Encoding and Attentive Modulation of Dynamic Motion Stimuli

Dissertation

zur Erlangung des Grades eines Doktors
der Naturwissenschaften

vorgelegt von

Heiko Stemmann
Institut für Hirnforschung
Fachbereich 2
Universität Bremen

Bremen, November 2008

1. Gutachter: Prof. Dr. Manfred Fahle
2. Gutachter: Prof. Dr. Michael Koch

1 Table of Contents

2 Introduction 1

2.1 The Receptive Field 1

2.2 Linear Receptive Field Models 2

2.3 Direction selectivity in Macaque visual area MT 3

2.4 Investigating the Encoding of Stimulus Features by Reconstruction 7

2.5 Attentive Modulation of Neural Response Behaviour 8

2.6 A network perspective on effects of attention on cortical motion processing 13

3 Overview of experiments 15

4 Encoding of dynamic visual stimuli by primate area MT neurons 16

4.1 Abstract 16

4.2 Introduction 16

4.3 Methods 16

4.4 Results 19

4.5 Discussion 21

5 Linearity of Macaque Area MT Neurons: Model-based Interpretation of Population Responses to Dynamic Motion Stimuli 23

5.1 Introduction 23

5.2 Materials and Methods 24

5.2.1 Recording and Surgical Techniques 24

5.2.2 Behavioural Task and Visual Stimulation 24

5.2.3 Data Analysis 26

5.2.4 Reconstruction of Movement Directions 31

5.2.5 Model Simulations 33

5.3 Results 34

5.3.1 Spike–Triggered Averages and Latencies 34

5.3.2 Tuning Curves 36

5.3.3 Linear Filter Estimates from Continuous and Discontinuous Stimulation 39

5.3.4 Stimulus Reconstruction 42

5.3.5 Comparison of MAP and ML Reconstructions 43

5.3.6 Temporal Stimulus Statistics and Reconstruction Accuracy 44

5.3.7 Coding Strategies for Dynamic Stimulus Trajectories 47

5.3.8 Time Integration of Stimulus Representation for Neural Populations 48

5.4 Discussion 50

5.4.1 Delays 50

5.4.2 Tuning Curves 51

5.4.3 Linear Filters 52

5.4.4 Reconstruction – Comparison of MAP and ML Estimates 52

5.4.5 Reconstruction – Stimulus Statistics 54
5.4.6 Reconstruction – Coding Strategies for Dynamic Stimulus Trajectories 54
5.4.7 Reconstruction – Time Course of Stimulus Representation for Neural Populations 55

5.5 Final Conclusions 55

6 Effects of Attention during Rapid Serial Visual Presentation on Macaque MT Neurons 56

6.1 Introduction 56

6.2 Methods 56
6.2.1 Animals and surgical procedures 56
6.2.2 Data acquisition and recording 56
6.2.3 Behavioural task and visual stimulation 57
6.2.4 Data Analysis 58

6.3 Results 59
6.3.1 Behaviour 59
6.3.2 Neural responses 59

6.4 Discussion and Conclusion 62

7 Imaging Cortical Networks of Spatial Attention in the Macaque 63

7.1 Introduction 63

7.2 Methods 64

7.3 Results 69

7.4 Discussion 76

7.5 Appendix 78
7.5.1 Behavioral Performance 78
7.5.2 Functional Brain Maps 81
7.5.3 Reaction times 83

8 Conclusion 85

9 Summary / Zusammenfassung 87

9.1 Summary 87

9.2 Zusammenfassung 88

10 Acknowledgements 90

11 References 92

2 Introduction

2.1 The Receptive Field

How much do we understand about processing in early visual areas of the brain? Several recent publications have asked this question (Carandini, Demb et al. 2005; Olshausen and Field 2005), showing that after decades of research our understanding is still not complete. Since the seminal findings of Hubel and Wiesel in cat (Hubel and Wiesel 1959) and primate primary visual cortex (Hubel and Wiesel 1968) much effort has been devoted to develop a model for neuronal processing within early stages of the visual system. By applying light spots of various sizes and shapes, these authors were the first to describe response properties of single neurons in the first stage of cortical visual processing. To describe their result, they used the concept of the *receptive field*, introduced 1938 for frog optic nerve fibers (Hartline 1938) and extended to cat retinal ganglion cells in 1953 (Kuffler 1953). The receptive field of a single neuron was defined as the area of the retina which must receive illumination in order to cause a discharge in a particular ganglion cell or nerve fiber. Kuffler (Kuffler 1953) showed that ganglion cells have concentric receptive fields, with an 'on' centre and an 'off' periphery, or vice versa. The 'on' and 'off' areas within a receptive field were found to be mutually antagonistic, i.e. a spot of light restricted to the centre of the receptive field was more effective than one covering the whole receptive field (Barlow, Fitzhugh et al. 1957). This definition of the receptive field only referred to the spatial extent and location of the area which modulates, in an excitatory or inhibitory fashion, the activity of the neural system under investigation.

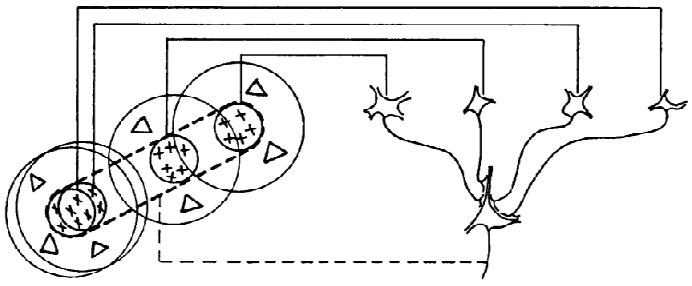

Figure 1: Scheme for explaining the organization of simple cell receptive fields. A large number of lateral geniculate cells, of which four are illustrated in the figure, have receptive fields with 'on' centers arranged along a straight line on the retina. All of these project upon the same single cortical cell. The receptive fields of the cortical cell will then have an elongated 'on' centre indicated by the interrupted lines in the receptive-field diagram to the left of the figure. (from (Hubel and Wiesel 1962)).

Hubel and Wiesel described additional determinants of cortical neural responses, namely form, size and orientation of the light stimulus. Responses to these attributes are formed by the spatial configuration of 'on' and 'off' regions within the receptive field (Hubel and Wiesel 1959). Hubel and Wiesel proposed a simple model of how these additional features maybe encoded, by supposing that afferent 'on' – 'off' – center geniculate cells are appropriately connected to cortical cells (see Figure 1). For example, computing a response profile over orientation space, i.e. plotting the averaged neural response to each orientation will result in an orientation *tuning curve*. In case of an orientation selective neuron this will result in a bell-shaped curve. Analogously, *tuning curves* for other features can be computed, e.g. directions, color, speed, etc..

2.2 Linear Receptive Field Models

The receptive field concept can be formalized. At the basis of most current models of neurons in the early visual system is the concept of a linear receptive field. Receptive fields act as linear filters. A linear filter describes the stimulus selectivity for a neuron: images that resemble the filter will produce a large response. For example, tuning for orientation in V1 can be described by filters elongated along one spatial axis (Hubel and Wiesel 1962). Basic models of neurons at the earliest stages of visual processing (retina, LGN, and V1 simple cells) typically include a single linear filter (Enroth-Cugell and Robson 1966; Movshon, Thompson et al. 1978). Models of neurons at later stages of processing (V1 complex cells and beyond) require multiple filters (Movshon, Thompson et al. 1978). The receptive field concept can be expected to include a third dimension: time. Taking into account all three dimensions, i.e. modelling the *spatiotemporal* receptive field, the full spatial layout of illumination in the recent past is specified.

Testing the linear receptive field experimentally led to the discovery of nonlinear phenomena, which cannot be explained by a linear account alone. For example, a suboptimal stimulus reduces the response of a visual neuron, but "firing rates" cannot go negative. There is a point at which neural activity is "clipped". In almost the same manner an optimal stimulus will increase the response of a neuron, but response rates cannot go infinitely high, because of the ~1 – 2 ms refractory period. These nonlinear phenomena can be modelled as a static nonlinearity, meaning that the linear response is passed through an input – output function that is invariant over time (Chichilnisky 2001). The combination of a linear receptive field and a static nonlinearity creates the linear-nonlinear (LN) model. This model predicts response rates but not actual spike times. Modeling spike responses by a Poisson process,

defined by rate (with equal mean and variance), leads to the LNP (linear-nonlinear-Poisson) model of spiking (Paninski, Pillow et al. 2004).

A common approach to characterize the receptive field properties of visual neurons is to apply a white noise stimulus and measure the activity of single neurons. In the next step the neural response, i.e. the action potentials (*spikes*) are reversely correlated with the stimulus. This approach is called spike-triggered average (STA), in which the spike-triggered stimulus distribution indicates the average feature in the stimulus space eliciting a spike (Chichilnisky 2001). When the system is linear, or approximately linear followed by a static nonlinearity, the linear filter is the "receptive field" and together with the static nonlinearity completely defines the system. White noise contains approximately equal energy over a range of temporal frequencies. While this flat energy spectrum is a nice feature for characterizing the receptive field, this spectrum differs from natural scenes, which contains less stimulus energy at higher temporal frequencies (Simoncelli and Olshausen 2001). The white noise approach can be applied to study motion processing, like for example in the direction of motion domain, by applying a random sequence of motion events (Borghuis, Perge et al. 2003). In this case the signal has a flat autocorrelation, i.e. there is no correlation between subsequently presented motion directions.

2.3 Direction selectivity in Macaque visual area MT

Extrastriate visual area MT (middle temporal area) was discovered by two different groups at roughly the same time. In 1971 Dubner and Zeki (Dubner and Zeki 1971) recorded electrophysiologically in anaesthetized macaques from a region situated on the posterior bank of the superior temporal sulcus (STS) which receives a direct projection from primary visual cortex (V1), particularly revealing the high proportion of direction-selective neurons. Around the same time, Allman and Kaas (Allman and Kaas 1971) discovered a large number of retinotopically organized maps in owl monkeys. One of these, which they named MT for middle temporal, contained neurons that responded better to drifting bars than to flashed spots. This region exhibited a dense myelination in the lower layers, which has also later been shown to be characteristic of the macaque motion area, which Zeki subsequently named V5. A number of subsequent studies indicated that MT was unique as a cortical area highly specialized for processing visual motion (Zeki 1974; Zeki 1980; Baker, Petersen et al. 1981; Van Essen, Maunsell et al. 1981; Maunsell and Van Essen 1983; Maunsell and Van Essen 1983; Felleman and Kaas 1984). These studies also showed that area MT is common to a number of different primate species. In the macaque brain area MT/V5 is located in the posterior bank of the dorsal part of the superior temporal sulcus (see Figure 2).

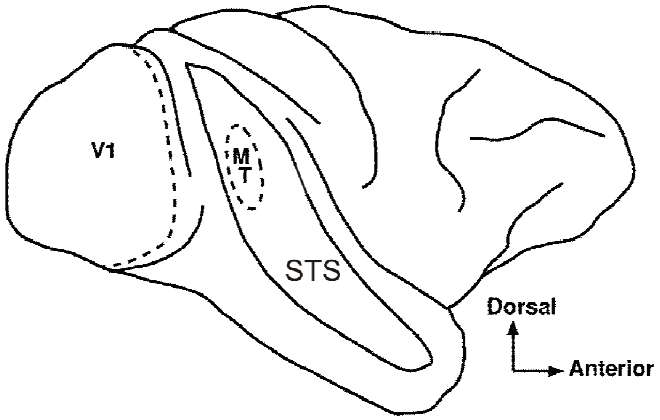

Figure 2: Schematic drawing of a lateral view on the right hemisphere of the macaque brain. Superior temporal sulcus (STS) has been opened to better illustrate the position of area MT/V5 lying in its dorsal part in the posterior bank. Primary visual cortex (V1) covers the whole occipital pole.

Area MT contains a complete representation of the contralateral visual hemifield (Gattass and Gross 1981; Van Essen, Maunsell et al. 1981). The representation is retinotopically organized, but the retinotopy is less regular than in striate cortex (Zeki 1974; Gattass and Gross 1981; Van Essen, Maunsell et al. 1981). The most characteristic feature of area MT is the high abundance of direction selective neurons (90% direction selective neurons are usually reported, (Zeki 1974; Albright 1984). The directional preference is evenly distributed along all directions, and there has been evidence for a columnar organization of direction selectivity in the initial publication by Dubner and Zeki (Dubner and Zeki 1971) but has only been shown definitively by Albright (Albright 1984). Malonek and colleagues (Malonek, Tootell et al. 1994), using optical imaging methods, confirmed the existence if direction columns and also revealed a columnar organisation for orientation. Eventually a columnar organization for disparity has been shown (DeAngelis and Newsome 1999) coexisting with the direction columns. There is also a clustering of neurons by speed preference but the clustering is not strictly columnar (Liu and Newsome 2003). Neurons in area MT lack colour selectivity (Zeki 1974; Maunsell and Van Essen 1983; Riecansky, Thiele et al. 2005), but modulatory influence of colour on neuronal sensitivity has been shown (Croner and Albright 1999).

Area MT has a rich set of interconnections with other cortical regions (see Figure 3) as well as with numerous subcortical structures. Cortical connections include V1, V2, V3, V3a, V4, V4t, MST (medial superior temporal area), FST (fundus superior temporal area), VIP (ventral intraparietal area), PO (parietooccipital area) and FEF (frontal eye fields).

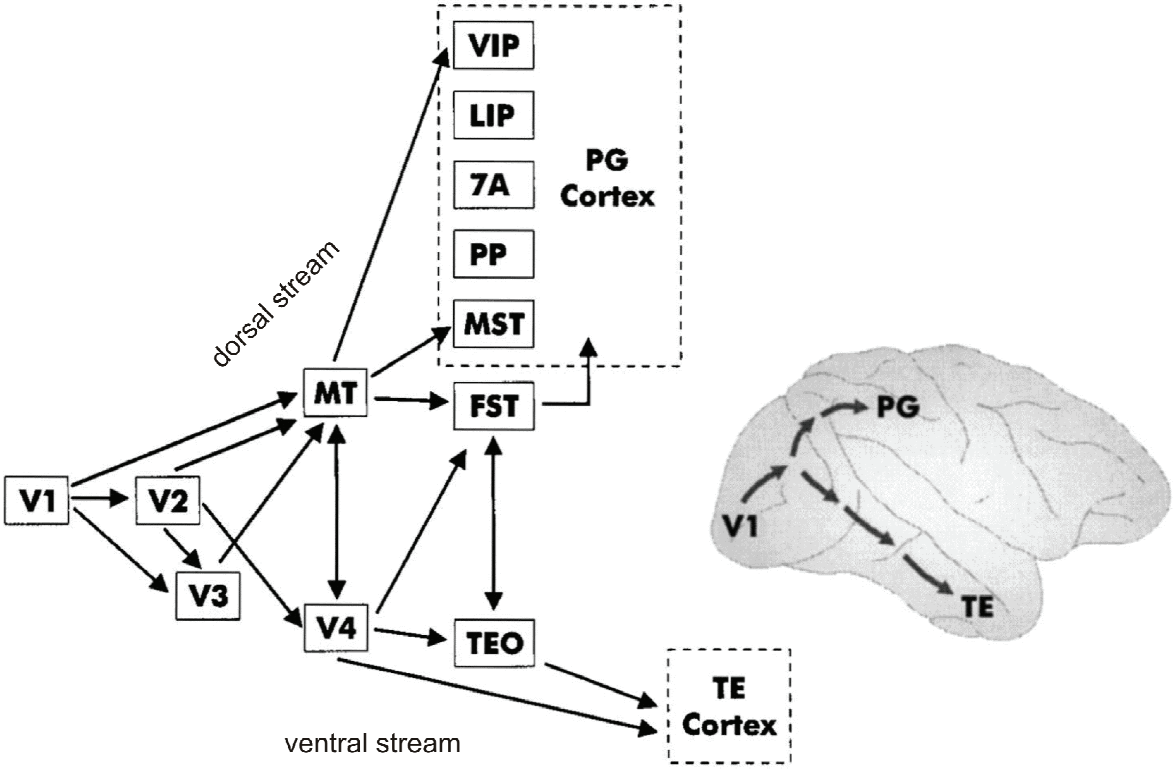

Figure 3: Simplified scheme of the cortical connections in the visual system of macaques. Two cortical streams originate from area V1. The dorsal (parietal) stream is involved in spatial vision, motion perception and visuo-motor coordination. The ventral (temporal) stream is engaged in object recognition and color perception. Arrows indicate major feed-forward connections between cortical areas. Cortical hierarchy increases from left to right. Abbreviations: MT, middle temporal: FST, fundus superior temporal; MST, medial superior temporal, PP, posterior parietal; LIP, lateral intraparietal; VIP, ventral intraparietal; PG, posterior parietal cortex; TE, inferior temporal cortex. (adapted from (Croner and Albright 1999))

Subcortical connections have been shown for the pulvinar, basal ganglia, superior colliculus, pretectum and the pons. The major input to area MT is provided by direction selective layer 4B and layer 6 neurons in primary visual cortex (Shipp and Zeki 1989). These receive predominantly input from magnocellular layer 4Cα (Hubel and Wiesel 1972). Input projections from area V2 origin are predominantly from the cytochrome-oxidase rich thick stripes (Shipp and Zeki 1985). There is an additional input from area V3 (Maunsell and van Essen 1983). Both sources, V2 thick stripe and V3, are innervated from a mostly separate population of layer 4B neurons in V1 (Livingstone and Hubel 1987; Felleman, Burkhalter et al. 1997; Sincich and Horton 2003) and contain direction selective neurons (Zeki 1978). However, it has been shown that response properties of area MT neurons are not entirely dependent on cortical input (Rodman, Gross et al. 1989). An additional afferent pathway to area MT passes through superior colliculus (Rodman, Gross et al. 1990) and probably pulvinar (Standage and Benevento 1983).

Area MT is the main input into the dorsal or posterior parietal processing stream (Ungerleider and Mishkin 1982), and its output target structures that are implicated in the analysis of optical flow (e.g. MST, VIP) and the generation of eye movements (e.g. LIP, FEF, SC, dorsolateral pons). The response properties of neurons in area MT and its connections implicate its involvement in motion processing. Numerous studies have been carried out to investigate the relationship of are MT neurons and motion perception and behaviour. It has been shown that motion sensitivity in area MT is closely correlated to perceptual sensitivity. Employing a motion direction discrimination task and concurrently recording single neurons in area MT of macaque monkeys, Newsome and colleagues were able to directly compare neuronal sensitivity with psychophysical performance of the monkey (see (Parker and Newsome 1998) for review). Remarkably, they discovered the neuronal sensitivity of single neurons to be as accurate as the monkey itself (Britten, Shadlen et al. 1992). Subsequent numerical simulations showed that the performance in a direction discrimination task can be modelled from responses of area MT neurons (Shadlen, Britten et al. 1996). Uka and DeAngelis (Uka and DeAngelis 2004) trained monkeys in a near/far - depth discrimination task and compared their performance with the sensitivity of single MT neurons. They found that single neurons were as sensitive as the observers. The behavioural paradigms in both studies demanded a relatively coarse judgement. Investigating neuronal sensitivity for a fine discrimination paradigm reveals sensitivity of individual neurons to be at best $2 - 3$ times worse than the observer as a whole when trying to predict behaviour based on the entire population of active neurons, while activity of the most sensitive neurons showed significant correlation with perceptual decisions (Purushothaman and Bradley 2005). This result suggests that performance of fine motion discrimination tasks critically depends on the activity of the most sensitive neurons.

Lesion studies in macaque monkeys also contributed to our understanding of the involvement of area MT in the perception of motion. Lesions of area MT result in a selective impairment of performance of motion perception indicated by elevated thresholds in perceptual tasks, e.g. motion detection, direction and speed discrimination (Newsome and Pare 1988; Cowey and Marcar 1992; Pasternak and Merigan 1994).

In chapter 4 the directional tuning in MT neurons during presentation of motion stimuli with two different dynamical behaviours is studied. It shows the feasibility of the reverse correlation approach to understand motion processing and the receptive field structure of area MT neurons. It could be shown that the underlying encoding mechanisms for motion direction in area MT are robust against different dynamic behaviours of the stimulus.

2.4 Investigating the Encoding of Stimulus Features by Reconstruction

Chapter 5 extends the idea of investigating the behaviour of MT neurons to different dynamic stimulus statistics by applying a reconstruction framework (Zhang, Ginzburg et al. 1998). Physical variables like orientation of a light bar or direction of motion are assumed to be coded as activity levels in populations of neurons. Thus physical variables cause specific neural activity patterns. If this is so, then it should be possible to reconstruct the physical variable from the neural activity pattern. This inverse problem (or decoding problem) is the problem the brain actually needs to solve. Reconstruction is useful in estimating how much information about the physical variables is present in the activity of a population of neurons. Second, by applying different reconstruction methods it provides insight into how the brain might use responses distributed over a population of neurons to solve computational problems such as visual object recognition. Two classes of reconstruction methods have been widely applied to population responses, namely probabilistic or Bayesian methods and basis function methods, i.e. population vector coding, optimal linear estimation or template matching (see for example (Zhang, Ginzburg et al. 1998) for a comparison of methods on decoding neural information from rat hippocampal place cells). Reconstruction methods have been applied to motor cortical activities during reaching tasks (Georgopoulos, Schwartz et al. 1986; Georgopoulos, Lurito et al. 1989; Schwartz 1994) using the population vector method, to disparity selective cells in visual cortex (Lehky and Sejnowski 1990) and to hippocampal place cells (Wilson and McNaughton 1993) using the template matching method.

Some authors claim that the ultimate test of our knowledge of the visual system is prediction: once we can predict the response of visual neurons to arbitrary stimuli, we can claim that we know what the visual system does (Carandini, Demb et al. 2005). The question to be answered in chapter 5 is: can we predict the neural response to a visual motion stimulus? Using two different stimulus statistics, a random walk, i.e. motion direction at any one moment is predictive for subsequent motion direction, and a white noise stimulus in the direction domain, it has been investigated whether it is possible to predict the neural responses of one stimulus statistics by simply applying the linear filters of a population of neurons obtained with the other statistics. This is an important question, because it directly impacts the way we think about the stimulus-response properties of cortical neurons. We apply a Bayesian method to determine the reliability of motion direction representation in area MT under different motion dynamics. Based on the direction *tuning curves* for a population of MT neurons assessed with a white noise direction stimulus, the task is to infer the direction of motion in each instance during presentation of a different dynamic motion

stimulus, directly extending the question in chapter 4 to what extend cortical representations of the outside world are sensitive to different stimulus statistics.

2.5 Attentive Modulation of Neural Response Behaviour

For most behaviour, visual information must be integrated with other, potentially non-visual information. Examples of these extraretinal signals are copies of the signals used to generate eye movements, so called efference copy or corollary discharge, to aid disambiguating retinal motion caused by the eye movement themselves from that to motion of an object. Other examples are eye position and attention.

Attention is a general term for cognitive processes which select meaningful perceptual information of our environment at the expense of other information classified to be less or not important at all. At any given moment a lot of information reaching our primary sense organs is completely irrelevant for our current goals and behaviour. Especially our visual system provides an overwhelming amount of information through the roughly 1.5 million axon fibres leading from the retina to the subsequent processing stages in the brain. Without any mechanism for sorting, evaluation and selection according to relevance, our visual system would be overcrowded with information that could not be processed in real time. Attention has been subject to scientific interest as early as the late nineteenth century, when William James published his monumental *Principles of Psychology* (James 1890) and devoted a whole chapter to the processes of attention. In that time investigation of cognitive processes was mainly based on introspection, i.e. self-observation and reporting of conscious inner thoughts and sensations. Since the middle of the twentieth century the process of perceptual selection has been extensively studied using all methods of cognitive neuroscience. I will discuss three fundamental questions about visual attention here, relevant to this thesis: (1) What is the locus of attentional selection? (2) What are the units of attentional selection? and (3) What are the neural correlates of attended versus unattended perceptual information?

Does selection occur early or late? According to the late selection theory, the whole visual information in the environment is perceptually analysed to a detailed level. Based on the results of this analysis, attention then selects relevant information for further analysis and response planning (Deutsch and Deutsch 1963). The early selection theory, in contrast, hypothesizes that incoming information is only processed rudimentarily, and attentional selection is prerequisite for object recognition and further perceptual analysis (Broadbent 1958). Modern imaging and electrophysiological techniques have been employed to investigate all stages of visual processing modulated by attention and to probe the relative timing of these effects to occur. A metabolic mapping study, for example, has shown

attention-dependent activity in macaque LGN (Vanduffel, Tootell et al. 2000), which is the earliest stage in the visual processing stream that receives feedback connections from subsequent processing stages. Attentional effects have also been described for the next stage in the processing hierarchy, V1. Motter (Motter 1993) found that many cells produced a higher response when the animals' attention was directed into the neuron's receptive field. Evidence for attentional modulation occurring in early stages of visual processing still leaves open the possibility that attentional effects occur at a temporally late step of processing thus reflecting feedback information of higher order stages. The latencies of attentional modulation reported in different studies vary in a large interval. Roelfsema and colleagues (Roelfsema, Lamme et al. 1998) reported attentional modulation to occur in V1 responses as late as 200 ms after stimulus onset, while Motter (Motter 1993) found V1 responses to be modulated directly beginning from stimulus onset. The major difference between these studies is that in the former study the subject was not informed about the location of the target object while in the latter the subject has been cued in advance to the location of the upcoming stimulus. Taken together the results implicate that neural responses to visual stimuli can be modulated by attention at very early stages and even during the first phase of processing. The results from Roelfsema (Roelfsema, Lamme et al. 1998) however show that this does not generalize for all sorts of tasks.

Moving up the processing hierarchy beyond V1, attention effects have been found to increase in magnitude, both for areas categorized to belong to the ventral and to the parietal pathway. Within the same experiments effects of attention have been shown to be stronger in V4 than in V1 (Haenny and Schiller 1988; McAdams and Maunsell 1999; Mehta, Ulbert et al. 2000), stronger in V2 than in V1 (Mehta, Ulbert et al. 2000), and stronger in area MST than in area MT (Treue and Maunsell 1996). These findings are in agreement with hierarchical accounts for attentional mechanisms where higher stages influence processing in lower stages in a top-down fashion according to the behavioural relevance of the visual information (Tsotsos, Culhane et al. 2001).

Does attention operate on locations, features or objects? What are the physical attributes, attention can operate on? There are three possible units selection can be based on: (1) The location in the visual field where a potentially relevant information resides can be enhanced. Many studies have been performed to investigate location-based mechanisms of selection by instructing subjects to focus their attention to a specific area in visual space in expense to the rest of the visual field. Consistent with the "spotlight" metaphor of attention (Eriksen and Eriksen 1974; Posner 1980), early studies by Moran and Desimone (Moran and Desimone

1985) showed neural enhancement for the attended stimulus over the unattended stimulus, when subjects were instructed to attend to stimuli at one location and ignore stimuli at another. The same has been shown in human subjects using positron emission tomography (PET) (Corbetta, Miezin et al. 1993). These early studies, however, were not able to conclusively show what has been selected, be it location, a specific feature or the whole object. Disambiguous evidence for space serving as a basic unit for attentional processes came from early psychophysical studies, showing that distracters closer to the target have greater perturbing influence than distracters farther away (Eriksen and Eriksen 1974). Subsequent neurophysiological (Connor, Gallant et al. 1996; Connor, Preddie et al. 1997) and neuro-imaging (Downing, Liu et al. 2001) studies confirmed this.

Can attention be allocated not only to a specific location in visual space but also to a particular feature – that is, enhancing the representation of perceptual information throughout the visual field that share a particular feature, like for example the same colour or the same shape? Evidence for feature-based attention mechanisms comes from a popular paradigm in visual psychophysics, visual search. In visual search experiments, targets and distracters differ by at least one feature. The efficiency of visual search depends on the number and type of distracters, being more efficient with low numbers of distracters and high dissimilarity between target and distracter. The number of items in a search array is called display size. The display size effect describes to which degree task performance depends on the display size. For feature search, searching for targets defined by a unique visual feature, such as colour, size or shape, the display size effect is roughly zero. This indicates that the target defining feature is enhanced over the whole visual field. Several psychophysical studies have shown that feature-based attention enhances behavioural performance across the visual field (Rossi and Paradiso 1995; Kumada 2001; Saenz, Buracas et al. 2003). Neural correlates of feature-based attention have been described in brain-imaging studies (Corbetta, Miezin et al. 1990; Wojciulik, Kanwisher et al. 1998; Saenz, Buracas et al. 2002). Neurophysiologic evidence for feature-based attention comes from single-cell recordings performed in area V4. Motter (Motter 1994; Motter 1994) trained monkeys to select a bar stimulus based on its color or luminance and then to discriminate the angular tilt of the selected stimulus. Most neurons were selectively enhanced when the colour or luminance in the receptive field matched the color or luminance of the cue. Robust effects of attention have also been shown in the dorsal pathway. Martinez-Trujillo and Treue (Treue and Martinez Trujillo 1999; Martinez-Trujillo and Treue 2004) examined attentional modulation in area MT by observing the neural responses to a behaviourally irrelevant stimulus within the receptive field while the animal

was instructed to attend to a moving stimulus outside the receptive field that could move either in the preferred or the null direction. Responses of neurons in area MT to the behaviourally irrelevant stimulus in the receptive field were modulated by the direction of motion that the animal was attending to. In another experiment these authors were able to show that feature-based and space-based attention mechanisms act independently and when determined separately in two experiments effects of both sum up approximately to the effect strength determined in a combined space- and feature-based attention experiment. Similar results have been found in area V4 (McAdams and Maunsell 2000).

