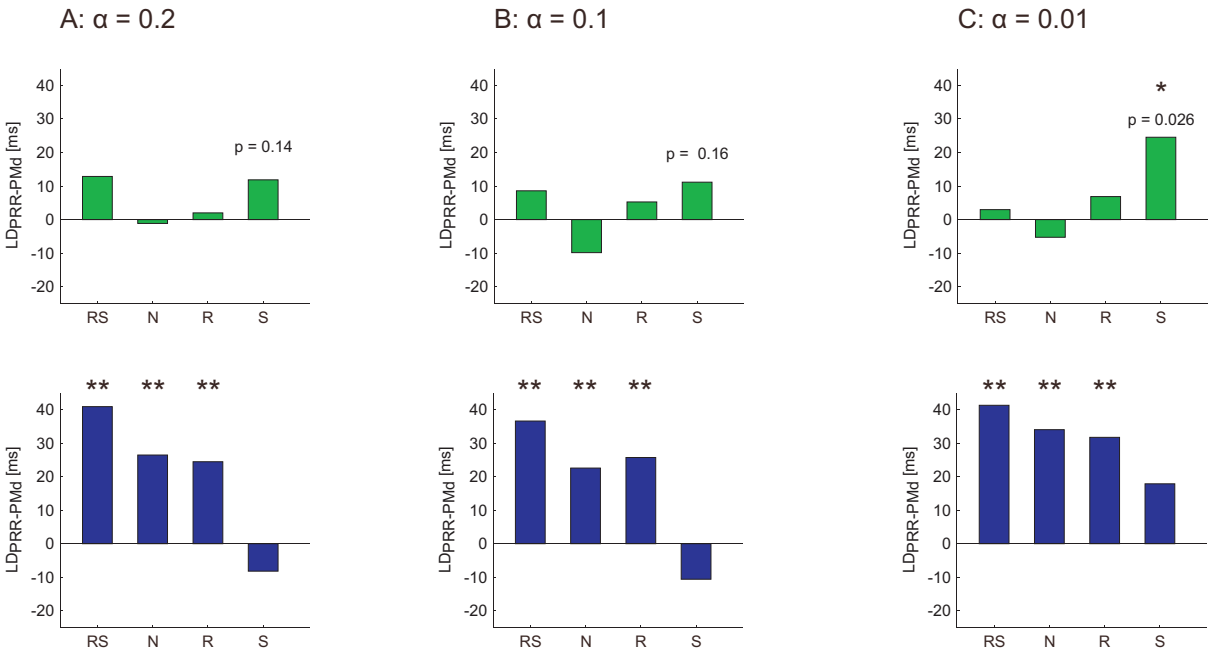

Figure S-1: Control for the effect of tuning significance thresholds on LD measures. Frontoparietal LDs are shown for pro- (top row) and anti-trials (bottom row) in the different cueing conditions. A-C: LD results when different significance thresholds were applied to the spatial tuning of each neuron in pro- and anti-reaches. A: LDs if significance of spatial tuning in single neurons and time windows was defined by $p<0.2$ (Kruskal-Wallis). B and C: Same as A but significance defined with $p < 0.1$ and 0.01, respectively.

Figure S-2

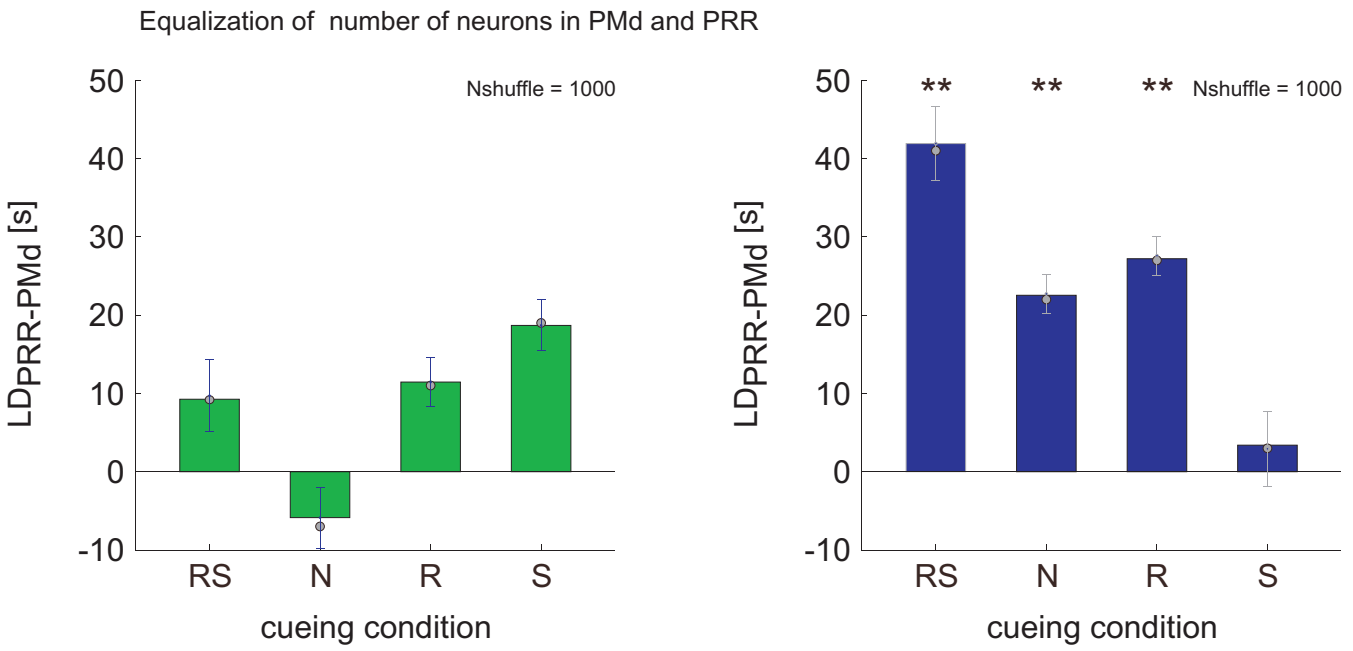

Figure S-2: Control for the effect of sample sizes in PMd and PRR on LD measures. Mean and standard deviation (1000 randomizations) of frontoparietal LDs for pro- (left panel) and anti-trials (right panel) in the different cueing conditions. In each randomization run a random sub-sample of the same number of neurons in PRR and PMd were taken. Asterisks indicate the level of significance (*: $p<0.05$, **: $p<0.01$). The dots indicate the original value with unequal number of neurons in PMd and PRR. The original LDs were not different from the LDs derived from balanced sample sizes ($p>0.4$).

Figure S-3

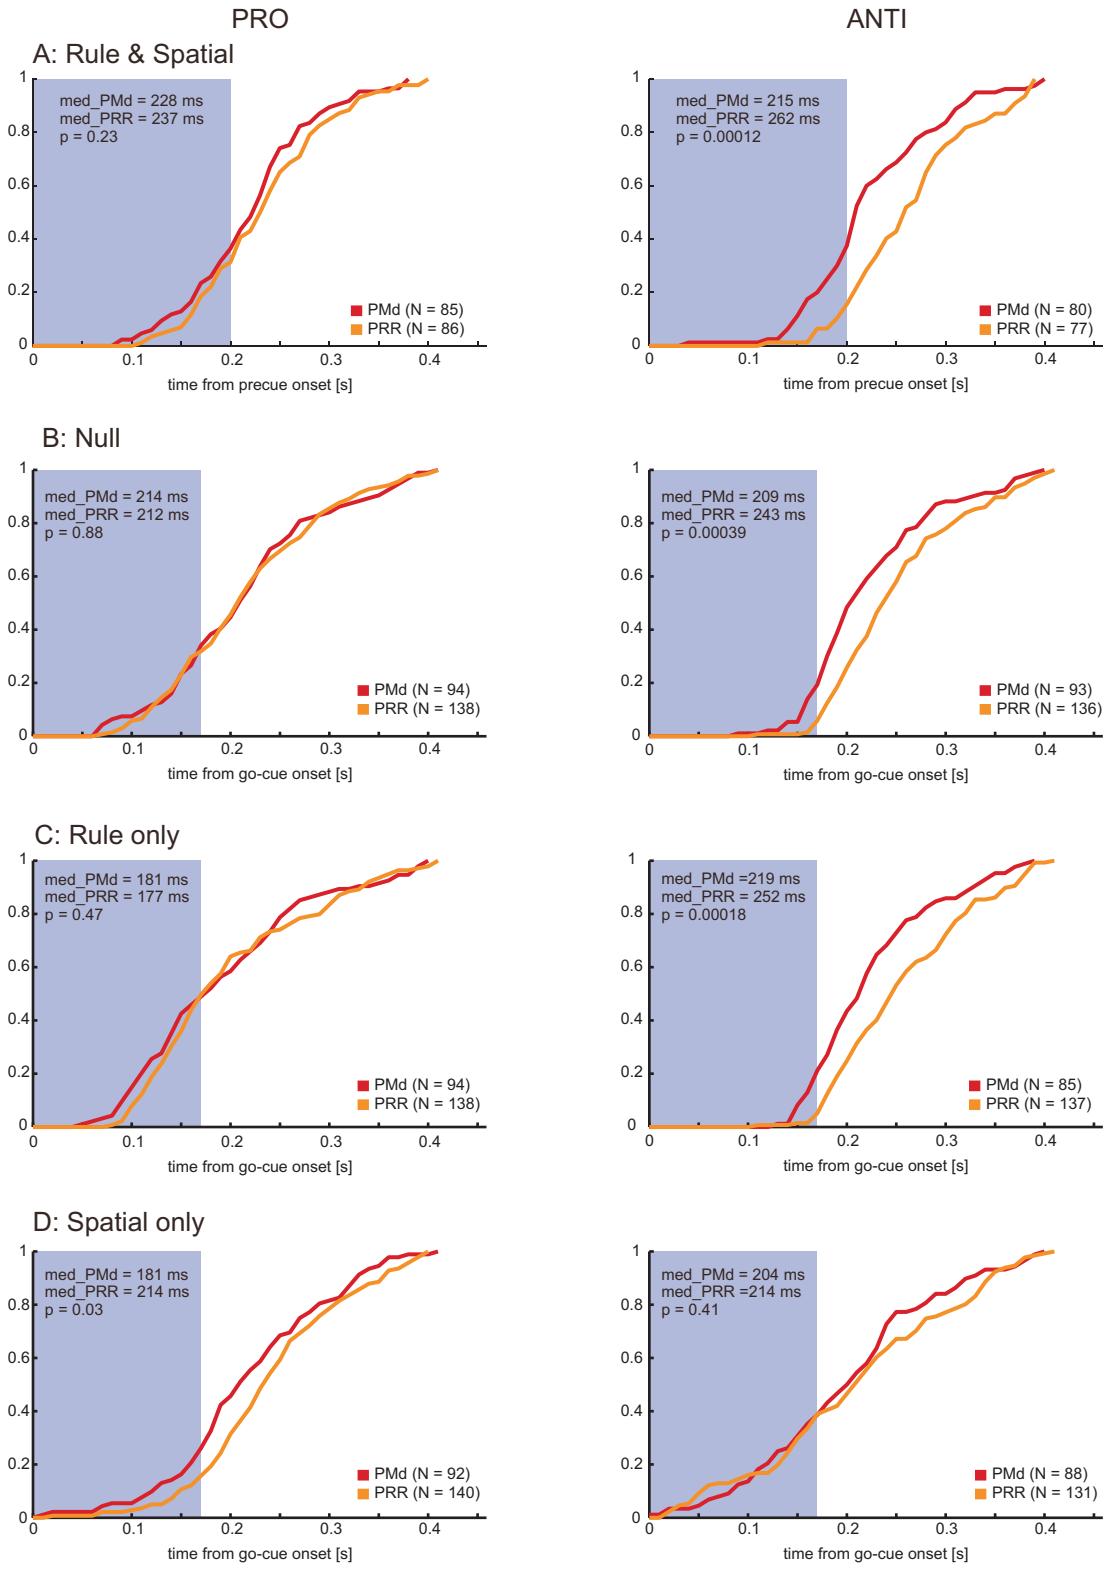

Figure S-3: Alternative neural latency measure. The curves show cumulative sums over the onset latencies of motor-related tuning. PMd (dark) and PRR (light) data are shown separately for pro- (left) and anti-trials (right) in the different cueing conditions (A-D). The motor-tuning onset latencies are derived for each single neuron within each condition. Additional to the cumulative sums the median onset latency and the p-value (ranksum-test) for the comparison between PMd and PRR onset latency distribution are provided. Also, the numbers of neurons in each area are provided, for which an onset latency could be computed.

Figure S-4

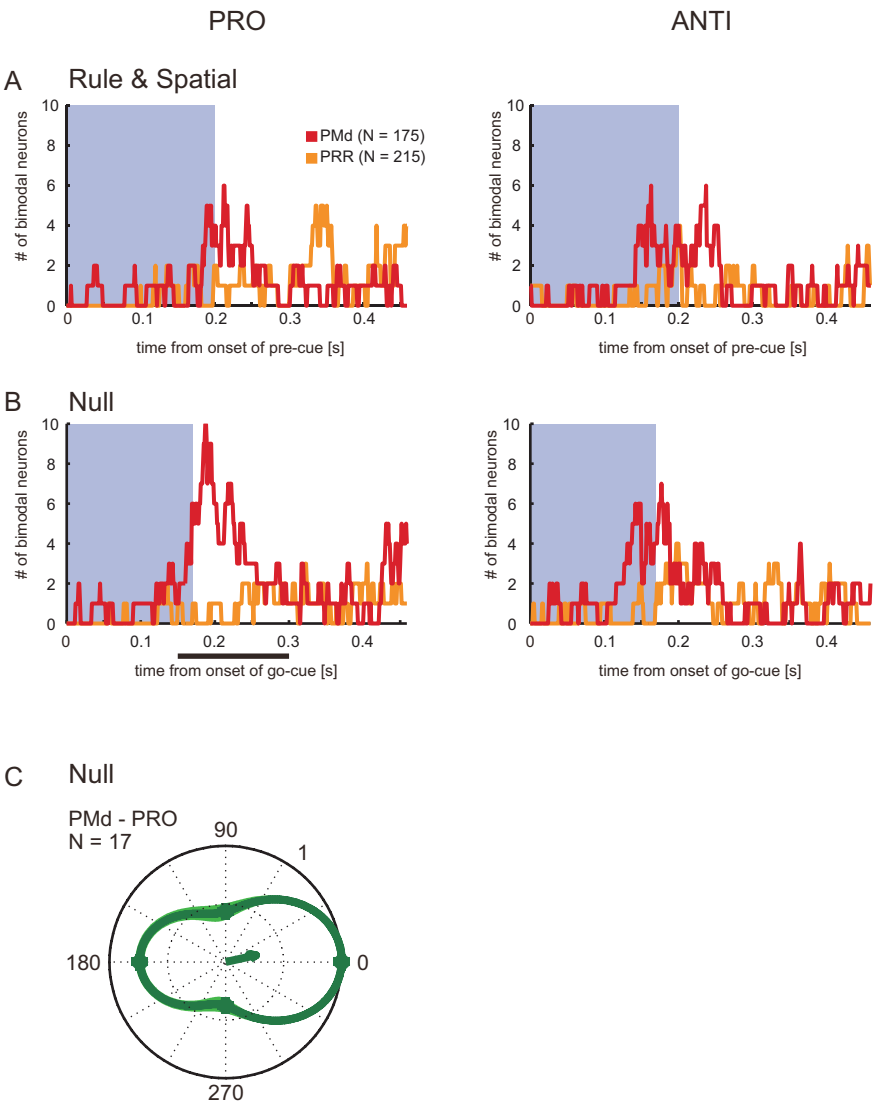

Figure S-4: Analysis of temporary bimodal tuning after cue presentation. A: Number of putatively bimodal neurons in PMd (dark) and PRR (light) and separate for pro- (left) and anti-trials (right) in the RS-condition after onset of the precue. B: Same as in A, but in the N condition after onset of the go-cue. C: Average normalized tuning of putatively bimodal PMd neurons in pro-trials in the N condition (the condition with the strongest indication for bimodal tuning). The tuning curve is bi-lobed, indicating bimodal tuning, even though not fully symmetrically. The time window, which was used for computing the directional tuning, is indicated by the black bar in B. Data are aligned and normalized to the average maximal response in pro-trials in the time window of analysis.

2.3 Choosing goals, not rules: Deciding among rule-based action plans

When a movement has to be planned towards a location that has to be inferred from a spatial stimulus (anti reach), some kind of remapping rule has to be applied. In the previous studies we analyzed situations in which the motor goal was unambiguously defined. Here we investigated a situation in which the motor goal is not definite but two equally probable potential motor goals remain. In this case planning a movement involves choosing the rule as well as computing a motor goal according to the chosen rule. It is unclear whether monkeys mentally construct all possible movement goals based on the remapping rules before selecting one of them, or whether they select a rule and after that only construct the associated motor goal for that rule. We could show that the frontoparietal reach areas can represent multiple motor goals at the same time and that these representations are biased by the monkey's choice preference. For this it is irrelevant if the motor goal was specified directly by a cue (pro reach) or by applying a remapping rule (anti reach). This means all potential motor goals are mentally constructed before one of them is selected and it supports the Bayesian brain theory.

Neuron
Article

Choosing Goals, Not Rules: Deciding among Rule-Based Action Plans

Christian Klaes,[1,2] Stephanie Westendorff,[1,2] Shubhodeep Chakrabarti,[2] and Alexander Gail[1,2,*]
[1]Bernstein Center for Computational Neuroscience, Göttingen, Germany
[2]German Primate Center, Kellnerweg 4, 37077 Göttingen, Germany
*Correspondence: agail@gwdg.de
DOI 10.1016/j.neuron.2011.02.053

SUMMARY

In natural situations, movements are often directed toward locations different from that of the evoking sensory stimulus. Movement goals must then be inferred from the sensory cue based on rules. When there is uncertainty about the rule that applies for a given cue, planning a movement involves both choosing the relevant rule and computing the movement goal based on that rule. Under these conditions, it is not clear whether primates compute multiple movement goals based on all possible rules before choosing an action, or whether they first choose a rule and then only represent the movement goal associated with that rule. Supporting the former hypothesis, we show that neurons in the frontoparietal reach areas of monkeys simultaneously represent two different rule-based movement goals, which are biased by the monkeys' choice preferences. Apparently, primates choose between multiple behavioral options by weighing against each other the movement goals associated with each option.

INTRODUCTION

When passing the ball to a player of his team, a soccer player can identify and select the proper target among many potential targets by the color of the jerseys. In this situation the physical targets are identical to potential targets of action (Figure 1A, left). However, when a striker is approaching the opponent goal, multiple alternative action goals have to be inferred from a single physical target (the goal keeper) via spatial transformation rules (Figure 1A, right). The striker might want to aim for the goal keeper, speculating that he or she will jump away, or for the opposite corner of the goal, hoping that the keeper stays. Recently, a lot has been learned on how primates represent and decide between multiple physical targets in target-selection tasks, and how different frontal and parietal cortical areas contribute to target valuation and selection (Sugrue et al., 2005; Gold and Shadlen, 2007; Churchland et al., 2008; Rangel et al., 2008; Andersen and Cui, 2009; Kable and Glimcher, 2009; Kim and Basso, 2010; Bisley and Goldberg, 2010; Cisek and Kalaska, 2010). Little is known, however, about decision processes in rule-selection tasks, which require choosing among goals based on a spatial transformation rule (Tremblay et al., 2002), and in which alternative goals might not be physically present as target stimuli, but have to be spatially inferred, like in the example of the striker.

In rule-selection experiments, alternative movements are conducted under identical spatial sensory conditions, but according to different context-defined transformation rules (Wise et al., 1996; Wallis and Miller, 2003). In antisaccade or antireach tasks (Figure 1A, right) a single visuospatial input is associated with two alternative movement goals: one that is directly cued by the sensory input (aim at the keeper), and another that has to be inferred from a spatial cue by applying a remapping rule (aim at the corner of the soccer goal opposite to the keeper) (Crammond and Kalaska, 1994; Shen and Alexander, 1997; Schlag-Rey et al., 1997; Everling et al., 1999; Zhang and Barash, 2004; Medendorp et al., 2005; Gail and Andersen, 2006). Two alternative decision processes are conceivable in such rule-selection tasks. The sensorimotor system could first choose among the alternative rules, and then only compute one sensorimotor transformation to encode the single motor goal that is associated with the selected rule (rule-selection hypothesis). Alternatively, the system could first compute all potential sensorimotor transformations, and then select among the multiple resulting motor-goal options (goal-selection hypothesis).

The difference between the rule- and goal-selection hypotheses should become obvious in areas of the brain that have "spatial competence" for movement planning, i.e., areas that exhibit spatially selective neural encoding of motor goal information. This is the case, for example, in the premotor cortex (Weinrich and Wise, 1982; Snyder et al., 1997; Crammond and Kalaska, 2000) and the posterior parietal cortex (Mountcastle et al., 1975; Snyder et al., 1997; Batista et al., 1999; Gail and Andersen, 2006). The rule-selection hypothesis predicts that such areas only encode one goal at a time, according to the preliminarily selected rule, but not multiple rule-based potential goals simultaneously (Figure 1B, left). The goal-selection hypothesis predicts that they simultaneously encode all alternative potential movement goals prior to the decision (Figure 1B, right). Therefore, the two hypotheses are distinguishable only at predecision stages, where the simultaneous existence of multiple, alternative, potential motor goals in a rule-selection task would favor the goal-selection hypothesis.

Evidence for potential motor goal encoding in spatial rule selection tasks, i.e., in situations like in the example of the striker, is lacking. Several areas of the brain have been thought

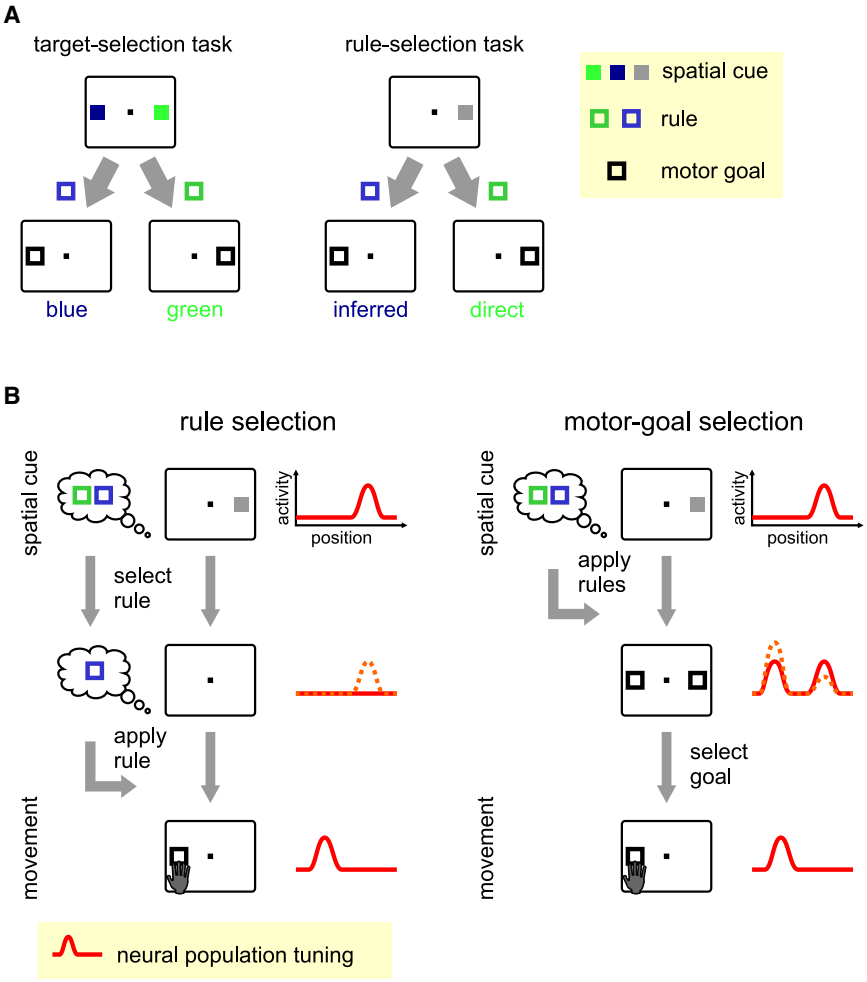

Figure 1. Target-Selection versus Rule-Selection Task

(A) In a target-selection task (left) two distinguishable (e.g. colored) spatial cues (targets) are presented, of which one should be chosen as movement goal according to a selection rule. Both potential targets have been spatially specified before the selection. In our rule-selection task only one spatial cue is provided, while the goal of the movement can be the position of the spatial cue (direct rule) or diametrically opposite to it (inferred rule).

(B) Neural rule- or goal-selection encoding in a rule-selection task. Decisions in ambiguous rule-selection tasks might follow a "rule selection" hypothesis (left columns). The rule will first be selected, and then be applied to the spatial cue to compute the single associated spatial motor goal. In spatial planning areas there will be either no spatial information encoded (solid red), or a memory of the spatial cue location (dotted orange) prior to the final selection. The "motor-goal selection" hypothesis (right columns) states that both alternative rules are applied to the spatial cue to create two competing spatial representations for the direct and the inferred motor goal, which may (dotted orange) or may not (solid red) be modulated by choice preferences of the subject. Note that the initial spatial encoding during the spatial cue presentation and the postdecision encoding of the final motor goal are identical in both hypotheses.

to encode multiple potential motor goals in space, but only in experiments involving selection among multiple physical targets (Basso and Wurtz, 1998; Cisek and Kalaska, 2005; Lau and Glimcher, 2008). However, in such tasks, multiple alternative spatial representations in the neural activity could be associated with multiple physical targets rather than motor goals. Therefore, target selection tasks are unsuitable for distinguishing between the rule- and the goal-selection hypotheses. We measured the spatial selectivity of neurons in monkey parietal and premotor cortex during reach planning in a novel rule-selection task (Figure 2). We show that two spatial, rule-based potential motor goals can be simultaneously encoded, supporting the goal-selection hypothesis.