Another possible unit attention mechanisms can operate on are whole objects. Early evidence for object-based attention came from psychophysical experiments carried out by Duncan (Duncan 1984). Using two spatially overlapping objects (a box with a line struck through it), it was found that two discrimination judgements that concern the same object can be made simultaneously without loss of accuracy, whereas two similar judgements concerning different objects cannot. Subsequent studies confirmed these early results of object-based attention mechanisms (Baylis and Driver 1993; Valdes-Sosa, Bobes et al. 1998; O'Craven, Downing et al. 1999; Davis, Driver et al. 2000; Mitchell, Stoner et al. 2003). Single cell electrophysiological data supporting the idea of object-based mechanisms is very rare and are confined to one work done in area V1 (Roelfsema, Lamme et al. 1998) and another work done in motion sensitive area MT (Wannig, Rodriguez et al. 2007). In the latter one monkeys were trained to attend to one of two transparent random-dot surfaces segregated by colour, which occupied the same region in space. The motion signal from the attended surface influenced the neural response more strongly than a physically identical motion signal originating from the unattended surface. This held true even when the segregation by colour was removed.

What are the neural correlates of attended versus unattended visual information? After having listed all these studies in favour or against early- versus late-selection, and for the various physical attributes being the operational unit of attention, there is one important question remaining: how exactly does attention modify the neural responses? Most studies trying to answer this question applied behavioural paradigms in combination with single cell electrophysiology, focusing on firing rates of single neurons. It has been shown in several single-unit studies that directing attention into the location of the receptive field leads to a response enhancement, i.e. higher firing rates, compared to directing attention somewhere else in the visual field for both ventral and dorsal pathway: V1: (Haenny and Schiller 1988; Motter 1993), V2 (Motter 1993; Luck, Chelazzi et al. 1997; Mehta, Ulbert et al. 2000), V3 (Nakamura and Colby 2000), V4 (Fischer and Boch 1985; Haenny, Maunsell et al. 1988;

Maunsell, Sclar et al. 1991; Motter 1993; Connor, Gallant et al. 1996; Connor, Preddie et al. 1997; McAdams and Maunsell 1999; Mehta, Ulbert et al. 2000), inferior temporal cortex (IT) (Richmond, Wurtz et al. 1983; Richmond and Sato 1987; Spitzer and Richmond 1991; Mehta, Ulbert et al. 2000), MT (Treue and Maunsell 1996; Treue and Martinez Trujillo 1999; Wegener, Freiwald et al. 2004; Wannig, Rodriguez et al. 2007), MST (Treue and Maunsell 1996; Treue and Martinez Trujillo 1999) and LIP (Bushnell, Goldberg et al. 1981; Colby, Duhamel et al. 1996; Gottlieb, Kusunoki et al. 1998).

Another potential mechanism for strengthening the neural representation of one stimulus over all other stimuli in the visual field is synchronization. Concerting the activity of all neurons involved in representation of the attended stimulus in a synchronous fashion would increase the impact of this representation to downstream processing stages and potential output structures like for example areas guiding motor responses or eye movements. An attentional mechanism based on synchronized cell assembly activity has first been proposed by Crick and Koch (Crick and Koch 1990) and subsequent modelling showed physiological plausibility (Niebur, Koch et al. 1993; Niebur and Koch 1994). The first study providing electrophysiological data was done in somatosensory cortex during performance of a visuo-tactile attention task (Steinmetz, Roy et al. 2000). Fries and colleagues (Fries, Reynolds et al. 2001) employing a visual spatial attention task, found increased synchronization of spikes with the LFP especially at high frequencies in the gamma-range (35 – 90 Hz). The same frequency range has been reported to play a major role in feature binding (Singer and Gray 1995). A similar enhancement in gamma power under attentional load has been reported in several EEG studies (Gruber, Muller et al. 1999; Muller, Gruber et al. 2000). Likewise Taylor and colleagues (Taylor, Mandon et al. 2005) reported a close coupling between attentional demand and oscillatory activity of local field potentials in the gamma-frequency range in area V4 of macaque monkeys performing a shape-tracking task.

Effects of attention on the response to moving stimuli have been described by Treue and Maunsell (Treue and Maunsell 1996) in area MT and MST, which were the first to describe attention effects in the parietal pathway. These authors used a single dot moving inside the receptive field along the cell's preferred and null direction and placed another dot outside the receptive field. The subjects were required to detect a velocity change on the target dot, which was instructed by a cue in the beginning of the trial. When attending to the dot inside the receptive field the responses of MT neurons were enhanced on average by 19%. This work has shown the attentional influence onto neural responses, namely tuning curves, in visual area MT on relatively long timescales (up to several seconds). Subsequent studies have

confirmed these early results by using a wide span of different behavioural tasks (Seidemann and Newsome 1999; Wegener, Freiwald et al. 2004; Wannig, Rodriguez et al. 2007; Busse, Katzner et al. 2008). In chapter 6 single unit recordings have been accomplished to investigate the influence of attentional state on the spatiotemporal receptive field of visual area MT neurons. By applying a similar stimulus design used in chapter 4 and 5 combined with an attention demanding behavioural task effects of attention on the response of neurons in area MT could be investigated on very fine time scales and with a high number of repetitions of each individual motion direction. Furthermore the motion direction space could be covered with very high resolution.

2.6 A network perspective on effects of attention on cortical motion processing

In chapter 7, a more systemic view is taken to investigate brain function. Using an only slightly modified behavioural paradigm described in chapter 6 the whole brain is investigated for attentional modulation by using functional magnetic resonance imaging (fMRI). Since its introduction in 1990 by Ogawa and colleagues (Ogawa, Lee et al. 1990; Ogawa, Lee et al. 1990) blood oxygenation level dependant (BOLD) functional magnetic brain imaging has been widely applied to investigate human brain function. In 1999 Logothetis and colleagues (Logothetis, Guggenberger et al. 1999) were the first using this method successfully in anaesthetized as well as in alert macaque monkeys. Subsequent studies in alert behaving monkeys proved the usefulness of this technique especially in performing comparative studies between humans and monkeys (Vanduffel, Fize et al. 2002; Tsao, Vanduffel et al. 2003). In addition to the opportunity to do comparative studies, fMRI allows to investigate brain function on large spatial scales up to covering the whole brain. This comes at the expanse of relatively poor temporal resolution in the range of several seconds. Functional magnetic resonance imaging in awake behaving macaques has been applied for investigation neural networks involved in motion processing (Vanduffel, Fize et al. 2001; Nelissen, Vanduffel et al. 2006)), in depth perception (Tsao, Vanduffel et al. 2003), in face processing (Tsao, Freiwald et al. 2003; Tsao, Schweers et al. 2008) and in the perception of colour (Conway, Moeller et al. 2007). The preeminent advantage of performing fMRI studies in macaque monkeys is the opportunity to target areas proven to be involved in performance of a task or involved in the representation of a specific feature with electrophysiological methods, i.e. single cell electrodes. Some of the above mentioned studies have already taken advantage of this opportunity (Tsao, Freiwald et al. 2006; Conway, Moeller et al. 2007).

Here the whole brain of macaque monkeys has been investigated for changes in cerebral blood oxygenation level while alternating spatial attention to two different positions in the visual field. Observing the whole brain allows tackling questions already raised in chapter 2.5 : (1) Does selection occur early in the visual cortex? Only few electrophysiological studies have shown attentional modulation for area V1 (Roelfsema, Lamme et al. 1998), but fMRI in humans has revealed strong modulation in area V1 during performance of an attention task (Huk and Heeger 2000). (2) What areas are involved in the top-down control of attention? In humans it has been proposed that the source of top-down biasing signals derives from a network of areas in frontal and parietal cortex (Kastner and Ungerleider 2000). In macaque monkey parietal area LIP has been considered to play a major role in controlling spatial mechanisms of visual attention (Gottlieb, Kusunoki et al. 1998). Conducting electric microstimulation in frontal area FEF (frontal eye field), Moore and Armstrong (Moore and Armstrong 2003) were able to observe neural effects in simultaneously recorded V4 neurons that resembled very well modulations due to spatial visual attention. V4 neurons exhibiting their receptive field location in the response field of the stimulated FEF neuron showed increased firing rates in response to electrical microstimulation. This result suggests an involvement of FEF in gain modulation of specific locations within the visual field. The same mechanism could be used to enhance the neural representation of one location in the visual field to the expense of all the others. (3) Are there potentially other areas involved? It could be possible that the dichotomy of the two processing streams, ventral and parietal, first proposed by Ungerleider and Mishkin (Ungerleider and Mishkin 1982), is not that strictly held when higher order cognitive functions, like for example attention, are involved. This might even hold true for paradigms perceptually clear-cut devoted to one or the other. In our case, a motion discrimination task involvement of mainly the parietal pathway would be expected. Functional magnetic resonance imaging nevertheless allows monitoring of the whole brain and investigation of areas not suspected to be involved in performance of this particular task.

3 Overview of experiments

The central question of my thesis was to investigate the neural mechanisms of motion representation in the macaque brain using different theoretical and experimental methods. These experiments and theoretical accounts are described in the following chapters.

The chapter "Encoding of dynamic visual stimuli by primate area MT neurons" has been published in Neurocomputing (Stemmann, Freiwald et al. 2005). Concept development has been done in co-operation with Winrich Freiwald and the Institute for Theoretical Neurophysics, Christian Eurich and Erich Schulzke. All the data acquisition and basic data analysis have been done by me. I created figures, and wrote the first draft of the manuscript. All further refinement of the manuscript have been done in close collaboration with Winrich Freiwald, Christian Eurich and Erich Schulzke.

The results presented in the chapter "Linearity of Macaque Area MT Neurons: Model-based Interpretation of Population Responses to Dynamic Motion Stimuli" are based on data, which have been acquired by me. I performed the basic data analysis and all further data analysis were done in close co-operation with Winrich Freiwald, Christian Eurich and Erich Schulzke.

Answering the question of how attention modulates the response properties of neurons in cortical area MT was subject of the experiments described in the chapter "Effects of Attention during Rapid Serial Visual Presentation on Macaque MT Neurons", which has been presented on the Society for Neuroscience Annual Meeting 2005, Washington, DC (Stemmann and Freiwald 2005). This project has been performed under supervision of Winrich Freiwald. I developed the concept for the behavioural paradigm, trained the monkeys and performed all data recording. I performed the data analysis and created all figures and wrote the manuscript.

The chapter "Imaging Cortical Networks of Spatial Attention in the Macaque" was aimed at describing the cortical network activated during a motion discrimination task. I performed the conceptual development in co-operation with Winrich Freiwald. Animal training, data acquisition, data analysis and preparation of figures has been done by me. I wrote the first manuscript and further refinement has been done in close co-operation with Winrich Freiwald. This work has been submitted to Neuron and is in the review process.

4 Encoding of dynamic visual stimuli by primate area MT neurons

4.1 Abstract

Neural stimulus selectivity is thought to be optimized for the representation of real-world stimuli. Neural coding properties, therefore, may adapt to different environments. Here, we address the question if tuning curves depend on the statistics of visual stimuli. This is done by studying the directional tuning of macaque area MT neurons exposed to dynamic motion stimuli of two different direction progression statistics. Despite an apparent difference of tuning curves across stimulus conditions, our results support the view that the underlying encoding system is robust and subject to only restricted malleability by stimulus statistics.

4.2 Introduction

Tuning properties in early visual cortex can be dynamic in several ways. Tuning curves can change within a few tens of milliseconds (Ringach, Hawken et al. 1997) and undergo substantial changes as the result of adaptation (Dragoi, Sharma et al. 2002) over the course of hundreds of milliseconds to seconds. Since many properties of early visual cortex can be understood in terms of an evolutionary optimization for an efficient representation of natural scene statistics (Olshausen and Field 1996), tuning dynamics and adaptivity may be hypothesized to optimize representations within different stimulus contexts. If this is so tuning properties ought to change with stimulus statistics.

In order to address this question, we designed a novel stimulus paradigm which mimics important aspects of natural scenes: subsequent stimulus states were determined both by a random component and a continuity requirement which dominates most of the motions we see. These stimuli were implemented as random walk trajectories in the motion direction domain, in which subsequent motion directions are correlated. Neural direction tuning obtained in this stimulus context was compared to tuning during stimulus sequences which only differed in the lack of any sequential correlations of motion directions.

4.3 Methods

We conducted extracellular recordings in area MT in two male macaque monkeys (*Macaca mulatta*). Surgical procedures, single-unit recording and data acquisition were standard, in short: prior to the experiments, the animals were surgically implanted with a head-holding device, a recording cylinder, and a scleral search coil. Surgical, animal care, and experimental procedures conformed to the guidelines of the National Institutes of Health for the care and

use of laboratory animals, the guidelines for the welfare of experimental animals issued by the Federal Government of Germany, and stipulations of local authorities. Recordings were taken with Tungsten microelectrodes, signals amplified, filtered (350-5000 Hz), digitized (sampled at 25 kHz) and stored on computer disk for offline analysis.

Monkeys performed a fixation task, foveating a small white fixation spot (0.13° x 0.13°) within a 3° x 3° window, while gaze direction was monitored with the indirect scleral eye-coil method. Visual stimuli were presented on a CRT monitor at a distance of either 57 cm or 86 cm at 100 Hz or 85 Hz refresh rate, respectively. After appearance of the fixation spot on the otherwise dark monitor and after the monkey started foveating the fixation spot, stimuli appeared inside a circular aperture covering the classical receptive field of the neurons under investigation. Trials in which the monkeys broke fixation before the end of stimulus presentation were discarded. After each successfully completed trial the animals were rewarded with a drop of juice.

Visual stimuli consisted of random-dot surfaces undergoing translational motion. In both stimulus paradigms the single dot diameter was 0.2°, the dot density 2 dots$^{\circ 2}$ visual angle, translation speed 7°/sec, dot lifetime was infinite, and motion coherence was 100%. Motion directions were updated at a rate of 50, 20 or 10 Hz. The length of individual motion sequences ranged from 3 to 5 seconds. In the first stimulus paradigm, an adoption of the approach in (Ringach, Hawken et al. 1997) to the motion direction domain (Borghuis, Perge et al. 2003), a stochastic sequence of motion directions was created (see Figure 4 A and B for illustration). The sequence was generated by pseudo-randomly selecting a new direction out of a set ζ of directions, sampled in steps of 30°. Selection from the set was done with replacement. The second stimulus paradigm also drew from ζ, but the difference between two subsequent motion directions was now determined by a Gaussian distribution with mean zero and standard deviation one, thus realizing a time-discrete random walk with fixed step size in direction space. The two stimulus paradigms will be referred to as *discontinuous* and *continuous*, respectively. Four different stimulus sequences have been generated, referred to as *trajectories*, with different starting directions (0°, 90°, 180° and 270°). Each stimulus sequence was repeatedly presented for 10 to 20 trials.

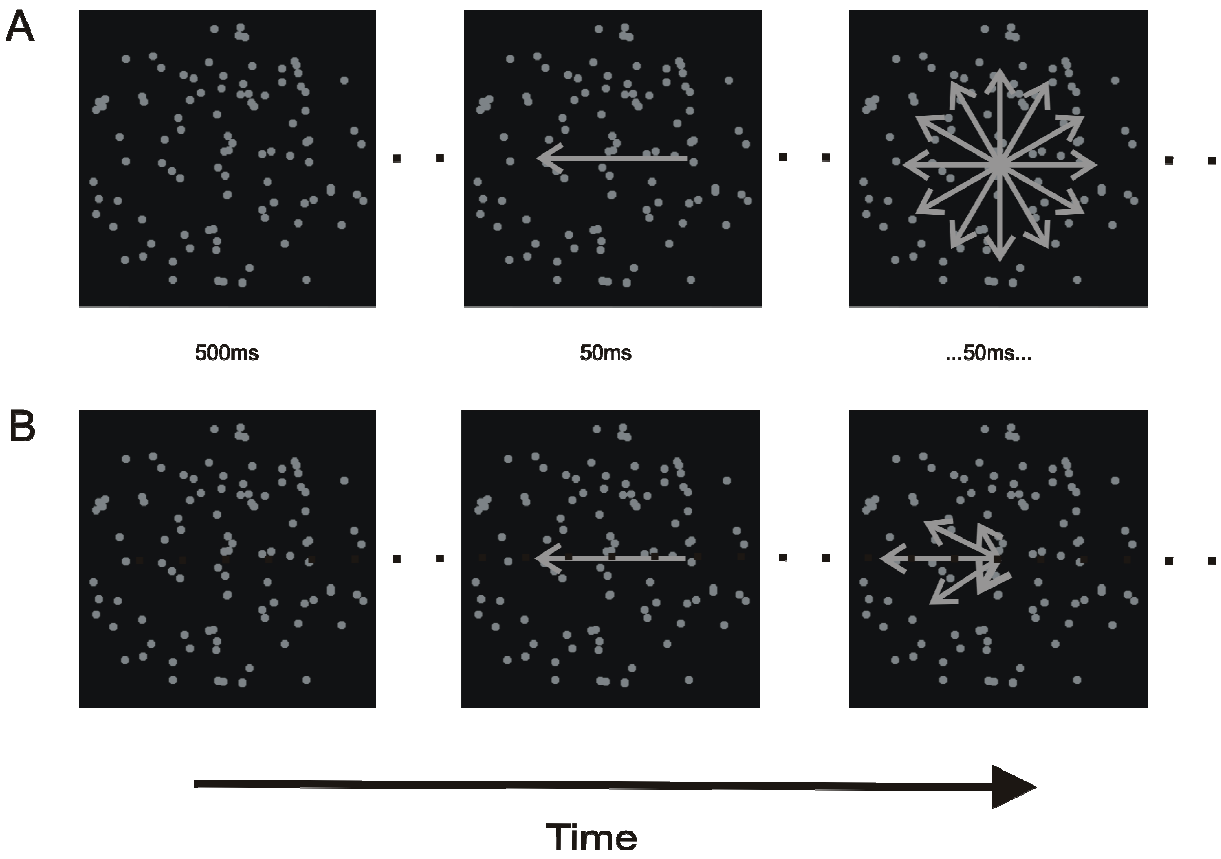

Figure 4: Sketches of the two different stimulus paradigms. In both paradigms, a trial starts with a 500 ms period of static presentation of the random dot stimulus followed by a translational movement for 50 ms (here: leftward, indicated by the gray arrow). The only, but crucial difference between both types of stimulation is indicated for the third stimulation phase by the set of arrows, whose lengths indicate choice probabilities. While in the first, "*discontinuous*", stimulus paradigm consecutive motion directions are chosen randomly out of the set ζ of directions with identical probabilities (**A**), in the second, "*continuous*", paradigm the consecutive directions are chosen from a Gaussian distribution centered on the current motion direction (**B**).

In order to assess the influence of stimulus statistics on neuronal responses we calculated estimates of tuning curves and optimal linear filters for the two stimulus paradigms. *Spike-Triggered-Averages* (STAs) were calculated for the discontinuous stimulation by reverse correlating the spike train with the stimulus sequence. As the discontinuous stimulus is uncorrelated, the STA closely approximates the optimal linear filter between stimulus and response. We defined the response *delay* of each neuron as the time difference between the occurrence of a spike and the maximum value of the time-dependent variance in the STA (see Figure 5A). *Tuning curves* were then computed from responses within a time window centered at this delay. In order to limit the influence of preceding and subsequent motion directions on the response, we divided the analysis window into five partitions of equal length, weighting the firing rates within these partitions according to different distributions (the uniform distribution corresponds to the "standard" tuning curves).

Direction indices were calculated as $DI = \frac{optTK - orthTK}{optTK + orthTK}$, where *optTK* (*orthTK*) indicates the tuning curve value at the (anti-)preferred direction.

4.4 Results

Figure 5 A and B show the spike triggered average of a typical cell and illustrate the procedure for delay estimation. When a spike of this cell occurred, it had, on average, most likely been preceded by a motion direction of about 300° approximately 100 ms earlier. The least likely stimulus event preceding the spike was a motion direction of about 120° at about the same delay. The tuning behaviour of the neuron depicted in Figure 5 is shown in Figure 6 A.

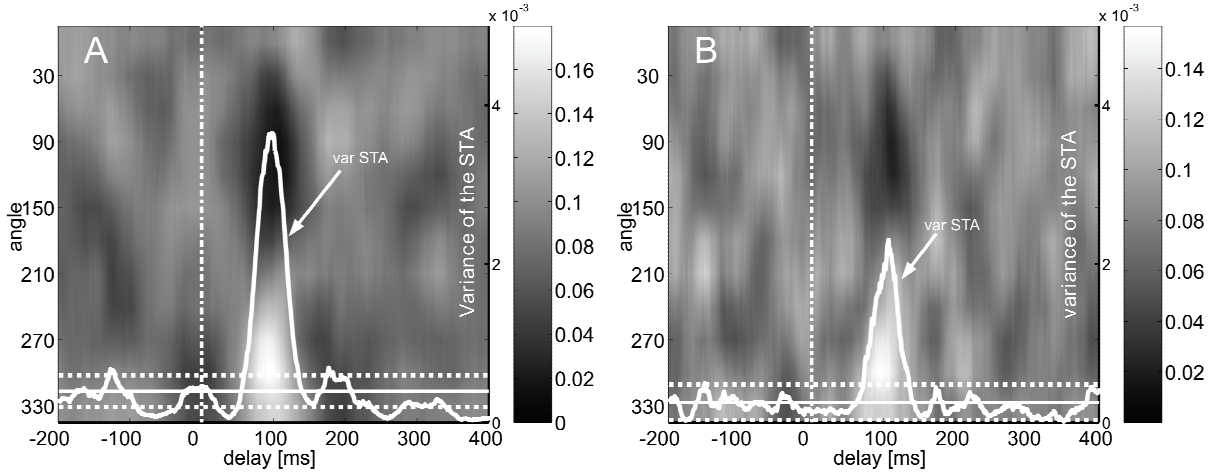

Figure 5: Spike triggered averages for an exemplary MT cell, obtained during discontinuous stimulation. (**A**) was obtained with 50 ms motion intervals, (**B**) with 20 ms motion intervals. The abscissa denotes the time before a spike event, the left ordinate denotes stimulus motion direction. A column within the spike triggered average for a fixed time delay τ represents the probability distribution (gray scaled) for the stimulus present at that delay before a spike. The white line shows the variance of the STA. The solid horizontal lines depict mean variance of the STA in the interval 200 ms to 100 ms before a spike; dashed lines indicate the twofold standard deviation; the dashed-dotted vertical line indicates the spike timestamp. The cell shows a characteristic response latency of approx. 100 ms for both interval lengths.

General activity and directional tuning differs in the two conditions. While activity levels did not vary systematically with stimulus statistics across the population of cells, directional tuning did. The histogram in Figure 6 B shows the distribution of directional indices for the whole population of 45 MT neurons. The two distributions' medians, *continuous* vs. *discontinuous*, differ significantly (paired Wilcoxon test: $p < 0.01$). Thus, when tuning curves in the two conditions are computed over the full width of one stimulus presentation window (here, 50 ms), directional tuning seems to be significantly more pronounced during random walk ("continuous") stimulation. However, this tuning difference between stimulus conditions depends on the choice of the analysis window (see Figure 6 C).

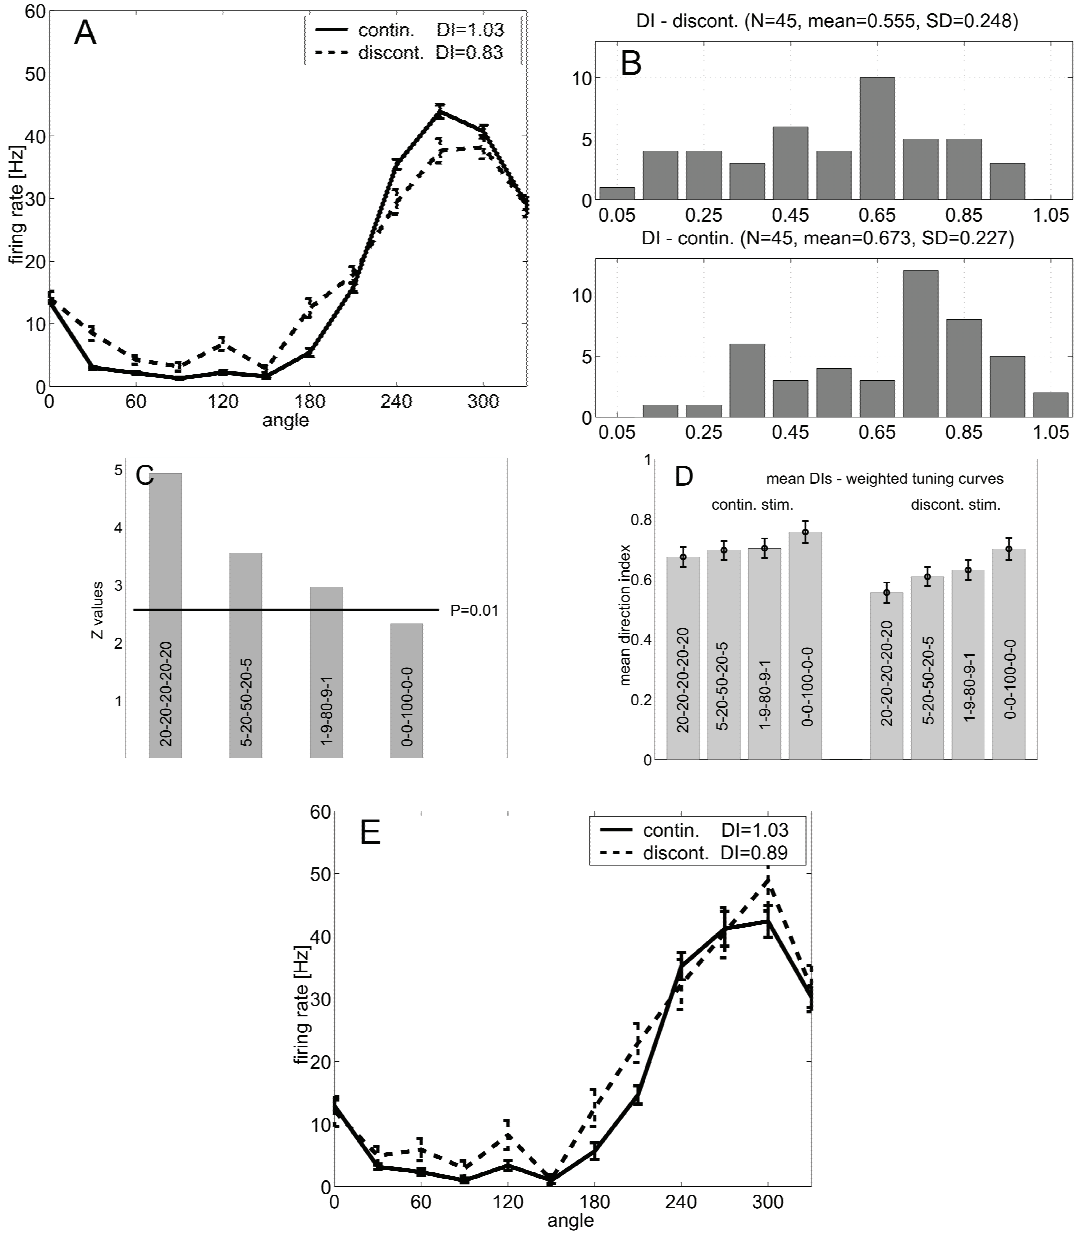

Figure 6: Emphasis of different partitions of analysis windows is crucial for the obtained tuning curves. (**A**) Tuning curves of the neuron in **Figure 5** for both stimulation paradigms, obtained with an analysis window of 50 ms (full stimulus presentation window) and taking into account the cell's response latency, obtained with STAs (see **Figure 5**). Solid line: continuous stimulation. Dotted line: discontinuous stimulation. Note that the continuous stimulation paradigm provides the sharper tuning curve. (**B**) Distribution of direction indices (DI) for all recorded neurons (n=45) and both stimulus paradigms. The mean DI obtained with the discontinuous paradigm (upper panel) is significantly lower ($p<0.01$, paired Wilcoxon test, see also (**C**)) than the mean DI obtained with the continuous paradigm (lower panel), indicating that tuning sharpness is reduced in the latter condition. DIs were obtained by using the full analysis window of 50 ms. (**C**) Significance of the difference between the mean DIs obtained with the two paradigms (ordinate) depends on the different weighting of the partitions of the 50 ms-analysis-windows (abscissa). The 50 ms-windows were divided in 5 parts of 10 ms. The numbers within the bars indicate the weighting of each of the windows´ partitions (see methods for details). With higher weighting of the central region, the significance level decreases, indicating that the tuning curves obtained with the two paradigms assimilate to each other. (**D**) Mean DIs of the population (ordinate) depend on the weighting of the windows´ partitions as done in (**C**) (abscissa). With higher weighting of the inner partitions, DIs increase both in the continuous as in the discontinuous paradigm, indicating sharper tuning curves. Note that the increase is stronger in case of the discontinuous stimulation, meaning that the differences between both paradigms decrease. (**E**) Tuning curves of the neuron in **Figure 5** for both stimulation paradigms, obtained with an analysis window of 10 ms, which is located in the center of the 50 ms presentation window, whereas the other partitions are discarded. Note the similarity of tuning curves in comparison to (**A**).

With increasing emphasis on the window's core region, a continuous decline of the significance level for the difference between the direction indices is observed. What are the reasons for this effect? As the analysis window shrinks, tuning curves are constructed with fewer spikes and, thus, the accuracy of determining the directional index may degrade, obliterating real tuning differences. A second possibility, however, is that tuning can in fact be more accurately determined with temporally restricted analysis windows, thereby minimizing the influence of responses to preceding or subsequent stimuli. Our data support the latter scenario. As is shown in Figure 6 D directional tuning improves during temporal focussing of the analysis window. This improvement is more pronounced for directional tuning during discontinuous than during continuous stimulation, thereby reducing the difference in directional tuning between the two conditions. This is illustrated in Figure 6 E for the example neuron, whose tuning curves are more similar when computed over a 10 ms than when computed over the full 50 ms (Figure 6 A) window.

4.5 Discussion

We have studied the malleability of neural tuning by stimulus statistics using as a model system directionally selective neurons from macaque area MT. When exposed to direction trajectories from two different statistics, neural tuning seems to be sharper during continuous stimulation. This effect can have two explanations. First, stimulus statistics genuinely alters the tuning properties of neurons in area MT. In the continuous stimulus paradigm, i.e., the random walk in movement direction, motion direction changes less abruptly than in the discontinuous stimulus paradigm and may better reflect physical motion we see in our natural environment. Neurons in MT may therefore adapt better to this stimulus statistics, resulting in tuning curves that are more selective. Second, since neurons in MT integrate stimuli over time as shown for example in Figure 5, a response at any point in time may result from the influence of more than a single stimulus. Due to the continuity constraint subsequent stimuli tend to be similar in case of continuous stimulation, but not in case of discontinuous stimulation. Therefore, sharper tuning can be expected during continuous stimulation - not as a result of a genuine change of tuning characteristics, but as the result of a direct effect of stimulus statistics on the analysis procedure. The evidence from area MT presented here lends support to the latter explanation. Directional filter properties of MT neurons appear robust and may not adapt to the different stimulus statistics used in our experiment. Since prolonged presentation of a single motion direction has been found to induce substantial adaptation (Kohn and Movshon 2004) and does constitute an important characteristic of stimulus

statistics, future research will need to identify the precise stimulus patterns to which the information processing system adapts and to which it does not. Our finding provides some constraints on the time-scale of such conditions and, furthermore, demonstrates that neural tuning curves obtained with a random stimulus sequence can be directly used to predict responses in quite different stimulus contexts.

5 Linearity of Macaque Area MT Neurons: Model-based Interpretation of Population Responses to Dynamic Motion Stimuli

5.1 Introduction

Information processing at virtually all levels of the visual system is adaptive: activity levels fall as stimuli fail to change. For example, cells in the retina adapt to the ambient light level, becoming less sensitive to light as its mean intensity increases (Shapley and Enroth-Cugell 1974), and neurons in infero-temporal cortex decrease their responses to repeated presentation of stimuli of identical shape (Miller, Gochin et al. 1991). Which aspects of a stimulus do neurons adapt to? Retinal ganglion cells have been shown to adapt to various aspects of the stimulus statistics, including image contrast, spatial scale (Smirnakis, Berry et al. 1997), orientation, spatial and temporal stimulus correlations (Hosoya, Baccus et al. 2005). Similarly, spatiotemporal receptive fields of neurons in primary sensory areas seem to be adapted to different stimulus statistics (Theunissen, David et al. 2001).

Adaptation properties of higher sensory areas in general are largely unknown, but work on motion-sensitive area MT has begun to determine the spectrum of adaptivity in one cortical area. Area MT is of special interest to the study of adaptation for two connected reasons. First, activity in area MT has been shown to be tightly linked to motion perception (Newsome, Britten et al. 1989; Salzman, Britten et al. 1990; Bradley, Chang et al. 1998; Croner and Albright 1999; Thiele, Dobkins et al. 2000; Dodd, Krug et al. 2001). Second, visual motion processing is strongly modified by adaptation, as has been shown in numerous psychophysical studies (e.g. (Mather, Verstraten et al. 1998). Thus, given the close link between MT activity and motion perception and the profound effects of adaptation on motion perception, MT activity should provide an assay to study the neural basis of motion adaptation. Indeed, recent studies have demonstrated adaptation of MT neurons to stimulus contrast (Kohn and Movshon 2003) and stimulus direction (Kohn and Movshon 2004). Not only can directional tuning of MT neurons adapt, but also its temporal window of integration over motion energy (Bair, Zohary et al. 2001).

In the present study, we asked whether MT neurons adapt to the temporal statistics of dynamic motion stimuli. We presented sequences of rapidly changing motion stimuli, which either changed direction randomly (Borghuis, Perge et al. 2003), or cohesively, such that a

motion direction at any one moment in time was predictive for the subsequent motion direction. We used this paradigm to test for changes of overall activity, direction selectivity and temporal integration of MT neurons in the awake, fixating macaque monkey.

5.2 Materials and Methods

5.2.1 Recording and Surgical Techniques

We conducted extracellular recordings in area MT in two male macaque monkeys (*Macaca mulatta*). Surgical procedures, single–unit recording and data acquisition were standard, in short: Prior to the experiments, the animals were surgically implanted with a head-holding device, a custom-designed recording cylinder, and the scleral search coil. Surgical, animal care, and experimental procedures conformed to the guidelines of the National Institutes of Health for the care and use of laboratory animals and were performed in accordance with the guidelines for the welfare of experimental animals issued by the Federal Government of Germany.

We used Tungsten microelectrodes with impedances of approximately 0.5 to 1 MΩ (World Precision Instruments), introduced to the cortex through a transdural guide tube positioned in a Teflon grid inside the recording cylinder.

Signals were amplified with conventional electrophysiological instrumentation (Multichannel Systems, Germany), filtered (350–5000 Hz), and displayed on an oscilloscope. Spike activity was digitized (sampled at 25 kHz) and stored on computer disk for offline analysis.

5.2.2 Behavioural Task and Visual Stimulation

The monkeys were alert, fixating their eyes and passively viewing during stimulus presentation. Eye position was monitored with the scleral search–coil method (Robinson 1963).

The monkeys viewed visual stimuli on a CRT monitor at a distance of either 57 cm or 86 cm. Stimuli were drawn with a commercially available graphics board (NVidia GeForce FX). The refresh rate of the monitor was either 100 Hz or 85 Hz.

At the beginning of each trial, a small, white fixation point (0.13×0.13°) appeared on the otherwise dark monitor. After the monkey started fixating the fixation spot, stimuli appeared inside a circular aperture, which was at least as large as the classical receptive field of the neurons under investigation, as determined by a mapping procedure with manually controllable gray bars. The animals were required to maintain fixation within a 3×3 deg

window centred on the fixation point. Trials were discarded in which the monkeys broke fixation before the end of stimulus presentation. After each successfully completed trial animals were rewarded with a small amount of liquid.

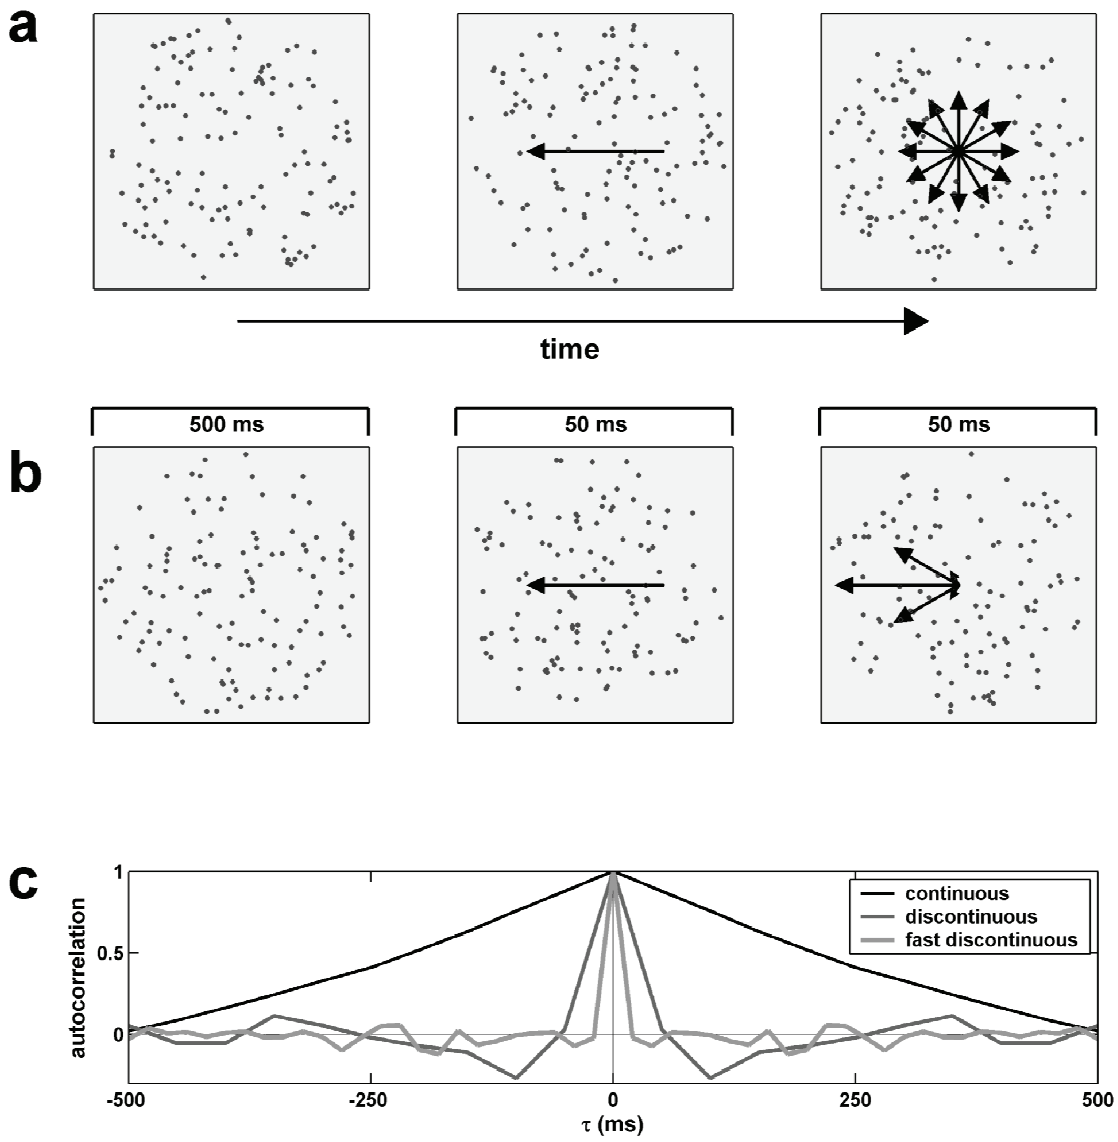

Figure 7: Visual stimulation with discontinuous (a) and continuous (b) direction sequences. A static random dot cloud (RDC) is shown for the first 500 ms of each trial (leftmost images). Then the RDC starts moving coherently in a direction. The direction for the next 50 ms is determined by stochastic processes. The scaled transition probabilities are depicted by the length of the arrows in the rightmost images: For discontinuous stimulation (a) the new direction is independent of the preceding direction (equal length of all arrows), while for continuous stimulation (b) there is a strong preference of the preceding direction. Figure (c) shows the autocorrelation functions of the three temporal statistics. The continuous stimulation exhibits a broad autocorrelation (black line). The discontinuous stimuli approach the delta function for decreasing single stimulus presentation time (dark grey line: 50 ms per direction, grey line: 20 ms).

Visual stimuli consisted of random dot surfaces with translational movements at 100% coherency and constant velocity. Apart from motion direction (see below) all stimulus parameters were constant and were not adjusted to the preferences of individual cells, that is:

dot density was 2 dots per square deg visual angle, dot size was 0.2°, velocity 7°/s and the dot lifetime was infinite. In the first stimulus paradigm (Figure 7a), which was adapted from (Ringach, Hawken et al. 1997), a stochastic sequence of motion directions was created by pseudo–randomly selecting a new direction out of a set ζ of directions, sampled in steps of 30 deg, either every 20 ms or every 50 ms (basic motion interval). The sequence was repeatedly presented for 10 to 30 trials for each cell. These stimulations are labelled '*fast–discontinuous*' and '*discontinuous*', respectively, throughout the remaining text. The second stimulus paradigm (Figure 7b) also draws from ζ, but with the difference that the direction changes from one time window to the next are given by a Gaussian probability distribution centred over the actual direction. This stimulus sequence, therefore, constitutes a random walk on the discretized motion direction axis. Four different stimulus sequences of this kind were generated, which we label '*continuous*'. Each started with a different direction (0 deg, 90 deg, 180 deg and 270 deg) in order to cover the full direction space.

The presentation of the different stimulus sequences was done in a pseudo–randomized order to ensure equal numbers of trials for each of three stimulus conditions continuous, discontinuous, and fast–discontinuous.

The autocorrelation function (Figure 7c) of the continuous stimulus consists of two parts: A high frequency component which results from the presentation time of 50 ms per direction and affects also the spike triggered average from discontinuous stimulus sequences, and a low frequency component which results from the stimulus–value dependent transition probability $P(\theta_t|\theta_{t-1})$ of the continuous stimulus sequence.

5.2.3 Data Analysis

The spike–sorting procedure was carried out in a custom–made program written in IDL (RSI, Boulder, CO). All other data analysis, trajectory reconstructions and simulations, were performed in MATLAB (MathWorks, Natick, MA).

Quantities written in bold lower case letters denote vector quantities throughout this chapter.. Upper case letters denote matrices or maximum values for indices.

Spike–triggered averages (STAs) (de Boer and Kuyper 1968) were computed as follows. First, stimuli at time t, $\mathbf{s}[t]$, were defined to be vectors of length K, where K denotes the number of possible directions (K=12 in our experiments). The vector entry corresponding to the actually presented stimulus direction is 1, while all other entries are set to 0. In this formulation $\mathbf{s}[t]$ can be interpreted as a probability distribution of directions at some time point t.

STAs $A[\theta_k, \tau]$ were now calculated as the average stimulus vectors **s** at time delay τ before a spike of the given neuron:

$$A[\theta_k, \tau] = \frac{1}{N} \sum_{i=1}^{N} s[\theta_k, t_i - \tau]. \quad \mathbf{1}$$

Time is discretized in steps of 1 ms. t_i denotes the time of the i–th action potential, N is the total number of action potentials, and k = 1,…,K is the index of the discrete set of directions θ_k. The elements of a STA column for fixed τ sum up to 1: $\Sigma_k A[\theta_k, \tau] = 1$. This definition of the STA extends the more common characterization (e.g., (Rieke, Warland et al. 1999) by a second dimension for direction in addition to time. STAs were calculated separately for continuous and discontinuous stimulation.

STA *variances* of the directions θ_k were determined as function of the delay τ according to

$$\sigma^2[\tau] = \left\langle \left(\langle A[\theta_k, \tau] \rangle_k - A[\theta_k, \tau] \right)^2 \right\rangle_k . \quad \mathbf{2}$$

The angle brackets $\langle \cdot \rangle_k$ denote averaging over motion directions θ_k. A high variance at some delay τ indicates that the stimulus probability distribution differs strongly from a flat probability distribution. This deviation is interpreted as relevant for spike creation and correlates with the strongest tuning of the cell (Borghuis, Perge et al. 2003).

The *latency* of a cell is defined to be the delay τ, which maximizes the variance $\sigma^2[\tau]$ determined from discontinuous stimulation (Mazer, Vinje et al. 2002). Latency values could also be calculated by using the STA from continuous or fast–discontinuous stimulation. Note that this definition of latency differs from the one commonly used, the delay between stimulus and response onsets.

We validated our latency–estimation procedure in a reconstruction of a discontinuous trajectory: In addition to the individual delays of the cells we introduced a global time delay τ_{add} for all neurons used for a reconstruction and determined trajectory estimation errors as a function of τ_{add}.

Tuning curves were determined for each neuron separately from continuous, discontinuous, and fast–discontinuous stimulation by averaging the firing rate responses to each direction (Stemmann, Freiwald et al. 2005).

The response rate to a single direction presentation was estimated as follows (Figure 8a): The spike train of a neuron was shifted relative to the stimulus sequence by the previously

determined latency τ_{lat}. The neural response rate was then defined as the number of spikes within the basic motion interval divided by the duration of the motion interval (20 ms for fast–discontinuous stimulation or 50 ms for the other two stimulation paradigms).

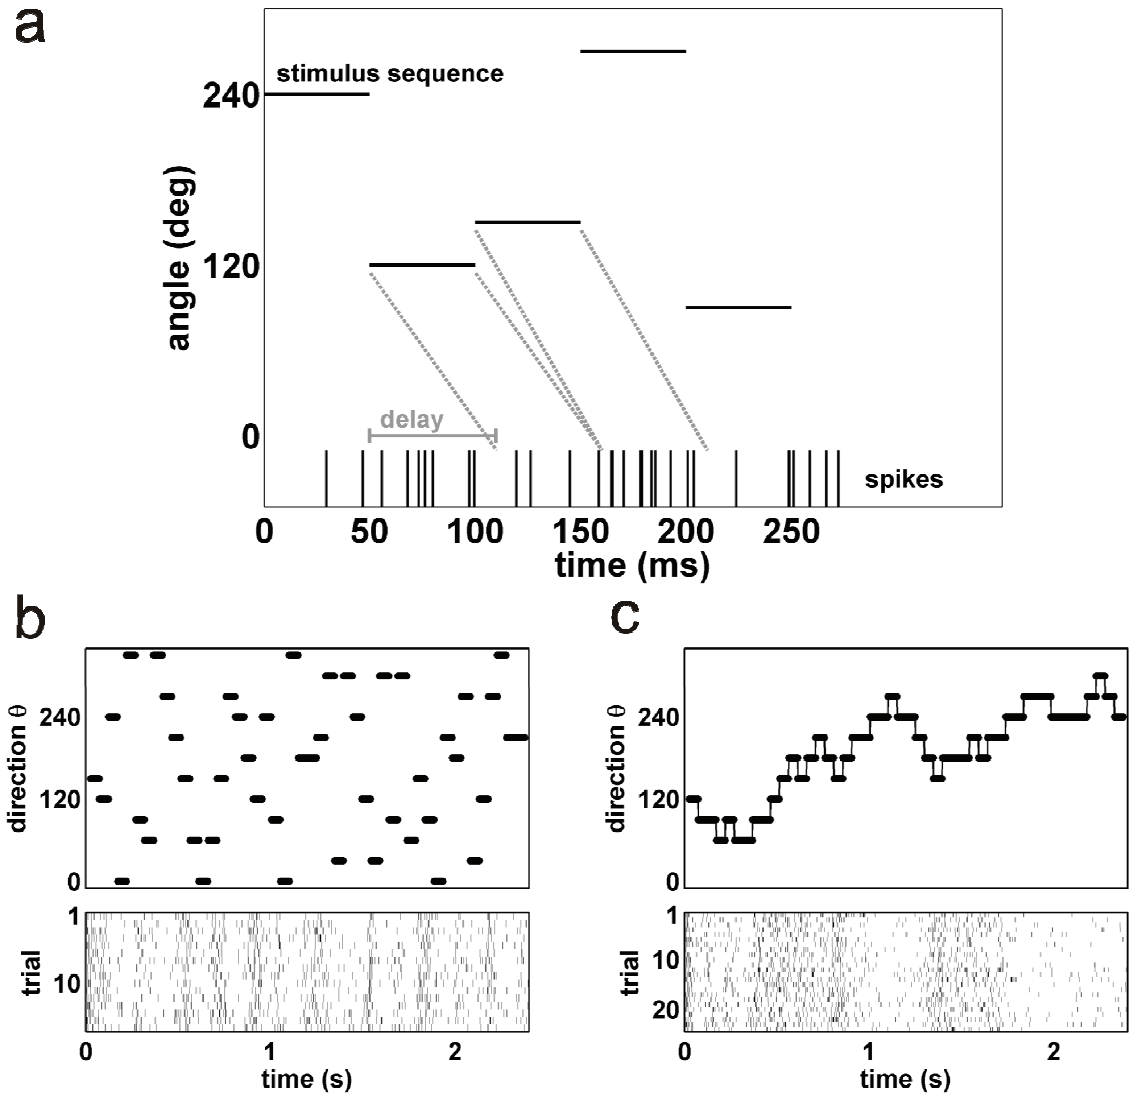

Figure 8: (**a**) Relating spike trains to dynamic stimuli and determining firing rate estimates as responses to a stimulus. The spike train time axis is shifted by the neural delay against the stimulus time axis. A firing rate in response to a stimulus presentation is obtained by dividing the spike count in the respective time bin by the duration.
(**b** and **c**) Examples of stimulus sequences and elicited spike trains for discontinuous (**b**) and continuous (**c**) stimuli.

In order to reduce influences of the preceding and the subsequent stimulus on response magnitude estimation, tuning curves for continuous and discontinuous stimulation (both with 50 ms basic motion interval duration) were estimated by calculating response rates with a variable windowing method. The presentation time of 50 ms was subdivided into five 10 ms windows, and the response magnitude determined for each window. Response magnitude within each basic motion interval was now computed by a weighted average of the responses

in each of the five windows. The procedure described above corresponds to weighting all five windows equally. The weighting concentrated on the central part of each motion interval. The weightings used for additional tuning curve estimations were (0.05, 0.2, 0.5, 0.2, 0.05), (0.01, 0.09. 0.80, 0.09, 0.01), and (0.0, 0.0, 1.0, 0.0, 0.0).

Tuning curves of the neural population were characterized quantitatively by Selectivity Indices and Direction Indices. The *Selectivity Index* (SI) of a neuron was calculated according to the following formula (Vogels and Orban 1994):

$$SI = \frac{\sqrt{\left(\sum_{k=1}^{K} f(\theta_k)\,\sin(\theta_k)\right)^2 + \left(\sum_{k=1}^{K} f(\theta_k)\,\cos(\theta_k)\right)^2}}{\sum_{k=1}^{K} f(\theta_k)},$$

3

where f denotes the tuning curve and the index $k = 1,\ldots,K$ the discrete set of directions. Values for SI are limited to the interval [0;1].

Direction indices (DI) were estimated by $DI = (f_{opt}-f_{null})/(f_{opt}+f_{null})$. f_{opt} is the tuning curve value for the preferred direction θ_{opt}, f_{null} the mean response to the null direction $\theta_{null} = \theta_{opt} - 180$ deg. Tuning curves for Selectivity and Direction Index estimation were Gaussian or von–Mises (Swindale 1998) fits to the original data.

Linear filters were calculated separately for continuous, discontinuous, and fast–discontinuous stimulation. Linear filters can be thought of as extensions of spike triggered averages, which aim to eliminate influences of stimulus autocorrelations on the spike–triggered averages. A linear filter F defines a linear relationship between a stimulus sequence s and a response r:

$$r[t] = r_0 + \sum_{k=1}^{K}\sum_{m=1}^{M} F[\theta_k, m]\; s[\theta_k, t-m+1],$$

4

where m and k are the indices for the time and direction domain, respectively. M is the length of the filter in time. In most cases filter time length was 200 ms, which corresponds to M = 201, as we discretized time in 1 ms bins. M determines the size of the time window in which the stimulus is assumed to influence the response r at some time t. The filter F can be determined by minimizing the time–averaged squared error $\langle (r_{est}[t]-r_{exp}[t])^2\rangle_t$. The estimated rate $r_{est}[t]$ is computed from equation (4). The solution to this problem is given by linear regression theory. By reformulating equation (4) to fit to matrix notation we get $\mathbf{r} = S\,\mathbf{f}$.

$\mathbf{r}=(r_1,\ldots,r_T)^T$ is the rate vector, its entries are the r[t] from equation (4), T denotes the total duration of a stimulus sequence. **f** is the filter F written as a column vector by rearrangement of the indices k and m, such that l=k+(m–1)K is the index of **f**. This reduction of dimensions introduces redundancy into the stimulus matrix $S=(\mathbf{s}_1,\ldots,\mathbf{s}_T)^T$: Each column vector $\mathbf{s}_t$ is now the alignment of the direction probability distributions at the times t, t–1,…,t–M+1 and its dimension is L=MK – see (Chichilnisky 2001) for an extensive description of the index rearrangement procedure. The vector **f** which minimizes the time–averaged squared error is given by

$$\mathbf{f}_{est}= Q_{ss}^{-1}\, Q_{rs,} \qquad \mathbf{5}$$

(Kay 1993), where $Q_{ss} = S^TS$ is the stimulus autocorrelation matrix and Q_{rs} the stimulus–response correlation. Q_{rs} is identical to the spike–triggered–average up to a shift and a scaling factor. After estimating **f** the two–dimensional filter matrix F is obtained by appropriately reverting the index rearrangement.

The inversion of the autocorrelation matrix $Q_{ss}=S^TS$ required in equation (5) is very sensitive to noise sources. Noise has the strongest influence on undersampled directions of the input space (Machens, Wehr et al. 2004), which is given by the rows of the stimulus matrix S. These undersampled directions can be identified by using the singular value decomposition $S=UDV^T$, so that the autocorrelation matrix can be written as $Q_{ss} = VD^2V^T$ (Press, Teukolsky et al. 1992). U and V are orthogonal matrices, while D is a diagonal matrix whose entries are the eigenvalues of S. The eigenvectors of the smallest eigenvalues correspond to the directions of the undersampled input space. The principal components regularization method (PCR) (Hastie, Tibshirani et al. 2001; Theunissen, David et al. 2001) exploits this characteristic and considers only contributions from the eigenvectors of the G largest eigenvalues by setting the inverses of the smallest eigenvalues to zero. G is a regularization parameter. The optimal value G_{opt} was determined by a combination of a cross validation procedure and variance considerations. The optimal value G_{cv} according to cross validation minimized the time averaged response error: $G_{cv} = \underset{G}{\text{argmin}} \mid \mathbf{r}_{est,G} - \mathbf{r}_{exp} \mid^2$. G_{cv} now defined a neighbourhood, in which G_{opt} was determined such that the variance of the respective linear filter estimate peaked at the same time as the spike–triggered average and the variance full width at half maximum was minimal. This combination of criteria yields robustness with respect to the noise sensitivity of the regularization method.

5.2.4 Reconstruction of Movement Directions

Reconstructions were performed within a Bayesian inference framework (Seung and Sompolinsky 1993; Zhang and Sejnowski 1999) that includes Maximum Likelihood as a special case. Our choice of the Bayesian method was motivated by several reasons: First, it allows to include prior information in the estimation process, which is particularly useful for estimating continuous stimulus trajectories, second, it fulfils optimality constraints within a probabilistic framework (statistically efficient estimator), and, finally, notable evidence supports the notion that Bayesian inference may be implemented in the brain (Ernst and Banks 2002; Kohn and Movshon 2004; Kording and Wolpert 2004; Rao 2004).

The response of a population of N neurons to a direction presentation θ_t at time step t is given by $\mathbf{m}=(m_1,\ldots,m_N)$, where m_j denotes the number of action potentials emitted by neuron j in the motion interval Δ (50 ms or 20 ms, depending on the stimulus trajectory). The rate response is the spike count divided by the length of the time window: $\mathbf{r}_t = \mathbf{m}_t / \Delta$.

The conditional probability of observing a rate response $\mathbf{r}_t$ to direction θ is given by

$$P(r_t \mid \theta_t) = \prod_{i=1}^{N} P(r_{i,t} \mid \theta_t) = \prod_{i=1}^{N} \exp(-f_i(\theta_t)\Delta) \frac{(f_i(\theta_t)\Delta)^{m_{i,t}}}{m_{i,t}!}.$$

6

where $f_i(\theta_t)$ denotes the tuning curve of the i–th neuron. Two assumptions were made in this formula: First, we assumed that neurons fire independently of each other (no noise correlations, (Wilke and Eurich 2002). This assumption is necessary because data from several trials was pooled. Second, we assumed Poissonian spiking statistics (Bair, Koch et al. 1994), O'Keefe, 1997 #232}. If $P(\mathbf{r}_t|\theta_t)$ is interpreted as a function of θ_t, it is also the *likelihood* $L(\theta_t,\mathbf{r}_t)$. The Maximum Likelihood (ML) inference chooses as its estimate the direction, that is most likely to elicit the observed response:

$$\theta_{t,ML} = \arg\max_{\theta_t} L(\theta_t).$$

7

If prior information is available about the spatio–temporal stimulus distribution in form of $P(\theta_t)$, the ML method can be extended to incorporate this additional information by Bayes' formula:

$$P(\theta_t \mid r_t) = \frac{P(r_t \mid \theta_t)\,P(\theta_t)}{P(r_t)}.$$

8

The denominator $P(\mathbf{r}_t)$ normalizes $P(\theta_t|\mathbf{r}_t)$ and needs not to be calculated explicitly. The maximum–a–posteriori (MAP) estimate θ_{MAP} is given by the direction, which maximizes the posterior probability $P(\theta_t|\mathbf{r}_t)$:

$$\theta_{t,MAP} = \arg\max_{\theta_t} P(\theta_t \mid r_t).$$

9

Hence, if the prior probability $P(\theta_t)$ is flat, the MAP method is identical to ML inference.