Potential motor goals can encode all alternative choices as defined by the task (options), or biased representations of all choices based on previous reward experience (preferences),

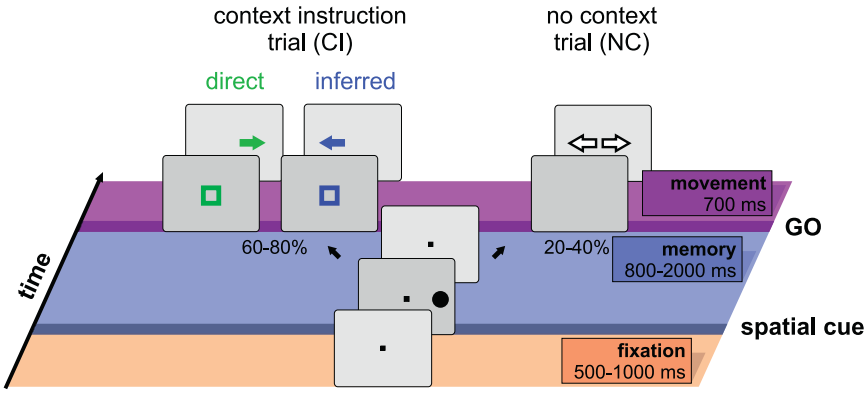

Figure 2. Rule-Selection Task with Instructed Delay

In our potential motor-goal (PMG) task a single spatial cue was presented at one of the four cardinal directions prior to a memory period. 60%–80% of the trials were context-instruction trials (PMG-CI) in which a contextual cue was shown in a second cue period ("GO"). This contextual cue instructed a direct reach toward the position of the preceding spatial cue (green) or an inferred reach toward the diametrically opposite direction (blue). In 20%–40% of the trials no contextual cue was shown (PMG-NC), and the monkeys were free to choose the goal either according to the direct or inferred transformation rule.

depending on which stage of the decision process they represent. So far, empirical evidence for preference encoding has been lacking for skeletomotor tasks, even in target selection experiments. Many previous oculomotor studies showed modulation of neural target responses by choice probability or some form of value assignment (preference encoding) in different brain areas of monkey (Basso and Wurtz, 1998; Dorris and Munoz, 1998; Platt and Glimcher, 1999; Sugrue et al., 2004; Dorris and Glimcher, 2004; Yang and Shadlen, 2007; Lau and Glimcher, 2008; Kim and Basso, 2010; Louie and Glimcher, 2010) and human (Hampton et al., 2006; Kable and Glimcher, 2007; Yanai et al., 2008; Wunderlich et al., 2009). Target-selection experiments using skeletomotor behavior, like reaching, showed encoding of freely selected targets in the parietal reach region (PRR) (Scherberger and Andersen, 2007; Pesaran et al., 2008), and potential motor goal encoding in the dorsal premotor cortex (PMd) of monkeys (Cisek and Kalaska, 2005), but the task designs in these studies did not allow dissociation of option versus preference encoding. We tested for the encoding of options versus preferences by using a mixture of instructed and free-choice trials in combination with different probabilistic reward schemes. We show that potential motor goal signals in monkey parietal and premotor cortex during reach planning represent mostly choice preferences, rather than options or preliminary selections.

RESULTS

Balanced Choices between Alternative Rule-Based Motor Goals

To distinguish between the rule- and goal-selection hypotheses we first tested if two potential rule-based motor goals can be encoded simultaneously, since this would provide evidence for the goal-selection hypothesis (Figure 1B, right). We designed a potential motor goal (PMG) task, in which subjects had to choose between two rule-based motor goals in each trial, and characterized the spatial selectivity of neural activity as a function of the spatial motor goal(s) during ambiguous reach planning.

Two male rhesus monkeys were trained to perform a memory-guided antireach task with instructed delay (Figure 2). A single spatial cue was combined with an optional contextual color cue. The contextual cue defined one of two spatial transformation rules according to which the spatial cue had to be mapped onto the associated motor goal. The reach goal could either be identical to the spatial cue (direct reach) or opposite to it (inferred reach). In each trial of this PMG task, both options were available in parallel to the subjects during reach planning, since the contextual cue was presented only at the end of the instructed delay, while the spatial cue was presented prior to the delay.

The PMG task consisted of two randomly interleaved trial types, either with context instruction at the end of the delay period (PMG-CI, 60%–80%), or without context instruction (no context, PMG-NC, 20%–40%). We used the ''free-choice'' PMG-NC trials to probe the subjects' behavioral choice preferences, and manipulated the subject's choice preferences by varying the reward schedules (see below).

The performance in PMG-CI trials was high for both monkeys (PMG-CI: 88 ± 1% [monkey A], 80 ± 1.4% [monkey S]). Most errors could be attributed to ocular fixation breaks, while improper choices in instructed trials, i.e., confusions of the direct and inferred reach goal, were rare in both monkeys (<2%). In the PMG-NC trials, both direct- and inferred-goal choices with proper ocular fixation and timing were considered correct, while reaches to any other directions were considered incorrect. In the first data set we are going to present, correct PMG-NC trials were rewarded according to a bias minimizing reward schedule (BMRS).

The BMRS was designed to reinforce balanced choice behavior, by taking the reward history of the monkey into account and reducing the reward probability if the behavior was biased. In the BMRS, random behavior with an equal amount of choices for either motor goal (on a short-term average across few trials) leads to 50% reward probability, while any consistent bias in choices leads to lower reward probabilities (see Experimental Procedures). With the BMRS, the direct choices (40 ± 0.1%, monkey A; 39.4 ± 2.5%, monkey S) and inferred choices (48.7 ± 0.1%, monkey A; 44.9 ± 2.6%, monkey S) were mostly balanced, with only a small bias in favor of inferred choices (Figure 3A). The overall balance between direct and inferred reach choices in PMG-NC trials suggests that the monkeys had close-to-equal preference for the two potential motor goals in BMRS sessions (= balanced data set).

Rule Selection versus Goal Selection—Potential Motor Goal Encoding

According to the goal-selection hypothesis, the planning of two equipotent alternative actions should lead to the neural encoding of both corresponding motor goal representations simultaneously. According to the rule-selection hypothesis, we would have to expect only one motor goal representation at a time despite balanced behavioral choices on average (Figure 1B). In the balanced choice condition, we recorded 145 (66 [A], 79 [S]) neurons in PRR, of which 97 (67%; 49 [A], 48 [S]) fulfilled the criteria to be tested for the encoding of potential motor goals (see Experimental Procedures). For the purpose of separating the rule-selection from the goal-selection hypothesis PMG-CI and PMG-NC trials were analyzed jointly, since the trial types are indistinguishable and unpredictable to the subjects prior to the optional contextual cue at the time of the GO signal.

Figure 3B shows an example neuron from PRR with a bimodal spatial selectivity profile from the balanced data set in the PMG task. We first tested the neurons spatial selectivity in two reference conditions. In the definite motor goal (DMG) task the monkeys were unambiguously instructed about the pending motor goal prior to memory period, i.e., the spatial and the contextual cue were shown at the beginning of the memory period (see Experimental Procedures). During such unambiguous planning in the DMG task, the neuron's responses reflected the unique downward motor goal in the ''direct'' (Figure 3B, left) and ''inferred'' (Figure 3B, center) context. This is indicated by the selectivity profiles for direct and inferred reaches that show the neural response as a function of the cue position, and that are shifted by 180° relative to each other (Figure 3B, bottom). Such motor-goal selectivity is characteristic for PRR (Gail and Andersen, 2006; Gail et al., 2009), and common to most directionally selective neurons of the current study (>80% across

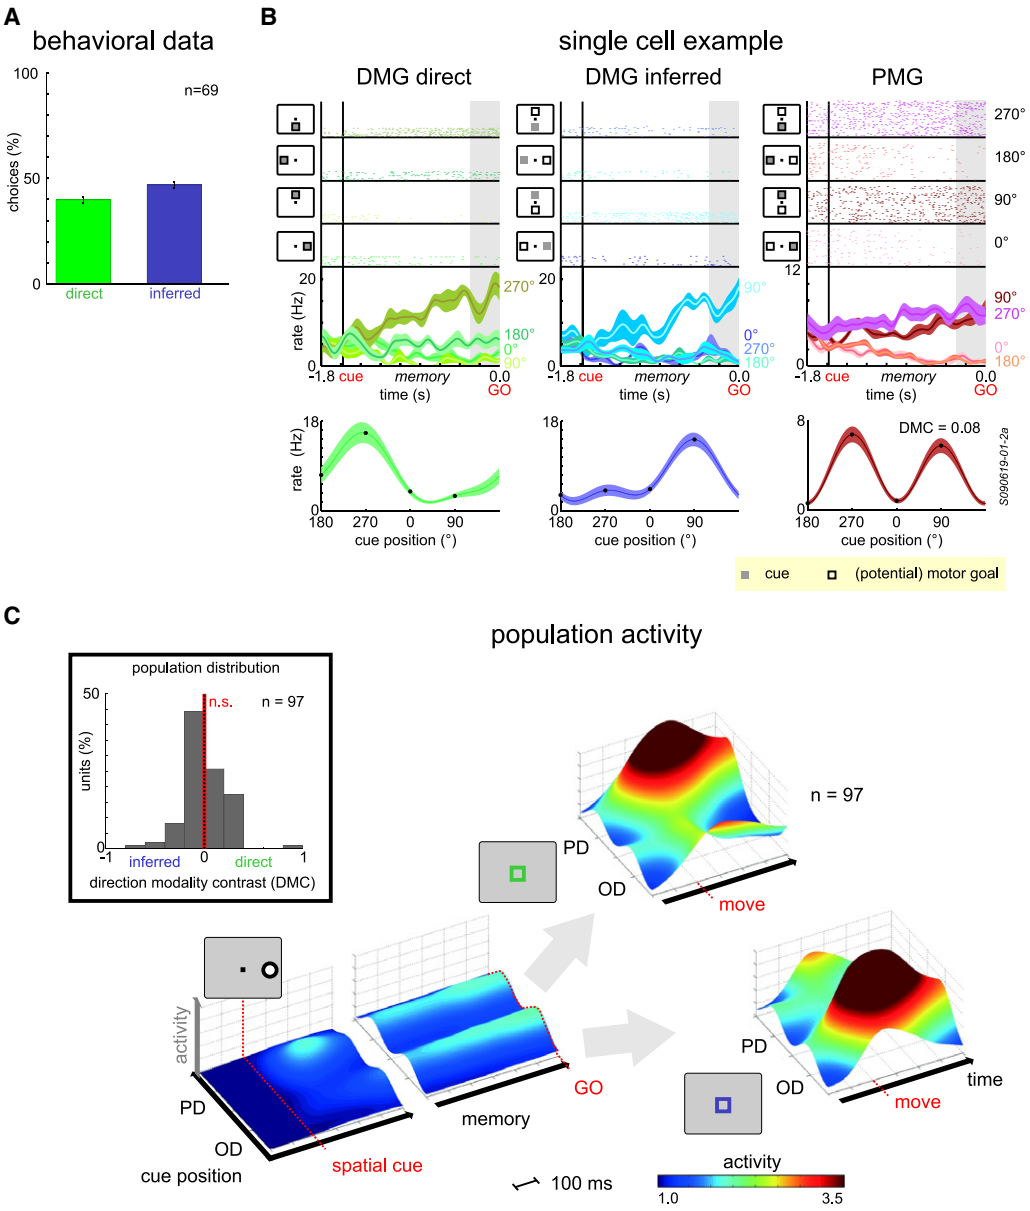

Figure 3. Neural Encoding of Potential Motor Goals in a Data Set with Balanced Choice Behavior

(A) Balanced percentage of direct (green) and inferred (blue) choices in PMG-NC trials with a bias minimizing reward schedule (BMRS; n = total number of behavioral testing days for both monkeys; error bars denote SEM).

(B) Response of an example PRR neuron. The top panels show the spike rasters in the definite motor goal (DMG) task with the direct rule (left column), the inferred rule (center) and in the PMG task (right), with average responses for each spatial cue position (0°, 90°, 180°, and 270°) below. Trials were aligned to the GO signal, while "cue" marks the average onset time of the spatial cue. Pictograms at the side of the spike rasters illustrate the spatial cue positions (filled squares) and the corresponding motor goal positions (open squares) according to the task. The bottom panels show the selectivity profiles of the neuron (average rate) as a function of cue position in the late memory period (shaded time window). Selectivity profiles were interpolated for illustrative purposes only. The shaded curves denote the SEM.

(C) Population results for the balanced data set in PRR. The average normalized activity of all eligible PRR neurons during the PMG-CI task is shown aligned to the spatial cue onset, the GO signal, and the movement onset (dotted red lines). Selectivity profiles were aligned to the neurons' preferred directions in DMG trials before averaging (PD, preferred direction; OD, opposite-to-preferred direction). Direct-cued and inferred-cued PMG-CI trials physically differ only at the time of the context instruction, hence, data are plotted jointly for the cue and memory periods. The inset shows the distribution of direction modality contrast (DMC) values for all eligible neurons in the late memory period.

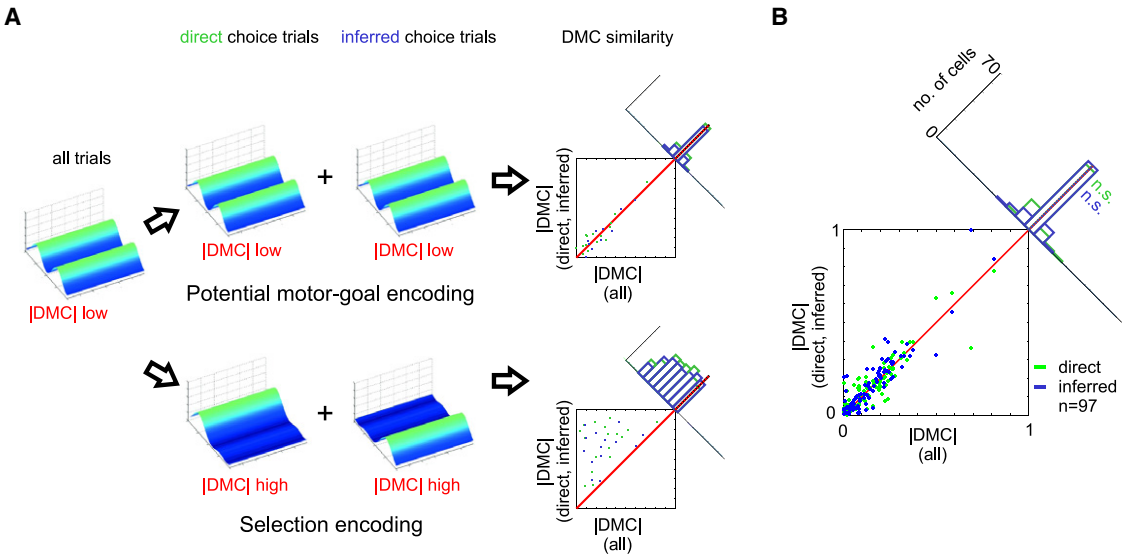

Figure 4. Encoding of Potential Motor Goals versus Preliminary Selections

(A) Schematic description of how to test the two alternative hypotheses with a choice-selective analysis of the free-choice trials (PMG-NC). Bimodal selectivity in the PMG trials could be the result of averaging across trials with alternating preliminary selection of either of the two potential motor goals (lower row), or reflect potential motor-goal encoding proper, independent of the later choice (top row). The similarity of the absolute DMC values, computed separately for direct-choice (green) and inferred-choice (blue) trials, with the choice-indifferent absolute DMC values, computed for all trials, allows to differentiate the two possibilities (see Figure S2 for a control).

(B) Choice-selective versus choice-indifferent absolute DMC values. The difference histogram shows that there is no significant deviation from the unity line, indicating encoding of two potential motor goals, and contradicting preliminary selection encoding.

data sets). Importantly, in the ambiguous PMG task (Figure 3B, right), the neuron was always most active if the previous spatial cue in a PMG task potentially indicated a downward (270°) reach, i.e., when it had appeared either at the upper (90°) or lower (270°) position. Since the spatial selectivity profile is plotted as a function of cue location, the bimodal activity profile with a peak separation of 180° indicates encoding of a single motor-goal direction in two different cue conditions, not two different motor-goal directions. Notably, the neuron was not active in trials with right-side (0°) or left-side (180°) cues, but only for those two directions (up and down) that were equally probable to instruct a downward motor goal.

The bimodal response profile of the example neuron in Figure 3B in the PMG task matched the prediction of the goal-selection hypothesis, and contradicts the rule-selection hypothesis. The bimodal profile mimicked the response pattern one would expect when averaging (*not* summing) the two response profiles in the DMG task. This means, the response pattern during planning of two equipotent alternative potential motor goals was an equally weighted linear combination of the response patterns during unambiguous planning of the two respective unique motor goals. In a model-based analysis we quantitatively confirmed this view (see Figures S1 and S4 available online).

Bimodal selectivity profiles dominated the balanced data set in PRR. The average population activity in the balanced data set shows two stable ridges of activity during the memory period (Figure 3C). Since the cue-position axis marks the location of the spatial cue relative to the preferred direction (PD) of each neuron (as measured in the DMG task), the two ridges indicate that on the population level the direct and inferred goals are represented simultaneously during ambiguous reach planning. For quantitative analysis we characterized the bimodal versus unimodal selectivity of each neuron with a direction modality contrast (DMC). Positive DMC indices indicate selectivity for the direct motor goal; negative values indicate selectivity for the inferred motor goal. Indices close to zero indicate symmetric bimodal tuning (*not* lack of tuning) since only directionally selective neurons were considered (see Experimental Procedures). The mean DMC of the balanced data set did not significantly deviate from zero ($m = 0.001$; standard error of the mean [SEM] = 0.021, $p > 0.05$), indicating that in the balanced data set most neurons had bimodal selectivity profiles (Figure 3C, inset).

Potential Motor Goals or Preliminary Selections? Choice-Selective Analysis in PRR

The existence of a bimodal neural selectivity pattern in the balanced data set is not sufficient to demonstrate potential motor goal encoding. The monkeys could have preliminarily selected one of the two potential motor goals during the memory period in every trial, and randomly switched their selection from trial to trial. Such switching would be obscured in PMG-CI trials due to the explicit context instruction at the time of the GO cue. The bimodal selectivity pattern revealed by the above analyses would denote an artifact of averaging across inhomogeneous sets of trials in this case (Figure 4A, bottom).

With a choice-selective analysis of the free-choice (PMG-NC) trials we can rule out this possibility. We can instead show that both potential motor goals were encoded independently of the monkey's later choices (Figure 4A, top). If the monkeys made preliminary selections in every trial then this selection should

be reflected in an unambiguous neural encoding of this preliminary selection in the late memory period immediately prior to the monkeys decision (preliminary selection encoding). We sorted PMG-NC trials according to the choice of the monkey, and computed DMC values of the neural activity in the late memory period separately for the trials in which the monkeys freely chose the direct and inferred goals. If the low DMC of a neuron with bimodal selectivity was the averaging result of two opposite unimodal selectivity profiles, one for direct-choice trials and the other for the inferred-choice trials, then a low absolute value of the original DMC would be attended with high absolute values of the two choice-selective DMC values for this neuron. This means, preliminary selection encoding would be indicated by a low similarity between the original and the choice-selective DMC values across neurons (Figure 4A, bottom). Vice versa, we can reject the selection hypothesis if a neuron in both choice-selective subsets of trials shows a bimodal selectivity pattern, i.e., when low absolute values of the original DMC is attended with low absolute choice-selective DMC values, resulting in a high similarity between original and choice-selective DMC across neurons (Figure 4A, top).

The balanced data set in PRR yielded bimodal selectivity in PMG-NC trials separately within direct-choice and within inferred-choice trials. The absolute choice-selective DMC values for direct- and inferred-choice trials were highly similar to the absolute original DMC values (Figure 4B). This can be seen by the fact that the average distance of the data points from the unity line did not significantly differ from zero neither for direct- ($p_d > 0.05$) nor inferred-choice ($p_i > 0.05$) trials. When—as a control—the method was applied to the DMG data set, in which we know that the monkeys had selected the motor goal already during the memory period, then the choice-selective and original DMCs were highly and significantly dissimilar ($p_d = 0.0012$, $p_i = 0.00067$; see Figure S2). Additional variance tests indicated that it is unlikely that the bimodal selectivity profiles were the consequence of rapid switching between two alternative preliminary selections within the time of a trial (Figure S2).

Taken together, the results from the choice-selective analysis of the balanced data set indicated genuine encoding of potential motor goals rather than alternating preliminary selections in PRR. This supports the motor-goal selection hypothesis and argues against the rule-selection hypothesis.

Manipulating Behavioral Choice Preferences

Depending on which stage of the decision process a brain area belongs to, encoding of multiple potential motor goals in that area could represent the multiple options offered to the subject (the "menu"), or the competing behavioral goals associated with these options and weighted with the subject's preference for either choice. Motor-goal options were defined solely by the task. In any PMG trial two motor-goal options (the direct and inferred motor goal) were valid during the memory period. Encoding of motor-goal options should lead to the representations of two potential motor goals during the memory period of all PMG trials, irrespective of any choice preferences of the monkeys. Motor-goal preferences were defined by the monkeys' average choice behavior in PMG-NC trials. Since the monkeys had close-to-equal choice preferences for direct and inferred motor goals in the balanced data set, the bimodal selectivity profiles are not suited to dissociate encoding of motor-goal options versus motor-goal preferences. If, on the other hand, the monkeys had a bias in favor of one of the two options, then encoding of motor-goal preferences should lead to neural activities in the memory period of PMG trials that reflect the relative probability of selecting either potential goal in the PMG-NC trials. By using different reward schedules we recorded two data sets, one with balanced choice behavior (see above), and one with strong behavioral choice bias, to dissociate the options and preference encoding hypotheses.

In the second data set, correct PMG-NC trials were rewarded according to an equal probability reward schedule (EPRS). With the EPRS, in which a 50% reward probability independent of the choice history was guaranteed (reward probability: 52 ± 5%; $p > 0.05$ [A], 50 ± 4%; $p > 0.05$ [S]), both monkeys showed a strong bias in favor of the inferred reach goal (Figure 5A), i.e., most reaches in PMG-NC trials were directed toward the inferred motor goal (85 ± 4.0% monkey A, 63 ± 4.1% monkey S), and only a small fraction toward the direct goal (2.4 ± 0.8% monkey A, 17.8 ± 3.4% monkey S). In the remaining PMG-NC trials (12.6% monkey A, 19.2% monkey S) the monkeys aborted the trial without reaching, or reached toward one of the orthogonal goals (<1%). This means that both monkeys had a preference for the inferred goal when the transformation rule was unknown, and when either goal selection was rewarded with equal probability in EPRS sessions (= biased data set).

We can only speculate about the reason for the intrinsic bias of both monkeys during the EPRS (Figure S3). The reason behind this behavior is not immediately relevant for the purpose of dissociating options encoding from preference encoding at the neural population level, though. It is sufficient to note that both monkeys consistently had a similarly strong bias over an extended period of time in the EPRS sessions, and little to no bias in the BMRS sessions.

Options versus Preference Encoding in PRR

If neurons encoded behavioral choice preferences then we would expect encoding of only the inferred motor goal in the PMG trials of the biased data set, in contrast to the encoding of both potential motor goals simultaneously as seen in the balanced data set. This should be true in the late memory period of all PMG trials, since PMG-CI trials are indistinguishable from the PMG-NC trials prior to the GO signal, and were randomly interleaved.

In PRR, the biased data set contained a total of 258 (159 [A], 99 [S]) recorded neurons. A total of 148 (57%) neurons (96 [A], 52 [S]) of the biased data set fulfilled the criterion for the analysis of potential motor-goal encoding.

The PRR example neuron in Figure 5B was recorded in the biased data set and was most active during planning of leftward (180°) reaches in direct-cued or inferred-cued DMG trials. In PMG trials the neuron was only highly active if the spatial cue was presented at the right side (0°), i.e., as if an inferred instruction had been given externally or had been selected internally. Such unimodal selectivity for the inferred goal dominated the biased data set in PRR. The average normalized population activity showed only a brief response increase when the cue

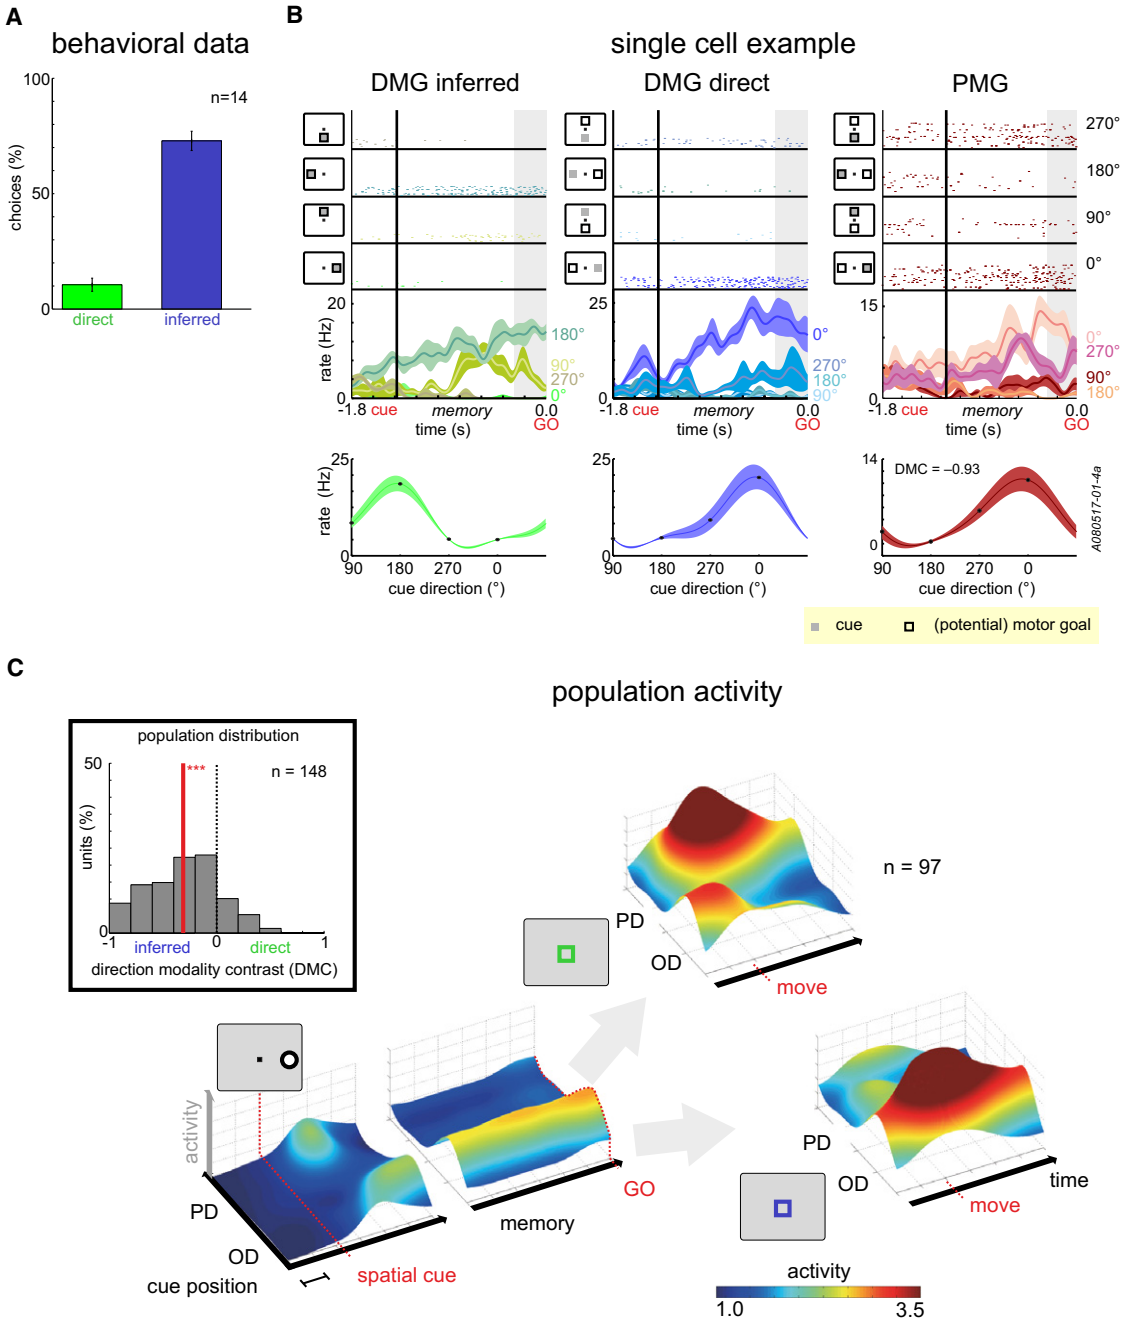

Figure 5. Neural Encoding of Motor Goal Preference in a Data Set with Biased Choice Behavior

(A) Biased percentage of direct (green) and inferred (blue) choices in PMG-NC trials with an equal probability reward schedule (EPRS). Same conventions apply as in Figure 3.

(B) Example PRR neuron from the biased data set. It showed motor-goal selectivity in the DMG task, as did the example in Figure 3. But in the PMG task it is only active if the spatial cue appears opposite to the neurons PD.

(C) Correspondingly, PRR population activity shows a strong representation at the opposite-to-cue position (OD) during the memory period. The DMC distribution is significantly biased for the inferred goal (mean DMC = −0.31; ***p < 0.001; rank-sum test).

matched the preferred direction (PD) of the neurons. This was followed by a high level of activity when the cue was opposite to the PD, corresponding to an encoding of the inferred goal throughout the memory period (Figure 5C). The mean DMC during the memory period of the biased data set was negative (m = −0.31; SEM = 0.028) and significantly different from zero (rank-sum test, p < 0.001) (Figure 5C, inset). This means that the behavioral preference was reflected in a significant bias of

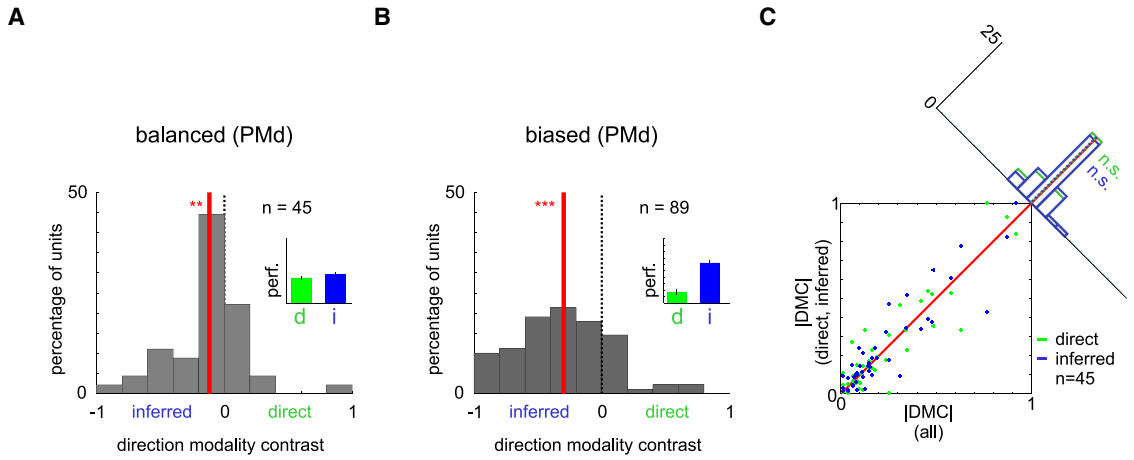

Figure 6. Potential Motor Goals in PMd

(A and B) Distribution of DMC values and the monkeys' choice-behavior (insets) of the balanced (A) and biased (B) data sets. Same conventions apply as in Figures 3 and 5 (***$p < 0.001$, **$p < 0.01$, rank-sum test).

(C) The choice-selective analysis of the DMC similarity in PMd. Same conventions apply as in Figure 4.

the neural directional selectivity in the population of PRR neurons. The inferred-goal neural preference is neither consistent with an unbiased equipotent encoding of the two task-defined motor goal options (options hypothesis), nor with an encoding of the previous instruction cue (visual memory), but it is consistent with the preference hypothesis.

Based on the observed inferred-goal selectivity in the biased data set alone, one could not dissociate preference encoding from preliminary selection encoding. But we can argue against the latter possibility based on the choice-independent bimodal response profiles in the choice-selective analysis of the balanced data set. Preliminary selection encoding would have had to reveal direct-goal neural selectivity in direct-choice trials, and inferred-goal selectivity in inferred-choice trials, which was not the case (see above).

Motor-Goal Encoding in Dorsal Premotor Cortex (PMd)

Another objective of our study was to compare parietal and premotor sensorimotor areas, which are well known to be involved in reach planning, while their role in reaching decisions is less clear (Cisek and Kalaska, 2005; Scherberger and Andersen, 2007). We conducted the same analyses for PMd as for PRR neurons. The biased data set contained 193 PMd neurons (118 monkey A, 75 monkey S), and the balanced data set 112 PMd neurons (monkey S). Of those, 46% fulfilled the criteria for the DMC analysis in the biased data set, and 40% in the balanced data set, which denote smaller fractions of neurons than in PRR (see above).

The analyses of potential motor-goal encoding in PMd revealed overall very similar results to PRR, but there were also differences. The distribution of DMC values in the biased data set of PMd (Figure 6B) revealed a significant bias in favor of the inferred motor goal ($m = -0.298$; SEM = 0.038, $p < 10^{-10}$, t test), as it was the case in PRR (Figure 5C, inset). In contrast to PRR (Figure 3C, inset), the DMC distribution in PMd (Figure 6A) also showed a significant remaining bias for inferred goals ($m = -0.11$; SEM = 0.05, $p = 0.004$) in the balanced data set. Note, though, that this bias in DMC values was significantly smaller ($p = 0.002$) than in the biased data set, which indicates that most neurons exhibited bimodal response profiles, while few had a weak bias for the inferred goal. Since the monkeys also had a small residual choice preference for the inferred goal (Figure 3A) this could mean that PMd is more strongly modulated by small choice preferences than PRR. The choice-selective analyses of the PMG-NC trials showed a high DMC similarity (Figure 6C), equivalent to PRR (Figure 4B). This, like in PRR, indicated that the bimodal directional selectivity was mostly *not* the consequence of preliminary selection encoding in combination with trial-by-trial switching of the behavioral choice. In summary, the PMd results are qualitatively very similar to PRR, suggesting similar encoding schemes in both areas. For a discussion of additional smaller differences between PRR and PMd as revealed by our model-based analyses and variance analyses see Figures S1 and S2.

Normalization of Multiple Motor-Goal Representations

Models of decision making often involve mutual competition between the neural representations of multiple coexisting alternative choices (Platt and Glimcher, 1999; Cisek, 2006). Such competition implies that the response of a neuron should be reduced when its preferred motor goal marks only one out of two equally valid behavioral options, compared to when the motor goal is unambiguously selected. The responses of the example neurons and the population activity plots in Figures 3 and 5 suggest that this is the case. The results indicate a halving of the neural response strength to each potential motor goal in the balanced PMG task compared to the corresponding unambiguous motor goal in the DMG task or biased PMG task. A quantitative analysis of the weight coefficients (scaling factors) in the model-based analysis confirmed this view (Figure S4). The reduced neural response strengths during the simultaneous presence of two alternative motor goals compared to a single

goal argues in favor of a competition between alternative motor goal representations.

DISCUSSION

The ability to plan multiple upcoming actions and decide among them is vital to an organism acting within a complex environment. We investigated how parietal and premotor reach planning areas encode the decision between different possible sensorimotor transformation rules that could be applied to a single visuospatial object. When monkeys were faced with two alternative spatial transformations, and chose them with equal preference, then two separate spatial motor goal representations coexisted in the frontoparietal reach network. This was the case despite the fact that only one goal was directly visually cued and the other had to be inferred from the visual cue by applying a remapping rule. Additionally, the parietal reach region (PRR) and the dorsal premotor cortex (PMd) predominantly encoded the variable choice preference between two potential motor goals. By using free-choice probe trials and two distinct reward schedules, we could rule out encoding of the monkeys' preliminary behavioral selections, as well as encoding of the task-defined choice options, during movement planning. Our results suggest that in rule-selection experiments the sensorimotor system first computes all potential motor goals associated with a currently valid set of potential transformation rules, weighs them according to the subject's choice preference, and then selects among these goals.

Deciding among Alternative Action Plans Rather than Transformation Rules or Targets

We showed that during movement planning two alternative potential reach goals can be represented simultaneously in PRR and PMd in a rule-selection task. In this task only one visuospatial target was presented at a time, allowing two alternative motor goals by applying two different mapping rules. Our results suggest that with preexisting knowledge about the visuospatial constraints of the task (knowing the spatial cue), and uncertainty about the to-be-applied rule (not knowing the context cue), the sensorimotor system constructs all remaining motor goal options, which are defined by the general context of the task, and are of subjective value to the monkey (see biased versus balanced condition below). We can reject the alternative rule-selection hypothesis according to which the monkeys in general would first select a rule, and then only compute the single associated motor plan. It is as if the sensorimotor system in a rule-selection task first creates all potential motor-goal representations and then applies the same computational decision algorithms as in a target-selection task.

The view that multiple spatial motor goal options can be simultaneously encoded prior to the decision in parietal and premotor areas is reminiscent of earlier saccadic target-selection experiments in the superior colliculus (Basso and Wurtz, 1998) and the lateral intraparietal area LIP (Platt and Glimcher, 1999; Sugrue et al., 2004; Dorris and Glimcher, 2004; Yang and Shadlen, 2007; Louie and Glimcher, 2010). They showed probabilistic, graded neural responses for preferred and nonpreferred targets, depending on saccadic choice probabilities or subjective values. Also, a study in PMd showed bimodal response profiles in a manual two-target selection task (Cisek and Kalaska, 2005). Our conclusions go beyond the previous findings, since these studies showed the coexistence of multiple spatial representations associated with alternative choices, but used target-selection tasks. We show that the simultaneous representation of mutually exclusive spatial motor goals in sensorimotor areas does not require the presentation of multiple alternative spatial physical target stimuli.

This simultaneous encoding of alternative competing motor goals is also fundamentally different from the representation of two sequential movement goals. Previous experiments showed that in the parietal cortex, during the planning of a multicomponent (double-step) movement, two neural populations were activated, each of which was selective for one of the single movement components (Medendorp et al., 2006; Baldauf et al., 2008). Double-step experiments do not induce a decision process between mutually exclusive action goals, and rather suggest that multiple components of a complex movement can be planned at once. Our finding of simultaneous encoding of alternative competing motor goals does complement previous observations in effector-selection experiments, which showed that alternative eye or hand movements to the same spatial target, instructed (Calton et al., 2002) or freely chosen (Cui and Andersen, 2007), can elicit simultaneous movement planning activity in LIP and PRR.

The advantage of the goal-selection scheme over the rule-selection scheme for decision making could be that—by computing all associated motor goal alternatives and their implicit action plans during the ambiguous state of planning—a more comprehensive cost-benefit calculation of each choice can be achieved. When the striker in our introductory example has to decide between aiming for the position of the goal keeper versus the opposite corner, then it is not enough to consider the likelihood of the keeper to jump or stay. Also the costs associated with the striker's action alternatives are relevant, e.g., the striker might be poor at aiming for right-side goals, or the ball might be in an immediate position that eases aiming for one corner but not the other. Our results imply that the decision process in our rule-selection experiment selected between competing motor-goal alternatives, not between different transformation rules or target stimuli, and that this competition likely happened in the sensorimotor areas that are involved in planning the respective movements.

Note, we do note rule out the possibility that in parallel a competition between the two potential rules takes place in rule-encoding frontal cortical areas (White and Wise, 1999; Wallis et al., 2001; Wallis and Miller, 2003; Genovesio et al., 2005). The rule-competition could then, in the extreme case, just be mirrored by probabilistic motor goal representations in downstream sensorimotor areas. Because of the observed response normalization in our data (see below), we believe that if at all there was a rule-competition in our task then it was paralleled by a goal-competition in the sensorimotor areas, which would make sense for economical reasons, as discussed in the previous paragraph (Cisek and Kalaska, 2010).

Preference versus Options Encoding

Potential motor-goal representations in our experiment depended on the preference of the monkeys, as defined by the

probability of behavioral choice of either action alternative. In our task design we cannot differentiate between choice probabilities and assigned subjective value (Sugrue et al., 2004; Samejima et al., 2005; Hampton et al., 2006; Kable and Glimcher, 2007; Lau and Glimcher, 2008; Wunderlich et al., 2009), as was attempted in a recent discounting experiment (Louie and Glimcher, 2010). Consequently, we speak more generally of preferences, as quantified by choice probabilities.

Simultaneous potential motor-goal encoding during reach planning had previously only been shown in PMd (Cisek and Kalaska, 2005). Since a dependence on the monkeys' choice preferences was not tested, it is unclear if this previous PMd data reflected preferences or task-defined motor-goal options (the menu, Padoa-Schioppa and Assad, 2006). The biased population tuning in the memory period of our biased data set contradicts options encoding, and suggests that potential motor-goal encoding predominantly reflected choice preferences in PMd.

In posterior parietal cortex, preference encoding between competing options has previously been shown in saccadic target-selection tasks (Platt and Glimcher, 1999; Sugrue et al., 2004; Dorris and Glimcher, 2004; Yang and Shadlen, 2007; Kable and Glimcher, 2007; Wunderlich et al., 2009; Louie and Glimcher, 2010). Corresponding data for skeletomotor movements, like reaching, and for rule-selection tasks in general is lacking. Previous target-selection tasks with reaching revealed post-GO-cue selection signals in PRR (Scherberger and Andersen, 2007; Pesaran et al., 2008), but no neural response modulations by choice preference was shown. Previous tasks with deterministic targets showed reward- or value-dependent modulations of the neural responses (Musallam et al., 2004; Iyer et al., 2010), but relative weighing of alternative options against each other was not tested.

Taken together, the principle of weighing alternative motor goal representations with behavioral choice preferences is not restricted to the saccade planning system, but can be found in the skeletomotor system as well, and neural implementations of this principle include not only parietal movement planning areas, but also areas in the frontal cortex, like PMd.

Competing Goal Representations

Models of decision making often imply mutual competition between the neural representations of multiple coexisting alternative choices (Platt and Glimcher, 1999; Cisek, 2006). In our experiment, this competition likely happened in the sensorimotor areas that we recorded from and that are involved in planning the respective movements, since we found reduced neural response strengths during the simultaneous representation of two alternative motor goals compared to a single goal (Cisek and Kalaska, 2005).

Conclusions

Our findings support the idea that reach decision making and movement planning, in tasks that require the selection of a spatial transformation rule, are integrative rather than sequential processes, mediated by overlapping action-specific neural populations in PRR and PMd (Scherberger and Andersen, 2007; Cisek, 2007; Andersen and Cui, 2009). The results provide evidence for competitive encoding of alternative potential reach plans in PRR and PMd, reflecting the monkeys' average choice preferences, but being independent of the immediate behavioral choice of the monkey. This is consistent with the idea that the brain utilizes probabilistic representations throughout all stages of the decision process until an action is finally required (Knill and Pouget, 2004). Importantly, our results suggest that in situations of uncertain choice of which transformation rule to apply, the sensorimotor system can construct all potential motor goal alternatives, and then select among these alternatives, once enough evidence for a proper choice is available, rather than preliminarily betting on one of the transformation rules and computing only the single corresponding motor plan. This strategy could denote a valuable and general principle in decision making, allowing a more comprehensive cost-benefit analysis that includes the consequential costs of the movements associated with each choice.

EXPERIMENTAL PROCEDURES

Potential Motor Goal Task with Context Instruction

In PMG-CI trials (Figure 2), one spatial and one contextual visual cue were presented to the subjects at different times during the trial (ViewSonic VX922 LCD screen; 5 ms off-on-off response time). The peripheral spatial cue was located at one of four possible positions (0°, 90°, 180°, and 270°) with an eccentricity of 9 cm (14.5° visual angle, VA) relative to the fixation point. The contextual cue consisted of a green (direct-cued) or blue (inferred-cued) frame around the central eye and hand fixation points. It instructed the subject to reach toward (direct, proreach) or to the position diametrically opposite of the spatial cue (inferred, antireach).

A trial was initiated by the monkey by fixating a small red square in the center of the screen (eye fixation tolerance: 2.0-3.0° VA; 224 Hz CCD camera, ET-49B, Thomas Recording, Giessen, Germany) and touching an adjacent white square of the same size (hand fixation tolerance: 4.0° VA, touch screen mounted directly in front of the video screen; IntelliTouch, ELO Systems, Menlo Park, CA). After a random period of 500–1000 ms (fixation period) the spatial cue was shown briefly for 200 ms. During the following 800–2000 ms (memory period) only the fixation squares were visible. The contextual cue was shown for 170 ms at the end of the memory period and the hand fixation square disappeared (GO signal). The monkey had to make a reach toward the instructed goal within a maximum of 700–1000 ms (movement period, 4.9° VA reach tolerance) and hold the goal position for 300–400 ms (feedback period). The monkey received visual feedback about the correct movement goal (filled circle of the same color as the contextual cue at the goal location) at the end of a correct trial.

Eye fixation had to be kept throughout the trial. Liquid reward and acoustic feedback indicated correct (high pitch tone, reward) or incorrect (low pitch tone, no reward) behavior. Correct choice of the instructed motor goal and fixation behavior were required for a PMG-CI trial to be considered correct. Only correct trials were used for the analysis.

Potential Motor Goal Task without Context Information

PMG-NC trials were similar to the PMG-CI trials, except that no contextual cue was shown at the end of the memory period. In those trials the monkey had to choose whether to reach to the direct or to the inferred goal. Until the end of the memory period PMG-CI and PMG-NC trials were indistinguishable. Only PMG-NC trials in which the monkey either reached for the direct or the inferred position were considered correct and were used for the analysis. Note that not all of the correct trials were rewarded. Reward depended on the used reward schedule (see below).

Definite Motor Goal Task

The DMG task differed from the PMG-CI trials only in the timing of the contextual cue. In the DMG task the spatial and the contextual cue were shown simultaneously at the beginning of the memory period. Only DMG trials with correct choices and ocular fixation were rewarded and analyzed.

Block Design and Randomization

The PMG and DMG tasks were presented in separate blocks. The DMG block consisted of typically ~100 trials, the PMG block of a minimum of ~300 trials. The order of the two tasks was variable across days.

PMG-NC and PMG-CI trials were randomly interleaved during PMG blocks. A PMG block contained 60%–80% (mean = 76%) PMG-CI trials and 20%–40% (mean = 24%) PMG-NC trials. In each task the four spatial cuing directions were randomly interleaved with equal probability. In PMG-CI trials and in the DMG task the direct-cued and inferred-cued trials were also randomly interleaved with equal probability.

Reward Schedules in PMG-NC Trials

We implemented two different reward schedules for PMG-NC trials.

One was the bias-minimizing reward schedule (BMRS). With a BMRS balanced behavior, i.e., 50% direct and 50% inferred reaches, leads to a 50% reward probability, while any biased choice behavior leads to lower reward probabilities. The BMRS algorithm takes the reward history of the monkeys into account and changes the probabilities for rewarding a direct or inferred reach in favor of the alternative that was chosen less often so far:

$$\begin{array}{l} p(R_d) = F(n_i - n_d) \\ p(R_i) = F(n_d - n_i) \end{array},$$

where n_i is the total number of rewarded inferred reaches and n_d is the total number of rewarded direct reaches. F was defined as

$$F(x) := \begin{cases} 1 & x>1 \\ 2/3 & x=1 \\ 1/2 & x=0 \\ 1/3 & x=-1 \\ 0 & x<-1 \end{cases}.$$

The second reward schedule was the equal-probability reward schedule (EPRS). In EPRS trials the monkeys were rewarded with 50% probability, no matter whether they reached for the direct or inferred goal, and regardless of the reward history. The reward probabilities for direct (R_d) or inferred (R_i) choices were

$$p(R_d) = p(R_i) = 0.5.$$

With the EPRS, the reward probability is independent of the behavioral strategy of the monkeys, as long as they chose between the two potential goals (see Figure S5 for data with 100% reward probability).

Biased and Balanced Condition

The recorded data was split into two distinct data sets. One data set contains only units that were recorded with the EPRS, before we trained the respective monkey with the BMRS. Since both monkeys showed a very similar choice bias during EPRS sessions (see Results), we refer to this data as the biased data set. The second data set contains only units recorded after we used BMRS and is referred to as balanced data set. Behavioral tests with the PMG-NC trials were conducted at the end of the neuronal recording period in the biased data set. Control experiments with simultaneous behavioral and neural recording of biased PMG-NC trials confirmed that results and conclusions are unaffected by this (see Figure S5).

Animal Preparation, Neural Recordings

Surgical procedures and neural recordings were described previously (Gail et al., 2009). Animal care and all experimental procedures were conducted in accordance with German laws governing animal care.

Neural Data Selection and Tuning Analyses

Extracellular recordings were conducted with up to five microelectrodes in parallel (“mini-matrix”; Thomas Recording, Giessen, Germany) on each chamber. Spike times and waveforms were recorded and subjected to additional offline sorting (Offline Sorter; Plexon).

All isolated units were tested for their directional selectivity (Kruskal-Wallis test; four groups of different spatial cue positions; sample sizes defined by the number of identical trial repetitions). Selectivity was tested independently for direct-cued and inferred-cued trials during the late memory period in the DMG task (average spike rate during the last 300 ms of the memory period, i.e., activity succeeding the precue with a time-lag of at least 500 ms, and immediately preceding the GO cue). The late memory period was chosen to extract movement planning activity without confounding effects of (1) immediate visual input from the cue stimuli; (2) transition phases from visual to motor-goal tuning (Gail and Andersen, 2006); or (3) visual and somatosensory input and motor-control signals related to movement initiation.

Only neurons that were significantly selective in direct-cued trials of the DMG task were used in the following analyses (Figure S6). For all analyses that involved PMG-CI or PMG-NC trials, we additionally required the neurons to be significantly directionally selective in the late memory period of PMG trials (Kruskal-Wallis, see above).

Population Activity

To visualize the temporal dynamics of spatial representations on a population level, we averaged the time-resolved spiking activity across all neurons that were directionally selective during the memory period of PMG trials. Before averaging, the directional selectivity profiles for each neuron were aligned relative to the interpolated preferred direction in the late memory period of the DMG task and normalized to the baseline level (average spike density in the 300 ms before spatial cue onset). The population activity was only used for illustrative purposes (see Figures 3C and 5C), not for quantitative statistical analyses.

Analysis of Bimodal Selectivity Profiles

We used a direction modality contrast (DMC) to quantify the bimodality of individual neuronal responses:

$$DMC = \frac{R_{MD} - R_{OD}}{R_{MD} + R_{OD}}.$$

R_{MD} is the mean firing rate of a neuron during the last 300 ms of the memory period of all PMG trials (PMG-CI and PMG-NC) at the same direction that evoked the maximum response (MD) in the DMG task. R_{OD} is the firing rate for trials in the opposite-to-maximum direction (OD). Since the MD is measured relative to the direction of the spatial cue in direct-cued trials of the DMG task, positive DMC indices indicate preferred selectivity for the direct motor goal (at the spatial cue location), whereas negative values indicate preferred selectivity for the inferred motor goal (opposite the spatial cue). Values around zero indicate symmetric bimodal selectivity, not lack of selectivity, since neurons without directional selectivity were removed from this analysis.

Choice-Selective Analyses

To differentiate between the selection and the preference hypotheses, we sorted the PMG-NC trials in the balanced data set according to the free choice of the monkey, and calculated the DMC separately for direct-choice and inferred-choice trials. That means if in a PMG-NC task the monkey reached toward a goal position as if the contextual instruction had been direct, the trial was labeled “direct choice” and if he reached toward a goal position as if the contextual instruction had been inferred, the trial was labeled “inferred choice.” The absolute choice-selective DMC values were then compared to the absolute original, choice-indifferent DMC values (average over all trials without sorting them according to the choice) in a similarity analysis (illustrated in Figure 4A). The DMG condition was used as a control for this similarity analysis (see Figure S2). To quantify the similarity between the choice-selective DMC values and the choice-indifferent DMC values, we calculated the distance from the unity line of the correlation plot, which is equivalent of calculating the difference between the choice-selective and choice-indifferent DMC values. We then used a t test to determine if the distribution of these differences was significantly deviating from zero.