If the stimulus trajectory under consideration is continuous, then the stimulus value at time t is close to the preceding one at time t–1. This implies that the posterior probability distribution $P(\theta_{t-1}|\mathbf{r}_{t-1})$ at time t–1 contains information about the stimulus at time t. Together with the knowledge about the dynamics of the underlying trajectory generation process $P(\theta_t|\theta_{t-1})$ we can write

$$P(\theta_t \mid r_t) \propto L(\theta_t) \int P(\theta_t \mid \theta_{t-1})\,P(\theta_{t-1} \mid r_{t-1})\,d\theta_{t-1}.$$

10

The integral over θ_{t-1} denotes the convolution of the estimate of the stimulus probability $P(\theta_{t-1}|\mathbf{r}_{t-1})$ at time t–1 with the process $P(\theta_t|\theta_{t-1})$, which generates the random walk, and can be used as a time–dependent prior probability distribution. The procedure of estimating $P(\theta_t|\mathbf{r}_t)$ with dynamical priors can be regarded as a predictor–corrector scheme, where the predictor is the extrapolation of information about the stimulus at time t–1 to time t, which is actualized (or: corrected) by the likelihood function $L(\theta_t)$ of the actual time step. It is known as "two–step reconstruction" (Zhang and Sejnowski 1999).

All stimulus sequences were reconstructed with ML. In addition, the continuous trajectories were estimated also with MAP.

Mean absolute reconstruction errors for a stimulus trajectory were determined as the average of the absolute errors over all direction presentations θ_i of the trajectory for a given set of neurons: $E_{abs} = \langle|\theta_i - \theta_{i,est}|\rangle_i$, where θ_i denotes the i–th direction in the stimulus sequence.

The chance level for E_{abs} denotes the mean absolute error that results from an estimator which selects its values by chance. These levels depend on the trajectory. For continuous trajectories chance level is 90 deg. As the discontinuous trajectories are pseudo–randomized sequences

such that every direction is presented the same number of times, chance levels depend on the number of repetitions of the same direction within one trial. The discontinuous trajectories with 50 ms presentation time are of length 48 (4 repetitions per direction) resulting in a chance level of 91.9. The fast–discontinuous trajectory consists of 12 repetitions per direction. In this case the chance level is 90.6. The longer the trajectory, the smaller is the change in probability. For totally randomized sequences chance level is 90 deg as e.g. for the continuous trajectory.

In contrast to the mean absolute error the *mean signed error* is defined by $E_{sign}(\Delta\theta) = \langle \theta_z - \theta_{z,est} \rangle_z$. The index z denotes a subset θ_z of all direction presentations θ_i in a trajectory, where the direction change from the preceding to the actual direction was $\Delta\theta$, independently from the absolute value of θ_z. $E_{sign}(\Delta\theta)$ may take negative values, as the averaging $\langle \cdot \rangle_z$ is not performed on the absolute values of the single estimation errors ($\theta_z - \theta_{z,est}$).

The neural population used for stimulus estimation consisted of pooled responses from different times (Georgopoulos, Schwartz et al. 1986; Recanzone 2000). In addition, repeated presentations of the same trajectory to single neurons were treated as independent neurons and added to our neural pool leading to an effective number N_{eff} of neurons.

5.2.5 Model Simulations

In order to interpret various results obtained from experimental data, we performed simulations with a population of virtual neurons (Figure 9).

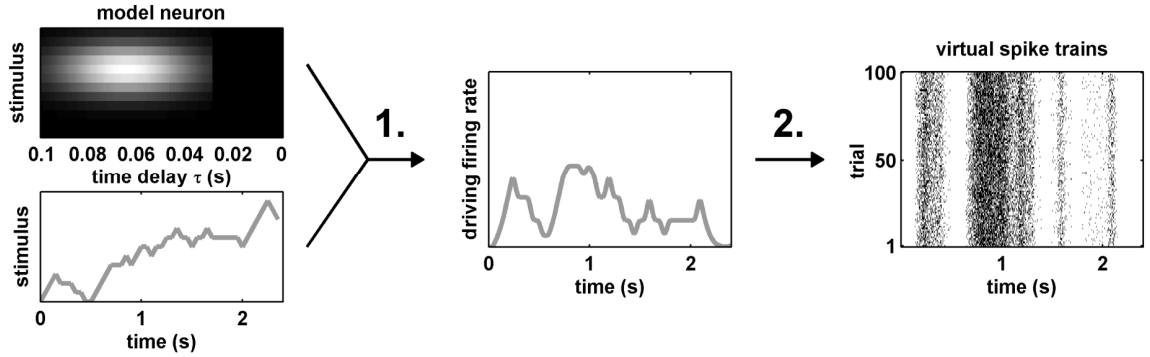

Figure 9: Obtaining virtual spike trains from a linear-nonlinear-Poisson (LNP) model neuron. The upper image on the left shows the direction-temporal receptive field of a virtual neuron. The color-coded spiking sensitivity is Gaussian shaped both in direction and in the time domain. The lower image on the left denotes a continuous trajectory. In step 1 the stimulus is convolved with a filter of the model neuron resulting in a time-dependent firing rate. If negative rate values exist, they are clipped to zero (rectification). In step 2 a Poisson process is applied 100 times to the driving firing rate resulting in 100 virtual spike trains.

100 model neurons were defined by a linear–nonlinear–Poisson (LNP) model (Simoncelli, Paninski et al. 2004). The linear filters representing the directio–temporal receptive field of the neurons were Gaussian both in direction and time domain.

Preferred direction, time delay, standard deviation in direction and time domain, baseline and stimulus–induced firing rates were chosen randomly but within physiologically plausible parameter ranges.

In contrast to the baseline firing rate, the mean stimulus–induced rate represents the informative part of the rate about the stimulus. Parameter ranges were: time delays τ=40–120 ms, standard deviation of the Gaussian for direction selectivity σ_θ=30–150 degrees, standard deviation of the Gaussian in the time domain σ_t=10–70 ms, baseline 3–17 Hz, mean stimulus–induced firing rates 10–30 Hz.

Responses were determined for each of the 100 virtual neurons to all stimulus trajectories used in the experiment. The linear part of the model consists in the convolution of a stimulus with the filter. The resulting mean firing rate was rectified (static nonlinearity) and this rectified rate used for the generation of 100 spike trains in response to each trajectory by repetitions of a Poisson process. The resulting spike trains were used for re–estimating time delays, tuning curves and linear filters for individual virtual neurons, and reconstructing continuous, discontinuous and fast–discontinuous trajectories from artificial populations with the identical procedures used for "real" spike data.

Hence, the model simulations provide us with results we would expect if the "real" population of MT neurons responded in a linear fashion with a static nonlinearity and Poissonian spike creation to the dynamic visual stimuli.

Throughout the results part we will compare the results obtained from experimental data with those from model simulations.

5.3 Results

5.3.1 Spike–Triggered Averages and Latencies

We recorded extracellularly from a total of 72 neurons of which 7 Neurons were discarded due to extremely low mean response rates (maximum of tuning curve <4 Hz). Figure 10a shows the spike–triggered average of a cell obtained from discontinuous stimulation with 50 ms duration of the basic motion interval. A spike of that neuron is most likely (probability p = 0.18) to be preceded by a motion direction of the random dot cloud of 300 deg at time point 100 ms before the action potential. The overlay denotes the variance of the underlying STA. To assess significantly high variance values we calculated mean and standard deviation from the variance at delays between –200 and –50 ms. This time period follows a spike and therefore cannot be relevant for action potential generation. Hence, the resulting mean and

standard deviation represent values for random probability distributions. For τ = 59–134 ms the variance of the STA deviates more than two standard deviations from random value (dashed–dotted lines in Figure 10a) and peaks at t = 97 ms. This value was taken as the latency of the neuron. The latency distribution of the experimental data set is depicted in Figure 10b. The complete bars correspond to the full distribution of the whole population of 65 neurons. Latencies of the neurons ranged from 40 ms to 160 ms with a mean of 104.3±25.3 (standard deviation, SD) ms. The darker bars denote the histogram of a neural subpopulation with prominently selective tuning. Latencies of the subpopulation were much more confined (80 ms to 110 ms, SD = 8.9 ms) and responsiveness was faster (mean = 94.4 ms). The measured latency distribution has larger values on average as found from response onsets in anesthetized macaque monkeys (87±45 (SD) ms in Raiguel et al. 1999 (Raiguel, Xiao et al. 1999). The higher latencies found here might be explained by the fact that we estimated optimal latencies (Mazer, Vinje et al. 2002) and not response onsets. We confirm the tendency to shorter latencies for more strongly tuned cells described in Raiguel et al. 1999 (Raiguel, Xiao et al. 1999).

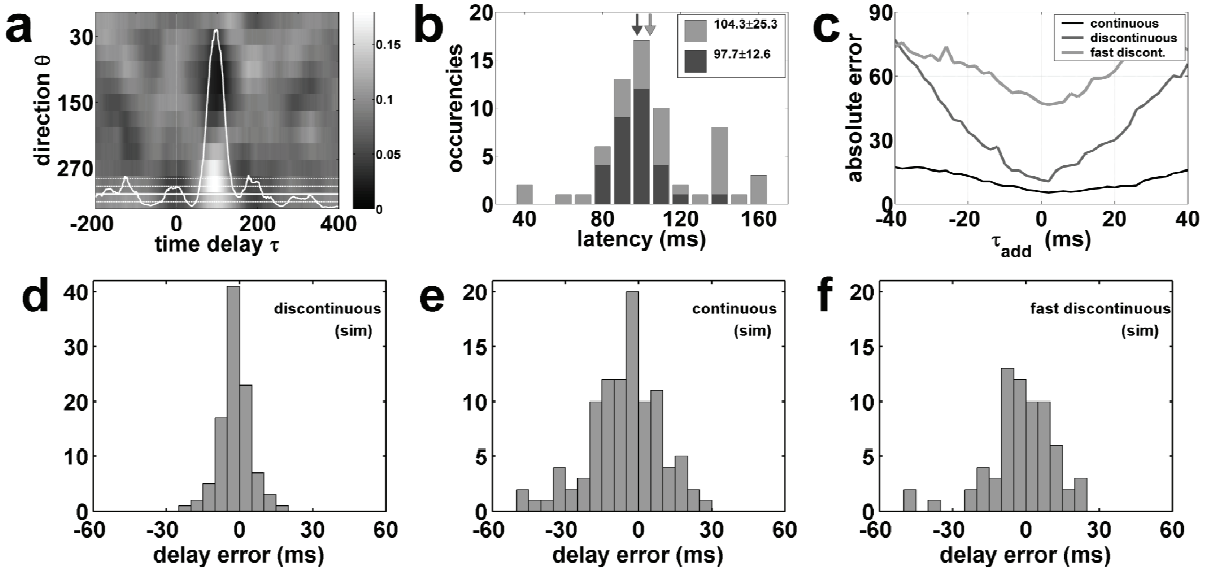

Figure 10: Latency-estimation from spike-triggered averages.
(**a**) Spike-triggered average of a MT neuron. The probability values for a stimulus being presented at a delay τ are gray–coded. The overlaid white curve denotes the variance of the STA (see text). Solid, dashed, and dashed-dotted lines denote the mean, one and two standard deviations of the STA variance, respectively.
(**b**) Histogram of measured latencies. The gray-shaded bars represent the complete population, while the black bars denote a subpopulation with the most prominent tuning. Average latencies are (104.3±25.3) ms and (94.4±8.9) ms, respectively.
(**c**) Errors of trajectory reconstruction with a set of neurons, whose latencies are globally shifted by some value τ_{add}. Reconstructions are best for τ_{add} =0 ms providing evidence that no systematic error was made in latency estimation.
(**d-f**) Differences between model neuron delays and delay estimates from discontinuous (**d**), continuous (**e**), and fast-discontinuous (**f**) stimuli. Discontinuous stimulation provided the most reliable estimates with 89% of estimates within 10 ms of the original delays. Average values with standard deviations of single values were -2±6.1 ms, -5.7±14.6 ms, and -2±13.7 ms, respectively.

We checked whether our method of latency estimation produced systematic errors by introducing additional latency shifts τ_{add}, constant for the entire neural population, and performing maximum–likelihood trajectory reconstructions for different values of τ_{add}. The additional delay was applied both to tuning curve and response firing rate estimation, which have to be determined separately for maximum likelihood reconstruction (see Methods). τ_{add} took values in the interval –40 to 40 ms in 2 ms steps. The result is shown in Figure 10c: For all three stimulation paradigms reconstructions perform best for $\tau_{add} = 0$ ms providing strong evidence that the latency estimation is not subjected to a bias.

As latency estimation from continuous or fast–discontinuous stimulation did not always match the results from discontinuous stimulation, we applied the latency estimation method to model neurons. Spike trains were generated artificially from a linear–nonlinear–Poisson model (see Methods) in response to the three stimulation paradigms. The data evaluation was performed in the same way as from experimental data. Here the true latencies of the virtual neurons are known and can be compared with estimates. Histograms of the differences between real and estimated latencies are plotted for a population of 100 virtual neurons in Figure 10 d–f for discontinuous (d), continuous (e), and fast–discontinuous (f) stimulation. The average differences are –2.0±6.1 (SD) ms, –5.7±14.6 ms, and –2.0±13.7 ms, respectively. The discontinuous stimulation provides the most accurate estimates. 89% of the estimates differed less than 10 ms from the original delay, documenting the reliability of the method. STAs from continuous stimulation are subject to strong blurring due to stimulus autocorrelations. This influence is reflected by the broad distribution of latency estimation errors (Figure 10e). Finally, spike–triggered averages from fast–discontinuous stimulation are often too noisy to reliably determining a single variance maximum (32 virtual neurons were discarded and not shown in the histogram in Figure 10f, because of badly conditioned variance plots).

5.3.2 Tuning Curves

Tuning curves are a frequently used measure for quantifying neural response selectivity to a stimulus feature. Usually, tuning curves are determined by measuring the response firing rate of a neuron to prolonged presentation of a stimulus value (Hubel and Wiesel 1959). In such cases, due to presentation lengths of usually one second or longer, the responses of a neuron can be attributed unambiguously to the single stimulus values. Here, however, the stimuli are dynamic and values change at rates of 20 or 50 Hz. Therefore, tuning curves must be

determined by averaging firing rate estimates from short stimulus presentations, where stimuli and spikes were related to each other as shown in Figure 8a.

Figure 11a shows the tuning curves of two neurons resulting from continuous (black lines) and discontinuous (grey lines) stimulation. For one neuron (solid lines, preferred direction approximately 120 deg) the tuning curve is much more modulated for continuous stimulation, while for the other neuron (dashed–dotted lines, preferred direction at around 300 deg) almost no difference in tuning is found. The preferred direction of each neuron is identical for both stimulation conditions.

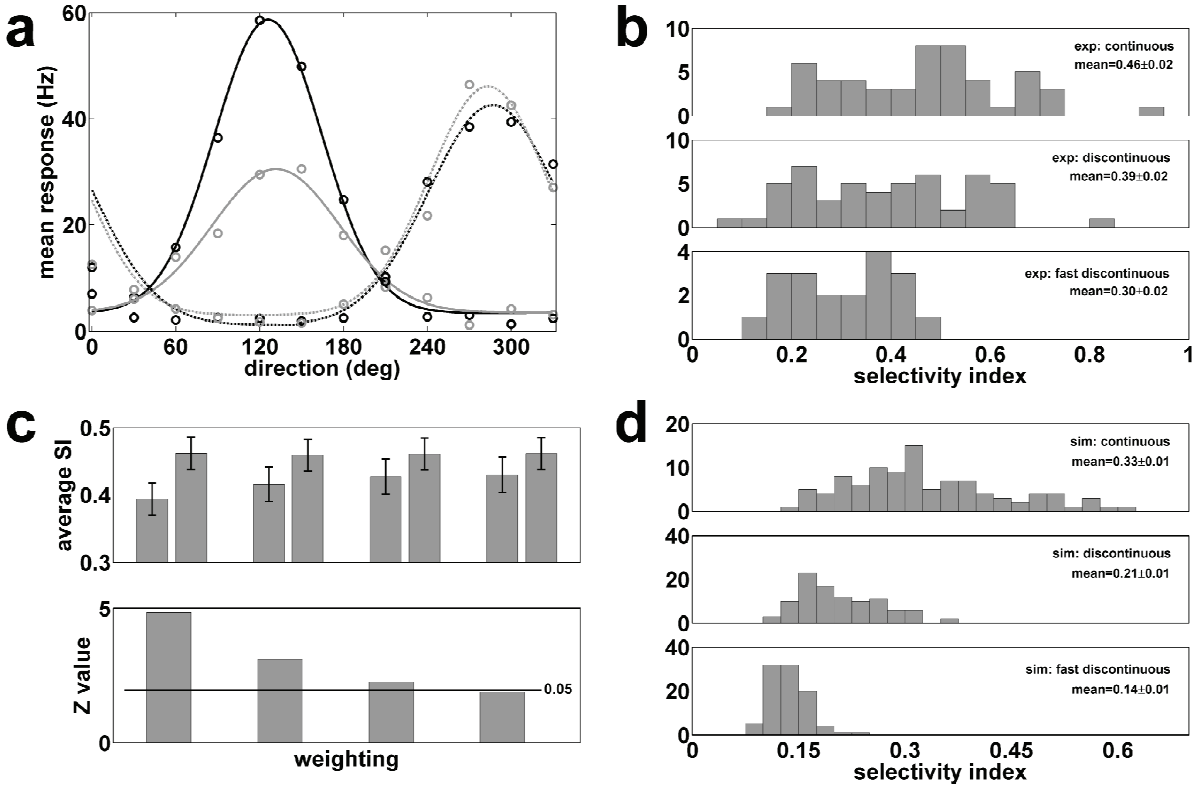

Figure 11: (**a**) Tuning curves of two neurons calculated from discontinuous (grey) and continuous (black) stimulus statistics.
(**b**) Selectivity index distribution for experimental data from continuous, discontinuous, and fast discontinuous stimulation (from top to bottom).
(**c**) Comparing the average selectivity indices from continuous and discontinuous stimulation for tuning curve sets with varying emphasis on the most centralized presentation time window. (**top**): Left bar of each group indicates SI for discontinuous stimulation, right bar is continuous stimulation (always 0.46±0.02). Weighting progressively shifts to the center time window from left group to right group. Differences in SI for both groups are decreasing with shift to the center. (**bottom**): Significance test for difference. Z values for a paired two-sided Wilcoxon test drop to values smaller than 1.95, which is the Z value for p = 0.05.
(**d**) Selectivity index distribution for model neuron from continuous, discontinuous, and fast discontinuous stimulation (from top to bottom).

Does the difference in tuning curve estimation as found for the first neuron reflect a response variability which is induced by the different temporal statistics of the stimuli, and maybe an adaptation to temporal statistics?

We addressed this question by comparing the selectivity index (SI, see methods) distribution of a neural population for continuous, discontinuous, and fast–discontinuous stimulation (Figure 11b). Average SI values with SD of the mean were 0.46±0.02, 0.39±0.02, and 0.30±0.02, respectively. The neural population was identical for the two 50 ms stimuli (N=51 neurons). 21 neurons from the population of 72 recorded neurons were excluded from this evaluation due to low firing rate values and failure of fitting a Gaussian or a von–Mises function to the tuning curve from experimental data. The population for the fast–discontinuous stimulation was a subpopulation of the one used for the 50 ms stimuli and consisted of 19 neurons.

A comparison of the average SIs across stimulus types reveals an increase of SIs with increasing width of stimulus autocorrelation functions (see Figure 7c). This comparison also holds if average SI values for continuous and discontinuous stimulation are calculated from the subpopulation of neurons used during fast–discontinuous stimulation. The respective mean SI values for the subpopulation were 0.44±0.04 and 0.37±0.04 (histograms not shown).

A paired two–sided Wilcoxon test comparing the SI distributions from continuous and discontinuous stimulation was highly significant ($p<0.00001$). The comparison of average direction indices for the three stimulus types exhibited an identical dependency on the stimulus statistics. The values were: 0.66±0.03, 0.57±0.03, and 0.43±0.05 (Stemmann, Freiwald et al. 2005). Therefore, the dependency of SI and DI values on the stimulus statistics suggests that the neural selectivity depends on stimulus dynamics.

However, these results do not fully answer the question, whether neurons have intrinsically different direction selectivity for stimuli with different temporal autocorrelation functions: Stimulus autocorrelations may influence tuning curve estimation. For dynamic stimuli, responses may be influenced by more than a single stimulus interval. In order to reduce or possibly eliminate the influence of temporally nearby stimuli we calculated tuning curves of the stimuli with 50 ms basic motion interval by dividing the 50 ms into five 10 ms bins.

We introduced different weighting schemes, which defined the contributions of the five time bins to the response rates. These weights were chosen such that they gave different emphasis on the central time bins.

In Figure 11c (upper panel) the mean selectivity indices for discontinuous (left bars) and continuous (right bars) stimuli are shown for different weightings. Error bars correspond to

the SD of the mean. As shown in the lower panel, the leftmost bars represent equal weighting w=(0.2, 0.2, 0.2, 0.2, 0.2). The other weightings were w=(0.05, 0.20, 0.50, 0.20, 0.05), w=(0.01, 0.09, 0.80, 0.09, 0.01), and w=(0.0, 0.0, 1.0, 0.0, 0.0). The last weighting corresponds to completely ignoring spikes which occur in the time windows 0–20 ms and 40–50 ms of every motion interval. Mean SI values for continuous stimulation were 0.46±0.02 (SD of the mean) for all weightings, while SI values for discontinuous stimulation were 0.39±0.02, 0.42±0.03, 0.43±0.03, and 0.43±0.03. This evaluation shows that the average selectivity indices from continuous and discontinuous stimulation approach each other, as we increase the influence of the time period of a single stimulus presentation which is potentially less influenced by other stimuli. A significance test confirms this result (Figure 11c, bottom panel). The Z values for a paired both–sided Wilcoxon test decrease: Z values were 4.84, 3.10, 2.25, and 1.87 (from left to right). The latter Z value is smaller than 1.95, which is the Z value for p=0.05. Hence, the null hypothesis of a single underlying distribution for selectivity indices from both continuous and discontinuous stimulation cannot be rejected for tuning curve sets from the most centralized weighting.

In addition, model simulations confirm the hypothesis that directional selectivity remained robust across different stimulus statistics. Figure 11d shows the selectivity index histograms for 100 neurons, which were modelled as linear filters with a static nonlinearity. The mean average SI values are 0.33±0.01, 0.21±0.01, and 0.14±0.01. Though the model neurons are modelled as linear filters with a static nonlinearity, the mean SI values exhibit the same trend as found in experimental data: The statistics with the larger autocorrelation function has the largest SI. Both analyses suggest that differences in tuning curve estimation from different temporal stimulus statistics are purely stimulus–induced and do not necessarily reflect neural nonlinearities or adaptation to stimulus statistics.

5.3.3 Linear Filter Estimates from Continuous and Discontinuous Stimulation

In the previous paragraph we showed that differences in neural tuning for different stimulus statistics could be fully explained by feed–forward stimulus influence without neural adaptation. Here we are interested in adaptational effects in the time domain. For that purpose we compare linear filters (see Methods 5.2.3) of neurons estimated from continuous and discontinuous stimulation. Filters are defined both in direction and time. We quantify the temporal properties of the filters by the full width at half maximum (FWHM) of the linear filter variance $\sigma^2(\tau)$, which is defined analogously to its counterpart calculated from the STA (see Methods 5.2.3).

As linear filter estimates from fast–discontinuous stimulation converged poorly, only two stimulus conditions (continuous and discontinuous) were compared. Figure 12 a and b show the linear filters Q_{rs} of an exemplary neuron to continuous (a) and discontinuous (b) stimulation, if we assume that the stimulus autocorrelation matrix is a unitary matrix. The difference between the so–defined linear filter Q_{rs} and the STA is, that the filter gives a direct relationship between stimulus and response, while the entries of the spike–triggered average denote the probabilities. Nevertheless Q_{rs} is a rescaled and shifted version of the STA (Dayan and Abbott 2001). The FWHMs of the filters shown in Figure 12a and b are >200 ms and 44 ms, respectively. The obvious influence of stimulus autocorrelation makes it impossible to directly compare the filters without deconvolution. In Figure 12c and d filters are shown after deconvolution with the true stimulus autocorrelation matrix (see Methods 5.2.3). The temporal widths now are 57 ms in (c) and 40 ms in (d). The influence of deconvolution is very strong for the continuous stimulation, resulting in a strong decrease in FWHM. For discontinuous stimulation the decrease is much smaller (as the decorrelation only needs to account for the 50 ms basic motion interval). Even after deconvolution we find a residual 17 ms difference of FWHMs. Is this difference due to an adaptation of neural properties to slowly and smoothly changing stimuli?

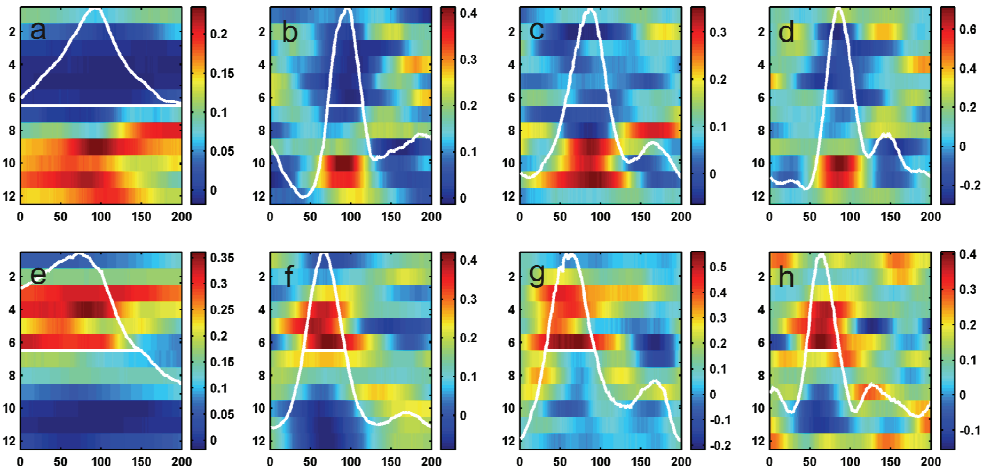

Figure 12: Linear filters of a neuron (top row). (**a**) shows the linear filter from continuous stimulation if the influence of stimulus autocorrelation is ignored. The effect of stimulus correlations is clearly visible and results in a FWHM larger than 200 ms. (**b**) shows the filter of the same neuron stimulated by discontinuous trajectories (FWHM = 44 ms). (**c**) and (**d**) show the filters after deconvolution and choice of an appropriate regularization parameter for denoising. FWHMs are 57 ms (40 ms), respectively. The y-axis for the variance is not shown as units are arbitrary and only the baseline var=0 is important. The baseline is always given by the bottom of each image.
Linear filters for a model neuron (bottom row). The original filter of the LNP model (not shown) had temporal width 34 ms and preferred direction was 150 deg. Temporal widths for the conditions continuous and discontinuous without autocorrelation (**e**) and (**f**), and continuous and discontinuous with autocorrelation (**g**) and (**h**), were >150 ms, 53 ms, 62 ms, and 42 ms, respectively. Deconvolution helps in reproducing the original filter shape for both stimulus conditions, while it fails to reproduce the preferred direction for the filter from continuous stimulation. Also, some influence from stimulus correlations is still present after deconvolution.

For resolving this question we compared the filters resulting from virtually stimulating a LNP model neuron with the two stimulation paradigms. This comparison indicates what we can expect from the filter estimation method if a neuron responds according to the same model to the stimuli. The original filter was defined by gaussian functions both in direction and in space domain. The delay was 70 ms, the preferred direction 150 deg. Standard deviations were σ_t=20 ms and σ_θ=60 deg for the two domains, respectively, resulting in a temporal FWHM of 34 ms. The filter was rescaled to produce 5 Hz background firing rate and an average stimulus–induced firing rate of 20 Hz, when applied to the stimuli.

Figure 12e and f show the re–estimated linear filters without deconvolution. The shapes and the temporal widths of >150 ms in (e) (continuous) and 53 ms in (f) (discontinuous) are qualitatively comparable to those of the "real" neuron. After deconvolution the widths are 62 ms and 42 ms, respectively. Two points can be made: First, the deconvolution procedure improves the estimation result for the original filter shape for discontinuous stimulation, even when it fails to reproduce the underlying filter width of 34 ms. Second, deconvolution also strongly improves the filter estimation from continuous stimulation, but is not capable of removing all influences of stimulus autocorrelation. A residual difference in temporal width of 20 ms separates the two stimulation conditions. These simulation results suggest that apparent differences of neural temporal integration time are due to stimulus and not neural properties. The simplest assumption about the neurons is that they do not adapt to stimulus statistics. Though neural filter for continuous and discontinuous stimulation are shaped quite differently in Figure 12 a–d, the hypothesis of a common underlying filter cannot be discarded for the exemplary neuron.