SUPPLEMENTAL INFORMATION

Supplemental Information includes six figures and can be found with this article online at doi:10.1016/j.neuron.2011.02.053.

ACKNOWLEDGMENTS

We thank Sina Plümer and Ludwig Ehrenreich for their help in data collection. This work was supported by the Federal Ministry for Education and Research (BMBF, Germany, grants 01GQ0433, 01GQ0814, 01GQ1005C), the Deutsche Forschungsgemeinschaft (DFG) Collaborative Research Centre 889 "Cellular Mechanisms of Sensory Processing," and the Alexander von Humboldt Foundation.

Accepted: February 23, 2011
Published: May 11, 2011

REFERENCES

Andersen, R.A., and Cui, H. (2009). Intention, action planning, and decision making in parietal-frontal circuits. Neuron *63*, 568–583.

Baldauf, D., Cui, H., and Andersen, R.A. (2008). The posterior parietal cortex encodes in parallel both goals for double-reach sequences. J. Neurosci. *28*, 10081–10089.

Basso, M.A., and Wurtz, R.H. (1998). Modulation of neuronal activity in superior colliculus by changes in target probability. J. Neurosci. *18*, 7519–7534.

Batista, A.P., Buneo, C.A., Snyder, L.H., and Andersen, R.A. (1999). Reach plans in eye-centered coordinates. Science *285*, 257–260.

Bisley, J.W., and Goldberg, M.E. (2010). Attention, intention, and priority in the parietal lobe. Annu. Rev. Neurosci. *33*, 1–21.

Calton, J.L., Dickinson, A.R., and Snyder, L.H. (2002). Non-spatial, motor-specific activation in posterior parietal cortex. Nat. Neurosci. *5*, 580–588.

Churchland, A.K., Kiani, R., and Shadlen, M.N. (2008). Decision-making with multiple alternatives. Nat. Neurosci. *11*, 693–702.

Cisek, P. (2006). Integrated neural processes for defining potential actions and deciding between them: A computational model. J. Neurosci. *26*, 9761–9770.

Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. Philos. Trans. R Soc. Lond. B Biol. Sci. *362*, 1585–1599.

Cisek, P., and Kalaska, J.F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: Specification of multiple direction choices and final selection of action. Neuron *45*, 801–814.

Cisek, P., and Kalaska, J.F. (2010). Neural mechanisms for interacting with a world full of action choices. Annu. Rev. Neurosci. *33*, 269–298.

Crammond, D.J., and Kalaska, J.F. (1994). Modulation of preparatory neuronal activity in dorsal premotor cortex due to stimulus-response compatibility. J. Neurophysiol. *71*, 1281–1284.

Crammond, D.J., and Kalaska, J.F. (2000). Prior information in motor and premotor cortex: Activity during the delay period and effect on pre-movement activity. J. Neurophysiol. *84*, 986–1005.

Cui, H., and Andersen, R.A. (2007). Posterior parietal cortex encodes autonomously selected motor plans. Neuron *56*, 552–559.

Dorris, M.C., and Glimcher, P.W. (2004). Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. Neuron *44*, 365–378.

Dorris, M.C., and Munoz, D.P. (1998). Saccadic probability influences motor preparation signals and time to saccadic initiation. J. Neurosci. *18*, 7015–7026.

Everling, S., Dorris, M.C., Klein, R.M., and Munoz, D.P. (1999). Role of primate superior colliculus in preparation and execution of anti-saccades and pro-saccades. J. Neurosci. *19*, 2740–2754.

Gail, A., and Andersen, R.A. (2006). Neural dynamics in monkey parietal reach region reflect context-specific sensorimotor transformations. J. Neurosci. *26*, 9376–9384.

Gail, A., Klaes, C., and Westendorff, S. (2009). Implementation of spatial transformation rules for goal-directed reaching via gain modulation in monkey parietal and premotor cortex. J. Neurosci. *29*, 9490–9499.

Genovesio, A., Brasted, P.J., Mitz, A.R., and Wise, S.P. (2005). Prefrontal cortex activity related to abstract response strategies. Neuron *47*, 307–320.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574.

Hampton, A.N., Bossaerts, P., and O'Doherty, J.P. (2006). The role of the ventromedial prefrontal cortex in abstract state-based inference during decision making in humans. J. Neurosci. *26*, 8360–8367.

Iyer, A., Lindner, A., Kagan, I., and Andersen, R.A. (2010). Motor preparatory activity in posterior parietal cortex is modulated by subjective absolute value. PLoS Biol. *8*, e1000444.

Kable, J.W., and Glimcher, P.W. (2007). The neural correlates of subjective value during intertemporal choice. Nat. Neurosci. *10*, 1625–1633.

Kable, J.W., and Glimcher, P.W. (2009). The neurobiology of decision: Consensus and controversy. Neuron *63*, 733–745.

Kim, B., and Basso, M.A. (2010). A probabilistic strategy for understanding action selection. J. Neurosci. *30*, 2340–2355.

Knill, D.C., and Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. Trends Neurosci. *27*, 712–719.

Lau, B., and Glimcher, P.W. (2008). Value representations in the primate striatum during matching behavior. Neuron *58*, 451–463.

Louie, K., and Glimcher, P.W. (2010). Separating value from choice: Delay discounting activity in the lateral intraparietal area. J. Neurosci. *30*, 5498–5507.

Medendorp, W.P., Goltz, H.C., and Vilis, T. (2005). Remapping the remembered target location for anti-saccades in human posterior parietal cortex. J. Neurophysiol. *94*, 734–740.

Medendorp, W.P., Goltz, H.C., and Vilis, T. (2006). Directional selectivity of BOLD activity in human posterior parietal cortex for memory-guided double-step saccades. J. Neurophysiol. *95*, 1645–1655.

Mountcastle, V.B., Lynch, J.C., Georgopoulos, A.P., Sakata, H., and Acuna, C. (1975). Posterior parietal association cortex of the monkey: Command functions for operations within extrapersonal space. J. Neurophysiol. *38*, 871–908.

Musallam, S., Corneil, B.D., Greger, B., Scherberger, H., and Andersen, R.A. (2004). Cognitive control signals for neural prosthetics. Science *305*, 258–262.

Padoa-Schioppa, C., and Assad, J.A. (2006). Neurons in the orbitofrontal cortex encode economic value. Nature *441*, 223–226.

Pesaran, B., Nelson, M.J., and Andersen, R.A. (2008). Free choice activates a decision circuit between frontal and parietal cortex. Nature *453*, 406–409.

Platt, M.L., and Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. Nature *400*, 233–238.

Rangel, A., Camerer, C.F., and Montague, P.R. (2008). A framework for studying the neurobiology of value-based decision making. Nat. Rev. Neurosci. *9*, 545–556.

Samejima, K., Ueda, Y., Doya, K., and Kimura, M. (2005). Representation of action-specific reward values in the striatum. Science *310*, 1337–1340.

Scherberger, H., and Andersen, R.A. (2007). Target selection signals for arm reaching in the posterior parietal cortex. J. Neurosci. *27*, 2001–2012.

Schlag-Rey, M., Amador, N., Sanchez, H., and Schlag, J. (1997). Antisaccade performance predicted by neuronal activity in the supplementary eye field. Nature *390*, 398–401.

Shen, L., and Alexander, G.E. (1997). Preferential representation of instructed target location versus limb trajectory in dorsal premotor area. J. Neurophysiol. *77*, 1195–1212.

Snyder, L.H., Batista, A.P., and Andersen, R.A. (1997). Coding of intention in the posterior parietal cortex. Nature *386*, 167–170.

Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2005). Choosing the greater of two goods: Neural currencies for valuation and decision making. Nat. Rev. Neurosci. *6*, 363–375.

Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. Science *304*, 1782–1787.

Tremblay, L., Gettner, S.N., and Olson, C.R. (2002). Neurons with object-centered spatial selectivity in Macaque SEF: Do they represent locations or rules? J. Neurophysiol. *87*, 333–350.

Wallis, J.D., and Miller, E.K. (2003). From rule to response: Neuronal processes in the premotor and prefrontal cortex. J. Neurophysiol. *90*, 1790–1806.

Wallis, J.D., Anderson, K.C., and Miller, E.K. (2001). Single neurons in prefrontal cortex encode abstract rules. Nature *411*, 953–956.

Weinrich, M., and Wise, S.P. (1982). The premotor cortex of the monkey. J. Neurosci. *2*, 1329–1345.

White, I.M., and Wise, S.P. (1999). Rule-dependent neuronal activity in the prefrontal cortex. Exp. Brain Res. *126*, 315–335.

Wise, S.P., di Pellegrino, G., and Boussaoud, D. (1996). The premotor cortex and nonstandard sensorimotor mapping. Can. J. Physiol. Pharmacol. *74*, 469–482.

Wunderlich, K., Rangel, A., and O'Doherty, J.P. (2009). Neural computations underlying action-based decision making in the human brain. Proc. Natl. Acad. Sci. USA *106*, 17199–17204.

Yanai, Y., Adamit, N., Israel, Z., Harel, R., and Prut, Y. (2008). Coordinate transformation is first completed downstream of primary motor cortex. J. Neurosci. *28*, 1728–1732.

Yang, T., and Shadlen, M.N. (2007). Probabilistic reasoning by neurons. Nature *447*, 1075–1080.

Zhang, M., and Barash, S. (2004). Persistent LIP activity in memory antisaccades: Working memory for a sensorimotor transformation. J. Neurophysiol. *91*, 1424–1441.

Supplemental information S1 – Model-based tuning analysis (related to Figures 3-6)

The analysis of the direction modality contrast (DMC) only addresses the spike rate difference between the direction with the maximum response (MD) and the opposite direction (OD) for each neuron. Other characteristics of spatial selectivity are not considered, like the exact preferred direction, which can lie between the probed directions, and gain-modulation effects between direct and inferred reaches, which have been shown to occur in both PRR and PMd (Gail et al. 2009). For example, a neuron which is tuned for the same reach direction in direct-cued and inferred-cued trials, but due to gain-modulation has a substantially different overall spike rate in the two contexts (e.g. 15 Hz in the direct context and 30 Hz in the inferred context), could show an asymmetric bimodal selectivity profile in the balanced PMG data. The DMC index in this case would be negative, indistinguishable from the DMC of a unimodal selectivity profile, and erroneously suggesting the representation of only a single spatial motor goal at the inferred position.

Therefore, as an alternative to the DMC, we modeled the to-be-expected selectivity profile in the PMG task (T^{PMG}) as a linear combination of the selectivity profiles in the direct-cued (d) and inferred-cued (i) DMG trials:

$$T^{PMG} = A(\alpha_d T_d^{DMG} + \alpha_i T_i^{DMG})$$

Each selectivity profile T consists of a 4-element vector with the neuron's responses to the four probed cue locations (left, right, up, down). With the model profiles we tested

three discrete alternative coding hypotheses. We fixed the linear coefficients and α_i to different preset values for each hypothesis and fitted *A* as the only free parameter on a least square basis. When both motor goals are encoded with equal strength, then the tuning in the PMG task should combine direct and inferred tuning of the DMG task with equal weights ($\alpha_d = \alpha_I = 0.5$; mean model). When a unique goal is encoded, then the PMG selectivity profile should either reflect the DMG selectivity profile of direct reach trials ($\alpha_d = 1$, $\alpha_i = 0$; direct model), or the inferred reach trials ($\alpha_d = 0$, $\alpha_i = 1$; inferred model), depending on which of these two alternatives was preferred or selected. We predicted the PMG selectivity profiles from the DMG profiles for these three models, and compared the mean square error (MSE) between the predicted and the measured profiles. We computed spike densities (Szucs, 1998) for a time-continuous representation of the neural response by convolving each spike train with a Gaussian kernel ($\sigma = 50$ ms) and averaging across trials. We then determined the best-fitting model for each neuron as a function of time. In each time bin the three models were fitted to the tuning profiles, separately in the PMG-CI and PMG-NC task. In the PMG-CI task the trials were sorted according to direct- and inferred-goal instruction. In the PMG-NC task, trials were sorted according to direct- and inferred-goal choice of the monkeys.

To quantify preference and selection encoding we defined a 'preference signal' and a 'selection signal' (Fig. S1b bottom). The preference signal was determined by the total number of neurons in the PMG-NC task of the balanced dataset, which best fit the mean model, irrespective of the choice of the monkey. That means if a neuron fits best with the mean model in direct- and inferred-choice trials for a particular point in time it contributed to the preference signal for that point in time. The selection signal was

likewise determined by the total number of neurons in the PMG-NC task of the balanced dataset, which best fit the direct model in direct-choice trials and the inferred model in inferred-choice trials.

Note that in the instructed PMG-CI trials we cannot determine the selection encoding during the memory period, since any preliminary selection signals would have to be overruled, and hence obscured, by the context instruction at the time of the 'GO'-cue. In non-instructed PMG-NC trials, on the other hand, preliminary selection signals could be detectable prior to the 'GO'-signal and would be predictive about the pending choice. This is why we only computed the preference and selection signals for the PMG-NC trials. After the 'GO'-signal we expect the selection signal to become dominant irrespective of the strength of the signal in the previous memory period, because the situation then is similar to a DMG task.

Note also that in the biased dataset the preference and selection signal can not be distinguished since the strong behavioral bias for inferred goals makes direct-goal choices extremely rare. This is why we analyzed the preference and selection signals only in the balanced dataset.

Model-based PRR results

In the PMG-CI task, most PRR neurons in the balanced condition (Fig. S1a; top left) were best fitted with the mean model (66.0 %), while inferred model (26.8 %) and direct model (7.2 %) were less prominent during the late memory period. In the biased dataset, the inferred model provided the best fit for most neurons (84.5 %), followed by the mean (12.8 %) and direct (2.7 %) models (Fig. S1a; bottom left). In Figure S1b (left) PMG-NC

trials were analyzed accordingly. Dominance of the mean model in the balanced dataset and dominance of the inferred model in the biased dataset together suggest predominant encoding of motor-goal preferences in PRR during movement planning.

In the balanced dataset preference and options encoding can not be distinguished. Similarly, we can not distinguish preference and selection encoding in the biased dataset. Our conclusion that the neural populations in our experiment mainly encode preferences are derived from the fact that the unimodal selectivity (inferred model) in the biased dataset rules out general options encoding, and the predominant bimodal selectivity in the balanced dataset (mean model) rules out general selection encoding. The latter dissociation can be best made with a choice-selective analysis of the PMG-NC trials. If the selection hypothesis was true, then the mean model (red curves) would always have to be less likely than the direct model in the trials with direct choices (green solid curve) or the inferred model in trials with inferred choices (blue dotted curve). The two curves at the bottom of the panels in Fig. S1b denote the fraction of neurons at each point in time which, according to the neuron's model fit, complied best with the preference or with the selection hypothesis. For the late memory period (shaded area) the percentage of neurons showing a selection signal was 2.9 %. For the neurons to be considered they had to best fit the direct model in direct-choice trials and the inferred model in the inferred-choice trials. The percentage of neurons that showed a preference signal in the same time window was 36.8 %. In this case to be considered the neurons had to best fit the mean model in direct- and inferred-choice trials. The preference signal was higher than the selection signal during the memory period, while selection encoding became dominant after the 'GO'-cue.

The fact that the direct model more often than the inferred model was the best fitting model in direct-choice trials, and the inferred model more often was the best model in inferred-choice trials, points to partial selection encoding. Yet, in hardly any neuron both conditions were simultaneously met, leading to the negligible fraction of truly selection encoding neurons during the memory period.

Model-based PMd results

PMd neurons in the balanced condition (Fig. S1a; top right) show a qualitatively similar result as those in PRR (mean model: 51.1 %; inferred model: 40.0 %; direct model: 8.9 %), when PMG-CI trials are considered. The same is true for the biased condition (inferred model: 67.4 %; mean model: 31.5 %; direct model: 1.1 %). Overall, the results suggest a predominant encoding of motor-goal preferences, like in PRR. Quantitatively, the difference in the fractions of neurons for which the mean or the inferred model provided the best fits was smaller in PMd than PRR. Also, in the PMG-NC trials (Fig. S1b, right), the monkeys' preliminary selections modulated a larger fraction of neurons in PMd, as can be seen from the stronger separation of the direct- (solid) and inferred-choice (dotted) curves and a reduced preference signal (21.1 %) compared to PRR. The percentage of neurons showing a selection signal during the late memory period was 7.6 %, which is more than double the amount compared to PRR. We did not observe such large differences between PMd and PRR in the similarity analyses presented in the main text (cp. Fig. 4b and Fig. 6c). Assuming that PMd neurons are only slightly more modulated by a preliminary selection than PRR neurons the different results could be the

consequence of the analysis method. The model based analysis classifies neurons by utilizing a winner-take-all mechanism, e.g. even if a model fit is only slightly better than another fit the assignment to a class is absolute. Therefore small modulatory differences between PMd and PRR can result in large differences in the number of neurons assigned to a particular class. The similarity analysis, on the other hand, does not enhance small differences, which explains why we cannot observe a similar difference between PMd and PRR in the similarity analyses presented in the main text (cp. Fig. 4 and Fig. 6c).

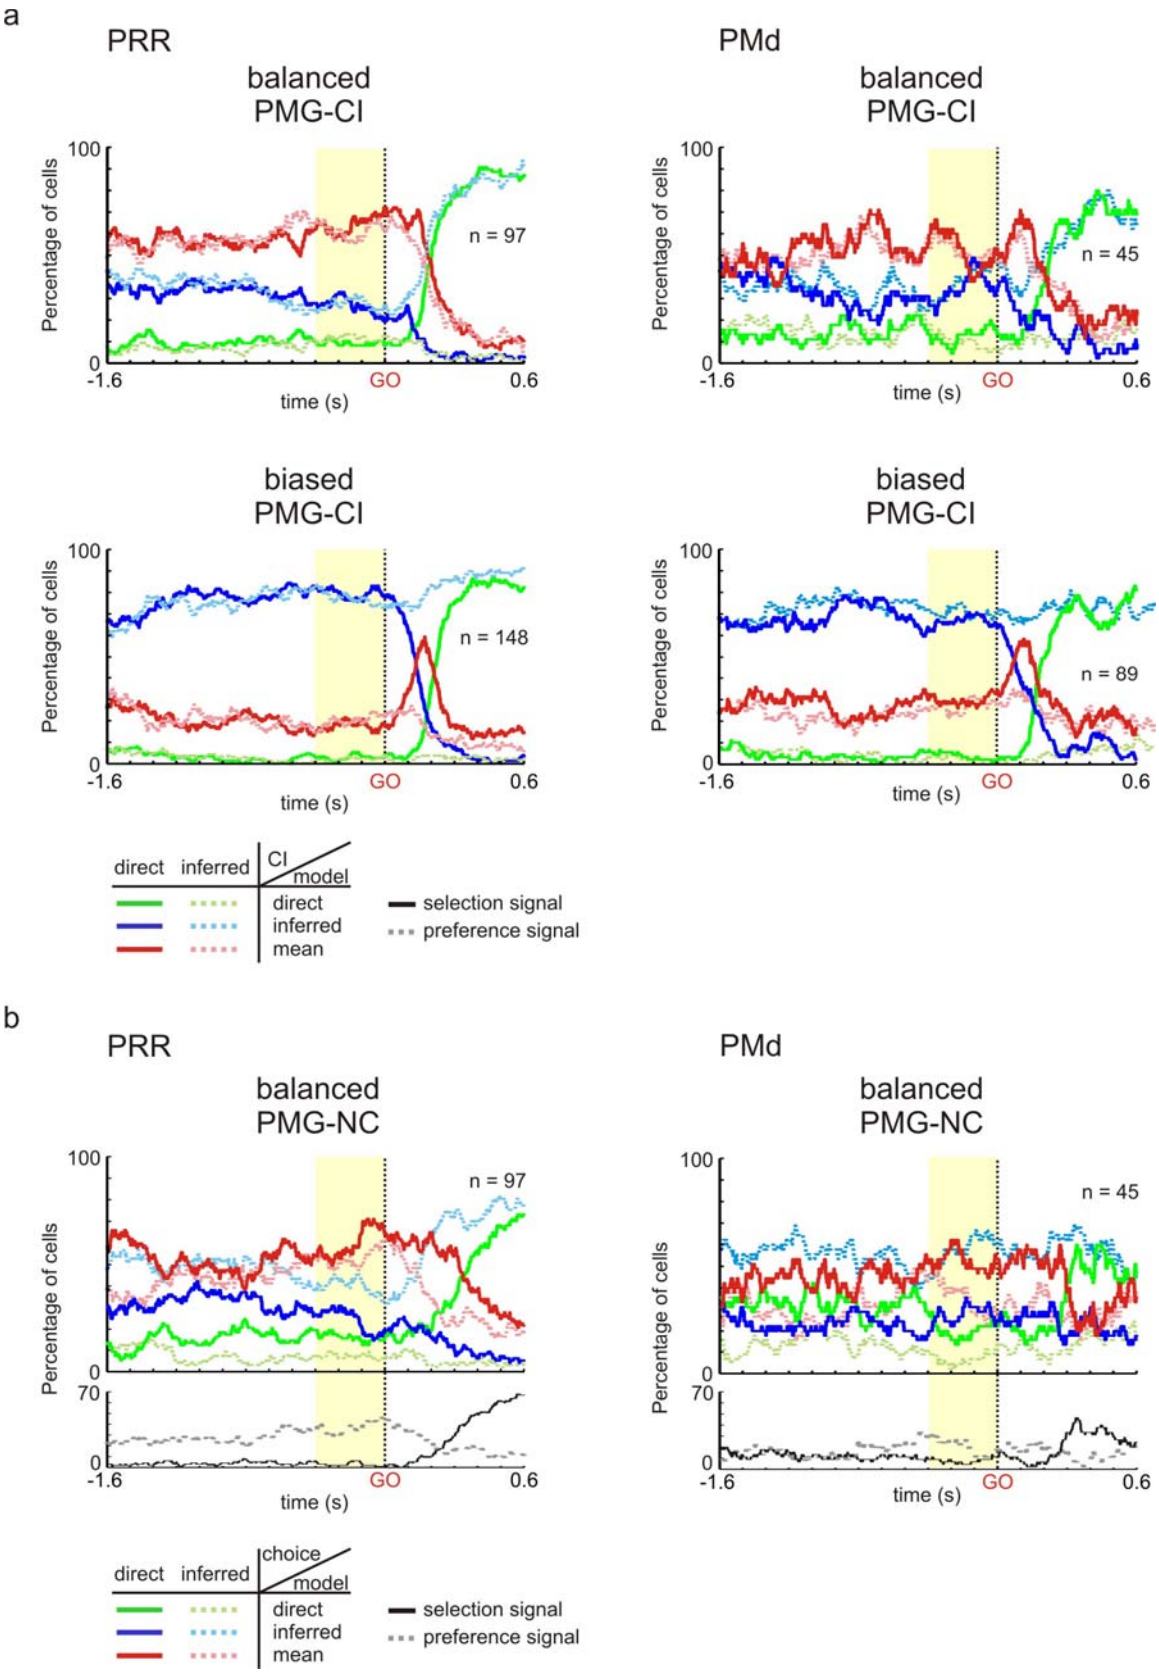
a
PRR
balanced
PMG-CI
PMd
balanced
PMG-CI
biased
PMG-CI
biased
PMG-CI
Percentage of cells
100
0
-1.6
GO
0.6
time (s)
n = 97
n = 45
n = 148
n = 89
direct
inferred
CI
model
direct
inferred
mean
selection signal
preference signal
b
PRR
balanced
PMG-NC
PMd
balanced
PMG-NC
70
choice
model

Fig. S1: Results of the model-based analysis for PRR (left) and PMd (right). (a) Model fits for the balanced (top) and biased (bottom) condition in the PMG-CI task. Curves depict the percentage of neurons which fitted best with the direct model (green), inferred model (blue), or mean model (red). Direct-instructed (solid) and inferred-instructed (dotted) trials are analyzed separately. The spike densities were aligned to the 'GO'-signal. The late memory period (shaded area) was used to calculate the mean percentages of neurons listed in the text (SI 1). (b) Same analysis for the PMG-NC task. The smaller graphs (bottom) show the selection (solid black) and preference (dotted grey) signals, as defined in the text.

Supplemental Information S2 – Control analyses for ruling out the selection hypothesis (related to figures 4b and 6c)

Sensitivity of choice-selective DMC similarity analysis

The DMC similarity analysis was used to determine if the monkeys used a trial-to-trial guess-and-switch behavioral strategy (4b and 6c), i.e. to test if the bimodal motor-goal representations could be an averaging effect across trials in which the monkeys randomly switched between planning a direct or inferred reach. To test the sensitivity of the DMC similarity analysis, we applied it to neuronal data of the DMG task. First, the mean firing rates for each reach direction in DMG without distinction between direct-cued and inferred-cued trials were used to calculate mean DMC values. The idea of this was to simulate a scenario in which a bimodal selectivity profile is artificially created as an artifact of averaging across direct- and inferred-choice trials. Second, all correct DMG trials were sorted according to the choice of the monkey (as instructed by the context cue prior to the memory period in DMG trials) to compute choice-selective DMC values. Since neurons in PRR and PMd typically showed motor-goal selectivity in the DMG task, choice-selective profiles were unimodal with high absolute values of the DMC (Fig. S2). As expected, the differences between the choice-selective and choice-indifferent DMC values (see Methods) are significantly different from zero for PRR (p_{direct}=0.0012; $p_{inferred}$=0.0007) and PMd (p_{direct}=0.0032; $p_{inferred}$=0.0018). This shows that our DMC correlation method is a sensitive measure, which would detect trial-to-trial guess-and-switch strategies if present.

Within-trial motor-goal switching

A bimodal selectivity profile of a neuron could also be the consequence of a rapid behavioral switching between the two possible motor goals during single trials. To rule this out, we computed the variance of the spike densities in trials with reaches to the direction of maximum response (MD) and to the opposite direction (OD) in the PMG task, and compared it with the variance in the MD and OD trials in the DMG task. The variance was computed across time (within the memory period) and across trials, without differentiating between direct-cued and inferred-cued trials. The variances in the DMG task was expected to be high, since MD and OD by definition induce different spike rate levels. If the switching hypothesis was true, then the variance in the PMG trials should be similarly high. If genuine potential motor goal encoding was true, then the variance in PMG trials should be smaller than in the DMG task, since MD and OD trials in PMG would induce similar spike rate levels. We tested our prediction by analyzing the normalized variance

$$Var_{norm} = \frac{\log(\mathrm{var}(spikedensity))}{\log(\mathrm{mean}(spikedensity))}$$

during the last 600 ms of the memory period. With the normalization we compensated for the fact that the variance in our data increased with increasing mean spike rate, like it has been shown for V1 and MT neurons (Snowden and Hess, 1992). We extended the analysis period to 600 ms, compared to the 300 ms for all other analyses, to take into account possible longer switch periods and to improve the statistical power.

For the balanced dataset in PRR we found that the variance in PMG was on average smaller than in DMG ($p=1.84 \times 10^{-7}$, paired Wilcoxon signed rank test; Fig. S2b). This difference in the spike rate variance argues against the hypothesis that the

bimodal selectivity profiles in PRR are mostly a consequence of averaging across short-term alternating choice-selective responses within or across trials. Instead it suggests that in PRR both potential motor goals are represented simultaneously within in each trial in the balanced dataset.

In PMd the variance in PMG trials was not lower than in DMG trials ($p>0.05$; Fig. S2d). This could be caused by slightly stronger trial-to-trial choice-selective signals, which the model-based analysis suggested to be present in PMd (see also Supplemental Information S1; Fig. S1b right). The stronger the choice-selective signal the higher the variance. This makes the variance measures in PMG and DMG more similar and reduces the distance to the unity line.

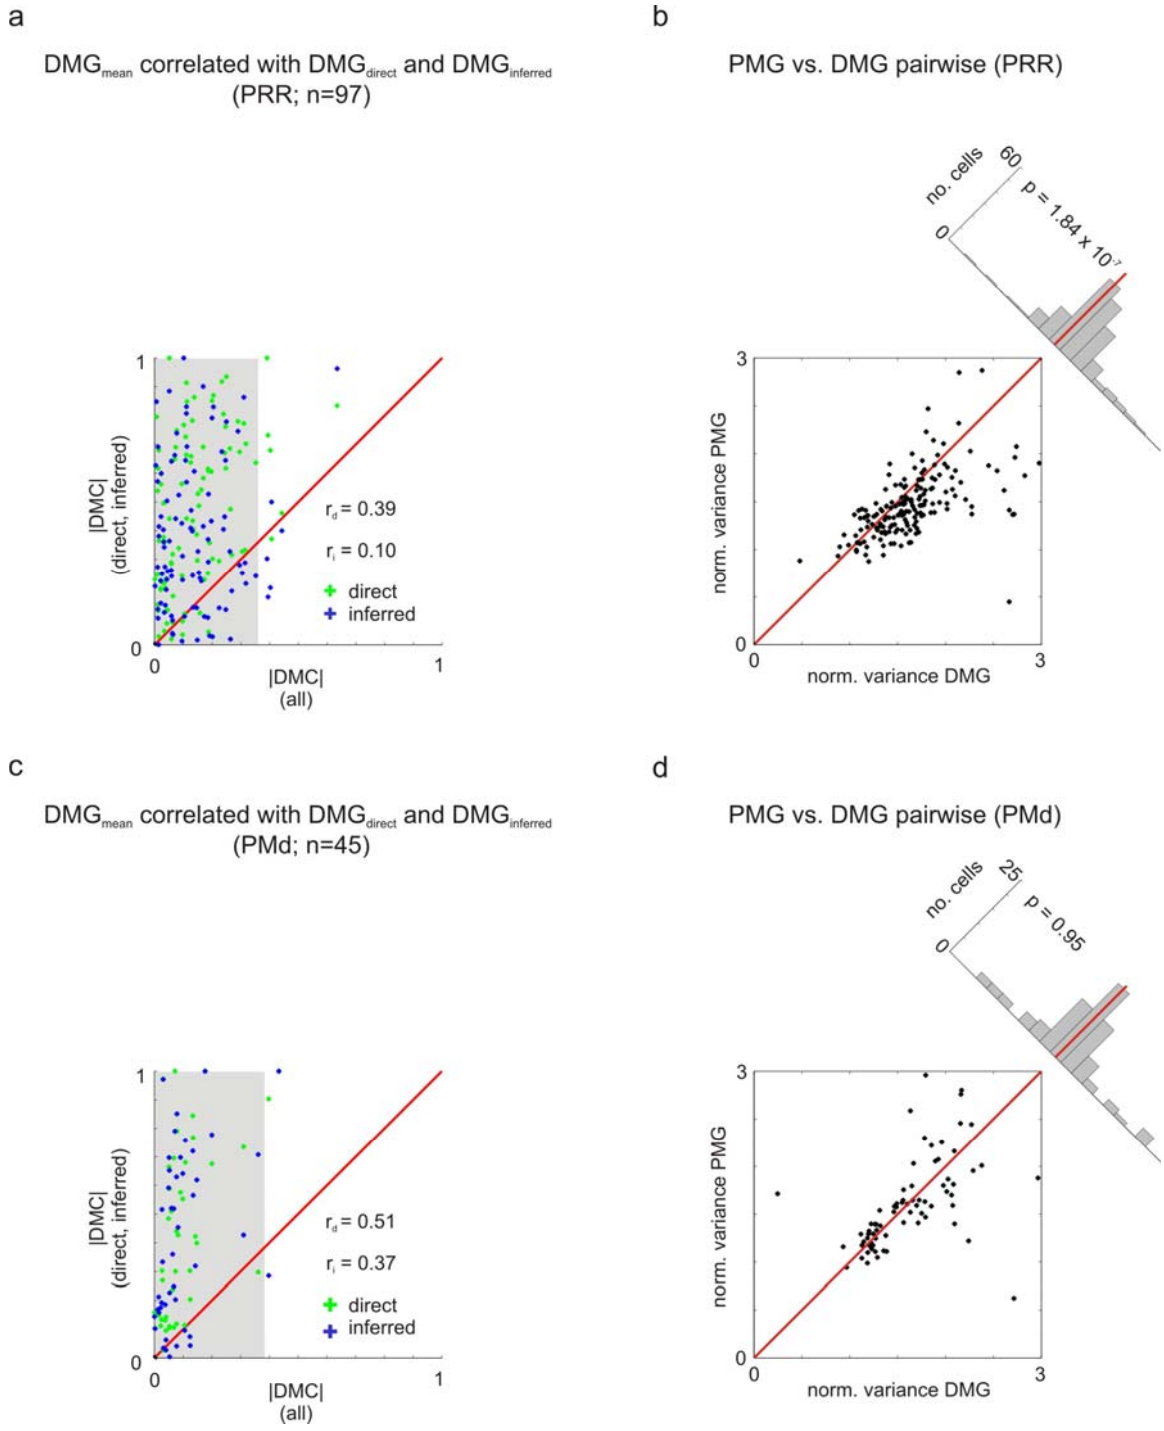

Figure S2: Control measures to rule out the selection encoding hypothesis. (a) DMC similarity analysis for the balanced dataset in the DMG task. The scatter of the data along the vertical axis demonstrates the sensitivity of this measure (see text for explanation). Same conventions as in Fig. 4b and 6c apply. (b) Within-trial motor-goal switching

analysis. The normalized variances for each neuron in MD and OD trials are compared between the DMG and PMG task (see text). Top right shows the distribution of the variance differences between DMG and PMG for all neurons. The p-value denotes the result of a paired Wilcoxon sign rank test.

Supplemental Information S3 – Source of the behavioral bias in the biased dataset (related to Fig. 5a)

Even though direct or inferred context instructions were equally likely in PMG-CI trials, and both choices were rewarded with equal amount and probability in PMG-NC trials, the monkeys' choice behavior and neuronal selectivity profiles showed a strong preference for the inferred reach goal in the biased dataset. At first glance, the preference for inferred reaches seems counter intuitive. In natural environments direct reaches are more common, and are easier to perform, since they do not require a spatial transformation.

One advantage of planning the less intuitive inferred reaches by default might be the preparation of the more difficult response alternative in order to comply with any constraint on reaction times after the 'GO'-signal. Alternatively, the more intensive training of the inferred task could have initially biased the monkeys. In the Bayesian sense, the frequency of exposure to inferred trials during training might have acted as prior probability which biased the decision (Körding and Wolpert, 2004; Diedrichsen et al., 2010). Since in the EPRS any choice strategy yields the same average reward, the monkeys did not have an incentive to deviate from any pre-existing behavioral bias. Even worse, any pre-existing bias should become enhanced since the monkeys, due to their bias, make the experience that more trials of the type for which they have a bias are being rewarded (simply because the monkey makes more of these trials).

Note, that our conclusions in the main manuscript are independent of the origin of the behavioral bias. Nevertheless, we believe that a combination of two of the above effects can be seen in our experiment. When, after several month of showing balanced behavior in the BMRS, one or our monkeys returned to the EPRS, he quickly developed a strong inferred-choice bias again (Fig. S3). This can hardly be explained with an imbalance between direct-cued and inferred-cued trials in the initial training stages. Rather, we interpret this as indication for a strategic advantage of planning inferred reaches in PMG trials. An initial small bias induced by this asymmetry might be self-enhanced via the Bayesian prior probability effect.

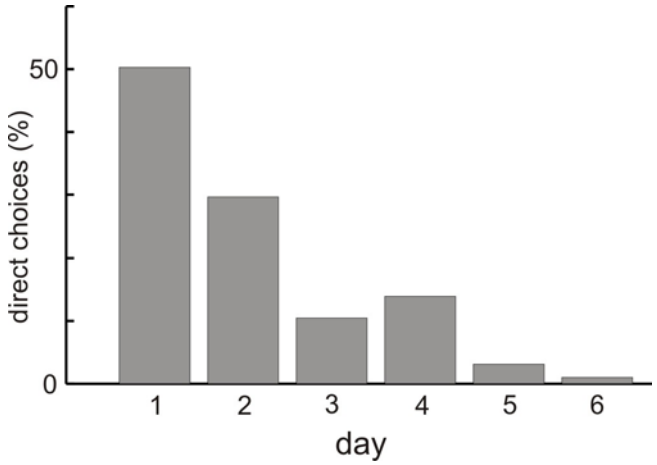

Figure S3: Re-occurrence of the inferred-choice bias in behavior when returning to the EPRS after several months of BMRS. The bars show the percentage of direct choices over six consecutive training days with an EPRS. Starting from a balanced behavior the monkey quickly developed a strong inferred-choice bias, even though direct-choices and inferred-choices were rewarded equally.

Supplemental Information S4 – Normalization and scaling factors (related to figures 3b and 5b)

The example neurons in Figures 3b and 5b, and the activity plots in Figures 3c and 5c show that the activity during the memory period is lower for two equally preferred motor goals than for a single preferred goal. To quantify this, we analyzed the distribution of the scaling factors *A* that were used for the model based analysis (see SI 1). The scaling factor was the fitted parameter in this analysis, and it quantifies the relative response strength of each neuron between the DMG and PMG tasks. For the inferred model (and also the direct model, which was never the dominant model, and therefore is not further considered) a scaling factor of 1 (one) means that the activity in the PMG task is not different from the DMG task. Instead, a scaling factor of 1 (one) in the mean model indicates halving of the activity in the DMG compared to the PMG task (see Experimental Procedures).

Figure S4 shows the distribution of scaling factors for the inferred model (S4a; biased dataset) and the mean model (S4b; balanced dataset). The mean scaling factors did not significantly deviate from 1 in either case, which indicates that the neural response strength on average is only half as strong when two equally preferred motor-goals are present, compared to a single one. This argues in favor of a mutual competition between multiple motor goals, as implemented in dynamic field models (Erlhagen and Schoner, 2002; Cisek, 2006) or other models of decision making (Averbeck and Seo, 2008; Eliades and Wang, 2008). This competition seems to be active even during reach planning, well before the final decision is enforced after the 'GO'-signal.

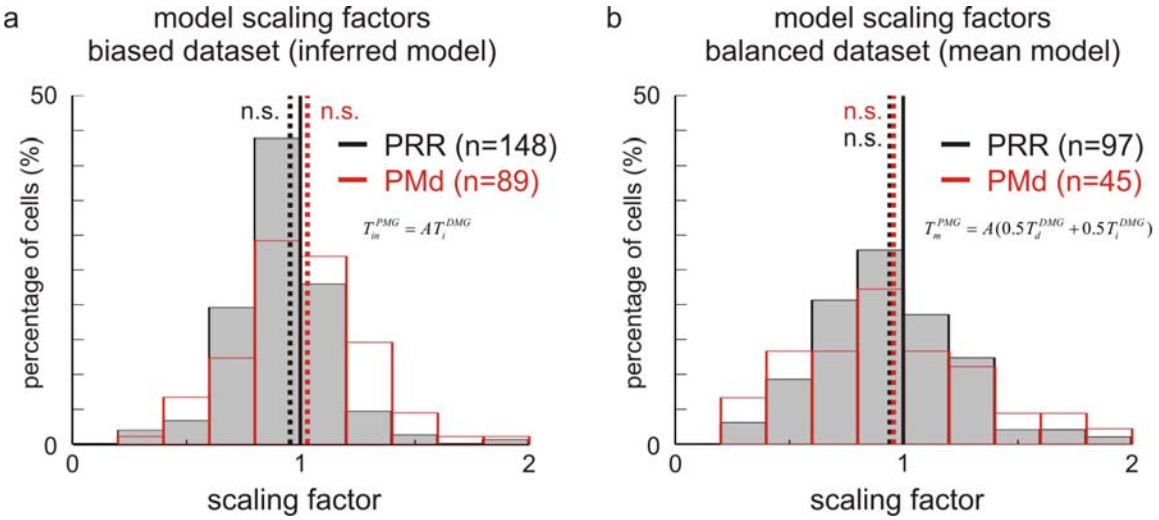

Figure S4: Scaling factors for model fit. (a) Distribution of the best-fit scaling factors (A in the linear model) of the inferred model for the biased dataset in PRR (grey; black outline) and PMd (blank; red outline). (b) Distribution of scaling factors A of the mean model for the balanced dataset (see Methods). The dotted lines indicate the mean values of the distributions, none of which significantly deviated from 1. This means that the neural response strengths to either of the two potential motor goals in the balanced data set were on average half as strong as the responses to a unique single motor goal in the biased dataset.

Supplemental Information S5 – Simultaneous behavioral and neural data collection with biased dataset (related to Figure 5)

In the main manuscript, behavioral testing of choice preferences with PMG-NC trials in the dataset with equal probability reward strategy (EPRS) was introduced at the end of the corresponding neural collection period. To ensure that independent behavioral and neural data collection causes no confounds on our interpretation, we also collected one biased dataset with the EPRS and strictly parallel neuronal and behavioral recordings (monkey A). This EPRS had a 100 % reward probability for direct and inferred reaches, and showed a strong behavioral bias for inferred choices (Fig. S2a), i.e. most reaches were directed towards the inferred motor goal (80 ± 4.6 %) and only a small fraction towards the direct goal (1.6 ± 1 %), as in the EPRS-50 dataset of the main manuscript. The equivalent choice bias between the EPRS-50 and EPRS-100 datasets is not surprising. Neither EPRS pushes a monkey to change its choice behavior, since the reward is independent of the pursued behavior. The distribution of direction modality contrast (DMC) values in the control experiment ($m=-0.27$; $p=8.9 \times 10^{-4}$; Fig. S2b) was qualitatively the same as in the biased dataset of the main manuscript (Fig. 4d). The model based analysis also yielded equivalent results between the main and control data (cp. Figs. S5c and S1a). This control experiment shows that simultaneous behavioral testing and neural recording in PMG-NC and PMD-CI trials yields identical results to subsequent testing and recording, and that the overall reward probability in an EPRS does not affect the results.

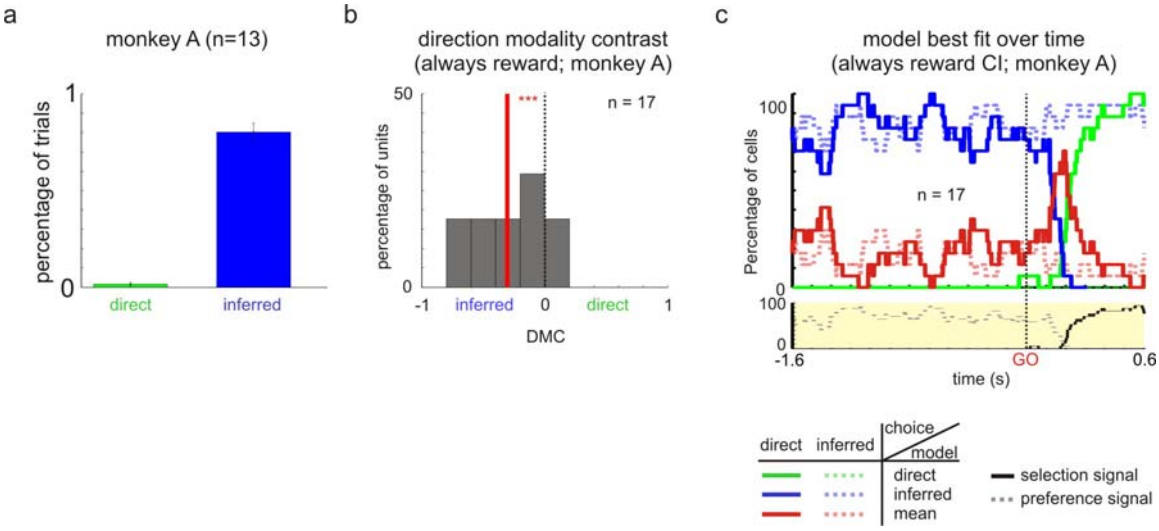

Figure S5: Preference encoding for a dataset with simultaneous recording of behavioral and neural recording during an EPRS in PRR of monkey A. Same conventions as in Figs. 5 and S1 apply. (a) The behavioral bias was the same when 100 % reward for either the direct or inferred motor goal was applied compared to when the 50 % EPRS was used (5a). The DMC values (b) and the model-based analysis (c) derived from neurons recorded in parallel with behavioral testing yielded the same results as for the subsequent testing (Fig. 5c).

Supplemental Information S6 – Alternative definitions of selection and options encoding (related to Methods)

For our analyses we had two distinct criteria that neurons had to fulfill to be included. First neurons had to be directional selective in the DMG task. This criterion was necessary because we only wanted to analyze neurons which were related to the reach task. Also, we needed a defined preferred direction to calculate the DMC values and selectivity profiles for the model based analyses, which we could only determine from neurons which were directionally selective in the DMG task. Second, neurons also had to be directionally selective in the PMG task. This restriction ensured that non-selective neurons were separated from bimodal neurons, which would both have a DMC value close to zero.

We also used two alternative neuron selection criteria. We did this in order not to exclude neurons which could have counted for the selection or option hypotheses, if these hypotheses were defined less strictly. First, we analyzed all neurons which were directionally selective in the memory period of DMG trials and after the 'GO'-signal in PMG trials, but not in the PMG memory period. These cells were only active when a definite motor goal was defined, but not when a preliminary (potential) motor goal existed, which could be interpreted as definite selection encoding (as opposed to a preliminary selection signal during the memory period of the PMG task). These criteria applied to only a small fraction of neurons in PRR (9 % biased; 4 % balanced) and PMd (3 % biased; 10 % balanced), which does not change our interpretations. Second, we analyzed all neurons which were not directionally selective in the DMG task and also not

in the PMG task after the 'GO'-signal, but in the PMG task during the memory period. These cells were only active when preliminary options existed, but not when a definite motor goal was defined, which could be interpreted as pure options encoding, which would only come into effect when more than one motor goal option is available. These criteria, too, applied only to small fractions of neurons in PRR (5 % biased; 2 % balanced) and PMd (10 % biased; 2 % balanced), which also does not change our interpretations of the main manuscript.

Reference List

Averbeck,B.B. and Seo,M. (2008). The Statistical Neuroanatomy of Frontal Networks in the Macaque. PLoS Computational Biology *4*, e1000050.

Cisek,P. (2006). Integrated Neural Processes for Defining Potential Actions and Deciding between Them: A Computational Model. Journal of Neuroscience *26*, 9761-9770.

Diedrichsen,J., White,O., Newman,D., and Lally,N. (2010). Use-Dependent and Error-Based Learning of Motor Behaviors. Journal of Neuroscience *30*, 5159-5166.

Eliades,S.J. and Wang,X. (2008). Chronic multi-electrode neural recording in free-roaming monkeys. J Neurosci Methods *172*, 201-214.

Erlhagen,W. and Schoner,G. (2002). Dynamic field theory of movement preparation. Psychological Review *109*, 545-572.

Körding,K.P. and Wolpert,D.M. (2004). Bayesian integration in sensorimotor learning. Nature *427*, 244-247.

Snowden,R.J. and Hess,R.F. (1992). Temporal frequency filters in the human peripheral visual field. Vision Research *32*, 61-72.

Szucs,A. (1998). Applications of the spike density function in analysis of neuronal firing patterns. J Neurosci Methods *81*, 159-167.

2.4 Sensorimotor learning biases choice behavior: A learning neural field model for decision making

In the final study presented in this work we developed a computational model of the decision making process in the frontoparietal network for reaching. Based on the electrophysiological findings of the previous study we constructed a multi-layer dynamic field model that was able to learn arbitrary spatial remapping rules. The model was trained to perform a task which required rule-based spatial remapping of a motor goal, which it learned by utilizing a hebbian style learning algorithm. The trained model developed activity patterns and neuronal tuning profiles consistent with the empirical data. We then examined how actions are planned in free-choice trials. Depending on input statistics and reward schedule, the model showed a either a preference bias towards one goal or an equal preference of the potential goals. These results are similar to the observations in monkeys performing the same task and provide a framework for the decision between different actions in ambiguous situations. The model successfully integrates visuomotor transformations, working memory and action selection into one framework for decision making.

Sensorimotor learning biases choice behavior: A learning neural field model for decision making

Christian Klaes[1*], Sebastian Schneegans[2*], Gregor Schöner[2], Alexander Gail[1]

[1] *Bernstein Center for Computational Neuroscience, German Primate Center – Leibniz Institute for Primate Research, Kellnerweg 4, 37077 Göttingen, Germany*

[2]*Institute for Neural Computation, Ruhr University Bochum, Universitätsstr. 150, 44780 Bochum, Germany*

**C.K and S.S. contributed equally to this work*

Corresponding author:

Alexander Gail, Ph.D.

Tel +49(0)551 3851 358

Fax +49(0)551 3851 452

E-mail: agail@gwdg.de

Short title:

A learning neural field model

Keywords:

Modeling, decision making, free-choice, sensorimotor transformation, spatial remapping

Manuscript

Pages: 35

Figures: 5

Words in summary: 250

Words in introduction: 500

Words in discussion: 1426

Total characters: 43.180

Acknowledgements:

This work was supported by the Federal Ministry for Education and Research (BMBF, Germany, grants 01GQ0433, 01GQ0814 and 01GQ0951) and the State of Lower Saxony (grant VWZN2563).

Summary

It has been proposed that the selection and specification of possible actions are not sequential operations, but that the decision for an action emerges from competition between different movement plans, which are specified and selected in parallel. For action choices based on ambiguous sensory input, this likely happens within frontoparietal sensorimotor areas. These areas have been shown capable of encoding alternative spatial motor goals in parallel during reach planning, and also show signatures of competitive value-based selection among these goals. Such sensorimotor decisions should not only be driven by the sensory evidence and reward expectancy in favor of the one or the other action choice, but also by the subject's learning history of the different sensorimotor associations. In previous computational models of competitive neural decision making in ambiguous situations the associations between input and corresponding motor output were always predefined. Such hard-wiring does not allow modeling of how decisions are influenced by the sensorimotor learning history or by changing reward contingencies. We present a dynamic neural field model (DNF) which can learn arbitrary sensorimotor associations with a reward-driven Hebbian learning algorithm, and which at the same time simulates the dynamics of action selection, as observed in monkey cortical recordings. With this adaptive model we investigated how input statistics and reward contingencies influence the model's choices when simulating recent empirical findings from decision experiments. The field model provides an integrated and dynamic account for the operations of sensorimotor integration, working memory, and action selection required for decision making in ambiguous situations.

Introduction

Sensory inputs usually does not automatically lead to unique action choices. Instead, the association of a specific sensory input with a specific action has to be learned from experience. The system should allow different actions dependent on behavioral context and remain adaptable by further experience. Hence, the learning mechanism that achieves these goals should be inherent to action selection, and as such to decision making.

While traditional psychological theories tended to view decision making as the outcome of a higher cognitive process which is separate from perception and action (Tversky and Kahneman, 1981), more recent neurophysiologically motivated ideas emphasize the integrative nature of sensorimotor processing and action selection (Cisek, 2006;Cisek and Kalaska, 2010;Andersen and Cui, 2009;Klaes et al., 2011). Several cortical areas form frontoparietal networks for making goal-directed saccades, like the lateral intraparietal area (LIP) and the frontal eye fields (FEF), or goal-directed reaches, like the parietal reach region (PRR) and the dorsal premotor cortex (PMd) (Wise et al., 1997;Burnod et al., 1999;Andersen and Buneo, 2002;Battaglia-Mayer et al., 2003). At the same time, neurons in these areas show signatures of valuation and selection of action, like modulation of the neural response with the subject's choice preference based on reward expectancy or other decision variables (Cisek and Kalaska, 2005;Sugrue et al., 2004;Gold and Shadlen, 2000;Platt and Glimcher, 1999;Musallam et al., 2004;Klaes et al., 2011;Pastor-Bernier and Cisek, 2011;Scherberger and Andersen, 2007;Pesaran et al., 2008;Kable and Glimcher, 2009;Kiani and Shadlen, 2009).