Comparisons between experimental and model data can only be made qualitatively as the "true" underlying filter or neural firing model of the experimental neuron is not known (if it exists). Therefore, simulation results could be modified by changing the nonlinearity in the model (which was the simplest possible: Rectification of the firing rate) or the filter width. In addition, filter estimation with deconvolution is very sensitive to noise. The critical point in the noise sensitivity is the inversion of the autocorrelation matrix, which is essential for the deconvolution procedure. Noise effects can be reduced by modifying the inverse of the autocorrelation matrix appropriately such that the influence of undersampled directions of the input space is suppressed. This was done by principal component regression, which ignores those directions, but introduces an additional parameter (the regression parameter G), which denotes the number of input directions which are not suppressed. The optimal parameter value G_{opt} was determined by minimizing the quadratic distance between the original firing rate and

the rate estimated with the filter with respect to the regularization parameter. This method robustly yielded a restricted range of values for G, but did not uniquely characterize the optimal value for G, which would be necessary for quantitative comparisons. The difficulty of quantitatively evaluating parameter values from linear filter estimation with regularization was also encountered in the estimation of auditory spectrotemporal receptive fields (Machens, Wehr et al. 2004).

5.3.4 Stimulus Reconstruction

In the previous paragraphs we focussed on single neuron properties resulting from different stimulus statistics. Now we turn to properties of neural populations and analyze the ability of stimulus trajectory reconstruction for different temporal statistics.

Stimulus value estimations were done with classical estimation theory using maximum likelihood (ML) inference and Bayesian estimation employing a dynamic prior, resulting in maximum–a–posteriori (MAP) estimates. ML and Bayesian inference, as used for stimulus estimation from neural responses (see Methods eq.6), is composed of two parts: The first part consists of the "knowledge" about neural response properties, which is deduced from a part of the data set. Here, we use the tuning curves as our knowledge about response selectivity. The second part, the test set, is the neural responses to a stimulus trajectory that is going to be reconstructed. For a subset of neurons, responses are available for all three stimulation conditions. This setup allows us to study reconstructions where tuning curves and response result from different stimulus statistics.

Figure 13a shows the MAP reconstruction of a continuous trajectory. The original sequence (dotted black line) consisted of 96 single stimulus presentations of 50 ms duration each, and trajectory length was 4.8 seconds. Reconstructions were performed with tuning curves from continuous (dark grey line, labelled *c50*) and discontinuous (grey line, labelled *d50*) stimulation. The mean absolute error averaged over all direction presentations was 2.8 deg and 8.4 deg, respectively. Responses of 34 physical neurons were used; by pooling the repeated presentations (trials) of the trajectory to the individual neurons we reached a population size of 390 effective neurons (see Methods). As shown before, the different stimulation conditions can result in different tuning curve heights (e.g. Figure 11a). This is the reason why the reconstruction with the set of tuning curves performs better. This effect is not due to the higher average selectivity index, because the reverse also holds: A discontinuous trajectory is reconstructed better with its own tuning curve set (data not shown). However, the reconstruction of the trajectory in Figure 13a with tuning curves from discontinuous

stimulation still works very reasonably. This result shows that the reconstruction is not confined to the class of stimuli, with which the neural statistics was determined.

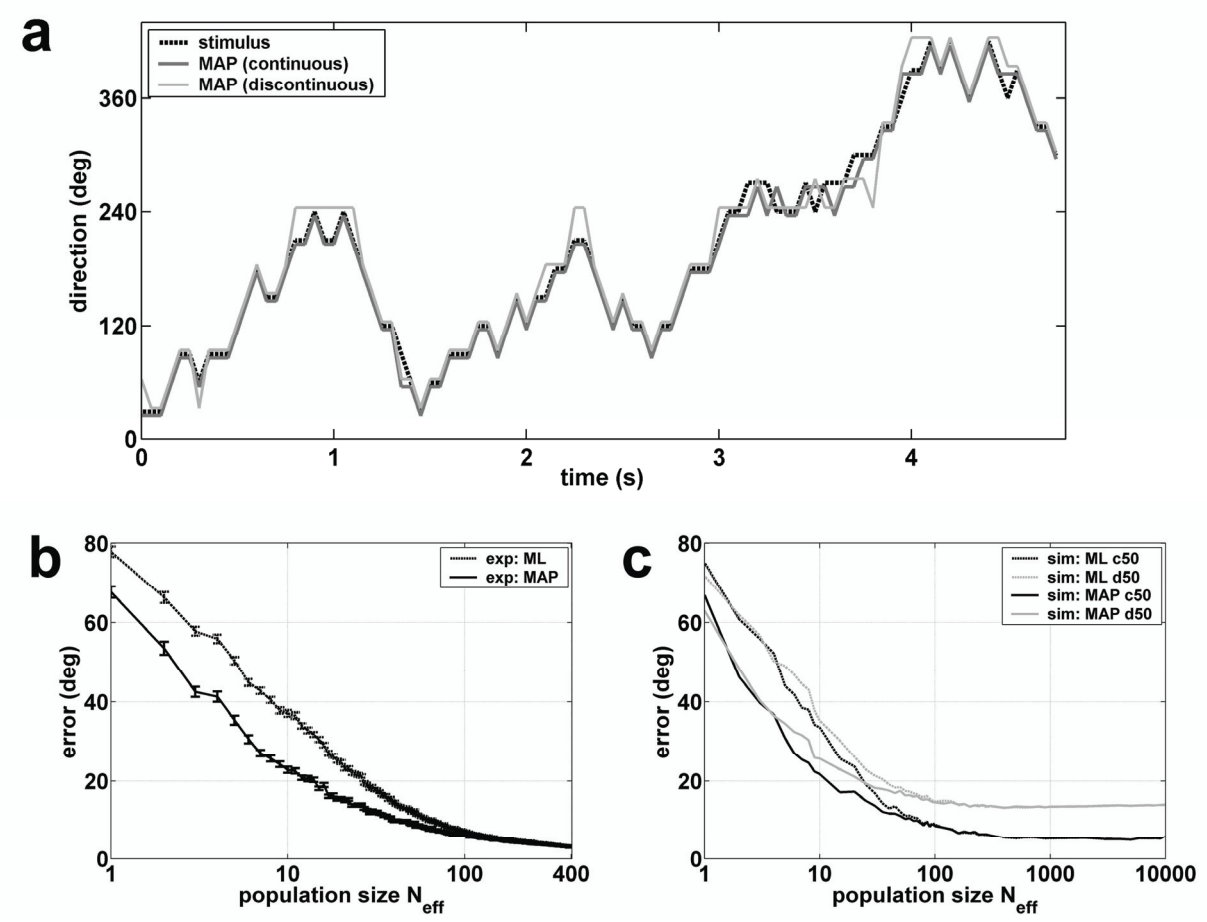

Figure 13: (**a**) Reconstruction of a continuous direction trajectory with a population of 390 effective neurons. The original stimulus is represented by the black dotted line; MAP estimates with tuning sets from continuous (discontinuous) stimulation are shown in dark grey (grey). Mean errors are 2.8 and 8.4 deg, respectively.
(**b**) Comparison of MAP and ML method for the reconstruction of the trajectory in (**a**) in dependency of the population size N_{eff}. For N_{eff}>100 the methods perform almost equally.
(**c**) A similar plot as (**b**) for a virtual population of LNP-model neurons. The legend denotes the reconstruction method and the tuning curve set used for reconstruction (c50: continuous, d50: discontinuous).

5.3.5 Comparison of MAP and ML Reconstructions

What advantage would a neural population have, if (in some way) it knew the statistics of the stimulus? In the case of a discontinuous stimulus sequence the population should rely on maximum likelihood estimation, since no information about a stimulus value can be inferred from observations of preceding stimulus states. However, if the stimulus is continuous knowledge about preceding stimulus values is informative. The population could use a modified form of the stimulus probability distribution of the preceding time–step as prior knowledge about the actual stimulus value. This information can be mixed with the new information about the actual stimulus value rendering the estimate more reliable.

Figure 13b shows how the incorporation of a prior affects reconstruction accuracy. The trajectory from Figure 13a is reconstructed with MAP and ML and with the same population of 390 effective neurons. The prior used for MAP reconstruction is tailored for the continuous stimulus: It is a convolution of the stimulus probability distribution estimate of the former time step with the probability distribution which underlies the generation of the random walk sequence (see Methods, eq.10). For each N_{eff} we randomly selected 100 subsets of N_{eff} neurons from the total of 390 effective neurons in a bootstrap procedure (Shao and Tu 1995). The average of the mean absolute errors resulting from the reconstructions of each of the 100 subsets is plotted as a function of N_{eff}. Chance level is 90 deg and is determined by taking random values as estimates. Error bars denote the standard deviation of the mean. For larger values of N_{eff} the methods perform almost equally, the errors are 6.1 deg and 7.1 deg, respectively, for N_{eff}=100. The importance of the prior is strongest for small neuron populations, and vanishes for population sizes of 100 and more neurons.

Figure 13c compares the two reconstruction methods for the virtual population of 100 LNP–model neurons (100 trials per trajectory). The continuous trajectory of Figure 13a is reconstructed with the tuning curve sets from continuous and discontinuous stimulation.

The simulation shows that for small values of N_{eff} (N_{eff}<10) the use of the prior compensates the use of tuning curves from different statistics (compare MAP *d50* and ML *c50* in Figure 13c). For large effective neuron numbers (N_{eff}>1000) the reconstruction error does not tend to zero, but saturates. The saturation value depends on the statistics used to determine the set of tuning curves. The training sets from discontinuous and fast–discontinuous stimulation are not able to approach the estimation performance from the genuine tuning set, which has the largest SI. Interestingly, the saturation value does not directly depend on the average SI value, as the error for the two tuning sets from discontinuous and fast–discontinuous stimulation converge, though their mean SI values differ considerably. What factors contribute most to the limits of reconstruction accuracy? We will address this question in the next section.

5.3.6 Temporal Stimulus Statistics and Reconstruction Accuracy

Here we analyze how well trajectories from different stimulus statistics can be reconstructed from neuronal responses by using tuning sets from the same statistics. We used a population with N_{eff}=195 effective neurons and compared results from maximum likelihood estimation. Figure 14a shows estimation errors for the three stimulus statistics as a function of N_{eff}. Estimation is most accurate for continuous stimulation, followed by discontinuous and worst for fast–discontinuous stimulation. Reconstruction errors for the full population are 4.7 deg, 23.8 deg, and 50.1 deg, respectively, where chance level is given by 90 deg. The same

reconstruction procedure applied to the virtual population of LNP neurons with N_{eff}=10000 reveals that reconstruction saturates (data not shown). Why is the reconstruction error for fast–discontinuous stimulation more than two times larger than for discontinuous stimulation, and why does it saturate for large numbers of neurons? A straightforward assumption is that a 20 ms spike train does not contain sufficient information to allow for a more precise reconstruction. In this case, stimulus reconstruction based on 20 ms spike train segments should saturate irrespective of the stimulus sequence used. We therefore reconstructed the two stimulus trajectories with 50 ms basic motion interval using responses from shorter time windows. Figure 14b shows the reconstruction error as a function of the time window size used for reconstruction. In this case, responses to a stimulus value were estimated by using the spikes which occur in a time window of length τ_0, where the window starts at the beginning of each stimulus presentation shifted by the delay of the neuron (see the inlay of Figure 14b). Spikes generated in the period τ_0–50 ms of each stimulus presentation are ignored and do not contribute to response estimates.

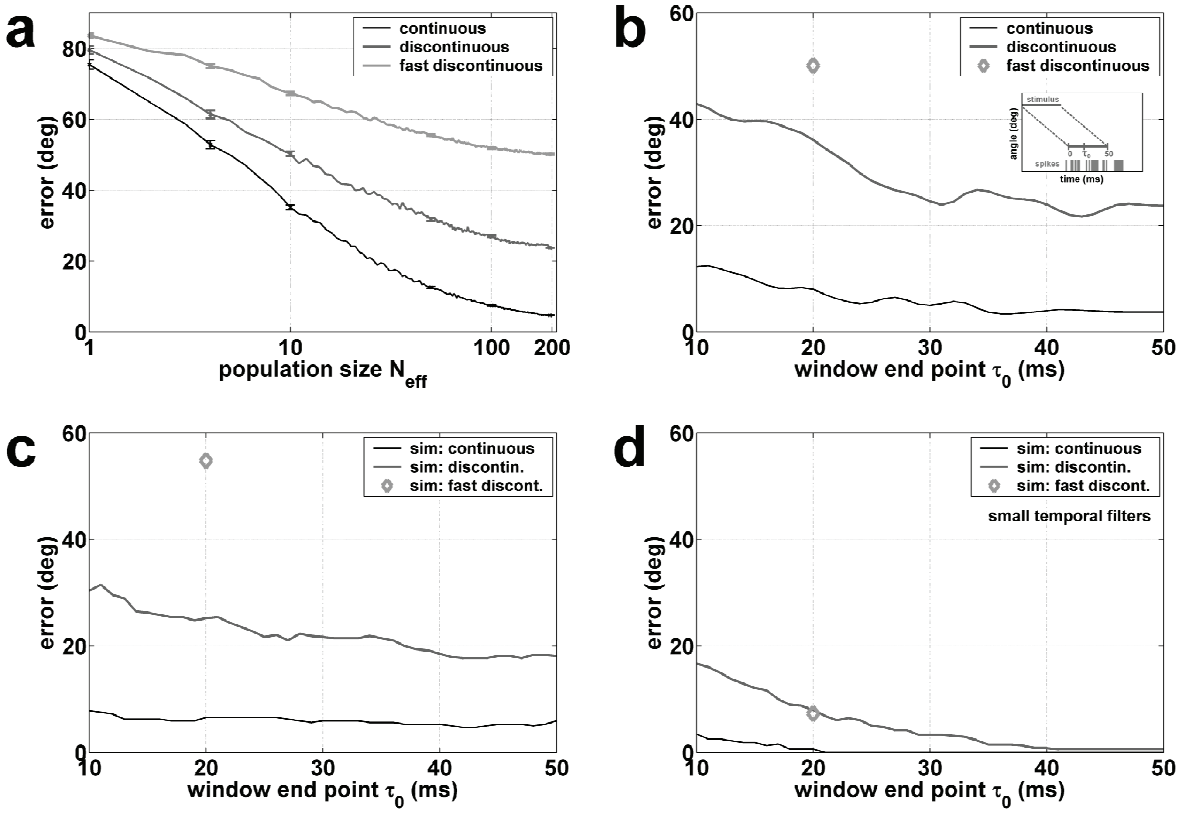

Figure 14: (**a**) Maximum likelihood reconstruction errors for continuous (black), discontinuous (dark grey), and fast-discontinuous (grey) trajectories as a function of the effective population size.
(**b**)-(**d**): Reconstruction of the 50 ms trajectories as a function of the time window end point τ_0. Only the first τ_0 ms of the 50 ms are used for rate estimation, as shown in the inset of (**b**). The figures represent experimental data (**b**), the reference LNP-neuron population (**c**), and a LNP-neuron population with small temporal filter widths (**d**). Reconstruction errors from fast-discontinuous stimulation are shown for comparison (grey dots). Curves in (**b-d**) are slighty smoothed (Gaussian filter, τ=1 ms).

Figure 14b shows that the reconstruction from responses taken from the first 10 ms of the stimulus presentation suffices to yield an estimate error of 12.3 deg in the case of continuous stimulation. In addition, reconstruction of the discontinuous trajectory from the first 20 ms of each stimulus presentation results in an error of 36.1 deg. Thus, shortness of spike train segments can only partly account for the limits of the representation of rapidly changing stimuli. Another possible reason could be the low average selectivity index of the tuning set from fast–discontinuous stimulation when used for reconstruction. However, the reconstruction of the fast–discontinuous stimulus is best with its generic tuning set. Reconstructions with tuning curves from continuous (discontinuous) stimulation yielded an error of 53.7 deg (54.2 deg).

In order to find the reason for the saturation effect in stimulus estimation we performed model simulations with a population of 100 LNP neurons and calculated stimulus estimates in the same way as for the experimental data. Results from the virtual population are shown in Figure 14c, which shares several characteristics with reconstructions from actual data (Figure 14b); the difference between reconstruction from fast–discontinuous stimulation and from the first 20 ms of the discontinuous stimulation is even more pronounced (errors are 54.8 deg and 25.2 deg, respectively). This similarity between experimental and modelling data suggests that we do not need to think of adaptive effects for explaining the different estimation errors (as we did before for tuning curves and linear filters), but may test the influence of parameters on neural responses.

The modelling approach allows to selectively changing parameter values of the virtual population. As the only difference between the discontinuous and the fast–discontinuous stimulus is one of time scale, we assumed that the reason for the discrepancy in reconstruction errors in Figure 14b and c is due to the length of the neural time integration. We therefore defined an additional population of LNP–neurons with shorter temporal filter widths σ_t. These were now drawn uniformly from the range 3–40 ms, while the former ones were taken from a gaussian with mean 40 ms and standard deviation 3 ms. The average time width now was 20.4 ms in comparison to the 39.8 ms in the reference population. With the shorter time integration of the linear filter the situation changes. Figure 14d shows that reconstruction of the fast–discontinuous stimulation is as good as from the first 20 ms of the discontinuous stimulation, the errors are 9.4 deg, and 9.2 deg. Furthermore, the reconstruction error of the continuous stimulus does not saturate with increasing neural pool size, but rather tends to zero. Therefore, the width of the filter in the time domain relative to stimulus update rate is the relevant

parameter which causes saturation of estimations. The temporal filter widths of the population of MT neurons are too large to account for a stimulus which switches direction at 50 Hz. By comparison of Figure 14b and Figure 14c we can assume that the average temporal filter width is somewhat smaller than the average filter width σ_t=39.8 ms of the LNP model population.

5.3.7 Coding Strategies for Dynamic Stimulus Trajectories

How do the individual parameters in our model influence reconstruction quality of dynamic trajectories? We tested the following parameters: Background firing rate, mean stimulus–induced rate, tuning width, and temporal filter width.

The reference model population described in the Methods section is the starting point of our variations. Each data point in Figure 15 is produced by a distinct neural population, which differed from the reference by the value of a single parameter only. The value of this parameter, averaged over the respective population, is the abscissa value. Each population produced 100 responses to each trajectory. Reconstructions were performed by using either 1 or 10 trials per neuron, resulting in effective population sizes of N_{eff}=100 and N_{eff}=1000. In order to avoid influences resulting from the selection of trials we extracted 10 trial sets in a bootstrap manner. The average decoding error from these 10 sets represents the ordinate of a data point. The error bars denote the standard deviations of the mean.

Figure 15 shows the influence of background firing rate (a), stimulus–induced firing rate (b), tuning width (c), and filter time width (d) on errors obtained from maximum likelihood estimations with N_{eff}=100 (solid lines) and N_{eff}=1000 (dashed lines). Note that N_{eff}=1000 is already within the saturation regime of Figure 13c and Figure 14. Baseline firing rate does not affect reconstruction accuracy (Figure 15a). In Figure 15b stimulus–induced firing rate only influences reconstruction accuracy before the population reaches saturation level. For N_{eff}=1000 (dotted lines) mean stimulus–induced activity is not a relevant parameter. Figure 15c shows that tuning width has an impact on estimation accuracy: Narrow tuning curves perform better than larger ones. The advantage gained from sharp tuning is demonstrated best by discontinuous stimulus reconstructions with the smaller populations (solid dark grey line). However, stimulation type–specific saturation levels seem to exist: In the case of continuous stimulation (dotted black line) narrow tuning width yields almost no advantage. Also, reconstruction errors during fast–discontinuous stimulation remain high even for the narrowest tuning widths. Again, as in the previous paragraph, the filter time width is the relevant parameter which sets lower bounds for estimation accuracy (Figure 15d). The error for all three stimulus types tends to zero for short integration windows, with the strongest

error decrease for fast–discontinuous stimulation from 56.9 deg (σ_t = 38.3 ms) to a residual error of 0.5 deg (σ_t = 2.8 ms).

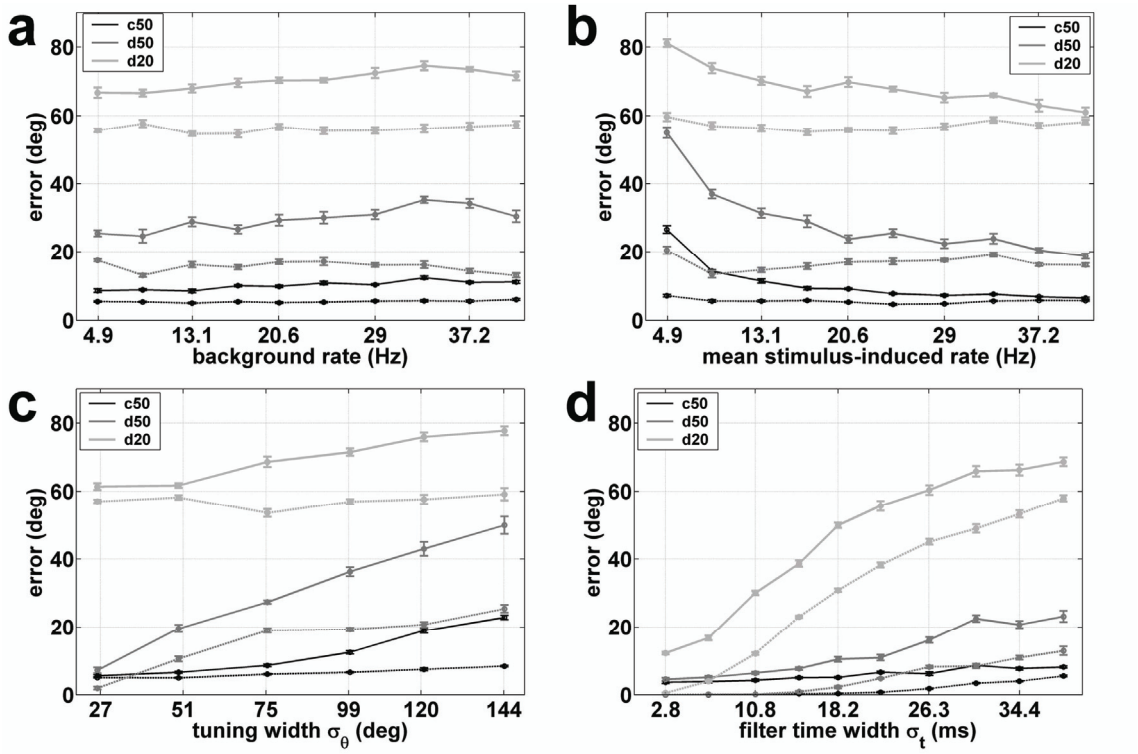

Figure 15: Influence of LNP model parameters on decoding accuracy of a virtual population. Starting from parameter ranges of a reference population we varied the ranges of one of the following four parameters: Background rate (**a**), mean stimulus-induced rate (**b**), tuning width (**c**), and temporal width (**d**). Each data point is computed from a population of 100 neurons with the abscissas denoting the average of the parameter under consideration. Solid (dashed) lines depict estimations using N_{eff}=100 (N_{eff}=1000) effective neurons, corresponding to 1 (10) trial(s) per neuron. N_{eff}=1000 represents saturation level for the reference population (see the blue line in Fig.7c). The baseline firing rate does not influence reconstruction. Mean stimulus-induced rate has some influence on representation accuracy only for population sizes below saturation level (N_{eff}=100, solid lines in b). Tuning width and time width have the strongest impact on reconstruction errors. The latter is mostly responsible for saturation, as can be seen from the dashed lines in (**d**).

5.3.8 Time Integration of Stimulus Representation for Neural Populations

Recently, Osborne et al. (Osborne, Bialek et al. 2004) published results about the time course of directional information in responses of single macaque MT neurons. In the single cell study, about 80 percent of the total information a MT neuron can give about stimulus direction was available 100 ms after response onset.

Here we exploit our finding that responses from MT neurons to dynamic stimuli can be successfully modelled by populations of LNP–neurons and tested the time course of stimulus representation for a *population* of neurons. We used the decoding accuracy as measure for stimulus representation instead of mutual information as in the aforementioned paper.

Figure 14b showed how noise from unfaithful responses is decreased as the integrating time window increases with τ_0. Due to finite temporal filter width the subsequent stimulus entered the estimate of the previous stimulus for larger τ_0. This effect hindered a further reduction of uncertainty for discontinuous trajectories with 50 ms motion interval. In order to avoid this restriction we stimulated a reference population of LNP neurons by a newly created discontinuous trajectory where the basic motion interval was extended to 300 ms.

Decoding was performed for three population sizes (N=1, 10, 100) by randomly selecting 100 subpopulations for each size from a pool of 200 LNP neurons.

Figure 16 depicts relative decoding errors for three different population sizes (N=1, 10, 100) as a function of increasing time window. Errors were determined by randomly selecting 100 subpopulations of a given size from a pool of 200 LNP neurons and averaging the estimation errors. Error bars denote the standard deviation of the mean. Relative errors are the rescaled versions of the stimulus estimation time courses, such that maximum errors are labelled with 1 and minima with 0. For a single neuron (N=1) the absolute estimation error reduced from 90 deg to 62.2±0.9 (SEM) deg for a time window of 300 ms length (N=10: 16.0±0.4 deg, N=100: 0.9±0.1 deg), indicating that single neuron performance is very poor, compared to larger populations. The curves in Figure 16 show that a given relative error level is reached faster for a large population. Therefore, large populations not only improve the overall representational accuracy (which is a trivial statement), but may also positively influence the speed with which relative information about the stimulus is available.

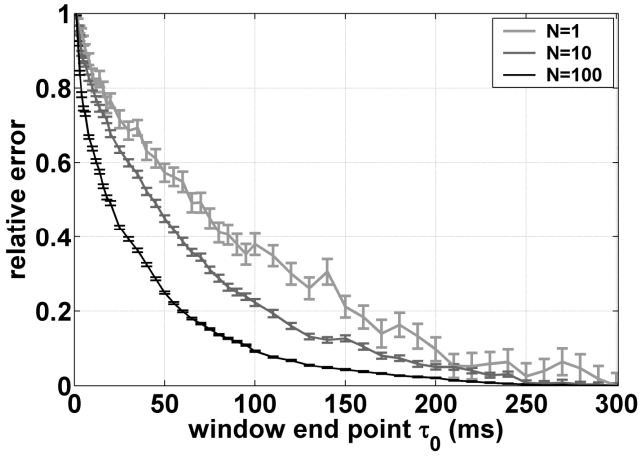

Figure 16: Time integration for LNP populations with N=1,10 and 100 neurons. A discontinuous trajectory with 300 ms per direction presentation was decoded using a population of neurons according to the reference model. Data was generated by selecting 100 subpopulations of a pool of 200 LNP neurons with a bootstrap method for each curve and averaging the results of the subpopulations. Error bars denote the standard deviation of the mean. The relative error is defined to be 1 for the maximum error of the respective curve and 0 for the minimum error. The relative error converges faster for larger populations.

5.4 Discussion

In the past, responses of MT neurons have been measured mostly for static stimulus conditions, where the stimulus parameter under consideration (mostly the orientation of a grating or the direction of a moving random dot cloud) is constant (Dubner and Zeki 1971). In that case static stimuli help eliminating the influence of neuronal time integration by averaging over time periods of several 100 ms. Temporal integration of single MT neurons has been considered (Osborne, Bialek et al. 2004), but only by using a static stimulus setup. Besides the examination of spike–triggered averages (Perge, Borghuis et al. 2005; Perge, Borghuis et al. 2005), only little is known about responses of MT neurons to dynamic stimuli.

We examined responses of a population of MT neurons to dynamic stimulus trajectories of random dot clouds in the direction domain with three different stimulus statistics, which we termed continuous, discontinuous, and fast discontinuous. The rates, with which stimulus values changed, were 20 Hz and 50 Hz, thus capturing a time scale, where neuronal time integration and stimulus dynamics are expected to interact with each other.
We compared delay, tuning curve, and linear filter estimates as well as reconstruction errors for the three stimulation types. Simulations with LNP–neuron populations were performed using the identical stimuli as in the experimental setup, and the results were used to interpret the experimental findings. Our setup required careful data analysis in order to avoid artefacts from rapidly changing stimulus values, which might suggest adaptational effects if data is evaluated naively.

5.4.1 Delays

Spike–triggered averages were determined as two–dimensional quantities, which, for a given delay, define the probability distribution on the stimulus values. Therefore, an entry of the STA denotes the probability of finding a stimulus value at a given delay before a spike.
We defined the latency of a neuron by the delay which maximized the variance calculated on the spike–triggered average from discontinuous stimulation. This latency estimation method was tested in two ways: We used experimental as well as simulated data for that purpose.
For the first test we introduced an additional global delay τ_{add} and reconstructed trajectories from the three stimulations with the maximum likelihood method. Errors as a function of τ_{add} showed that our delay definition provided the strongest correlation between the data and its

estimates. Correlations between the electrophysiological data and tuning properties were also used in Smith et al. (Smith, Majaj et al. 2005) to estimate latencies from white noise stimulation in the direction domain.