Previous models of decision making do not utiliyze learning in decision tasks with ambiguous choice situations. Conceptually, they either do not implement neural-inspired mechanisms of sensorimotor mapping, like the various threshold-models of decision making (see Gold and Shadlen (Gold and Shadlen, 2007) for review), are limited to solve predefined target-selection tasks (Cisek, 2006) or allowed sensorimotor association learning, but did not investigate decision making in ambiguous situations (Brozovic et al., 2007;Salinas, 2004;Fusi et al., 2007).

We present a dynamic neural field (DNF) model that aims to capture the processes of sensory integration, working memory formation and action selection in a context-dependent mapping task. The model implements a reward-driven Hebbian learning

mechanism that allows it to acquire simple associative sensorimotor mappings from experience. In our tasks, the model selects from a continuum of 'behavioral' options through an integrated competition process between potential action plans. With this framework, which reflects the conceptual idea of integrated sensorimotor and decision processing, we aim to provide explanations for previously observed choice behavior (Klaes et al., 2011) and make predictions on mutual interdependencies, which had not been reported before:

Hypothesis I: General arbitrary mapping rules, which associate spatial cue locations with arbitrary motor goals like in an anti-reach, can be learned through a reward-driven Hebbian mechanism by means of local associations.

Hypothesis II: The same continuous reward-driven learning mechanism which allows the acquisition of the mapping rules also adapts the choice behavior to changing probabilistic reward contingencies in free-choice trials.

Hypothesis III: Input statistics independent of the reward schedule should also bias the decision process in free-choice scenarios.

Methods:

Task Structure

A delayed-reach task that requires a context-dependent remapping of the target location was used to investigate the formation of reach plans in monkeys in an electrophysiological study (Klaes et al., 2011), and is emulated here for the model. Different variants of the task were employed, but they all share the same basic structure (Fig. 1c): Two cues are presented to the subject, a spatial and a contextual cue, that together determine the rewarded goal location for a reach movement. The spatial cue can be located at one of four equally spaced positions on a circle around a central fixation point, providing directional information. The contextual cue is presented as a color stimulus near the central fixation point, and determines a mapping rule for the trial. The mapping rule is either 'direct' (indicated by a green cue), meaning that the rewarded motor goal is located at the same position as the spatial cue, or 'inferred' (blue cue), which means the rewarded goal is located at the position diametrically opposite to the spatial cue location. The reach movement has to be executed after a delay upon a ‘go’-signal. The different task variants are described in detail at the end of this section.

Neurodynamic Model

We use an architecture consisting of multiple DNFs to capture the neural processes underlying cue perception, working memory for target location, movement plan formation and movement initiation. While the system architecture is meant to coarsely reflect the processing pathways in the cortex, and its structure is inspired by a previous model that explicitly aimed to capture activation patterns in specific cortical areas (Cisek, 2006), we refrain from claiming a strict correspondence between the elements of the model and distinct regions of the cortex. The single DNFs that comprise the architecture are defined based on functional considerations, and an individual field may reflect representations and functions that are distributed over different cortical areas (Erlhagen and Schöner, 2002).

The model learns the direct/inferred reach task through a reward-driven Hebbian learning mechanism, but is strongly pre-structured so that it can perform elementary behaviors without any training: It initially supports a direct mapping from a stimulus location to a

reach plan, has the capacity to form working memories of stimulus locations, and selects between competing reach plans through lateral interactions within dynamic representations. This pre-structuring simplifies the learning problem compared to most artificial neural network models that typically begin with weak random connectivity, but also allows a more accurate emulation of the task learning in monkeys. The learning procedure in the electrophysiological study presupposes the monkeys to spontaneously produce certain behaviors, like reaching to a visual stimulus. To adequately capture the learning process and specifically the effects of trial and reward statistics, we have to start with a similarly pre-structured system.

Dynamic neural fields

DNFs describe the neural activation patterns underlying behavior through the evolution of continuous activity distributions over time, emphasizing the role of attractor states and instabilities (Erlhagen and Schöner, 2002;Schneegans and Schöner, 2008). DNFs are based on the concept of population coding, in which a value along a certain feature dimension (e.g. the location of a visual stimulus or a parameter of a planned movement) is represented through the distribution of activity within a population of neurons. These neurons have different tuning functions that sample the underlying feature space (Averbeck et al., 2006;deCharms and Zador, 2000). Abstracting from the discrete spiking neurons, DNFs directly describe the activity distributions over the underlying feature space (Wilson and Cowan, 1973;Erlhagen et al., 1999;Erlhagen et al., 1999;Deco et al., 2008;Amari, 1977). This activity distribution evolves continuously in time under the influence of external input and lateral interactions, governed by a differential equation of the form

$$\tau\, \dot{u}(x,t) = -u(x,t) + h + S(x,t) + \int w(x,x') f\big(u(x',t)\big) dx' + \Theta.$$

Here, $u(x,t)$ is the activity at time t for a position x along the underlying feature dimension, $\dot{u}(x,t)$ is its rate of change over time, which is scaled with a time constant τ, and h is the (negative) global resting level for the field activity. Any point in the field receives external input $S(x,t)$ as well as endogenous input from other parts of the field, and is affected by additive noise Θ representing unspecific input and spontaneous activity. The lateral interactions are described by an interaction kernel $w(x,x')$, typically

consisting of a local excitatory and a long-range inhibitory component. This reflects the pattern of mutual excitation between neurons with similar tuning curves and inhibition between those with dissimilar tuning curves found throughout the neocortex. The interactions are only triggered if the activity at a field position exceeds the threshold (set arbitrarily at zero in our models) of a sigmoid output function,

$$f(u(x,t)) = \frac{1}{1 + \exp\big(-\beta u(x,t)\big)}.$$

The pattern of lateral interactions promotes the formation of localized bumps of activity, or peaks, as attractor states of the field dynamics (Fig. 1b).

These peaks serve as units of representation, yielding a certain value along the underlying feature dimension through their position in the field. Depending on the interaction parameters, different dynamic regimes can be achieved: With moderately strong, local interactions, the peaks can provide a representation of the current inputs that is stabilized against fluctuations, with a sharp transition from a sub-threshold state to a peak when a stimulus is first detected. For stronger self-excitation (balanced by sufficient inhibition), peaks may become self-sustained in the absence of input, yielding a model of working memory (similar to the implementation with spiking neurons described by Wang (Wang, 2001)). If strong global inhibition is present in a field, a competitive regime is created in which only a single peak can form, generating a selection decision that is stabilized over time.

For numerical simulations, the conceptually continuous field dynamics are discretized in space and time. To perform comparisons with electrophysiological data, the field output at one point in the field is equated to the firing rate of neurons with a corresponding tuning function.

Model architecture

The dynamic model for context-dependent reaching consists of a set of interconnected DNFs and discrete nodes that can be organized into four levels: Perception (spatial and context input fields), memory and association (association field), movement planning (decision field) and movement initiation (motor field) which are shown in Figure 1a.

A direct pathway from perception to movement planning and initiation implements a sensorimotor mapping for visually guided reach movements that is functional prior to any learning. The direct pathway comprises three DNFs defined over a one-dimensional space in angular coordinates (with circular boundary conditions). This space represents either the location of the spatial cue in the task (as its direction from the central fixation point) or the direction of a reach movement. The first of the three fields is the *spatial input field*, which features relatively weak, local interactions to form a stabilized representation of a currently presented spatial cue. It projects in a topological fashion onto the *decision field*. This field has moderate local excitatory and global inhibitory interactions, producing a soft competition behavior between different regions of the field. While these interactions promote the concentration of activity in a single region, they still allow multiple activity peaks to exist simultaneously if they are driven by multiple localized inputs. The decision field in turn reciprocally projects to the *motor field* in a topological manner. The motor field itself features stronger self-excitation and global inhibition, producing a strong selection behavior that only allows a single stabilized activity peak to exist at a time. It is held at a low resting level during most of the time, so that it cannot form a peak from the decision field's input alone. Only when the 'go'-signal is given in the task, the activity of the motor field is globally boosted and an activity peak can form, simulating a 'gating' mechanism for movement initialization. Similar 'gating' mechanisms have been described for saccade generation (Kobayashi and Isa, 2002).

There is an additional indirect pathway from the spatial input to the decision field, which runs through the *association field*. This field spans two dimensions, one of them corresponds to the angular spatial representation used in the other fields, the other one is initially not associated with a specific feature and instead provides redundancy in the existing representation that allows further specialization through learning. The association field receives input from the spatial input field that is topologic along the spatial dimension and unspecific along the other dimension, such that a single peak in the spatial perception induces a ridge of activity in the association field. The lateral interactions will then induce the formation of a localized peak from this ridge input.

Strong lateral interactions make these peaks self-sustained, such that they can serve as working memory of the spatial cue location after the cue itself is gone.

A second input to the association field comes from a set of two *contextual neurons (context input)*. These provide a simple symbolic representation of the color cue indicating the mapping rule for the current trial. The projection to the association field is initially weak and unspecific, with randomly initialized weights from both neurons to the whole field. These connections are modified during the learning phase as detailed below. The association field itself provides a second input to the decision field. The projection is initially topologic along the spatial dimension, so that it supports a delayed reach movement to the memorized location of a previously presented spatial cue, but it is likewise subject to learning.

Relationship to neurophysiology

The concept of representing the movement plan and the motor command for the reach movement through a population code over possible movement directions, as used in the decision and motor field in the model, is consistent with previous analyses of neural activity patterns during reach movements (Georgopoulos et al. 1986, Bastian et al. 2003). To draw the connection to the results of the experimental study, we compare the activity in the decision field to electrophysiological data from PRR (Klaes et al., 2011). We do not claim a one-to-one relationship between the field and the cortical area, however, since similar representations can also be found in PMd (Cisek and Kalaska, 2005). The function of the association field might likewise be distributed over parts of PRR and PMd (with possible contributions from the prefrontal cortex). Different regions in this field develop a combined selectivity for reach direction and context input that is consistent with a gain modulation by context described for neurons in these areas (Gail and Andersen, 2006). The motor field can be equated more directly to parts of the primary motor cortex (M1) or alternatively caudal parts of PMd, which both show strong correlations with the final motor goal. The two sources of input, spatial input field and contextual input field, resemble the visual input coming from visual extrastriate areas (Blatt et al., 1990;Johnson et al., 1996) and presumably top-down, rule-based influences from prefrontal cortex (Wise et al., 1997), respectively.

Learning

The projections from the context neurons to the association field and from the association field to the decision field are adapted according to a reward-driven Hebbian learning rule (Suri and Schultz, 1999;Fusi et al., 2007;Vasilaki et al., 2009;Stringer et al., 2007). We use two variants of the basic Hebbian rule that incorporate an implicit limit of weight growth, the 'instar' and 'outstar' learning rule in the formulation of Marshall (1995). These rules have successfully been used in topographical dynamic neural networks that are comparable to DNFs (Grossberg and Williamson, 2001;Kalarickal and Marshall, 2002). We further adapted them to be used in a reward-dependent manner: As in the original rules, we will change the weights to make them more similar to the output pattern of one group of neurons if the reward is positive, but in addition we will also make the weights more dissimilar if the reward is negative. The connection weights from the context neurons to the association field are updated according to the reward-dependent instar rule:

$$\Delta w_{ac}(\vec{x}, i) = \eta(r) \cdot f\big(u_a(\vec{x})\big) \cdot [g_c(i) - w_{ac}(\vec{x}, i)]$$

$$g_c = \begin{cases} f(u_c) & \text{for } r > 0 \\ \dfrac{\left|f\big(u_c(i)\big)\right|}{\left|1 - f\big(u_c(i)\big)\right|}\big(1 - f(u_c)\big) & \text{for } r < 0 \end{cases}$$

Here, $w_{ac}(\vec{x}, i)$ is the weight from the context node *i* to position $\vec{x}$ in the association field, Δw_{ac} is the change of that weight in one trial, $f\big(u_a(\vec{x})\big)$ is the association field output at that position and $f\big(u_c(i)\big)$ the output of context node *i*. The learning rate $\eta(r)$ depends on the reward *r* for that trial, and is larger if the reward is positive. The normalization in the case of negative reward is introduced to ensure that the overall weight changes are comparable for successful and fail trials.

. It is positive for successful trials and negative for failed trials. The weights from the association field to the decision field are adapted according to the reward-dependent outstar rule:

$$\Delta w_{da}(\vec{y}, \vec{x}) = \eta(r) \cdot [g_d(\vec{y}) - w_{da}(\vec{y}, \vec{x})] \cdot f\big(u_a(\vec{x})\big)$$

$$g_d = \begin{cases} f(u_d) & \text{for } r > 0 \\ \frac{|f(u_d)|}{|1 - f(u_d)|}\big(1 - f(u_d)\big) & \text{for } r < 0 \end{cases}$$

Analogously, $w_{da}(\vec{y}, \vec{x})$ is the weight from position $\vec{x}$ in the association field to position $\vec{y}$ in the reach planning field, Δw_{da} is the change of that weight, and $f\big(u_d(\vec{y})\big)$ is the output of the reach planning field at position $\vec{y}$.

The learning rule is applied only at the time when the reward for one trial is given after a sufficiently strong peak has formed in the motor field, indicating the initialization of a reach movement. The direction of the planned reach, given by the position of the activity peak in the field, is compared to the rewarded goal location according to the task requirements. The trial is considered a success, with a reward value of +1, if the reach direction falls within a tolerance window (±9º) around the desired goal direction, and a failure with a reward value of -1 otherwise.

Task variations

Inferred-reach (IR) training task: This task was used to train the model, which is initially only able to execute direct reaches, to perform both direct and inferred reaches depending on the context (Fig. 1d). At the beginning of an IR training trial a contextual cue and a spatial cue at one of four possible locations were presented. Our approach was then to briefly show a second spatial cue at the end of the memory period to indicate the correct goal location according to the remapping rule. Over many trials the salience of the second spatial stimulus was reduced until it eventually disappeared. By this the first spatial stimulus location was associated with a movement to the diametrically opposite location. For the training of the model the saliency is steadily reduced from 1 to 0 within 1000 trials. Trials with inferred and direct contextual cues were intermixed in the training so that the direct contextual cue was associated with the direct mapping rule (80% inferred and 20% direct if not indicated otherwise). Before any of the other task variants was applied to a model it was first trained with this task, to be able to perform context dependent direct/inferred reaches (IR trained model).

Definite Motor-Goal task (DMG): This task was used as a control condition to test if the model properly learned to make 'direct' and 'inferred' choices depending on the contextual cue after IR training. In the DMG task the spatial and the contextual cue were presented simultaneously at the beginning of the memory period (Fig. 1; left). The cues were only presented briefly and the relative presentation times, compared to the memory period, were chosen to be equivalent to the physiological study. The two mapping rules and four locations (see above) were presented with equal probability. Learning rate was set to zero in this control condition.

Potential motor-goal task with context instruction (PMG-CI): This task was used to examine the ongoing decision making process in situations with incomplete information. It was a variation of the DMG task in which the spatial and the contextual cues were separated in time. First only the spatial cue was presented. Therefore two potentially rewarded motor goals remained equally possible throughout the memory period, either at the location of the spatial cue or diametrically opposite to it. At the end of the memory period when the contextual cue was presented this ambiguity was resolved and only one rewarded target remained. During testing in this task the learning rate was set to zero.

Potential motor-goal task with no context instruction (PMG-NC): In the PMG-CI trials, the ambiguity that was created by not showing a contextual cue, was resolved at the end of the memory period. To test the free-choice behavior of the model we used a task in which no context instruction was shown at all. In this case two different reward schedules decided about which trials were rewarded and which not (see below). When the model was training PMG-NC trials, these were interspersed with PMG-CI trials (40:60 ratio).

Reward schedules

In the free-choice condition two algorithms determined which of the two potential motor goals would be rewarded. In the *equal probability reward schedule* (EPRS) both potential locations were rewarded with equal probability, irrespective of the choices of the model. From a game theoretical perspective the situation is equivalent to a matching pennies game in which the computer's strategy corresponds to the Nash equilibrium. With this

reward schedule we did not expect any change of preference in the subjects since the expected reward is independent of the subject's choice and no behavioral strategy leads to more than 50 % reward. In the *bias minimizing reward schedule* (BMRS) the success history was taken into account to decide which motor goal would be rewarded. Any behavioral bias for one of the motor goals was punished by lowering the probability of reward for that goal, so that the behavioral strategy that yields the highest reward (=50%) is one in which both motor goals are chosen with equal probability.

Monkey behavioral and electrophysiological data

The monkey behavioral and neuronal data and analyses that we refer to in this study are taken from a previous electrophysiological study and are described in detail elsewhere (Klaes et al., 2011). The population and single cell data that is presented comes from the parietal reach region (PRR) of two monkeys. Previously unpublished behavioral data is presented from these monkeys to test one of the predictions of the model (see Results: Field size and learning limitations).

Results:

Learning an arbitrary mapping rule

When the IR training task is applied to the adaptive DNF model, it acquires the formerly unknown spatial mapping rule in addition to the default mapping by forming individual associations of the possible stimulus combinations with the rewarded reach directions. These associations are formed in the following way: When the spatial cue is presented at the beginning of the trial, it induces an activity peak at the corresponding location in the association field and at the same time in the decision field (via the direct pathway). The association field peak remains self-sustained after the input disappears, and keeps supporting the activity in the decision field. The simultaneously presented contextual cue activates the corresponding context neuron, which likewise retains this activation through the neuron's self-excitation.

When the second spatial cue appears that indicates the desired reach goal, it also drives activity in the decision field – for inferred trials, at a location shifted by 180º from the initial cue location – and overrides the initial reach plan. The memory peak in the association field remains largely unchanged, as it is strongly stabilized by the field interactions and suppresses the formation of new peaks. The 'go'-signal given together with the second spatial cue leads to the execution of the currently selected reach plan. If the selected reach movement matches the goal location, the projections from the active regions in the association field to those in the decision field are strengthened (and others are weakened), eventually leading to the development of projections that implement the inferred reach. At the same time, specific connections from the context neurons to the association field form: Through the reinforcement of initial random variations in the peak positions, different areas in the association field develop selectivity for one of the context inputs.

The resulting connection patterns are illustrated in Figure 2. While connectivity from the context neurons to the association field is initially weak and unspecific (Fig. 2a, left), at the end of the training separate areas with a preference for either the direct or inferred context signal have formed around all four spatial cue locations in the association field (Fig. 2b, left). For those locations that show preference for the inferred context signal, the original topographic projections to the decision field (Fig 2a, right) have shifted by 180º

to implement the inferred reach response (Fig 2b, right). Spatial positions which have not been trained (i.e. spatial cues never appeared at those locations) do neither show a shift of projections nor a sensitivity for one of the contexts, even though there is a slight local spill over from trained locations due to the width of the interaction kernels.

The IR trained model was then tested in the DMG task, and reached a performance of 100% (n = 4000). This confirms that the training succeeded: The model can perform both the direct and inferred reach and flexibly switch between the two mappings depending on the context signal.

Input statistics during learning bias movement plan formation

One of the most surprising findings in the electrophysiological study was the observation of a strong bias in the trained monkeys to prepare and execute the inferred reach if the contextual cue was presented only after the memory period (in PMG-CI trials) or omitted completely (PMG-NC trials). We hypothesize that this bias arises from the input statistics during the early training of the monkeys, where the number of inferred reach trials substantially exceeded the number of direct reach trials. We use the adaptive DNF model to show how the reward-dependent Hebbian type learning rule can reproduce this bias in the decision process as an effect of the input statistics during training.

The IR trained model is capable of performing the PMG-CI task, in which the context cue is given only at the end of the memory period, with high reliability without additional training (88.0% correct; n = 2000). In the model, the initial presentation of the spatial cue induces the formation of a sustained activity peak in the association field. In the absence of context input, the separate regions with different context preferences do not influence the activity distribution in the DNF, and the peak typically spans both of them. The subsequent presentation of a context input changes the attractor states of the DNF, and the activity peak shifts towards the region that has a preference for the given context. The projections from that region to the decision field then select the appropriate action.

When the model is trained with a ratio of 80% inferred trials to emulate the intense inferred reach training procedure in the electrophysiological study, it develops a similar bias to prepare the inferred reach in PMG trials. In addition, the time course of activity in the decision field qualitative reproduces the recorded neural activity in PRR (Fig. 3a):

Following the presentation of the spatial cue, activity first forms at the direction of the direct reach (as a result of the direct pathway in the model), but is quickly replaced by a reach plan for the inferred movement. This plan is retained if the inferred context cue is given at the end of the memory period, or switches back to the direct reach if the direct context cue is presented.

The cause for this bias in the model is that the regions in the association field that have developed a preference for the inferred context are substantially larger than those for the direct context, as a result of the inferred reach being performed more frequently during training (see Fig. 2b). When the activity peak in this field forms before the presentation of the context cue, it typically covers a larger area of the inferred-context region, resulting in a stronger projection to the opposite reach direction in the decision field. The competitive interactions in the decision field then suppress the activity for the weakly excited direct reach direction.

Action selection in free-choice trials

For both the electrophysiological study and the model, we used the PMG-NC trials to test whether the bias in the observed movement plan also carries over to a biased response behavior in the absence of the context cue. PMG-NC trials and PMG-CI trials were interspersed (60% PMG-CI; 40% PMG-NC), and the reward schedule in NC trials was independent of choice (EPRS, see Methods). In 85% (±2%) of the PMG-NC trials with EPRS the monkeys had chosen the inferred motor goal. The adaptive DNF model selected the inferred motor goal in 96.6% of PMG-NC trials (n=1000), when it had been trained with 80 % inferred trials.

To further confirm our hypothesis that the input statistics during training directly influence the choice behavior, we systematically varied the ratio of direct to inferred trials during the IR training in the model and tested the resulting choice behavior and spatial response profiles in the decision layer with the PMG-NC task, when no contextual cue was presented at the end of the trial (Fig. 4). For testing the free-choice behavior we set the learning rate to zero in the PMG-NC trials. The number of inferred choices increases continuously with the ratio of inferred trials during IR training (Fig. 5a), in an approximately sigmoidal fashion (logistic function fit: m = 0.578, beta = 22.105; MSE =

$8.02 * 10^{-6}$). The sigmoidal curve is not centered at 50% inferred training trials, but at a slightly higher ratio, which is an effect of the direct reaches being the default action before training and the direct pathway from spatial input field to decision field. The difference in the underlying activation strength for inferred versus direct goal representations in the decision layer during the memory period also increases steadily and approximately linearly (linear fit: m = 16.1, b = -9.02, MSE = 0.62) with the fraction of inferred trials during IR learning (Fig. 5b).

We note that this result is indeed an effect of the input statistics, not of the expected reward for different choices: Even in training sets with 100% success rate for both direct and inferred reaches (due to a sufficiently strong second spatial cue guiding the action selection), the described biases still developed in the model.

How the reward schedule changes choice preferences

In the electrophysiological study, the monkeys were subsequently trained with a reward schedule which takes the success history into account and discourages biased behavior in order to balance the behavioral choices (BMRS, see Methods) in order to test whether they would develop a balanced movement plan and represent multiple possible reach directions simultaneously. We emulate this experiment in the model, thereby testing whether the learning mechanism can adapt the model's choice behavior in response to a new reward schedule. For this additional training, the weights that had developed in the previous testing phase with PMG-NC trials were taken as starting conditions. Similar to the monkeys in the electrophysiological study, the model developed a balanced choice behavior under the new reward schedule (49% direct trials, n = 2000). Correspondingly, during the memory period the two potential motor goals are equally represented in the decision field (Fig. 3c), as had been the case in the electrophysiological data (Fig. 3d). This confirms – as we expected based on the reward-dependent learning rule that is used – that the DNF model is capable of adapting its choice behavior to increase its overall reward probability, in a fashion that is consistent with the experimental data.

Field size and learning limitations

The metrics of the representations and interactions in the DNF model, specifically the width of the excitatory local interaction kernel relative to the size of the spatial dimension, limits its capability to learn an arbitrary number of associations. If the number of associations to be learned is increased without changing these metrics, the activated regions in the layers start overlapping. This leads to a strong susceptibility of the movement plan representation in the decision layer to small noise variations, which are amplified by the competitive interactions in the field, and ultimately produces an increased number of wrong decisions. Figure 5a shows a confusion matrix of the instructed directions and the choices of the model for four trained directions, where almost all trials are directed correctly. If the number of associations is increased to eight directions in two contexts, using the same model parameters, then errors occur in which mostly neighboring directions are confused (Fig. 5b). We found a similar pattern of errors, albeit with a slightly higher overall error rate, in the behavioral data of monkeys (Fig. 5c and d, previously unpublished).

We note that the observed errors in the experimental data and the model are not the result of a limited precision or specificity in the general movement plan formation: The monkeys have been able to reliably perform delayed direct reaches to arbitrary cued locations in their workspace, and the model can reproduce this behavior before training. The observed limitations therefore appear to be a result of the formation of new associations during IR training.

Discussion:

Neurophysiological data suggests that decision making and movement planning share an overlapping neural substrate in frontoparietal sensorimotor areas (Cisek, 2006;Klaes et al., 2011;Pastor-Bernier and Cisek, 2011;Glimcher, 2003;Gold and Shadlen, 2007;Andersen and Cui, 2009). In this case, the competitive interaction between alternative action plans for the purpose of selecting between them, i.e. the decision making, and the learning of sensorimotor associations should not be independent processes. With the current simulation study we could show how a DNF model, which

can learn new spatial remapping tasks via a reward-driven Hebbian learning rule, is at the same time sensitive to changing reward schedules and to input statistics in its free-choice decision behavior. In a free-choice task with two distinct reward schedules the learning model adapted its choice behavior, similar to what was previously observed in monkeys. With an equal probability reward schedule, i.e. when the reward was independent of the choice, the model easily showed biased behavior. In this case, the input statistics during the initial learning of the different sensorimotor mappings determined the behavioral choice preferences. The long-term learning history and the short-term reward history could both be critical variables which decide about behavioral biases in free-choice tasks and which exert their influence on choice behavior via the same neural mechanisms.

Learning arbitrary remapping rules through local associations

The architecture of the adaptive DNF model does not support any explicit representations of the rule that is indicated by the contextual cue (direct vs. inferred reach), but learns this rule implicitly by forming associations of individual stimulus combinations with the rewarded motor response. The sensitivity of the movement plan formation to the context cue is established by the specialization of attractor states in the association field, which develops regions with different context preferences. The trained system then combines spatial and context input in a fashion that is analogous to gain modulation mechanisms proposed for this type of task in previous models (Brozovic et al., 2007;Salinas, 2004). The model produces the same behavior that would be expected from the application of an explicit rule through these local associations, including the formation of consistent biases for a certain context.

Our prediction of capacity limitations due to limited neuronal resources supports this mechanism of local associations and is consistent with observations from the electrophysiological study (Klaes et al., 2011). When we increased the number of possible spatial cue locations to eight the error rates increased substantially, leading us to subsequently use only four directions for the task in the experiments (unpublished observation).

Since most of the time neighboring locations are confused, we can assume that the spatial configuration plays an important role in the learning process. Our model produces a

similar error pattern and in addition provides a possible explanation for it. If new associations are learned through the formation of additional attractor states in the DNFs, the number of possible associations is limited by the metrics of the representation. If this limit is exceeded, neighboring attractor states that reflect spatially contiguous associations are no longer sufficiently separated. Accordingly, if learning of spatial associations takes place in topographically organized neuronal maps, analogous neuronal mechanisms may account for the errors in the monkey behavior.

Integration of learning and action selection

With the adaptive DNF model, we aim to achieve two goals in a single neural architecture. On the one hand, we provide a process model of movement plan formation and action selection in a context-dependent mapping task that can reproduce neurophysiological data. It is in this respect similar to another recent modeling study of decision making in the fronto-parietal cortex (Cisek, 2006), but extends that approach to allow the selection of locations that were not explicitly spatially cued as motor goals.

On the other hand, the model also incorporates a learning mechanism that allows it to acquire new mapping rules, and to do so in a training procedure that closely emulates the training of the monkeys for the same task. This realistic reward-dependent learning differentiates it from previous models of this task (Brozovic et al., 2007;Salinas, 2004). Other theoretical accounts that focus on the learning process deal only with a small number (typically just two) of possible response choices, represented by discrete nodes (Vasilaki et al., 2009;Fusi et al., 2007;Urbanczik and Senn, 2009;Rigotti et al., 2010). They are therefore less suited to capture the process of action selection from a continuous space of motor acts in the fronto-parietal network. Furthermore, they generally do not investigate behavioral biases and free choice tasks (but see Soltani and Wang (2010) for a theoretical study investigating effects of prior probabilities).

Influence of reward contingencies

We could show that the application of different rewarding schedules in free-choice trials can influence the neuronal representations during the memory period and by that also influence the behavior. Using a reward-driven Hebbian learning algorithm enables the

model to adapt to changes in the reward schedule in a manner similar to what is called the ‘matching law’ (Herrnstein and Loveland, 1975;Sugrue et al., 2004;Soltani and Wang, 2006). In fact the model is very sensitive to the reward history and does not necessarily develop equal representations of the two potential motor goals after learning PMG-NC trials with a BMRS. If the ratio of PMG-NC trials is very high, it can even happen that the model 'unlearns' the trained mapping relationship because all weights are more or less reset to an indifferent state (data not shown). A similar observation, that errors can cancel the learned mapping, has been made in a model by Fusi et al. (2007). For this reason it is necessary to also present regular instructed trials along with the free-choice trials. Reducing the learning rates after initial learning of multiple associations would slow such unlearning process but also decrease the sensitivity to changing reward schedules. Similarly, learning in a free-choice task without forcing balanced behavior (i.e. EPRS reward schedule) can easily lead to a biased behavior, due to small imbalances in the probabilities of direct or inferred trials which can self-enhance the probability of the same choice in later trials. This is especially true if the reward probability is high (e.g. 100%) in which case an initially randomly choosen option will be more likely to be choosen again. Such a behavioral bias in free-choice trials is evident in our electrophysiological study (Klaes et al., 2011) and has also been reported in other studies (Barraclough et al., 2004;Scherberger and Andersen, 2007). Biased and balanced behavior in free-choice trials of the model were achieved by the same underlying learning mechanism which enabled the model to learn the arbitrary mapping in the first place. This could also be the case in cortical sensorimotor areas of the parietal or frontal cortex.

Influence of input statistics

Our results show that the input statistics during learning of new sensorimotor associations bias the decision in later free-choice trials. Even if the model is perfectly able to solve the task when all necessary information is provided, its free-choice behavior can be biased. Humans rely on prior probabilities if they have to base their decision on lacking or ambiguous evidence (Carpenter and Williams, 1995;Redding and Wallace, 2000;Sharma et al., 2003;Körding and Wolpert, 2004). From a Bayesian point of view, the activity distribution in our model during the memory period of PMG trials can be interpreted as a

representation of the prior distribution. Since the probabilities for direct and inferred trials were equal in the final task in the electrophysiological study we assume that the inferred bias was acquired during the training of the task, when more inferred than direct trials were presented (unpublished observation).

Conclusions and predictions

Our model successfully integrates sensorimotor processing and working memory formation with decision making and action selection. The reward-driven Hebbian learning mechanism we use to enable the learning of context dependent remappings is sufficient to also explain its susceptibility for probabilistic reward contingencies and input statistics. Most importantly we could reproduce the electrophysiological results from a previous study (Klaes et al., 2011), which showed a similar dependency on reward contingencies. Since continuous reward-driven neuronal weight adaptations change the behavior in free-choice trials, we can also predict that similar manipulations of the reward schedule will be able to produce any ratio of biased behavior. This could also be the source of matching behavior in foraging tasks. From this we can also predict that most initial biased behaviors in different kinds of tasks are a result of the systems history and former exposure to the task and that the same learning mechanism is responsible for learning of arbitrary mappings and biases in free-choice trials.

Figures:

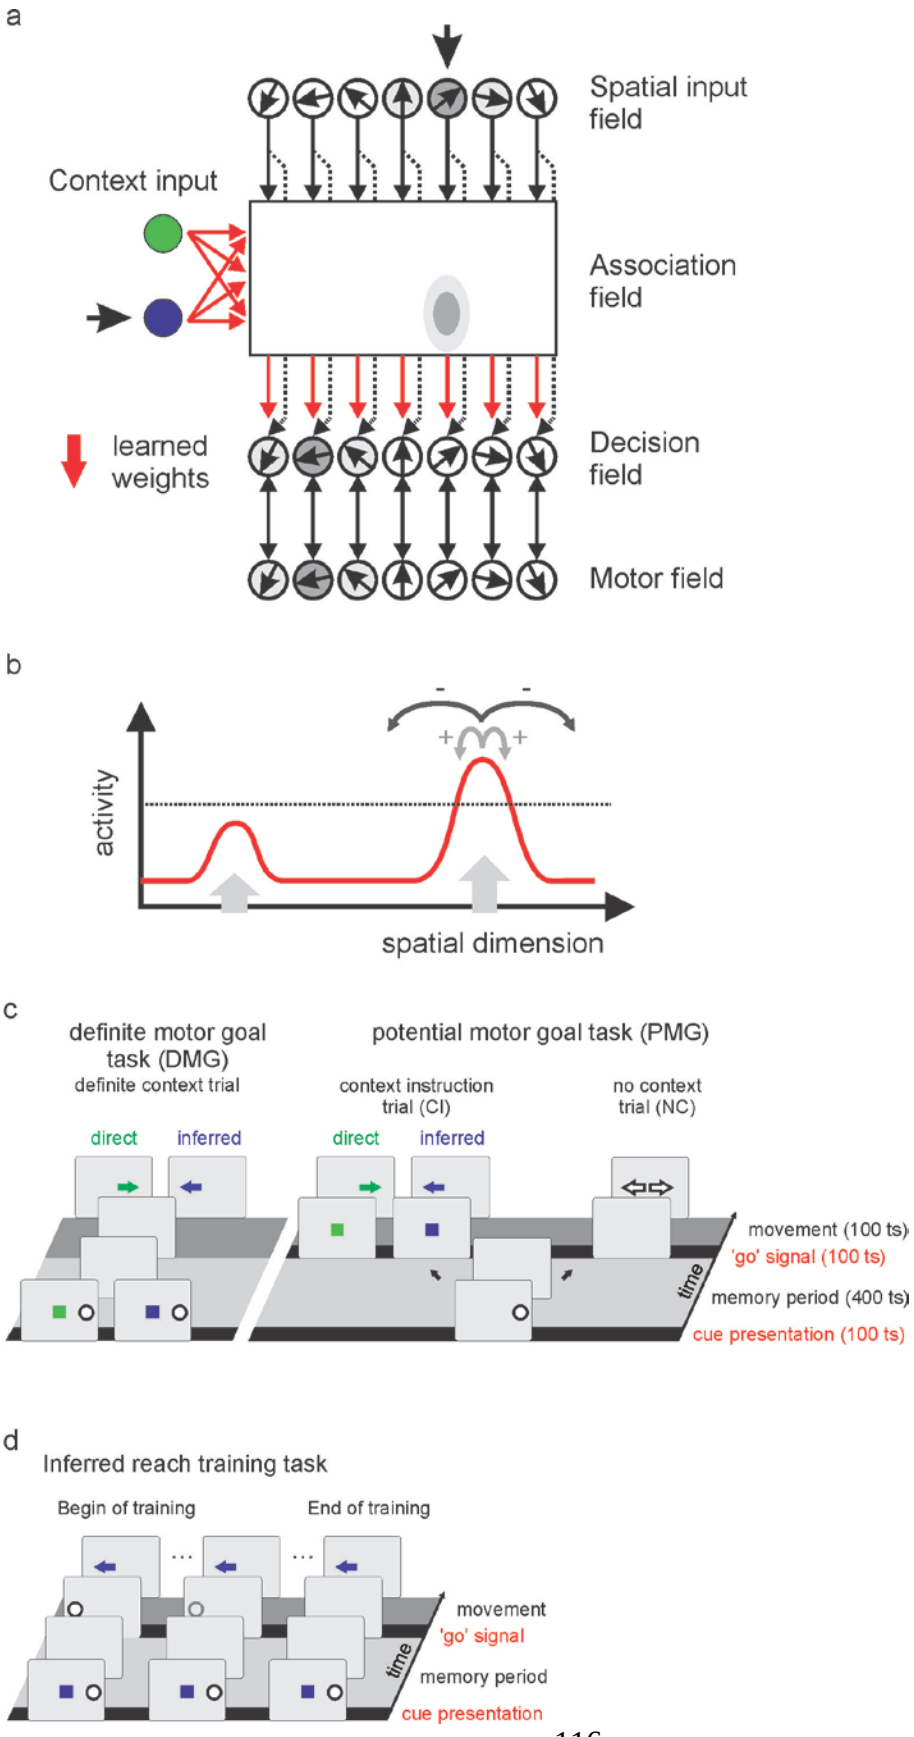
a
Spatial input field
Context input
Association field
learned weights
Decision field
Motor field
b
activity
spatial dimension
c
definite motor goal task (DMG)
definite context trial
direct
inferred
potential motor goal task (PMG)
context instruction trial (CI)
direct
inferred
no context trial (NC)
movement (100 ts)
'go' signal (100 ts)
time
memory period (400 ts)
cue presentation (100 ts)
d
Inferred reach training task
Begin of training
End of training
movement
'go' signal
time
memory period
cue presentation

Figure 1: Neural field model. (a) The model consists of several neural fields. The spatial, reach and motor fields are all one dimensional and topographically represent spatial positions (indicated by the arrows within the circles). The association field is two-dimensional. One dimension represents spatial positions (horizontal), the other one has initially no meaning associated to it, but develops a specificity for contextual information during training (vertical). The spatial field receives visual input from the location of the spatial cue. It projects in a topographically ordered manner to the decision field and to the association field, with the latter receiving additional input from the contextual cue. The decision field and the motor field are reciprocally connected. Only the connections between the contextual input and the association field and those between the association field and the decision field are affected by learning. In the beginning the former are initialized with small random weights, the latter show a topographic organization analogous to the other connections between the fields. The illustration shows a schematic snapshot of a trained network when a blue contextual cue and a spatial cue at the position indicated by the black arrow above the spatial field is presented (activity in the fields is indicated by shaded circles). This causes activation in the association field and through learned weights between the association field and the decision field a remapped representation of the motor goal in the decision field. This is then further propagated to the motor field where the selected action is evaluated. (b) Exemplary illustration of the activation in the decision field. External input from other fields (indicated by grey arrows at the bottom) localized activity (red), which is modulated by lateral interactions in the field. The interactions consist of local excitatory connections (light grey arrows marked with a '+') and surrounding inhibitory connections (dark grey arrows marked with a '-'), and implement a soft competition between distant field regions. Activity above a threshold value (dotted line) is influencing the motor field. (c) Illustration of the tasks that model and monkeys had to perform. In the beginning either a single spatial cue (PMG task) or a spatial and a contextual cue (DMG task) were presented indicated by a white circle (spatial cue) and a colored rectangle (contextual cue). During the memory period no cue was shown. The 'go'-signal indicated the subject to move towards the cued location, either towards the same location as the spatial cue (direct trial; green) or towards the diametrically opposite location (inferred trial; blue). In one part of the PMG

trials the contextual cue was presented at the end of the memory period (PMG-CI), and in another part no contextual cue was shown at all (PMG-NC) and a free choice had to be made (see Methods). (d) Illustration of the inferred reach training task, in which a second spatial cue is shown at the end of the memory period which indicates the rewarded goal position. This cue is gradually faded out over many trials during the training.

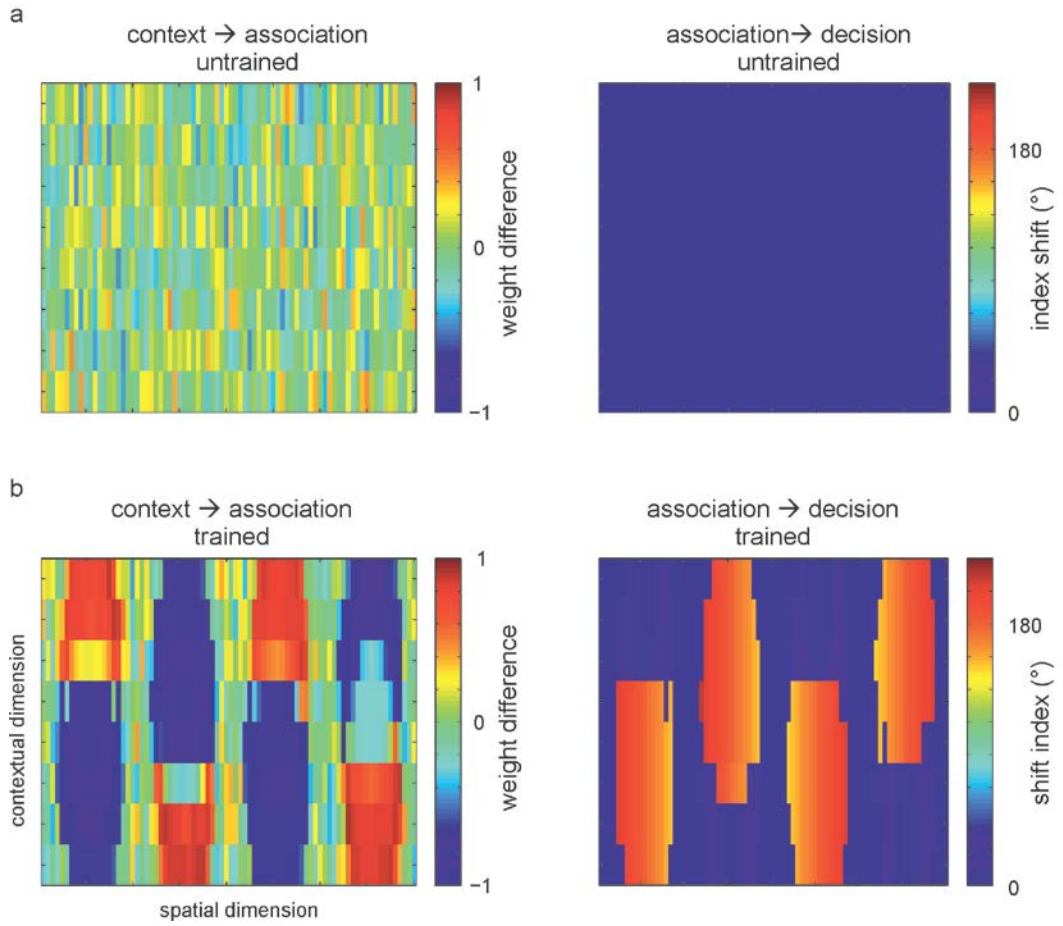

Figure 2: Weight changes during IR training. (a, c) Weight difference matrix from the context input to the association field. Each point in the field is calculated as the difference of the weights from the direct context input neuron and the inferred context input neuron to the association field. In the initial network weight differences are randomly distributed around 0 without any spatial pattern (a). After IR training distinct areas sensitive for direct or inferred context input evolve (c). The topographic connectivity from the association field to the decision field, which is initially established (b), also changes after

the training. In the beginning each point in the same spatial column preferably connects to the corresponding spatial position in the decision field (index shift = 0º). After IR training those areas which prefer the inferred context input preferably connect to the opposite spatial position in the decision field, corresponding to an index shift of about 180º (d).

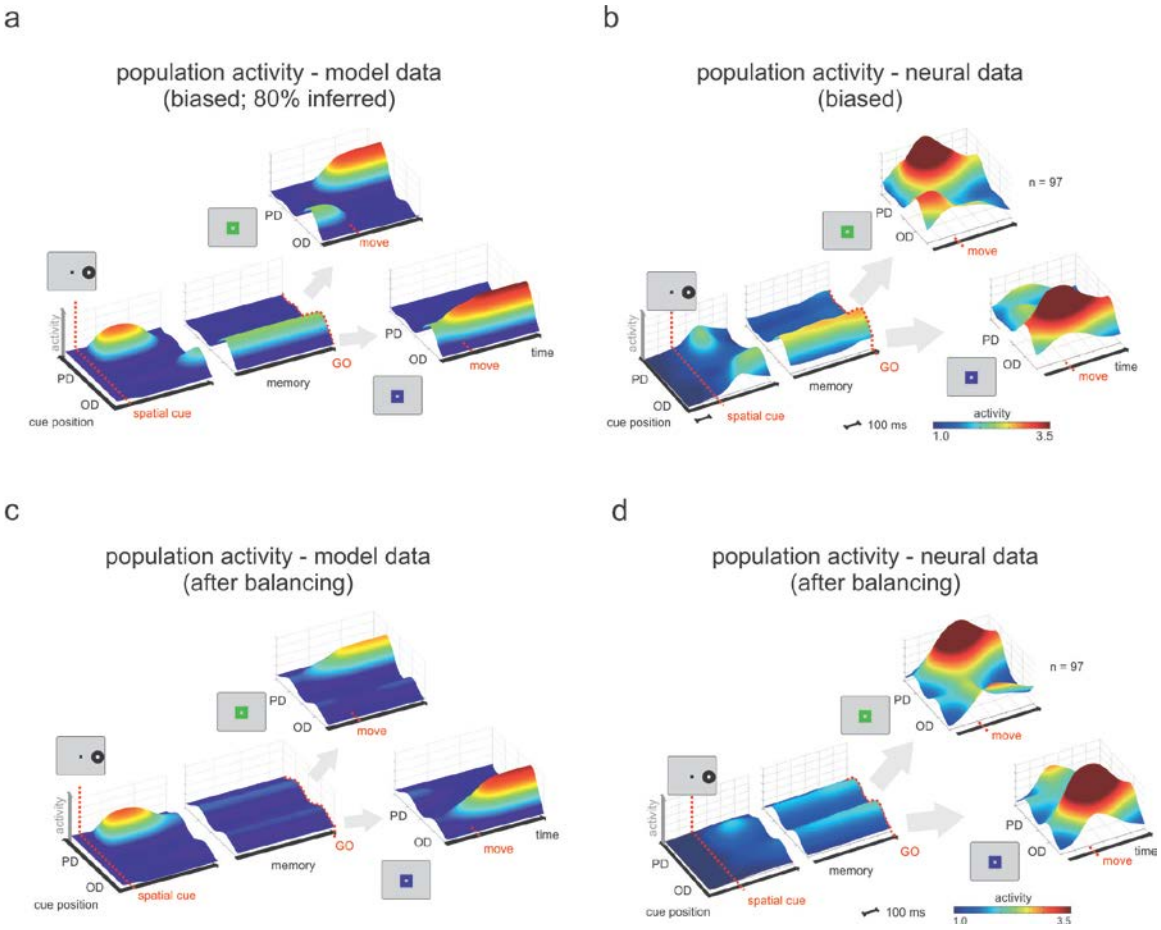

Figure 3: Comparison of the averaged and normalized population activity from the model reach layer (a, c) and from electrophysiological recordings in PRR (b, d) during the PMG task. Prior to averaging and normalizing, the real and model neurons' selectivity profiles were aligned according to their preferred directions in DMG trials (PD: preferred direction, OD: opposite-to-preferred direction. The averaged and normalized activity of real neurons during the PMG task in the biased (b) and balanced (d) datasets is shown for three epochs, aligned to cue onset, 'go'-signal, and movement onset, since the length of the epochs was variable. The model neurons were aligned accordingly even though the epochs had fixed lengths. It can be seen that during the memory period in the model and in the real data plots, only one activity ridge is stable throughout the memory period, before a bias minimizing reward schedule (BMRS; see Methods) was introduced (a, b). After introduction of the BMRS two stable ridges with a lower activity remain during the memory period (c, d).

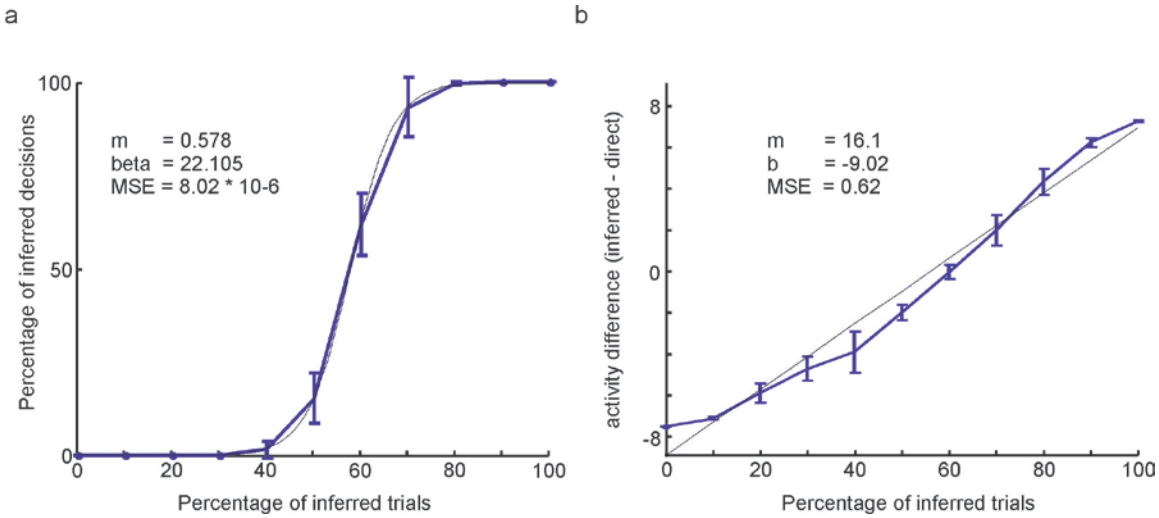

Figure 4: Influence of the input statistics on the behavior and neuronal representation during the memory period. The behavioral bias for inferred reaches in the free-choice trials depends on the percentage of inferred trials during IR training and rises continuously in a sigmoidal fashion (logistic fit function; black curve) (a). The difference of the mean neuronal activity at the preferred and opposite-to-preferred position during the memory period shows an approximately linear (linear fit; black curve) increase when the number of inferred trials is increased (b).

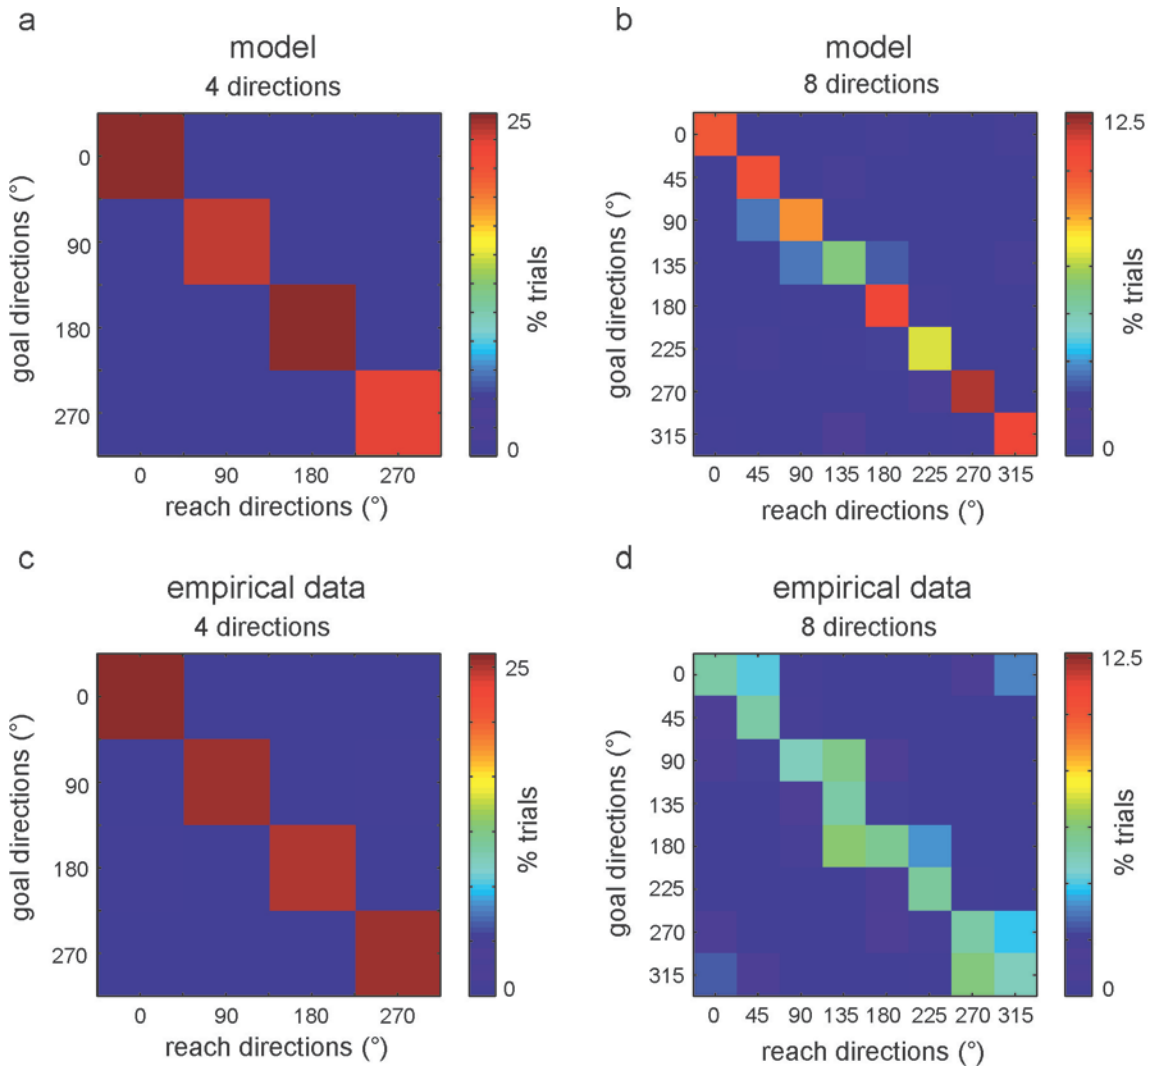

Figure 5: Confusion matrix for different numbers of possible motor goals. A model trained and tested with four possible spatial cue locations selects the correct reach direction in almost all trials, indicated by the high number of trials that fall on the diagonal (a). If the number of possible motor goals is increased to eight, performance drops but most miss reaches are targeted to nearby locations (b). In monkey experiments we can observe a similar error pattern if we increase the number of possible motor goals (c, d).