In the second test we used our method to estimate latencies from LNP neurons, where the true values are known. Histograms for the three stimulus conditions showed that discontinuous stimulation latencies were very accurate and more reliable than estimates from the other two stimulations: Blurring of the spike–triggered average through stimulus autocorrelations led to less accurate estimates from continuous stimulation, while noise resulting from the fast change of stimulus values defied estimates for fast–discontinuous stimulation. Therefore the stimulation with intermediate values for noise and width of the autocorrelation function gave the best estimates.

Raiguel et al. (Raiguel, Xiao et al. 1999) found an average of 87±25 (SD) ms in their latency measurements of MT neurons in anesthetized macaques. This values differs about 17 ms from the average value of 104.3±25.3 (SD) ms found in this study. It is speculative to attribute this difference to anaesthesia, while this study was performed on awake, behaving macaque monkeys. It is more likely to be due to the fact that latencies in the aforementioned paper were determined as response onsets for stimulus presentations of about 500 ms length. Mazer et al. (Mazer, Vinje et al. 2002) defined response onsets as estimates at the rising flank of the spike–triggered average variance, and "optimal" latencies by the maximum variance of the STA, as we did. For V1 neurons, they found a difference of 17 ms between latency values of the two different definitions. This result is in good agreement with the found difference.

Also, we found a tendency to shorter latencies for cells with stronger tuning. This result confirms a finding of Raiguel et al. (Raiguel, Xiao et al. 1999).

5.4.2 Tuning Curves

Tuning curves were estimated in order to test whether they differ for different stimulus conditions on time scales up to some seconds and whether differences can be attributed to adaptation of the neural system. Differences were found but could not be attributed to adaptation, but more likely result from feedforward influence of the different stimulation conditions. The first hint for this interpretation is given by calculating SI value distributions for continuous and discontinuous stimulation in dependency of the weighting of 10 ms time bins within the 50 ms basic motion interval. The more the weighting emphasized the central bins of each motion interval, the less significant was the two–sided Wilcoxon test. Finally, for the most centralized weighting the null hypothesis that both SI distributions are from an identical distribution could not be rejected with alpha=0.05. The second hint comes from the

virtual LNP neuron population. Though the model does not include any adaptation large differences are found in the SI value distribution. Therefore, our data can be completely explained without assuming adaptational effects in the direction domain.

5.4.3 Linear Filters

As stimuli differed in their temporal properties we tested whether temporal properties of neuronal responses adapted to the stimulus statistics. For this purpose we estimated optimal linear filters for the three stimulus types and focussed on the temporal aspect of the filters. The use of the inverse of the autocorrelation matrix is essential for eliminating stimulus correlations from filter estimates. However, a straightforward linear regression approach failed because parts of the input space are undersampled and were therefore very susceptible to noise. The principal component regularization is a well–known method to deal with this noise (Hastie, Tibshirani et al. 2001). It worked well for the stimuli with 50 ms presentation time in mainly eliminating autocorrelation influences, while it failed to produce faithful estimates for the fast stimulation. The reduction of noise is at the expense of the introduction of an additional parameter, the regularization parameter. The temporal widths of filter estimates depended strongly on the regularization parameter and it was impossible to determine a unique parameter value from cross–correlation methods.

Therefore, results are subject to interpretation during the data evaluation process. Note, that Machens et al. (Machens, Wehr et al. 2004) encountered the same problem with a similar noise reduction method for auditory neurons. On this basis the regularized filters from experimental and simulated data show identical characteristics in the comparison of the continuous and the discontinuous case suggesting that the differences found for temporal filter widths are not due to an adaptation in the time domain.

5.4.4 Reconstruction – Comparison of MAP and ML Estimates

Bayesian methods allow to integrate prior knowledge in the estimation of stimuli or combine information about stimuli in a statistically optimal fashion. For example, it has been shown in a task which employed two modalities that the brain optimally combines visual and haptic information (Ernst and Banks 2002), as if the brain combined probability distributions to produce estimates according to the combined ML function. In addition, experiments from Körding and Wolpert (Kording and Wolpert 2004) suggest that the brain has internal representations about the statistical distribution of the stimulus.

In this study we looked at the advantage a population of neurons could draw from prior knowledge. We compared ML and MAP reconstructions of a continuous trajectory where the MAP estimation employs a Bayesian framework which uses the stimulus–creating probability distribution as a dynamic prior. In comparison, ML evaluation can be interpreted as a special case of the Bayesian approach, where a flat prior is used, therefore ignoring any prior knowledge about the stimulus.

The impact of the prior is strongest, if the number of neurons available for stimulus encoding is very low. In our case, for population sizes of 100 and more neurons the knowledge about the true prior does not yield further information. Our results suggest that – if only applied for estimation of stimulus parameters – feedforward strategies and ML evaluation are sufficient. The situation may change if cognitive processes are involved, where a prior might play a fundamental role in the computational paradigm.

A prior can provide additional information for technical applications as online decoding; where information from a multi-electrode array must be converted in short time to a stimulus estimate (Chapin, Moxon et al. 1999) and pooling from several repetitions of the same stimulus is not possible. In that case a prior could be given by external knowledge about the statistics of stimulus value changes. Another case, where external knowledge about the stimulus improved decoding quality was the reconstruction of position of a foraging rat from hippocampal responses (Zhang, Ginzburg et al. 1998).

An interesting feature in Figure 13c is the saturation of reconstruction errors for both ML and MAP. The Cramer–Rao theorem in estimation theory (Kay 1993) states that the variance of any unbiased estimator is bound by the Cramer–Rao lower bound (CRLB), hence setting a limit on decoding accuracy. The ML estimator e.g. is able to attain this limit and is therefore called statistically efficient. Much work in theoretical neuroscience has been devoted to the optimal estimation of single stimulus values by a population of neurons employing the Fisher Information (FI), which is the inverse of the CRLB (Paradiso 1988; Seung and Sompolinsky 1993; Eurich and Wilke 2000; Wilke and Eurich 2002). Under some conditions (N uncorrelated neurons, gaussian tuning curves, poissonian spiking statistics) the FI for a population is proportional to the density of neurons in the stimulus space (Zhang and Sejnowski 1999). As N tends to infinity, the FI also tends to infinity and the CRLB approaches zero. However, in Figure 13c we find that the error does not tend to zero for $N>1000$.

In our simulations we had conditions that lead to a vanishing CRLB. Especially, the population was uncorrelated, where correlations are thought to limit coding accuracy (Shadlen

and Newsome 1994), but see also Shamir and Sompolinsky 2004 (Shamir and Sompolinsky 2004). Below we will discuss that the interaction of finite stimulus duration and finite temporal filter widths is responsible for the saturation effect by introducing a bias into the reconstruction. The saturation levels of ML reconstruction of the continuous trajectory with different tuning sets and N_{eff}=10000 are given by 5.9 deg and 14.1 deg for continuous and discontinuous sets, respectively. Therefore, the use of the tuning set which results form a different stimulus type leads to an additional bias in stimulus estimation.

5.4.5 Reconstruction – Stimulus Statistics

The temporal dynamics of stimulus trajectories had notable influence on decoding accuracy (Figure 14a). In our three examples we found that the broader the autocorrelation functions the better are the estimates. Interestingly, differences in estimation errors persist also for large values of N. We addressed the question why the fast–discontinuous estimates perform far worse than their discontinuous equivalents. By reconstructing the 50 ms stimuli as a function of increasing time window in Figure 14b it turned out that the 20 ms time window by itself contains enough information for a more accurate stimulus representation. Furthermore, the comparison with reconstructions from simulated data (Figure 14c) suggested that adaptational effects are not responsible for restricted reconstruction ability. Figure 14d shows that the temporal filter width is responsible for various saturation levels. For populations with short time integration, reconstructions of discontinuous and fast–discontinuous stimuli become virtually identical for window lengths of 20 ms. In addition, the error for the continuous stimuli tends to zero, showing that also the continuous stimulus is subject to a finite bias. The latter reconstruction demonstrates the unbiasedness of the ML estimator.

5.4.6 Reconstruction – Coding Strategies for Dynamic Stimulus Trajectories

We used our simulations approach to assess the influence of LNP neuron parameters on estimation of dynamic trajectories. Much work has been devoted in the last years to the search for favourable parameter values or coding strategies of neuronal populations for optimal single stimulus coding. Features like tuning width, maximal firing rate, background activity, stimulus dimensionality, and noise correlations have been examined (see e.g. (Eurich and Wilke 2000; Dayan and Abbott 2001; Wilke and Eurich 2002). Here we study the temporal filter characteristics which have not been considered previously due to the static stimulus estimation framework used in the literature.

Temporal filter width turned out to be most important, because it introduces a bias in the estimations. Another influential parameter we considered was the tuning width. We found that

sharp tuning improves decoding quality. This result is in agreement with theory for one–dimensional stimuli. In contrast to intuition and the theoretical findings, background activity and mean stimulus–induced rate played almost no role in our estimations. The small influence of background activity for a smaller neural population in Figure 15a suggest the following explanation: The other parameters, mainly temporal filter width, introduce a bias which acts as lower limit and which is already reached for the baseline values examined in this work. The identical argumentation also holds for the mean stimulus–induced rate. This result obviously depends on the population size and the integration speed.

5.4.7 Reconstruction – Time Course of Stimulus Representation for Neural Populations

The modelling approach was also used to determine the time course of estimation errors, when single stimulus values are coded by a population of neurons.

For that purpose a discontinuous trajectory with 300 ms basic motion interval was decoded by virtual populations with N=1, 10, and 100 neurons. Besides the rather trivial statement that coding accuracy improves with increasing population size we find that relative decoding errors decrease faster for larger populations. This result can be compared to the time course of mutual information of single units in MT recently measured by Osborne et al. (Osborne, Bialek et al. 2004). Here, the relative decoding error is used as equivalent to mutual information. As population size influences the convergence speed of relative errors it might not be sufficient to compare single neuron results with psychophysical speed thresholds. Therefore, it might be necessary to look at the time course of information resulting from neural populations in order to explain the speed with which the brain performs computations.

5.5 Final Conclusions

We found no evidence supporting the idea that neurons in area MT adapt to statistics of dynamical stimuli on a time scale of seconds, neither in direction nor in time domain. Furthermore, comparisons showed that experimentally recorded responses from neurons were very similar to those obtained by a population of LNP neurons, which apply a linear stage, a rectification, and a Poisson process as firing mechanism.

Is the description of neurons via models which employ linear filters justified? The linear filter approach is useful because of the simplicity of the model and by its ability to capture a good deal of neural responses in many cases (Touryan, Lau et al. 2002). But there exists evidence that temporal filter properties depend on parameters such as speed, spatial frequency and contrast of a sinusoidal grating (Bair and Movshon 2004). In this work, however, the quasi–linear approach reproduced results from experimental data.

6 Effects of Attention during Rapid Serial Visual Presentation on Macaque MT Neurons

6.1 Introduction

Two influential recent studies (McAdams and Maunsell 1999); (Treue and Maunsell 1996) have found that in both areas MT and V4, attention enhances response magnitude of a given cell similarly to all motion directions (respectively grating orientations). That is, attention scales the tuning curve vertically by a constant factor without further affecting its shape.
Earlier findings suggested that task difficulty could affect how attention modulates neural tuning (Spitzer, Desimone et al. 1988). Here we address the question of whether gain modulation can account for attention effects under conditions which require close monitoring and discrimination of constantly changing motion directions. We used a rapid serial visual presentation (RSVP) paradigm to study directional tuning in macaque area MT.

6.2 Methods

6.2.1 Animals and surgical procedures

Two male macaque monkeys (*Macaca Mulatta*) have been trained to perform an attention demanding motion discrimination task. Monkeys were 6 and 8 years old. Both animals have been employed in different attention paradigms before. For the purpose of this project they have been retrained over a period of two to three months. Surgical procedures have been described in more detail elsewhere (Wegener, Freiwald et al. 2004), see also chapter 4.3). All animal procedures conformed to the guidelines of the National Institutes of Health for the care and use of laboratory animals, the guideline for the welfare of experimental animals issued by the Federal Government of Germany and were approved by local authorities.

6.2.2 Data acquisition and recording

Single unit recordings were obtained with Tungsten microelectrodes from WPI (Sarasota, Florida). Electrodes were advanced with a hydraulic microdrive (Narishige, Tokyo, Japan). Signals were amplified and filtered (350-500Hz) with recording hardware from Multichannel Systems (Reutlingen, Germany). A custom made recording system digitized the analog signal sampled at 25 kHz and stored it on computer disk for offline analysis. Recordings were done in area MT.

6.2.3 Behavioural task and visual stimulation

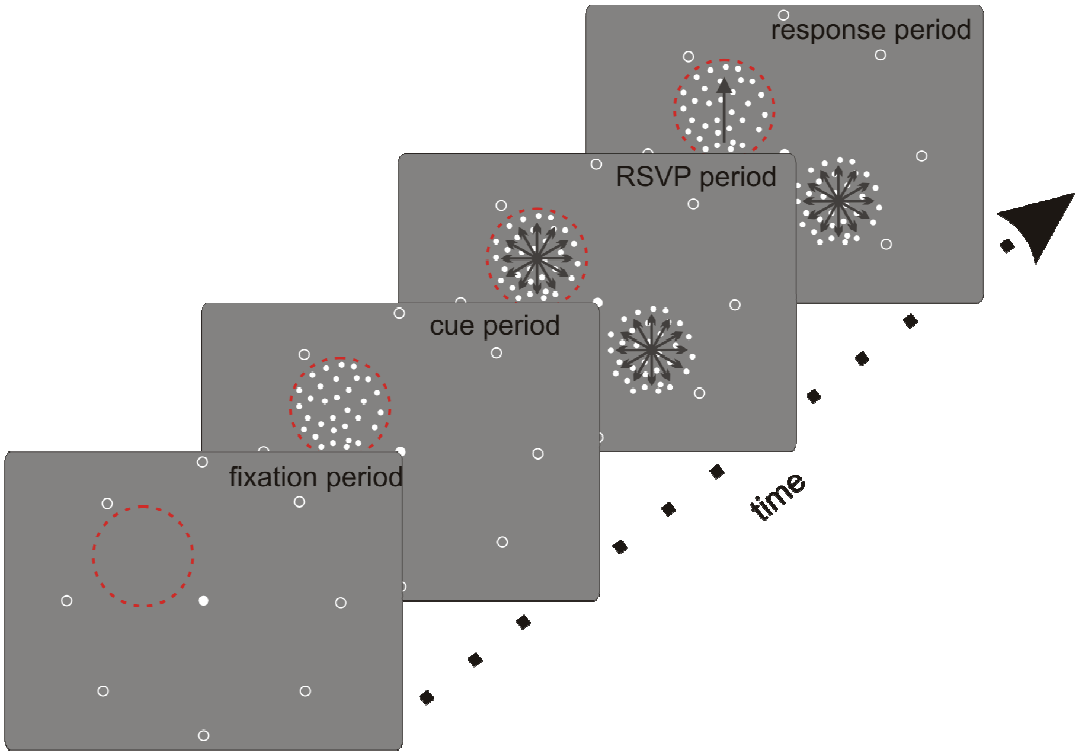

Figure 17: Stimulus configuration and behavioural paradigm. Monkeys started a trial by fixation to the fixation spot. 500 ms later the cue, a static dot field appeared at the location to which the monkey was required to pay attention to for 500 ms. After another 500 ms without stimulus, two random dot surfaces appeared moving into different motion directions. Direction of motion for each surface changed randomly in discrete steps of 30° (RSVP period). Monkeys had to track the target stimulus until the translation direction ceased to change (response period), while ignoring the distracter stimulus in the other hemifield. Red dotted circle represents receptive field border. See text for further explanations.

During foveation of a fixation spot, two spatially separated surfaces were presented, one within the receptive field of the neuron under investigation, the second at identical eccentricity, rotated 180° around the fixation spot. Random dots of each surface had a density of 5 dots/$^{\circ 2}$ and infinite lifetime. During translation, either 80% or 90% of the dots moved coherently towards a specified direction, and 20% or 10% moved in random directions.

After an initial fixation period of 500ms the monkey was instructed which surface would be behaviourally relevant by a static presentation of one surface. After another delay of 500ms both surfaces appeared. During this period random dot surfaces were changing their translational direction every 50ms, 60ms or 100ms in multiples of 30°. The monkeys had to track the target stimulus for 20(40)-80 translational phases until the translation direction ceased changing for up to 400ms followed again by rapid direction changes (see Figure 17). Monkeys were required to detect the prolonged translational movement and additionally discriminate its direction, which had to be indicated by a saccade to one of 8 saccade target

dots presented at 9.5° vis. angle in the periphery during the whole stimulation period. A trial was rated as successful if the animal initiated a response within 800ms after the beginning of the prolonged translational movement and reached the appropriate saccade target 500ms afterwards. A successful trial was rewarded with a drop of liquid.

6.2.4 Data Analysis

Similar to the analysis methods described in chapter 4.3 and chapter 5.2.3 tuning curves and spike–triggered averages have been computed. Spike-triggered averages have been generated by reverse correlating the spike train with the stimulus sequence presented in the receptive fields of the neurons (see Figure 18 sketching this procedure for one direction). The reverse correlation matrix was normalized to get the spike triggered averages in units of probability. Response delays of neurons were assessed by looking for the maximum value of the time-dependent variance over motion directions in the spike-triggered average. Tuning curves were then computed from responses sampled in a window centred on this delay. Note that this delay is specific to each individual neuron.

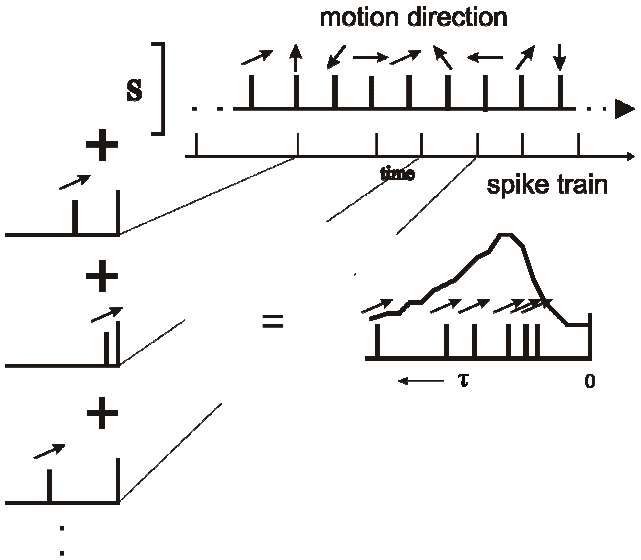

Figure 18: Spike-triggered average. The neural response is cross correlated with the motion impulse sequence. The sketch depicts the reverse correlogram for motion in left-upward direction.

For each neuron spike-triggered averages and tuning curves for attended and unattended condition were computed. Tuning was assessed for both the rapid serial visual presentation period and for the behavioural relevant prolonged translation period. Tuning curves subsequently were fitted with a von Mises function ($M(\theta) = A\exp\{\kappa[\cos(\theta-\varphi)-1]\}$), where A denotes the peak amplitude for the preferred direction φ and κ is the width parameter, using the Matlab Optimization Toolbox (The Mathworks Inc., Natich, MA, USA). The von Mises distribution is also known as the circular normal distribution, so it may be thought as a circular analogue of the normal distribution. Fitted parameters were used to compute

attentional indices ($(P-A)/(P+A)$, where P is response to attended stimulus in the receptive field and A is response to unattended stimulus in the receptive field).

6.3 Results

6.3.1 Behaviour

Both animals performed the task with approximately 70% correct trials, far above the chance level of 12.5%. Broke fixation errors which are unrelated to any stimulus behaviour were handled separately (approx. 30% of all trials in both animals). This paradigm thus is highly demanding not only requiring detection of an event but also discrimination of movement direction, while setting the whole visual system under a high perceptual load due to its multiple direction changes throughout the whole stimulation period.

6.3.2 Neural responses

We recorded from 99 direction selective neurons in area MT in two monkeys (62 + 37 cells, respectively). Here we report on attention effects on firing rates in response to rapidly changing translational direction of motion. We computed motion direction reverse-correlograms (spike triggered average, (STA), procedure explained exemplarily for one direction of motion in Figure 18). An exemplary neuronal response is shown in Figure 19. The probability values for a direction value at delay τ are colour coded for both attention conditions. Left graph shows the STA when attention is paid to the random dot surface inside the receptive field, right graph when attention is directed outside the receptive field.

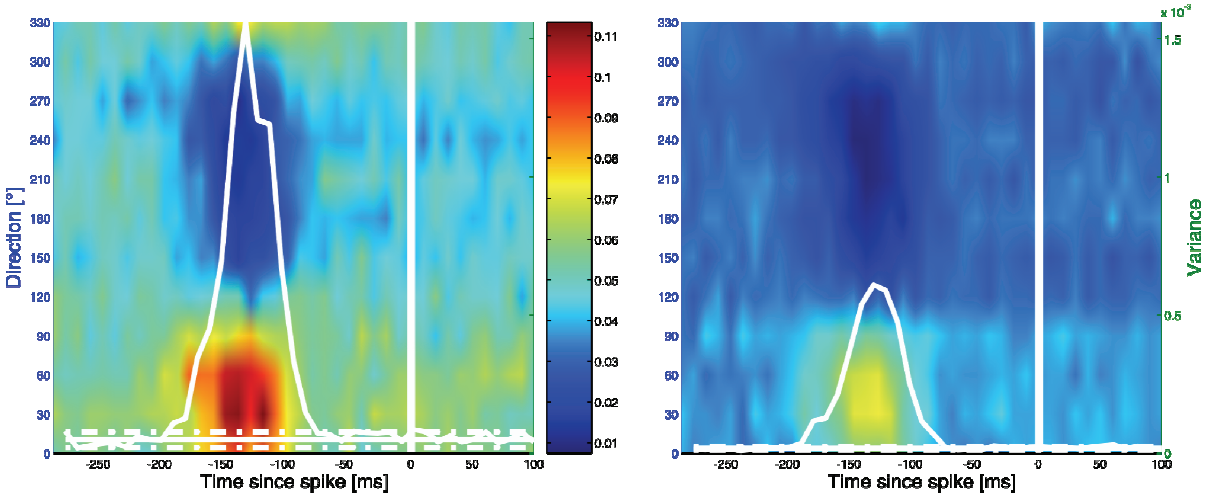

Figure 19: Spike triggered average for a single neuron in the attended condition (left) and for the unattended condition (right). During attention the overall activity is increased and the variance peaks higher. The probability value for the preferred direction is increased during trials in which the behaviourally relevant stimulus was inside the receptive field.

The overlaid white curve denotes STA variance across motion direction for each delay value τ. Solid and dashed-dotted horizontal lines denote the mean and plus minus two standard deviations of STA variance, respectively. A high variance at some delay τ indicates a significant relationship between motion direction and spiking activity. The latency of a cell is defined as that τ which maximizes the variance (Mazer, Vinje et al. 2002).

Tuning curves were determined for each neuron separately from both attention conditions by averaging the firing rate responses to each direction (Stemmann, Freiwald et al. 2005). The response to a single direction presentation was estimated as follows: The spike train was shifted relative to the stimulus sequence by the previously determined latency τ, which is calculated for both conditions together for this purpose. The response rate was then defined as the number of spikes within the basic motion interval normalized to number of presentation and duration. Figure 20 shows the tuning curves for the same neuron displayed in Figure 19. The red graph depicts the tuning curve for the attended condition.

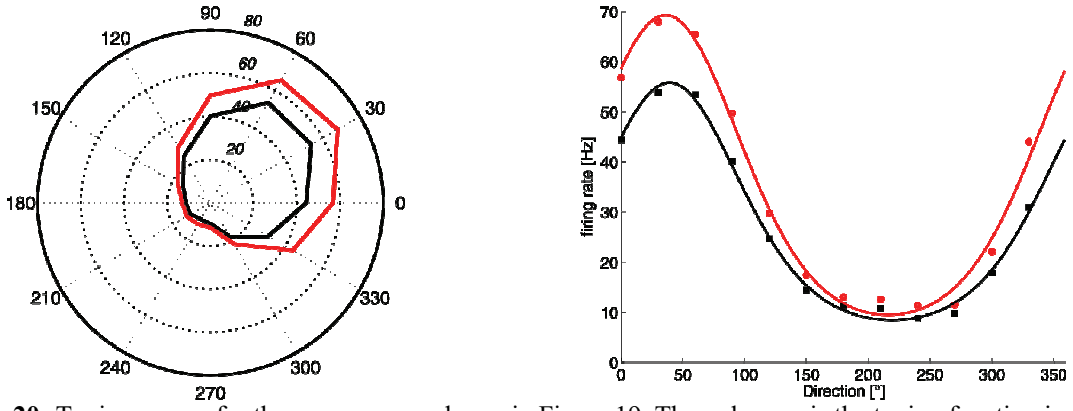

Figure 20: Tuning curves for the same neuron shown in Figure 19. The red curve is the tuning function in the attended condition, the black one in the non-attended condition. Curves in the right graph were fitted with a von Mises function ($M(\theta) = A\exp\{\kappa[\cos(\theta-\varphi)-1]\}$, where A is the amplitude at the preferred direction φ and κ is a width parameter).

This neuron exhibited a higher response in terms of firing rates for the attended condition compared to a situation when attention is directed somewhere else in the visual field. The response is stronger for all motion direction. In order to assess the attentional effect for the whole population of MT neurons we calculated an attention index that examined the attentional effect on the firing rate to the preferred direction of each individual neuron. Figure 21 shows the attentional effect size for the whole population of area MT neurons in response to the rapid serial visual presentation period of the stimulus.

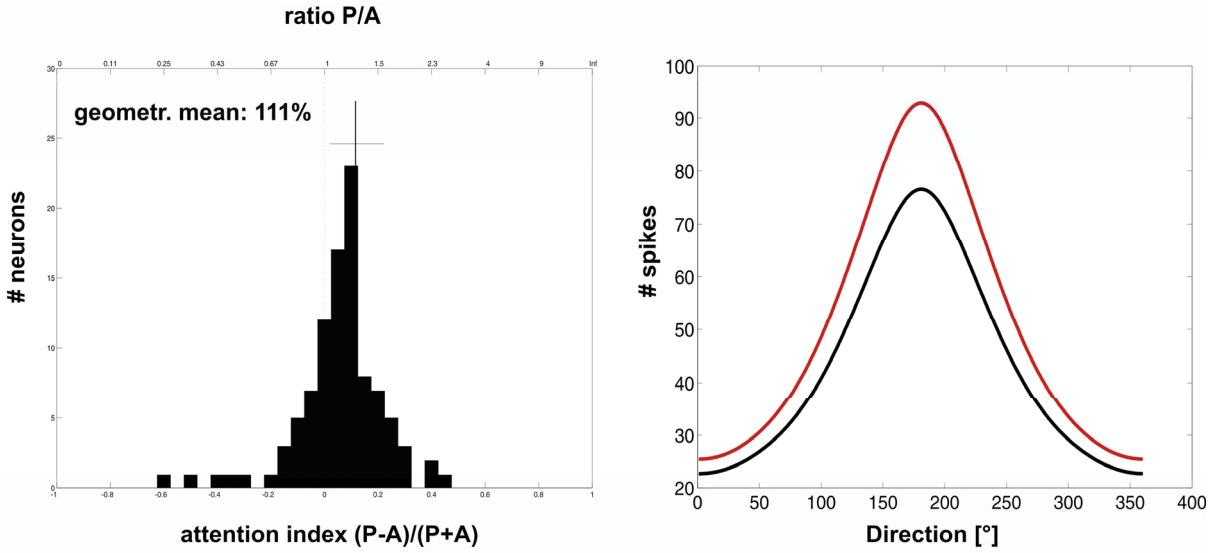

Figure 21: Attentional effect size over the population of MT neurons during the rapid visual presentation period. **Left:** Distribution of the neurons' attention indices. The top scale shows the corresponding fiiring rate rations. The histogram is shifted rightwards, with an average attention index of 0.055 (geom.. mean), indicating an attention dependent average increase in firing rate by 11% (amplitude A of von Mises function). The horizontal cross bar indicates the 99% confidence interval (Wilcoxon signed rank test). **Right:** Population average of tuning curves (based on von Mises fits). Red curve: population tuning during attended condition; black curve: population tuning during non-attended condition. Population tuning curves differ in amplitude over the whole range of motion directions.

There is a significant shift in the histogram indicating an average increase in firing rate by 11% for the preferred direction. The right graph shows the population tuning curve for both conditions. Neural response is enhanced over the whole range of motion directions.

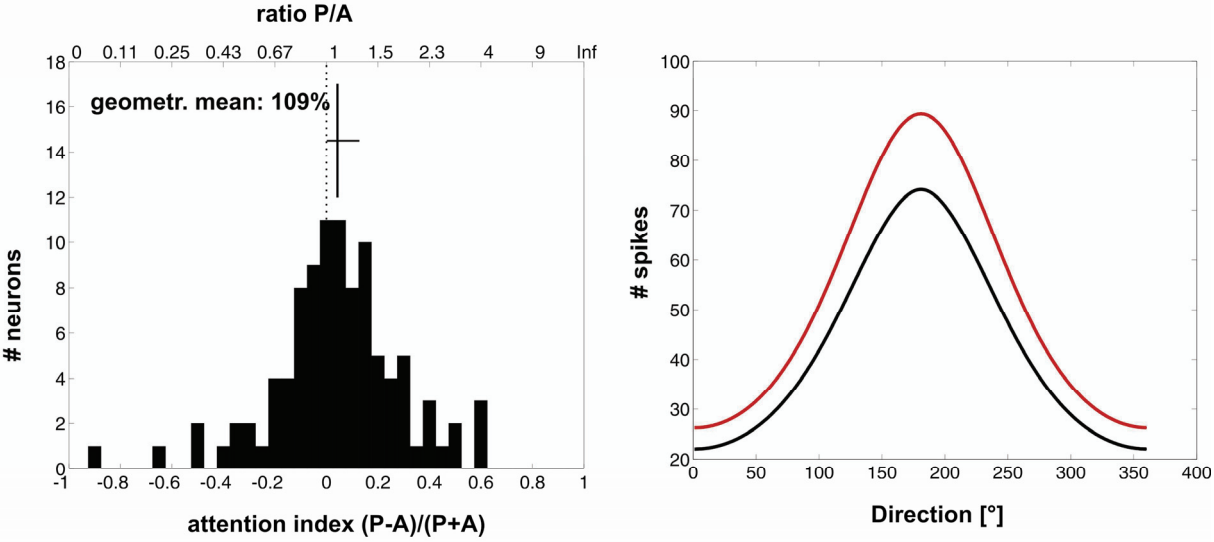

Figure 22: Attentional effect size calculated from the repsonses to the prolonged translational movement the monkeys were requiered to respond to. The attentional modulations are comparable to the modulations found for the responses during the rapidly changing period (Figure 21). Conventions are the same as in Figure 21.

Figure 22 shows the results for the tuning curves resulting from analysing the neural responses to the behavioural relevant movement. Neural response is 9% increased in the

attended condition. The population tuning curve is again positively modulated over the whole range of motion directions.

6.4 Discussion and Conclusion

The rapid serial visual presentation paradigm has been effectively combined with an attention demanding task in macaque monkeys. The stimulus regime allowed acquiring spike-triggered averages and tuning curves in the presence and absence of attention. Basic effects of firing rates have been reproduced. Treue and colleagues (Treue and Maunsell 1996; Treue and Martinez Trujillo 1999) were the first to report attention effects in area MT. In their first study (Treue and Maunsell 1996) using a single squared dot moving in the cells' preferred motion axis they reported attentional increases in firing rate of about 19%. Subsequent studies using random dot surfaces reported increases of 10% (Treue and Martinez Trujillo 1999). The result of this study matches the latter results from Treue very well. Whether the discrepancy between their early and late reports is related to the use of different visual stimuli needs further investigation.

Introducing different levels of noise by employing random dot stimuli might lead to an overall decrease in effect size and could also require additional control mechanisms provided by other areas of the brain. To investigate the interaction of area MT and all other cortical areas involved in the performance of a motion discrimination task, we employed functional magnetic resonance imaging techniques. The next chapter provides an overview of the whole cortical network activated during performance of a slightly, for fMRI purposes, adapted behavioural paradigm.

7 Imaging Cortical Networks of Spatial Attention in the Macaque

7.1 Introduction

The visual system is not a rigid, passive analyzer of incoming sensory information. Rather some information deemed important for current behavioral goals can be selected at the expense of other, task-irrelevant information (Chun and Wolfe 2001). This active selection process, attention, constitutes a critical link between the brain regions determining internal cognitive sets and the brain regions processing information about the outer world. Therefore, understanding attention requires us to know the entire network of brain areas involved (Kanwisher and Wojciulik 2000; Kastner and Ungerleider 2000) and, ultimately, how interactions between areas within this network (Buschman and Miller 2007; Saalmann, Pigarev et al. 2007) cause changes in information processing (Moore and Armstrong 2003).

Attention-induced changes of information processing have been studied mostly in the macaque monkey, and are now well understood at the single-cell level (Treue 2001; Reynolds and Chelazzi 2004). Yet at the level of cortical areas and networks, our picture of attention mechanisms in the macaque is incomplete (Heeger, Gandhi et al. 2001), because a global survey of brain areas showing attentional modulation has never been undertaken. The advent of fMRI in awake macaque monkeys (Vanduffel, Fize et al. 2001; Tsao, Vanduffel et al. 2003), provides a technique to systematically identify the areas subserving attentive stimulus processing and determine whether any key components have been missed by previous studies that focused on single areas.

Here, we used fMRI in macaque monkeys performing a spatial attention task (Figure 23A). Monkeys had to track the rapid motion sequence of one of two random dot surfaces (Figure 23B). When motion direction ceased changing, monkeys had to saccade in the corresponding direction (movie of stimulus available online http://www.neuro.uni-bremen.de/~stemmann/paradigma1_StemmannFreiwald.avi). Task design emphasized endogenous top-down spatial attention, because (i) the task-relevant event was not a feature change that could capture attention exogenously (Muhlenen, Rempel et al. 2005), (ii) the rapid serial visual presentation (RSVP) design put the visual system under a high perceptual load (Suppl. Text S1), (iii) attention had to be continuously deployed over seconds, and (iv) this attentional deployment was dissociated from saccade preparation.

Imaging activity across the entire brain during performance of this attention-demanding motion discrimination task allowed us to directly compare activations across cortical areas in order to address three fundamental questions about the mechanisms of attention: first, what is the locus of attentional selection? High perceptual load has been proposed to lead to early visual selection (Lavie and Tsal 1994), and thus attention effects are expected to occur as early as in area V1. On the other hand, if competition between the neural ensembles representing target and distracter (Desimone and Duncan 1995) or the match of stimulus size and receptive field size (Luck and Hillyard 2000; Hopf, Luck et al. 2006) are the main determinants of attentional modulation, attention effects would be expected at processing stages with larger receptive fields.

Second, is attentional modulation in this task specific to the motion-processing pathway? Theories of feature-based attention posit that attention can selectively recruit cortical areas (Corbetta, Miezin et al. 1990; O'Craven, Rosen et al. 1997) or even specific cell groups (Uka and DeAngelis 2004) whose feature selectivities are most informative about the task-relevant features. Thus areas like V1, V2, V3, and dorsal stream areas MT, MST, FST and VIP that contain large fractions of direction-selective cells (Zeki 1974; Baizer 1982; Desimone and Ungerleider 1986; Orban, Kennedy et al. 1986; Galletti, Battaglini et al. 1990; Colby, Duhamel et al. 1993) should be modulated by attention.

Third, which areas exert attentional control? A fronto-parietal network of areas is thought to control the "spotlight of attention" (Kastner and Ungerleider 2000). In the macaque, both the lateral intraparietal area, LIP (Gottlieb, Kusunoki et al. 1998; Bisley and Goldberg 2003), and the frontal eye fields, FEF (Thompson, Bichot et al. 2005), have been suggested to contain salience maps and should thus belong to the network responsible for guiding attention to task-relevant regions.

7.2 Methods

All animal procedures conformed to the NIH Guide for Use and Care of Laboratory Animals, regulations for the welfare of experimental animals issued by the Federal Government of Germany, and stipulations of local Bremen authorities.

Subjects and Surgery. Two male rhesus monkeys (6-10kg) were used in this study. Prior to training and MR scanning, each monkey was implanted with a MR-compatible plastic head post (Ultem, General Electric Plastics) attached to the skull by ceramic screws (zirconium oxide, Thomas Recording) and dental cement (Grip Cement, Caulk, Dentsply International).

All procedures followed standard anesthetic, aseptic, and postoperative treatment protocols, described in detail in (Wegener, Freiwald et al. 2004).

Visual Stimulation and Tasks. The main task both monkeys performed was a motion-tracking task in which macaques had to foveate a central fixation spot (0.25°) and covertly pay attention to one of two random dot surfaces (RDSs), positioned on the horizontal meridian to the left and right of the central fixation spot (Figure 23A). RDSs were presented in a circular aperture 6° of visual angle in diameter and positioned 5° from the fixation spot on the horizontal meridian. Dot density of each surface was 5 dots per square degree of visual angle, translation velocity 7°/s. Eye position of the animals was monitored by an infrared pupil tracking system (ETL-200, ISCAN Inc., Burlington, MA). After an initial fixation period, one RDS was cued as the behaviorally relevant stimulus (the target) by a short line (0.35° x 0.06°) to the side of the fixation spot, pointing to this surface. RDSs randomly changed translation direction every 60 or 50 ms (monkey Q and M, respectively) in random multiples of 20° (Figure 23B), until the translation direction ceased changing for up to 600 ms (the prolonged motion event, PME), to be followed again by rapid direction changes. The PME occurred at a random point in time after at least 20 and at most 60 brief motion events, independently in target and distracter stimulus sequence. While keeping fixation in a 2° x 2° fixation window, monkeys had to pay attention to the target motion sequence to detect the PME and discriminate its direction. Monkeys had to report the direction of the PME by a saccade to one out of 8 peripheral saccade target (ST) dots positioned 8.5° from the fixation spot (Figure 23A). A trial was rated successful if the animal initiated a response within 800ms after PME onset, and if the saccade reached the correct target directly in less than 500ms. Successful completion of a trial was rewarded with a drop of water or juice.

Both monkeys were also trained on a passive fixation task, in which they were rewarded for maintaining fixation within the 2° x 2° central fixation window. During the fixation task various stimuli could be presented. The fixation task was used, without further visual stimulation, as a control for the attention task and for localizer scans (meridian mapping, center-periphery mapping, and motion localizer, see 7.5.2).

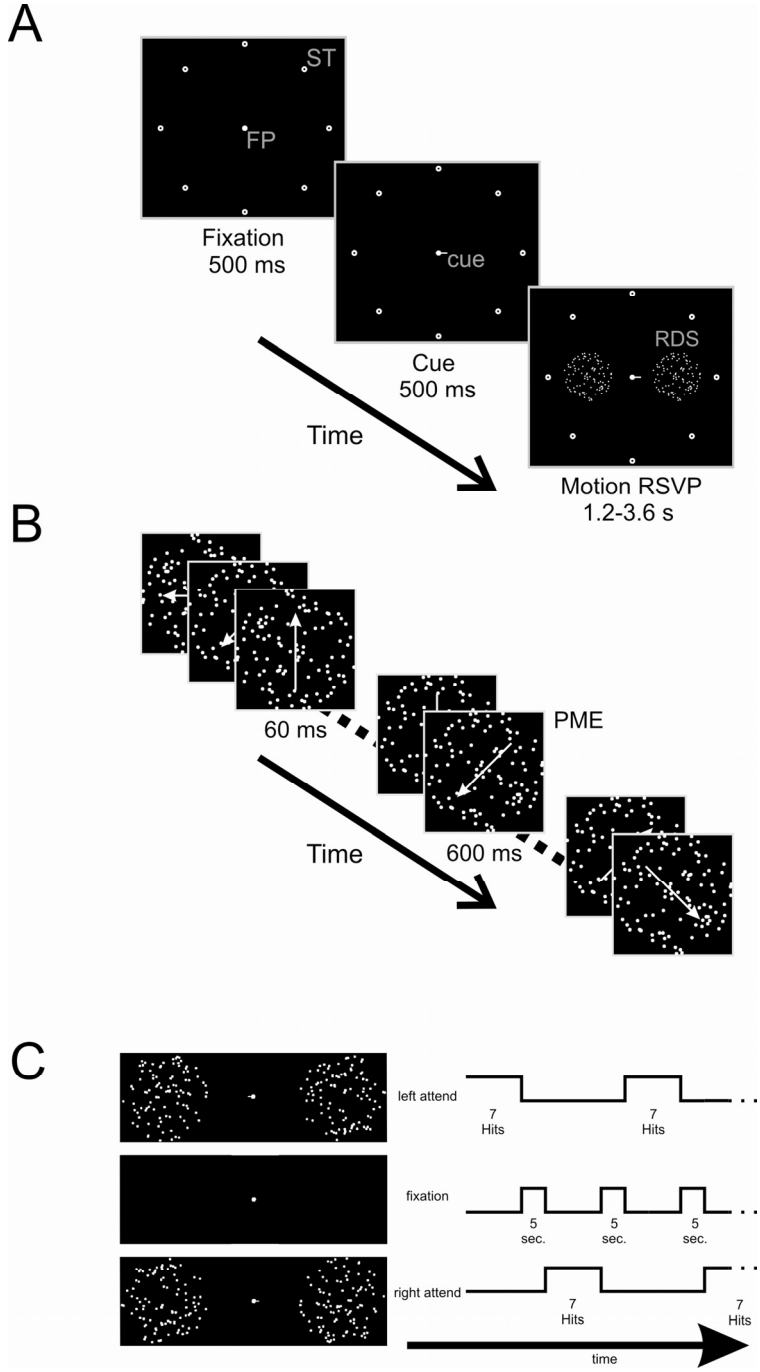

Figure 23: Design of attentive motion-discrimination task. (**A**) Stimulus configuration. Monkeys initiated a trial by foveating the fixation spot (FP). After a 500 ms delay, a horizontal bar (the spatial cue) occurred originating from the fixation spot, pointing either to the left or the right indicating which side attention had to be deployed to. After another 500 ms, two random dot surfaces (RDSs) appeared to the left and the right of the fixation spot at equidistant positions. (**B**) Event sequence of one of the RDSs. While both RDSs were changing their motion direction every 60 ms, the monkeys had to track the target stimulus for 20 to 60 direction changes until the translation direction ceased changing for 600 ms (prolonged motion event, PME) followed again by rapid direction changes. Monkeys were required to respond to the target surface PME by a saccade to the corresponding saccade target (ST). (**C**) Temporal sequence of behavioural conditions (attentive motion-discrimination and passive fixation) during scanning. Monkeys alternated between performing the attentive motion-discrimination task, paying attention to the left side (top), passive fixation condition (center) and attentive motion-discrimination task, paying attention to the right side (bottom). The sequence of behavioural conditions was 'LFRFLFR…', with 'L' 'attend left', 'F' 'fixate', and 'R' 'attend right'. During the 'attend left' and 'attend right' blocks, monkeys had to complete seven trials successfully ('hits'). Thus, duration of these blocks was variable, on average ~30 seconds. Passive fixation blocks separated the two attention conditions and lasted 5 seconds.

Meridian mapping: the retinotopic mapping procedure to define early visual areas was composed of two alternating 30 second blocks (V and H), separated by 30 second fixation periods. During block type V a vertical black-and-white checkerboard wedge (20° width) was shown with 1 Hz contrast reversals, during block type H, a horizontal wedge.

Center periphery mapping: in order to effectively separate activations resulting from foveal stimuli (fixation spot and bar cue) from activations resulting from the random dot surfaces, we modified the meridian mapping stimulus to alternate not between blocks of horizontal and vertical wedges, but between a foveal and a peripheral aperture displaying the checkerboard pattern, contrast inverting at 1Hz frequency. The diameter of the foveal aperture was 1°, the diameter of the peripheral aperture was 6.5° (slightly larger than the RDS diameter in motion discrimination task).

Motion mapping: To localize motion-responsive areas we used a motion stimulus adopted from (Nelissen, Vanduffel et al. 2006). Blocks of 30 seconds of blank screen were alternated with blocks of different motion types of random dot patterns (50 % white 0.2° dots on a black background, dot density was ~10 dots per square deg). The different types were: (1) slow translation with a velocity of 1°/s, (2) fast translation with a velocity of 8°/s, (3) rotation, (4) expansion and contraction, (5) static presentation, (6) opponent motion with a stripe width of 4°, and (7) opponent motion with a stripe width 1°. For the first four motion types motion direction changed with a frequency of 1 Hz to avoid adaptation.

Saccade mapping: in order to map saccade-related activity and localize the FEF, we had both monkeys perform a guided saccade task: in alternating 30 second blocks, monkeys were required to either maintain fixation on a central target, or make saccades every 1.5 seconds to a new random location within a grid of 3 x 3 possible saccade targets (grid spacing 8.5°).

Finally, monkeys were trained on an attentive motion-detection task. The temporal sequence of events was identical to that of the attentive motion-discrimination task. The detection task differed from the discrimination task in that the brief motion events were of completely incoherent motion, with all dots of the RDSs moving in random directions. The prolonged motion event was of low coherence (10%). Monkeys had to respond to the detection of the PME by saccading onto the target surface. Thus, the task allowed monkeys to plan the saccade to the target surface as they were paying attention to it. All other task requirements were identical to the attentive discrimination task.

Stimulus presentation and behavioral monitoring and synchronization to scanner were done by custom written software, run on a PC under Windows XP. During scanning, stimuli were displayed at 75 Hz with a resolution of 1280 x 1024 pixels, using a video beamer (JVC DLA - G15E) on a back projection screen, 49 cm in front of the monkey's eyes.

Scanning Procedures. After performance of the attentive motion-discrimination task reached asymptote (after 6-8 months training), each monkey was scanned in a horizontal 3T MR head-scanner (Allegra, Siemens, Erlangen, Germany). A radial surface coil was positioned immediately over the monkey's head.

Each experiment consisted of 10-15 functional scans of 6-9 min duration each. Functional time series consisted of gradient-echoplanar whole-brain images: repetition time (TR) 1.5s or 3.0s; echo time (TE) 30ms, 1.56 x 1.56 x 2 mm^3 or 1.5 x 1.5 x 1.5 mm^3 voxel size (24 transversal slices or 32 coronal slices). The effective field of view consisted of the entire macaque brain or was covered sequentially in independent sessions by partially overlapping slice positioning. In addition, for each subject, a high-resolution anatomical volume (3D-MPRAGE, 256 x 256 matrix, 128 slices, 0.5 x 0.5 x 0.5 mm^3 voxel size) was obtained in the ketamine-medetomidine anesthetized monkey. This anatomical MR data was used to generate inflated and flattened cortical representations for each subject using Freesurfer software, (http://surfer.nmr.mgh.harvard.edu).

For all functional imaging of monkey M, a contrast agent, ferumoxtran-10 (Sinerem, Guerbet, Paris; concentration: 21 mg Fe/ml in saline; dosage: 8 mg Fe/kg), was injected into the femoral vein prior to each scan session. Sinerem is the same contrast agent as MION, produced under a different name (dextran-coated iron oxide agent; MION (Vanduffel, Fize et al. 2001)). To confirm independence of the main result of the use of a contrast agent, we measured the BOLD response (without contrast agent) in monkey Q for scans of the attentive motion-discrimination task. All other functional scans in monkey Q used Sinerem.

Data Analysis. Functional data were analyzed using a block design in SPM5 (statistical parametric mapping, Wellcome Department of Imaging Neuroscience; London, UK). Freesurfer was used for registering functional to anatomical volumes and for surface flattening. Scans that showed unacceptable levels of movement artifacts or during which performance was low were discarded. Functional data were motion-corrected and spatially smoothed with a Gaussian kernel of 2mm full-width-at-half-maximum. Realignment

parameters were included as covariates of no interest in the general linear model (Friston, Holmes et al. 1994).

For each functional contrast, significance was assessed by t-scores, displayed as a statistical parametric map. Strength of activation was determined by the mean GLM beta values (scaled to percent signal change).

Boundaries of retinotopic visual areas were determined by meridian mapping (Sereno, Dale et al. 1995). Boundaries of areas inside the superior temporal and intraparietal sulci were determined by mapping with a motion localizer aided by known anatomical landmarks. To identify the brain regions activated by target and distracter stimulus, "peripheral activation zones" were defined by the contrast peripheral vs. central stimulation of the center-periphery mapping data. The intersection of visual cortical area with the peripheral activation zone defined the ROIs for which attentional modulation was assessed. Frontal eye-fields were defined by the saccade vs. no saccade contrast of the guided saccade task.

7.3 Results

The task was demanding for both monkeys, but they performed much better than chance. They detected the prolonged target motion event in 70% (monkey Q) and 51% (monkey M) of the trials; and when they detected this event, they generated a saccade in the correct direction in 97% and 86% of these trials (chance level = 12.5%, details in 7.5.1).

During scanning, blocks of attention task trials were separated by periods of passive fixation (Figure 23C). Comparing these two conditions allowed identification of task-related activations. Performance of the attention task activated cortical regions in occipital, temporal, parietal and, to a lesser extent, frontal lobes more than fixation (Figure 24A, Figure 25). Activity in somato-sensory and primary motor cortex (blue areas in Figure 24A and Figure 25) was reduced during task performance, likely due to less frequent licking during attention than fixation periods.

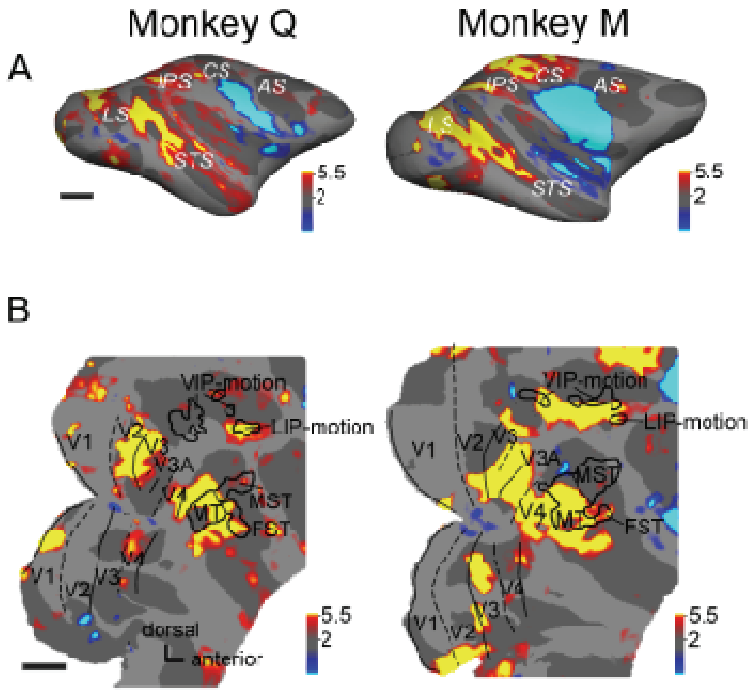

Figure 24: Task-related activition. (**A**) Statistical parametric maps of the contrast 'task vs. fixation' overlaid on the inflated right hemispheres of monkey Q (left) and monkey M (right). Yellow- and red coloured regions were significantly more activated by performance of the attentive motion discrimination task than by passive fixation at $P < 0.005$, corrected for multiple comparisons, while cyan-blue regions were significantly less active during active task performance. Scale bar is 1 cm. LS, lunate sulcus; IPS, intraparietal sulcus; STS, superior temporal sulcus; CS, central sulcus; AS, arcuate sulcus. (**B**) Same parametric maps overlaid on flattened posterior hemispheres. Dashed and solid black lines mark vertical and horizontal meridians, respectively. Motion-sensitive areas (black outlines) in the STS, MT, MST, and FST, and in the intraparietal sulcus, VIP and LIP, were mapped with a motion localizer (see Methods, and 7.5.2). FST, fundus of the STS; MT, middle temporal area; MST, medial superior temporal area.

In order to assign task-related activity to specific visual areas, we performed four localizer experiments. First, we charted retinotopic visual areas using meridian mapping with a checkerboard stimulus (Sereno, Dale et al. 1995; Vanduffel, Fize et al. 2002) (Figure 32A and Figure 33A). Second, we defined motion-sensitive regions by comparing activity to moving versus static random dots (Figure 32B and Figure 33B). Third, we scanned monkeys while they performed a guided saccade task to identify regions involved in saccade generation (Figure 32C and Figure 33C). Finally, we devised an eccentricity localizer that differentiated regions responding to the random dot surfaces from regions responding to fixation spot and spatial cue (Figure 32D and Figure 33D). With these four localizers we could identify cortical areas and assign task-performance related activity to them. Figure 24B shows areas activated by the attention task compared to fixation on a flattened map of the posterior half of the brain,

with visual area boundaries determined by retinotopy and motion localizers superimposed. We found task-related modulation in occipital areas V1, V2, V3, V3A, V4d, V4t, in superior temporal sulcus (STS) areas MT, FST, MST (weak), in the dorsal part of posterior inferotemporal cortex (PITd (Hikosaka 1997)), in intraparietal areas VIP and LIP, and in the FEF. Thus, attention modulated areas at all levels of the visual hierarchy, areas of the dorsal and areas of the ventral stream, and including areas like V4 and MT in which receptive field sizes matched the size of the target stimulus, as well as areas with much smaller receptive fields, e.g. V1 and V2, or larger ones, e.g., PITd.

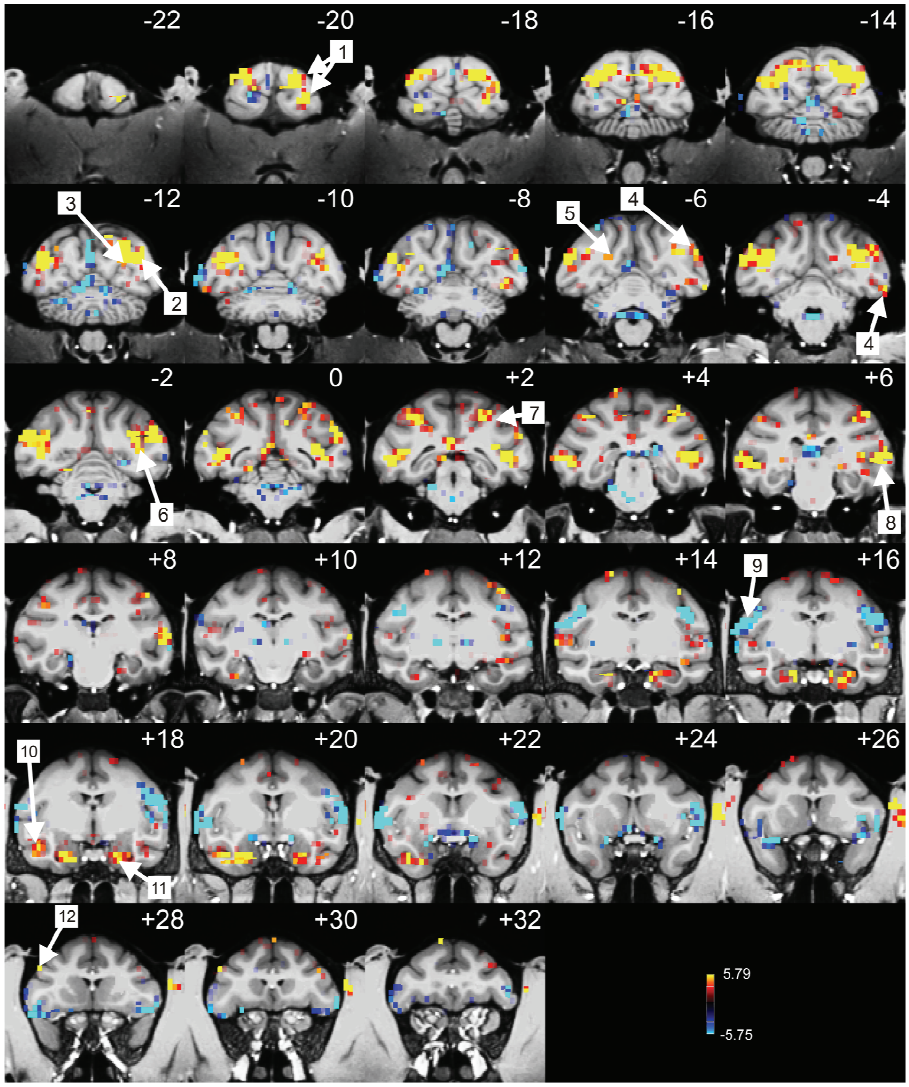

Figure 25: Visual activation in monkey Q. Statistical parametric maps, contrasting periods of task performance with periods of fixation, projected onto coronal slices of a high-resolution anatomy. Cyan and blue regions indicate less activity during task performance while red and yellow regions indicate higher activity. Numbers depict cortical areas: 1, area V1; 2, area V2; 3, area V3; 4, area V4; 5, area V3a; 6, area MT; 7, area LIP; 8, PITd; 9, somatosensory areas 1+2; 10, area TE; 11, amygdala; 12, FEF.

These differences in activation during attentive task performance, compared to passive fixation, could be due to differences in sensory stimulation, attention requirements, eye movements (saccades), or interactions between these components. To find the subset of cortical areas whose activation depends on the spatial direction of attention, we contrasted the two spatial attention conditions ‘attend contralateral’ and ‘attend ipsilateral’. These spatial attention conditions were dissociated from saccade planning, since saccades to any of the eight different targets were generated equally frequently in both conditions. In V1, the attended surface induced significantly larger activity than the distracter (Figure 26 & Figure 27). Similarly, the target activated extrastriate areas V2, V3, V3A, V4d, and V4t more than the distracter did. Spatial effects of attention were similarly strong in motion-selective area MT, but weak or absent in neighboring area MST. In motion-sensitive area FST, which, like MST, is intimately connected with MT (Ungerleider and Desimone 1986; Boussaoud, Ungerleider et al. 1990), the effect of spatial attention was inconsistent across monkeys (strong in M, weak in Q, Figure 27). Most surprisingly, we found strong effects of spatial attention in both animals further anterior in the STS in inferotemporal area PITd. Within the parietal lobe, we found attentional modulation of comparable magnitude and consistent across animals in area LIP only. Area 7a showed no effect of attention, area VIP a moderate one in one animal (monkey M) and not in the other. Within the frontal lobes, in which task performance led to only weak modulation (Figure 24A), we found spatial attention to modulate the FEF.