Reference List

Amari S (1977) Dynamics of Pattern Formation in Lateral-Inhibition Type Neural Fields. Biol Cybern 27:77-87.

Andersen RA, Buneo CA (2002) Intentional maps in posterior parietal cortex. Annu Rev Neurosci 25:189-220.

Andersen RA, Cui H (2009) Intention, Action Planning, and Decision Making in Parietal-Frontal Circuits. Neuron 63:568-583.

Averbeck BB, Sohn JW, Lee D (2006) Activity in prefrontal cortex during dynamic selection of action sequences. Nat Neurosci 9:276-282.

Barraclough DJ, Conroy ML, Lee D (2004) Prefrontal cortex and decision making in a mixed-strategy game. Nat Neurosci 7:404-410.

Battaglia-Mayer A, Caminiti R, Lacquaniti F, Zago M (2003) Multiple levels of representation of reaching in the parieto-frontal network. Cereb Cortex 13:1009-1022.

Blatt GJ, Andersen RA, Stoner GR (1990) Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. J Comp Neurol 299:421-445.

Brozovic M, Gail A, Andersen RA (2007) Gain Mechanisms for Contextually Guided Visuomotor Transformations. J Neurosci 27:10588-10596.

Burnod Y, Baraduc P, Battaglia-Mayer A, Guigon E, Koechlin E, Ferraina S, Lacquaniti F, Caminiti R (1999) Parieto-frontal coding of reaching: an integrated framework. Exp Brain Res 129:325-346.

Carpenter RHS, Williams MLL (1995) Neural computation of log likelihood in control of saccadic eye movements. Nature 377:59-62.

Cisek P (2006) Integrated Neural Processes for Defining Potential Actions and Deciding between Them: A Computational Model. J Neurosci 26:9761-9770.

Cisek P, Kalaska JF (2005) Neural Correlates of Reaching Decisions in Dorsal Premotor Cortex: Specification of Multiple Direction Choices and Final Selection of Action. Neuron 45:801-814.

Cisek P, Kalaska JF (2010) Neural Mechanisms for Interacting with a World Full of Action Choices. Annu Rev Neurosci 33:269-298.

deCharms RC, Zador A (2000) Neural representation and the cortical code. Annu Rev Neurosci 23:613-647.

Deco G, Jirsa VK, Robinson PA, Breakspear M, Friston K (2008) The dynamic brain: from spiking neurons to neural masses and cortical fields. PLoS Comput Biol 4.

Erlhagen W, Bastian A, Jancke D, Riehle A, Schöner G (1999) The distribution of neuronal population activation (DPA) as a tool to study interaction and integration in cortical representations. J Neurosci Methods 94:53-66.

Erlhagen W, Schöner G (2002) Dynamic field theory of movement preparation. Psychol Rev 109:545-572.

Fusi S, Asaad WF, Miller EK, Wang XJ (2007) A Neural Circuit Model of Flexible Sensorimotor Mapping: Learning and Forgetting on Multiple Timescales. Neuron 54:319-333.

Gail A, Andersen RA (2006) Neural Dynamics in Monkey Parietal Reach Region Reflect Context-Specific Sensorimotor Transformations. J Neurosci 26:9376-9384.

Glimcher PW (2003) The neurobiology of visual-saccadic decision making. Annu Rev Neurosci 26:133-179.

Gold JI, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. Nature 404:390-394.

Gold JI, Shadlen MN (2007) The Neural Basis of Decision Making. Annu Rev Neurosci 30:535-574.

Grossberg S, Williamson JR (2001) A neural model of how horizontal and interlaminar connections of visual cortex develop into adult circuits that carry out perceptual grouping and learning. Cereb Cortex 11:37-58.

Herrnstein RJ, Loveland DH (1975) Maximizing and matching on concurrent ratio schedules. J Exp Anal Behav 24:107-116.

Johnson PB, Ferraina S, Bianchi L, Caminiti R (1996) Cortical networks for visual reaching: Physiological and anatomical organization of frontal and parietal lobe arm regions. Cereb Cortex 6:102-119.

Kable JW, Glimcher PW (2009) The Neurobiology of Decision: Consensus and Controversy. Neuron 63:733-745.

Kalarickal GJ, Marshall JA (2002) Rearrangement of receptive field topography after intracortical and peripheral stimulation: the role of plasticity in inhibitory pathways. Network: Computation in Neural Systems 13:1-40.

Kiani R, Shadlen MN (2009) Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. Science 324:759-764.

Klaes C, Westendorff S, Gail A (2011) Choosing goals, not rules: Deciding among rule-based action plans. Neuron 70:536-548.

Kobayashi Y, Isa T (2002) Sensory-motor gating and cognitive control by the brainstem cholinergic system. Neural Netw 15:731-741.

Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. Nature 427:244-247.

Marshall JA (1995) Adaptive Perceptual Pattern Recognition by Self-Organizing Neural Networks: Context, Uncertainty, Multiplicity, and Scale. Neural Netw 8:335-362.

Musallam S, Corneil BD, Greger B, Scherberger H, Andersen RA (2004) Cognitive Control Signals for Neural Prosthetics. Science 305:258-262.

Pastor-Bernier A, Cisek P (2011) Neural Correlates of Biased Competition in Premotor Cortex. J Neurosci 31:7083-7088.

Pesaran B, Nelson MJ, Andersen RA (2008) Free choice activates a decision circuit between frontal and parietal cortex. Nature 453:406-409.

Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. Nature 400:233-238.

Redding GM, Wallace B (2000) Prism exposure aftereffects and direct effects for different movement and feedback times. J Mot Behav 32:83-99.

Rigotti M, Ben Dayan Rubin D, Morrison SE, Salzman CD, Fusi S (2010) Attractor concretion as a mechanism for the formation of context representations. Neuroimage 52:833-847.

Salinas E (2004) Fast Remapping of Sensory Stimuli onto Motor Actions on the Basis of Contextual Modulation. J Neurosci 24:1113-1118.

Scherberger H, Andersen RA (2007) Target Selection Signals for Arm Reaching in the Posterior Parietal Cortex. J Neurosci 27:2001-2012.

Schneegans S, Schöner G (2008) Dynamic field theory as a framework for understanding embodied cognition. In: Handbook of cognitive science: an embodied approach pp 241-269. Elsevier Books.

Sharma J, Dragoi V, Tenenbaum JB, Miller EK, Sur M. (2003) V1 Neurons Signal Acquisition of an Internal Representation of Stimulus Location. Science 13:1758-1763.

Soltani A, Wang X-J (2006) A Biophysically Based Neural Model of Matching Law Behavior: Melioration by Stochastic Synapses. J Neurosci 26:3731-3744.

Soltani A, Wang X-J (2010) Synaptic computation underlying probabilistic inference. Nat Neurosci 13:112-119.

Stringer SM, Rolls ET, Taylor P (2007) Learning movement sequences with a delayed reward signal in a hierarchical model of motor function. Neural Netw 20:172-181.

Sugrue LP, Corrado GS, Newsome WT (2004) Matching Behavior and the Representation of Value in the Parietal Cortex. Science 304:1782-1787.

Suri RE, Schultz W (1999) A neural network model with dopamine-like reinforcement signal that learns a spatial delayed response task. Neuroscience 91:871-890.

Tversky A, Kahneman D (1981) The framing of decisions and the psychology of choice. Science 211:453-458.

Urbanczik R, Senn W (2009) Reinforcement learning in populations of spiking neurons. Nat Neurosci 12:250-252.

Vasilaki E, Fusi S, Wang X-J, Senn W (2009) Learning flexible sensori-motor mappings in a complex network. Biol Cybern 100:147-158.

Wang X-J (2001) Synaptic reverberation underlying mnemonic persistent activity. Trends Neurosci 24:455-463.

Wilson HR, Cowan JD (1973) A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. Kybernetik 13:55-80.

Wise SP, Boussaoud D, Johnson PB, Caminiti R (1997) Premotor and Parietal Cortex: Corticocortical Connectivity and Combinatorial Computations. Annu Rev Neurosci 20:25-42.

3 Summary

In summary the studies presented in this thesis combine different aspects of reach planning in the macaque frontoparietal reach network. First we investigated the underlying mechanisms for arbitrary visuomotor transformations. We found evidence that gain modulation could be the mechanism behind these transformations. Then we asked the question how PRR and PMd interact and in which area the motor plans emerge first. Our results show that in cases that require a remapping a motor goal first emerges in PMd. To understand the decision making process in situations with ambiguous motor goals we analyzed the neuronal activity in free choice trials. We could show that potential motor goals are constructed before a decision is made, irrespective of the goals to be directly cued or inferred from a cue. These findings helped to develop a computational model for decision making in the frontoparietal reach network. The model could reproduce many finding from the electrophysiological studies and also helped to get a better understanding of the underlying mechanisms. This model can now be developed further and provide predictions for future experiments. We can now ask how the input statistics and reward history affect decision making processes in detail. Altogether these contributions support the importance of PRR and PMd for the planning of reach movements and for decision making.

4 Bibliography

Amador,N., Schlag-Rey,M., and Schlag,J. (1998). Primate Antisaccades. I.áBehavioral Characteristics. Journal of Neurophysiology *80*, 1775-1786.

Andersen,R.A., Bracewell,R.M., Barash,S., Gnadt,J.W., and Fogassi,L. (1990). Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. Journal of Neuroscience *10*, 1176-1196.

Andersen,R.A. and Buneo,C.A. (2002). Intentional maps in posterior parietal cortex. Annu.Rev.Neurosci *25*, 189-220.

Andersen,R.A. and Mountcastle,V.B. (1983). The influence of the angle of gaze upon the excitability of the light- sensitive neurons of the posterior parietal cortex. Journal of Neuroscience *3*, 532-548.

Batista,A.P., Buneo,C.A., Snyder,L.H., and Andersen,R.A. (1999). Reach plans in eye-centered coordinates. Science *285*, 257-260.

Batista,A.P., Santhanam,G., Yu,B.M., Ryu,S.I., Afshar,A., and Shenoy,K.V. (2007). Reference frames for reach planning in macaque dorsal premotor cortex. Journal of Neurophysiology *98*, 966-983.

Bell,A.H., Everling,S., and Munoz,D.P. (2000). Influence of Stimulus Eccentricity and Direction on Characteristics of Pro- and Antisaccades in Non-Human Primates. Journal of Neurophysiology *84*, 2595-2604.

Blatt,G.J., Andersen,R.A., and Stoner,G.R. (1990). Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. J Comp Neurol. *299*, 421-445.

Boussaoud,D., Jouffrais,C., and Bremmer,F. (1998). Eye position effects on the neuronal activity of dorsal premotor cortex in the macaque monkey. J Neurophysiol *80*, 1132-1150.

Brozovic,M., Gail,A., and Andersen,R.A. (2007). Gain Mechanisms for Contextually Guided Visuomotor Transformations. Journal of Neuroscience *27*, 10588-10596.

Buneo,C.A., Jarvis,M.R., Batista,A.P., and Andersen,R.A. (2002). Direct visuomotor transformations for reaching. Nature *416*, 632-636.

Burnod,Y., Baraduc,P., Battaglia-Mayer,A., Guigon,E., Koechlin,E., Ferraina,S., Lacquaniti,F., and Caminiti,R. (1999). Parieto-frontal coding of reaching: an integrated framework. Experimental Brain Research *129*, 325-346.

Chang,S.W.C., Papadimitriou,C., and Snyder,L.H. (2009). Using a Compound Gain Field to Compute a Reach Plan. Neuron *64*, 744-755.

Cisek,P. (2006). Integrated Neural Processes for Defining Potential Actions and Deciding between Them: A Computational Model. Journal of Neuroscience *26*, 9761-9770.

Cisek,P. and Kalaska,J.F. (2005). Neural Correlates of Reaching Decisions in Dorsal Premotor Cortex: Specification of Multiple Direction Choices and Final Selection of Action. Neuron *45*, 801-814.

Cohen,Y.E. and Andersen,R.A. (2002). A common reference frame for movement plans in the posterior parietal cortex. Nature Reviews Neuroscience *3*, 553-562.

Colby,C.L., Gattass,R., Olson,C.R., and Gross,C.G. (1988). Topographical Organization of Cortical Afferents to Extrastriate Visual Area Po in the Macaque - A Dual Tracer Study. Journal of Comparative Neurology *269*, 392-413.

Connor,C.E., Gallant,J.L., Preddie,D.C., and Van Essen,D.C. (1996). Responses in area V4 depend on the spatial relationship between stimulus and attention. Journal of Neurophysiology *75*, 1306-1308.

Crammond,D.J. and Kalaska,J.F. (1994). Modulation of preparatory neuronal activity in dorsal premotor cortex due to stimulus-response compatibility. Journal of Neurophysiology *71*, 1281-1284.

Crammond,D.J. and Kalaska,J.F. (1996). Differential relation of discharge in primary motor cortex and premotor cortex to movements versus actively maintained postures during a reaching task. Exp Brain Res *108*, 45-61.

Crammond,D.J. and Kalaska,J.F. (2000). Prior Information in Motor and Premotor Cortex: Activity During the Delay Period and Effect on Pre-Movement Activity. Journal of Neurophysiology *84*, 986-1005.

Gail,A. and Andersen,R.A. (2006). Neural Dynamics in Monkey Parietal Reach Region Reflect Context-Specific Sensorimotor Transformations. Journal of Neuroscience *26*, 9376-9384.

Glimcher,P.W. and Rustichini,A. (2004). Neuroeconomics: The Consilience of Brain and Decision. Science *306*, 447-452.

Gold,J.I. and Shadlen,M.N. (2000). Representation of a perceptual decision in developing oculomotor commands. Nature *404*, 390-394.

Goodale,M.A. and Milner,A.D. (1992). Separate Visual Pathways for Perception and Action. Trends in Neurosciences *15*, 20-25.

Grol,M.J., de Lange,F.P., Verstraten,F.A.J., Passingham,R.E., and Toni,I. (2006). Cerebral Changes during Performance of Overlearned Arbitrary Visuomotor Associations. Journal of Neuroscience *26*, 117-125.

Gross,C.G. (1992). Representation of visual stimuli in inferior temporal cortex. Philos.Trans.R.Soc.Lond B Biol.Sci. *335*, 3-10.

Hallett,P.E. (1978). Primary and secondary saccades to goals defined by instructions. Vision Research *18*, 1279-1296.

Johnson,P.B., Ferraina,S., Bianchi,L., and Caminiti,R. (1996). Cortical networks for visual reaching: Physiological and anatomical organization of frontal and parietal lobe arm regions. Cerebral Cortex *6*, 102-119.

Knill,D.C. and Pouget,A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. Trends in Neurosciences *27*, 712-719.

Kurata,K. (1991). Corticocortical inputs to the dorsal and ventral aspects of the premotor cortex of macaque monkeys. Neuroscience Research *12*, 263-280.

Kurata,K. (1994). Information processing for motor control in primate premotor cortex. Behavioural Brain Research *61*, 135-142.

Logothetis,N.K., Pauls,J., and Poggio,T. (1995). Shape representation in the inferior temporal cortex of monkeys. Curr Biol *5*, 552-563.

Marconi,B., Genovesio,A., Battaglia-Mayer,A., Ferraina,S., Squatrito,S., Molinari,M., Lacquaniti,F., and Caminiti,R. (2001). Eye-hand coordination during reaching. I. Anatomical relationships between parietal and frontal cortex. Cerebral Cortex *11*, 513-527.

Miller,E.K. (2000). The prefontral cortex and cognitive control. Nature Reviews Neuroscience *1*, 59-65.

Mitz,A.R., Godschalk,M., and Wise,S.P. (1991). Learning-dependent neuronal activity in the premotor cortex: activity during the acquisition of conditional motor associations. J Neurosci *11*, 1855-1872.

Mountcastle,V.B., Lynch,J.C., Georgopoulos,A.P., Sakata,H., and Acuna,C. (1975). Posterior parietal association cortex of the monkey: command functions for operations within extrapersonal space. Journal of Neurophysiology *38*, 871-908.

Muhammad,R., Wallis,J.D., and Miller,E.K. (2006). A Comparison of Abstract Rules in the Prefrontal Cortex, Premotor Cortex, Inferior Temporal Cortex, and Striatum. Journal of Cognitive Neuroscience *18*, 974-989.

Pesaran,B., Nelson,M.J., and Andersen,R.A. (2006). Dorsal Premotor Neurons Encode the Relative Position of the Hand, Eye, and Goal during Reach Planning. Neuron *51*, 125-134.

Platt,M.L. and Glimcher,P.W. (1999). Neural correlates of decision variables in parietal cortex. Nature *400*, 233-238.

Pouget,A. and Sejnowski,T. (1997). Spatial Transformations in the Parietal Cortex Using Basis Functions. Journal of Cognitive Neuroscience *9*, 222-237.

Preuss,T.M., Stepniewska,I., and Kaas,J.H. (1996). Movement representation in the dorsal and ventral premotor areas of owl monkeys: A microstimulation study. Journal of Comparative Neurology *371*, 649-676.

Salinas,E. and Thier,P. (2000). Gain Modulation: A Major Computational Principle of the Central Nervous System. Neuron *27*, 15-21.

Snyder,L.H., Batista,A.P., and Andersen,R.A. (1997). Coding of intention in the posterior parietal cortex. Nature *386*, 167-170.

Snyder,L.H., Batista,A.P., and Andersen,R.A. (2000). Intention-related activity in the posterior parietal cortex: a review. Vision Research *40*, 1433-1441.

Stoet,G. and Snyder,L.H. (2004). Single Neurons in Posterior Parietal Cortex of Monkeys Encode Cognitive Set. Neuron *42*, 1003-1012.

Sugrue,L.P., Corrado,G.S., and Newsome,W.T. (2004). Matching Behavior and the Representation of Value in the Parietal Cortex. Science *304*, 1782-1787.

Toni,I. and Passingham,R.E. (1999). Prefrontal-basal ganglia pathways are involved in the learning of arbitrary visuomotor associations: a PET study. Experimental Brain Research *127*, 19-32.

Toni,I., Rushworth,M.F.S., and Passingham,R.E. (2001). Neural correlates of visuomotor associations - Spatial rules compared with arbitrary rules. Experimental Brain Research *141*, 359-369.

Ungerleider,L.G. and Mishkin,M. (1982). The two cortical visual systems. In Analysis of Visual Behavior, D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield, eds. MIT Press), pp. 549-586.

Van Opstal,A.J., Hepp,K., Suzuki,Y., and Henn,V. (1995). Influence of eye position on activity in monkey superior colliculus. J Neurophysiol *74*, 1593-1610.

Wallis,J.D., Anderson,K.C., and Miller,E.K. (2001). Single neurons in prefrontal cortex encode abstract rules. Nature *411*, 953-956.

Wallis,J.D. and Miller,E.K. (2003). From Rule to Response: Neuronal Processes in the Premotor and Prefrontal Cortex. Journal of Neurophysiology *90*, 1790-1806.

Weinrich,M. and Wise,S.P. (1982). The premotor cortex of the monkey. J Neurosci *2*, 1329-1345.

White,I.M. and Wise,S.P. (1999). Rule-dependent neuronal activity in the prefrontal cortex. Experimental Brain Research *126*, 315-335.

Wise,S.P., Boussaoud,D., Johnson,P.B., and Caminiti,R. (1997). Premotor and Parietal Cortex: Corticocortical Connectivity and Combinatorial Computations. Annual Review of Neuroscience *20*, 25-42.

Wise,S.P. and Murray,E.A. (2000). Arbitrary associations between antecedents and actions. Trends in Neurosciences *23*, 271-276.

Zhang,J., Riehle,A., Requin,J., and Kornblum,S. (1997). Dynamics of single neuron activity in monkey primary motor cortex related to sensorimotor transformation. J Neurosci *17*, 2227-2246.

Zipser,D. and Andersen,R.A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. Nature *331*, 679-684.

Curriculum Vitae

Christian Klaes

Scientific Education

since 11/2006	Graduate student, German Primate Center, Göttingen, Germany Advisor: Dr. Alexander Gail
05/2003-05/2004	Study of neurocomputing (Neuroinformatik) at the Ruhr-Universität-Bochum Qualification: certificate Thesis: Design and implementation of behavioural modules for a mobile, autonomous robot Supervisor: Gregor Schöner
08/1997-04/2003	Study of biology at the Ruhr-Universität-Bochum Qualification: University degree in biology Thesis: Perceptual learning in humans: analysis of the effects of rTMS on the visual recognition ability in humans
08/1987-05/1996	Abitur, Walram Gymnasium, Menden, Germany

Previous working experience

12/2004-02/2006	Practical training at T-Systems in the field of electronic toll collection as a systems integrator, Mülheim, Germany
07/2001-04/2003	Scientific assistant at the department of neurocomputing, Ruhr-Universität-Bochum, Germany

Journal publications

Klaes, C., Westendorff, S., Chakrabarti, S., Gail, A (2011). Choosing goals, not rules: Deciding among rule-based action plans. Neuron 70, 536-548.

Westendorff, S., Klaes, C., and Gail, A. (2010). The Cortical Timeline for Deciding on Reach Motor Goals. Journal of Neuroscience *30*, 5426-5436.

Gail, A., Klaes, C., and Westendorff, S. (2009). Implementation of Spatial Transformation Rules for Goal-Directed Reaching via Gain Modulation in Monkey Parietal and Premotor Cortex. Journal of Neuroscience *29*, 9490-9499.

Schneider, S., Igel, C., Klaes, C., Dinse, H. R., Wiemer, J.C. (2004). Evolutionary Adaptation of Nonlinear Dynamical Systems in Computational Neuroscience. Genetic Programming and evolvable machines *5*, 215-227

Published abstracts

Klaes, C., Ragert, P., Jancke, D., Tegenthoff, M., Dinse, H.R. (2003) rTMS induced improvement of human orientation discrimination. Soc Neurosci Abstr 29: 911.22

Dinse, H.R., Klaes, C., Ragert, P., Tegenthoff, M. (2005) rTMS induced improvement of human visual orientation discrimination. Proceedings of the 6th Meeting of the German Neuroscience Society/30th Göttingen Neurobiology Conference 2005 Neuroforum Suppl 1, 179A.

Klaes, C., Westendorff, S., Gail, A. (2008) Constructing reach-goal representations in the fronto-parietal reach network. Poster at the meeting of the Federation of European Neuroscience Societies, Geneva, Swiss

Westendorff, S., Klaes, C., Gail, A. (2008) Motor-goals for reaches are encoded earlier in premotor than parietal cortex. Poster at the meeting of the Federation of European Neuroscience Societies, Geneva, Swiss

Klaes, C., Westendorff, S., Gail, A. (2008) Mapping rules have complementary effects on PRR and PMd motor goal representations. Poster at the meeting of the Society for Neuroscience, Washington, DC

Klaes, C., Westendorff, S., Gail, A. (2009) Gain modulation of reach related neurons in PRR and PMd of monkeys. Göttingen Meeting of the German Neuroscience Society, Göttingen, Germany

Klaes, C., Westendorff, S., Gail, A. (2009) Contextual gain modulations. Poster at the Primate Neurobiology Meeting 2009, Göttingen, Germany

Klaes, C., Westendorff, S., Gail, A. (2009) Gain modulation of reach related neurons in the parietal reach region and dorsal premotor cortex of monkeys. Poster at the Neurizons conference, Göttingen, Germany

Klaes, C., Westendorff, S., Gail, A. (2009) Parietal reach region represents visually cued and spatially inferred potential motor goals simultaneously. Poster at the meeting of the Society for Neuroscience, Chicago, IL

Westendorff, S., Klaes, C., Gail, A. (2009) Spatially inferred, but not directly cued reach goals are represented earlier in premotor than posterior parietal cortex. Poster at the meeting of the Society for Neuroscience, Chicago, IL

Klaes, C., Schneegans, S., Schöner, G., Gail, A. (2010) Decision making in the Posterior Parietal Cortex: A learning neural field model. Poster at the Bernstein

Conference on Computational Neuroscience 2010, Berlin, Germany

Klaes, C., Schneegans, S., Schöner, G., Gail, A. (2010) A learning neural field model of decision making. Poster at the meeting of the Society for Neuroscience, San Diego, CA

Klaes, C., Schneegans, S., Schöner, G., Gail, A. (2011) A learning neural field model of decision making. Poster at the Computational Neuroscience (CNS) meeting, Stockholm, Sweden