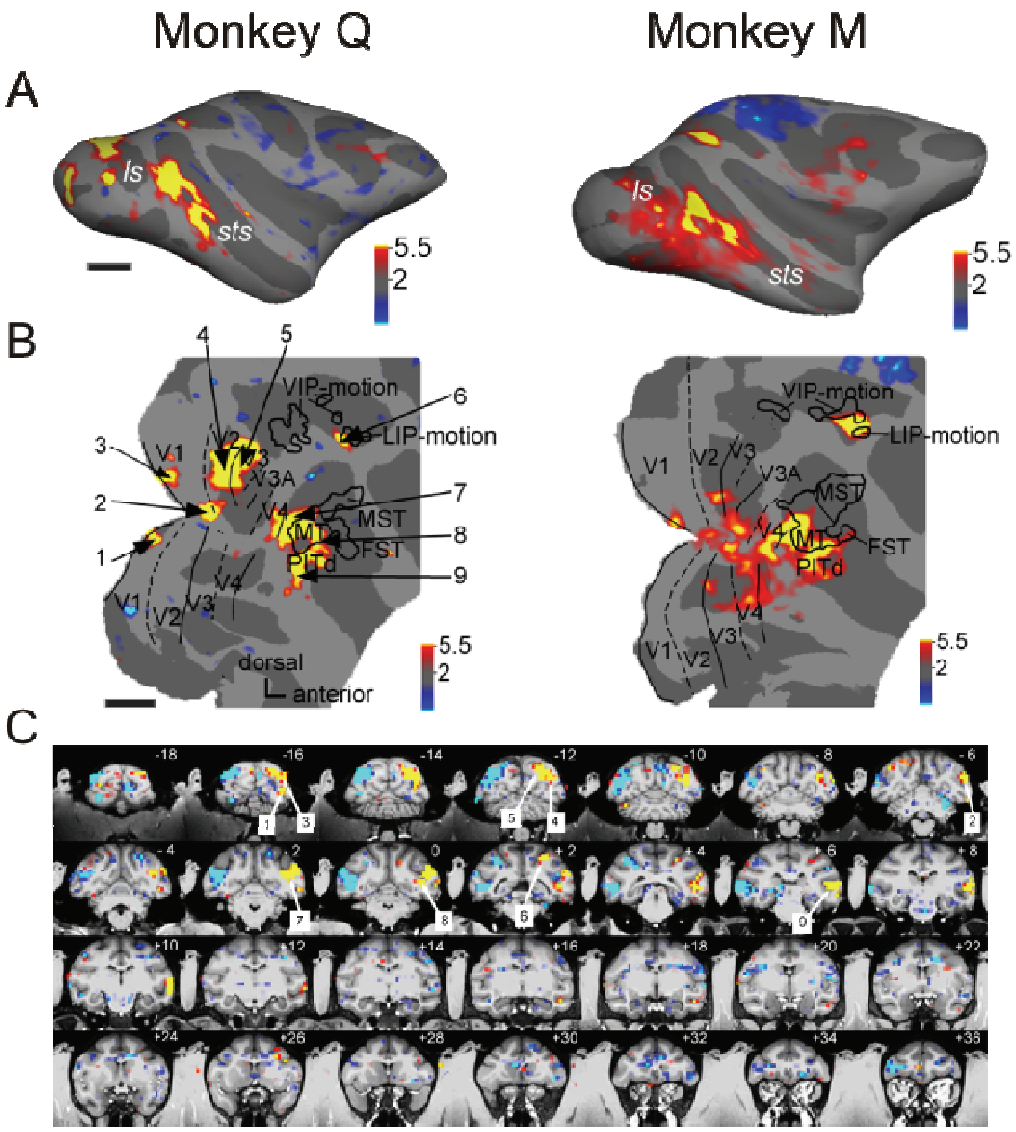

Figure 26: Spatial attentional modulation. (**A**) Statistical parametric maps for the contrast 'attend contralateral vs. attend ipsilateral'. Conventions as in Fig. 2A. (**B**) Same parametric maps overlaid on flattened posterior hemispheres. Conventions as in Fig. 2B. Left: numbers point to regions of significant activation shown on coronal slices in (**C**): 1, area V1 lower hemifield; 2, foveal representation; 3, area V1 upper hemifield; 4, area V2; 5, area V3; 6: posterior parietal area LIP; 7: V4t; 8: area MT; 9: area PITd) (C) Parametric maps on coronal slices of high-resolution anatomy, left hemisphere on the left. Cyan and blue indicate higher activity for right than left side target, yellow and red indicate higher activity for left than right target.

The patterns of task-related activation (Figure 24) and of spatial attentional modulation (Figure 26 & Figure 27) of posterior parietal and prefrontal cortical areas suggest LIP is the prime source of control for sustained attention. We wondered whether the relatively weak FEF activation in this task, was caused by the dissociation of saccade generation from the preceding, sustained deployment of spatial attention required for correct task performance. Overtraining of a task with this requirement may, in the course of months, reduce the influence of attentional control by the FEF, a structure which may support only brief deviations of the focus of spatial attention from the direction of saccade preparation (Juan, Shorter-Jacobi et al. 2004). We therefore trained monkeys on a second attention task in which

the directions of sustained attention and saccade planning coincided. Monkeys had to detect the occurrence of a 600 ms low-coherence motion event that interrupted an incoherent motion sequence and report that event by a saccade onto the target surface (Figure 28A). After two (monkey Q) and three weeks (monkey M) of training on this new task, monkeys were scanned while performing blocks of trials of the new motion detection task interleaved with blocks of the old motion discrimination task (Figure 28B). The two attention tasks differed in three major ways: attention saccade contingency, feature discrimination vs. detection, and feature strength (coherent vs. incoherent motion). Comparing activity across these two tasks allowed us to address two further questions about the nature of the attention effects we had seen: is the involvement of ventral stream area PITd in attentive motion processing specific to the process of feature discrimination? Does the involvement of dorsal-stream motion areas MT, MST, FST and VIP depend on the strength of motion signals?

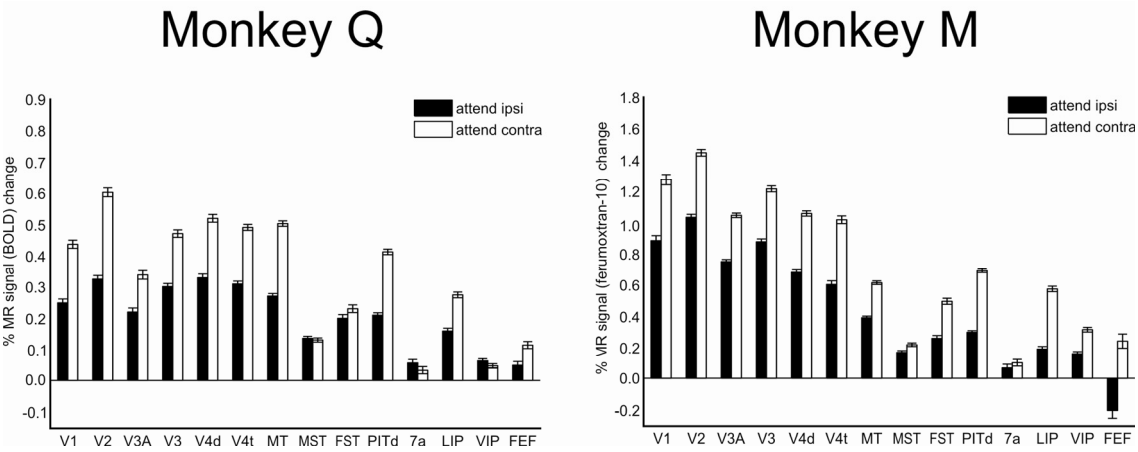

Figure 27: Bargraph representation of the percent signal change in selected cortical areas. Data is merged over both hemispheres for each subject separately (Q left, M right). Percent MR signal change is calculated for each ROI relative to the baseline preceding each block.

After training on the motion detection task, both monkeys reacted much more slowly than before (417 ms vs. 285 ms in monkey Q, and 366 ms vs. 296 ms in monkey M, respectively, during discrimination), but reaction times during detection and discrimination task were similar (see 7.5.3). This massive increase in response latency likely reflects the heightened requirements for cognitive control in a behavioral setting requiring two different kinds of attentive sensory-motor-mapping as well as the necessity, in the detection task, to integrate sensory evidence for the occurrence of the low coherence motion over time (Shadlen and Newsome 1996). We found area MT to be activated and modulated by spatial attention in both tasks, while motion-selective areas MST, FST and VIP were much more weakly, if at all, activated and modulated by the attention condition (Figure 29). In contrast, area PITd was

robustly activated and modulated by spatial attention in both tasks. Areas LIP and FEF were activated to a similar degree in both tasks, with the attentional modulation of area LIP larger than that of the FEF. Thus, the pattern of areas recruited by spatial attention depended neither on specific task requirements nor on the strength of motion signals.

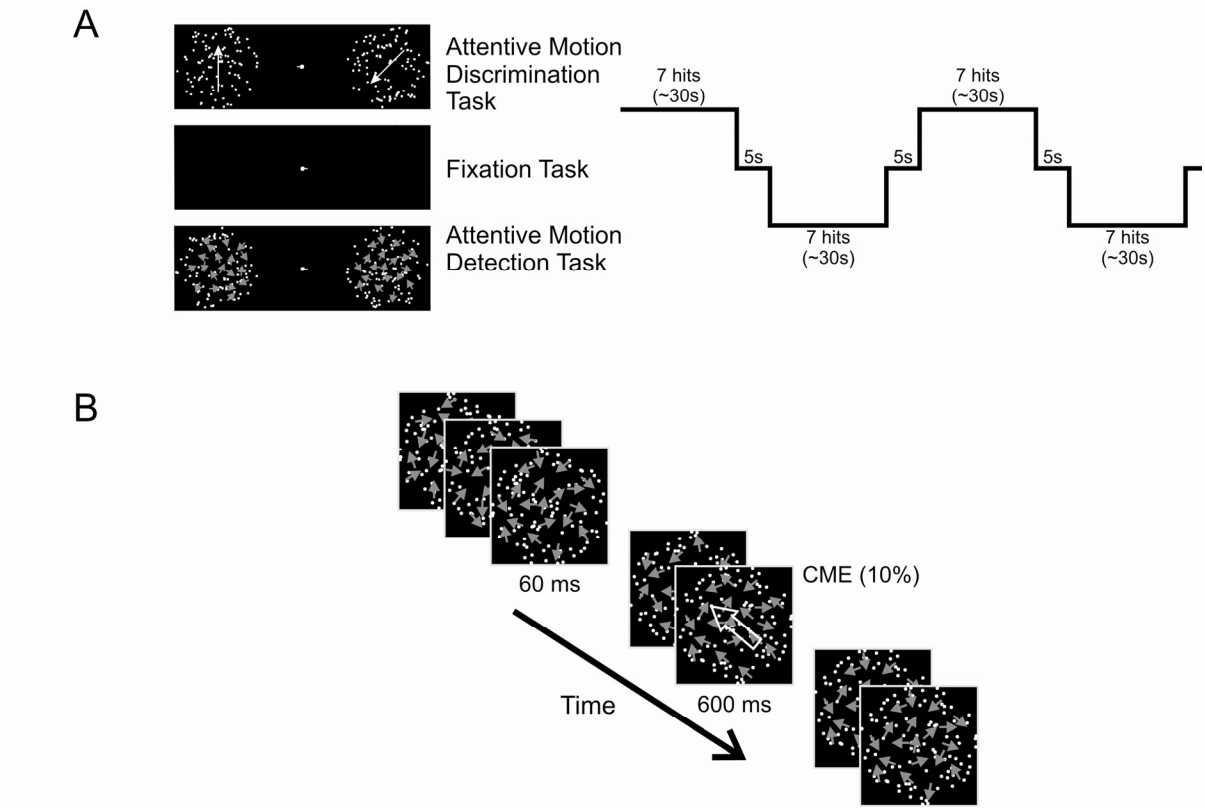

Figure 28: Design of attentive motion-detection task. (**A**) Temporal sequence of behavioural conditions (attentive motion-discrimination, passive fixation, and attentive motion-detection) during scanning. Monkeys alternated between performing the attentive motion-discrimination task, passive fixation condition (center) and the attentive motion-detection task (bottom). The sequence of behavioural conditions was 'IFEFIFE…', with 'I' 'discriminate', 'F' 'fixate', and 'E' 'detect'. During attentive discrimination and attentive detection monkeys had to complete seven trials successfully ('hits'). Thus, duration of these blocks was variable, on average ~30 seconds. Passive fixation blocks separated the two attention conditions and lasted 5 seconds. (**B**) Event sequence of one of the RDSs. The sequence of events was identical to that in the motion-discrimination task (Fig. 1B), but motion during the brief 60ms periods was incoherent, while motion during the task-relevant prolonged motion event was 10% coherent (CME, coherent motion event). Monkeys were required to respond to the target surface CME by a saccade onto the target surface, irrespective of motion direction.

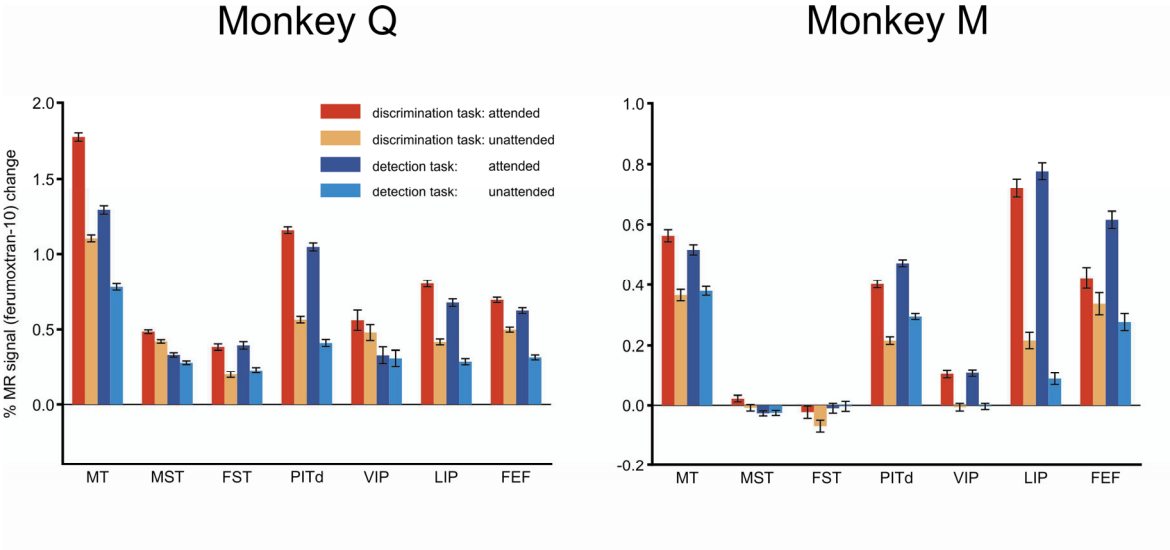

Figure 29: Attentional modulation in motion discrimination and motion detection tasks. Percentage MR signal change is shown for the subset of cortical areas covered by the slice prescription used for this experiment.

7.4 Discussion

The composition of the network of areas subserving spatial attention has four striking characteristics. First, within fronto-parietal regions, the network is confined to areas LIP and FEF. Both areas have been proposed to contain salience maps (Gottlieb, Kusunoki et al. 1998; Bisley and Goldberg 2003; Thompson, Bichot et al. 2005) that would be important for directing attention. Second, attention modulated cortical activity as early as primary visual cortex. This finding agrees with reports in human subjects (Huk and Heeger 2000), ruling out the possibility that the sites of attentional modulation in man and macaque differ systematically (Heeger, Gandhi et al. 2001). Furthermore, this finding agrees with the theoretical prediction that high perceptual load should lead to early attentional selection (Lavie and Tsal 1994).

Third, we found many areas containing direction-selective cells to be modulated by attention. This is predicted by feature-based accounts of attention (Corbetta, Miezin et al. 1990; Beauchamp, Cox et al. 1997; O'Craven, Rosen et al. 1997), because attention was directed to motion. Yet attention effects in motion-sensitive areas MST, FST and VIP, located at high levels of the visual cortical hierarchy (Felleman and Van Essen 1991), were inconsistent across animals, weak, or absent. This is surprising in light of several earlier reports on an increase in the strength of attention effects with hierarchical rank (Motter 1993; Treue and Maunsell 1996; Mehta, Ulbert et al. 2000; Cook and Maunsell 2002). Thus the strength of the attentional modulation of neural activity measured by fMRI (Logothetis, Pauls et al. 2001; Niessing, Ebisch et al. 2005) does not necessarily grow with increasing hierarchical rank. Rather the site of attentional selection can be sharply focused, even separating areas with

close anatomical links like MT and MST (Ungerleider and Desimone 1986; Boussaoud, Ungerleider et al. 1990).

Fourth, we did find attentional modulation in a high-level cortical area, but one hitherto not suspected to be involved in attentive motion processing. This inferotemporal cortical area, PITd, belongs to the ventral stream. PITd cells have large receptive fields often encompassing parts of both visual hemifields (Kobatake and Tanaka 1994; Hikosaka 1997). Thus the two random dot surfaces in our paradigm could have competed for processing resources in this area, and attention could have biased this competition in favor of the target (Desimone and Duncan 1995; Duncan, Humphreys et al. 1997). Neighboring areas V4 and TEO have been suggested to exert a spatial attentional filtering function (De Weerd, Peralta et al. 1999). PITd may also implement an attentional filter, yet it will take electrophysiological recordings targeted to this area to reveal the precise role PITd neurons play in attentive motion processing.

The grouping of cortical areas into two processing streams (Ungerleider and Mishkin 1982) is a basic organization principle of the primate visual system. Different functional characterizations of the two streams have been given, one associating the dorsal stream with motion processing and the ventral stream with color and shape processing (Livingstone and Hubel 1988), a second one emphasizing the representation of stimulus location in dorsal, and stimulus identity in ventral stream (Ungerleider and Mishkin 1982), and a third one postulating the dorsal stream to serve vision for action and the ventral stream vision for perception (Goodale and Milner 1992). In our two tasks, motion processing was clearly important, while shape and color were not. Thus, if feature domain determined the locus of attentional selection, we should have found strong attentional modulation of the dorsal, not the ventral path. Since we found the opposite, feature domain was not the determinant for the locus of attentional selection. In both tasks, the spatial location of the target stimulus had to be selected, which explains activation and attentional modulation of LIP (Colby and Goldberg 1999), yet precise stimulus localization was not necessary. In contrast, stimulus identity, determined neither by color nor shape, but by motion, was of central relevance. Thus this motion within a static stimulus constituted an object attribute very much like color or shape (Ferrera, Rudolph et al. 1994). And this object attribute did not directly guide action, but rather it had to be brought into a perceptual format useful for task performance (Mirabella, Bertini et al. 2007). These are the characteristics of vision for perception, proposed to be a ventral stream function (Mirabella, Bertini et al. 2007). Thus our results suggest that it is *how* a visual feature is processed rather than *what kind* of feature that determines the locus of

attentional selection. For understanding the mechanisms by which attention selects visual features, area PITd, likely, is key.

7.5 Appendix

7.5.1 Behavioral Performance

7.5.1.1 Attentive Motion Discrimination Task

The attentive motion discrimination task (Figure 23) was demanding for both monkeys, Q and M. But both performed far better than chance. When tested outside the scanner, both monkeys completed more than 60% of trials correctly (Q=66.1%, M=64.5%, 14 days, 27461 trials monkey Q, 27 days, 44587 trials monkey M, Figure 30A and B, below). Only rarely were saccades generated to the wrong saccade target at the right time after the target prolonged motion event (PME), i.e. within 800ms after the start of the event (‘wrong’ errors, Q=1.7%, M=4.7%). Responses related to the PME on the distractor surface occurred in 6.6% (Q) and 2.9% (M) of trials (‘distractor’ in Figure 30). Thus, behavioral evidence suggests that both monkeys paid close attention to the cued side. By far the most frequent error category in both monkeys were responses that occurred before the target PME (‘early’ errors, Q=24.0%, M=27.2%), while only few trials were completed with too late a saccade after the PME (‘late’ errors, Q=0.3%, M=0.5%) or without any response (‘ignore’ errors, Q=1.3%, M=0.2%). Thus, both animals followed a strategy that erred on making too quick responses. Combining the different error sources, performance can be expressed in the following way: monkeys detected the PME in 67.7% (monkey Q) and 69.2% (monkey M) of the trials, respectively. When they detected this event, they generated a saccade in the correct direction in 97.5% and 93.3% of these trials (chance level=12.5%). Monkey Q, by these measures, performed better than monkey M.

During scanning, both monkeys continued to perform better than chance, but performance was substantially lower than outside the scanner due to a marked increase of ‘early’ errors (Figure 30C and D, below, hits: Q=66.7%, M=43.7%, 15 days, early errors: Q=25.6%, M=46.5%, 5644 trials monkey Q, 11534 trials monkey M). Other sources of errors were similarly low as outside the scanner (‘wrong’: Q=2.2%, M=7.5%, ‘distractor’: Q=3.6%, M=2.1%, ‘late’: Q=0.2%, M=0.3%, ‘ignore’: Q=0.4%, M<0.1%). Monkeys detected the PME in 70.3% (monkey Q) and 51.2% (monkey M) of the trials, respectively. When they detected this event, they generated a saccade in the correct direction in 96.7% and 85.6% of these trials. Thus during scanning monkey Q performed better than monkey M.

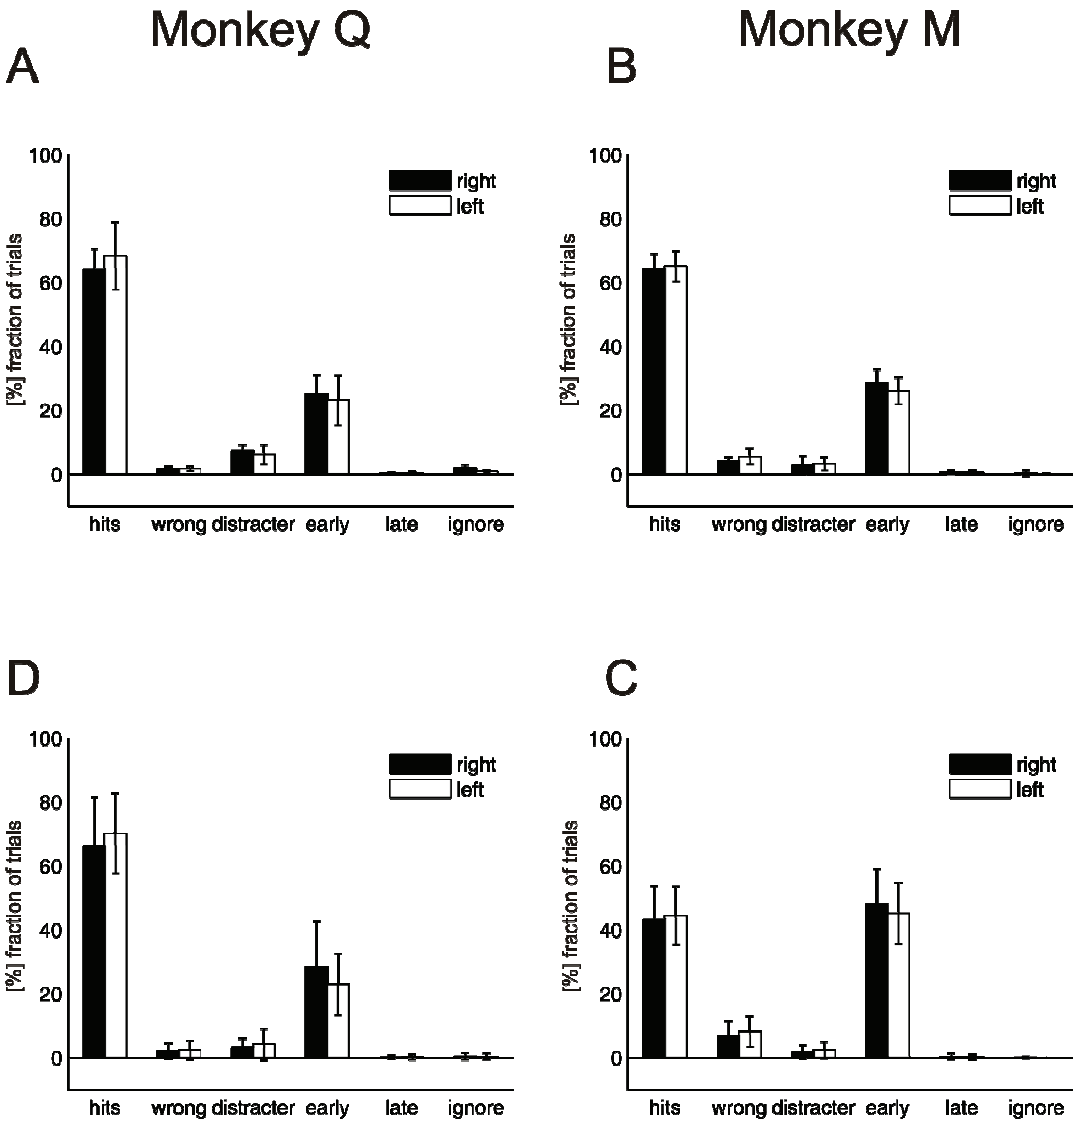

Figure 30: Performance of monkeys Q (left) and M (right) during training outside the scanner (**A** and **B**) and inside the scanner during acquisition of functional MR data (**C** and **D**). Error bars indicate the standard deviation from the mean. Black and white bars represent fraction of trials for the different positions of the behavioural relevant surface, to the right or the left from the fixation spot respectively**.**

7.5.1.2 Attentive Motion Detection Task

After monkeys had been trained on the attentive motion detection task, they performed this task in alternating blocks with the attentive discrimination task, while they were being scanned (3 days, 2589 trials monkey Q, 4 days, 3624 trials monkey M). Performance of both tasks is detailed in Figure 31. Monkey Q performed better than 90% in both tasks, monkey M's performance was 50.2% in the motion discrimination task and 72.9% in the motion detection task. As before, the main sources of error were responses made before the PME occurred on the target surface (Q=6.0%, M=31.4%). Thus, both monkeys performed better in the motion discrimination task than they had before (Figure 30).

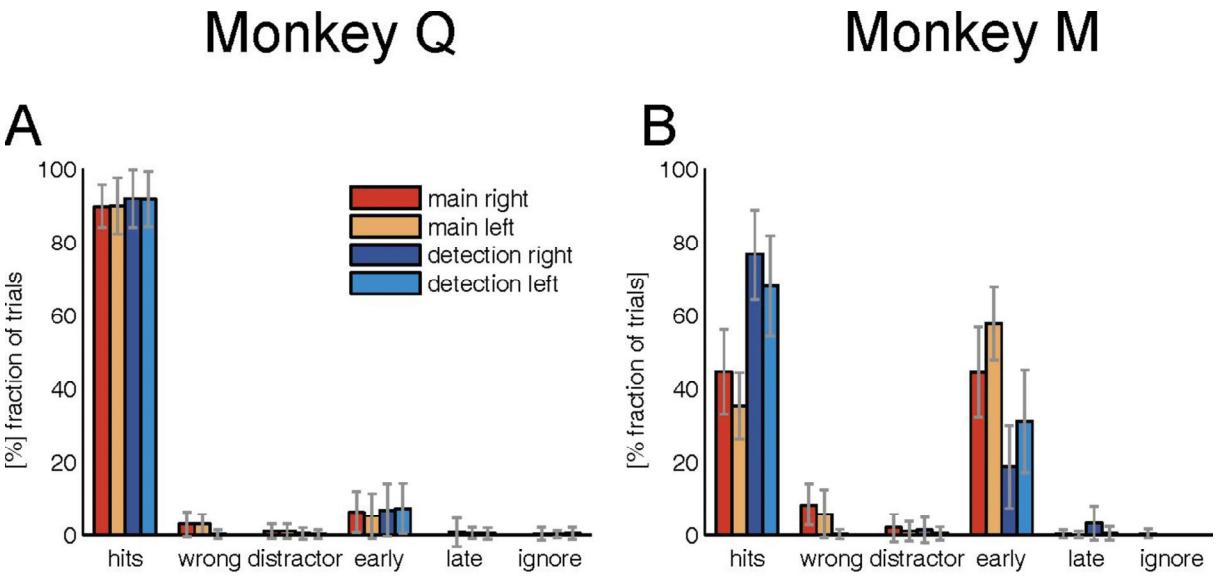

Figure 31: Performance of monkeys Q (left) and M (right) during performance of the attentive motion detection task inside the scanner (**A** and **B**). Error bars indicate the standard deviation from the mean. Red and orange bars represent values for right and left position of the behavioural relevant surface performing the original discrimination task, blue and light blue for performing the attentive motion detection task.

7.5.2 Functional Brain Maps

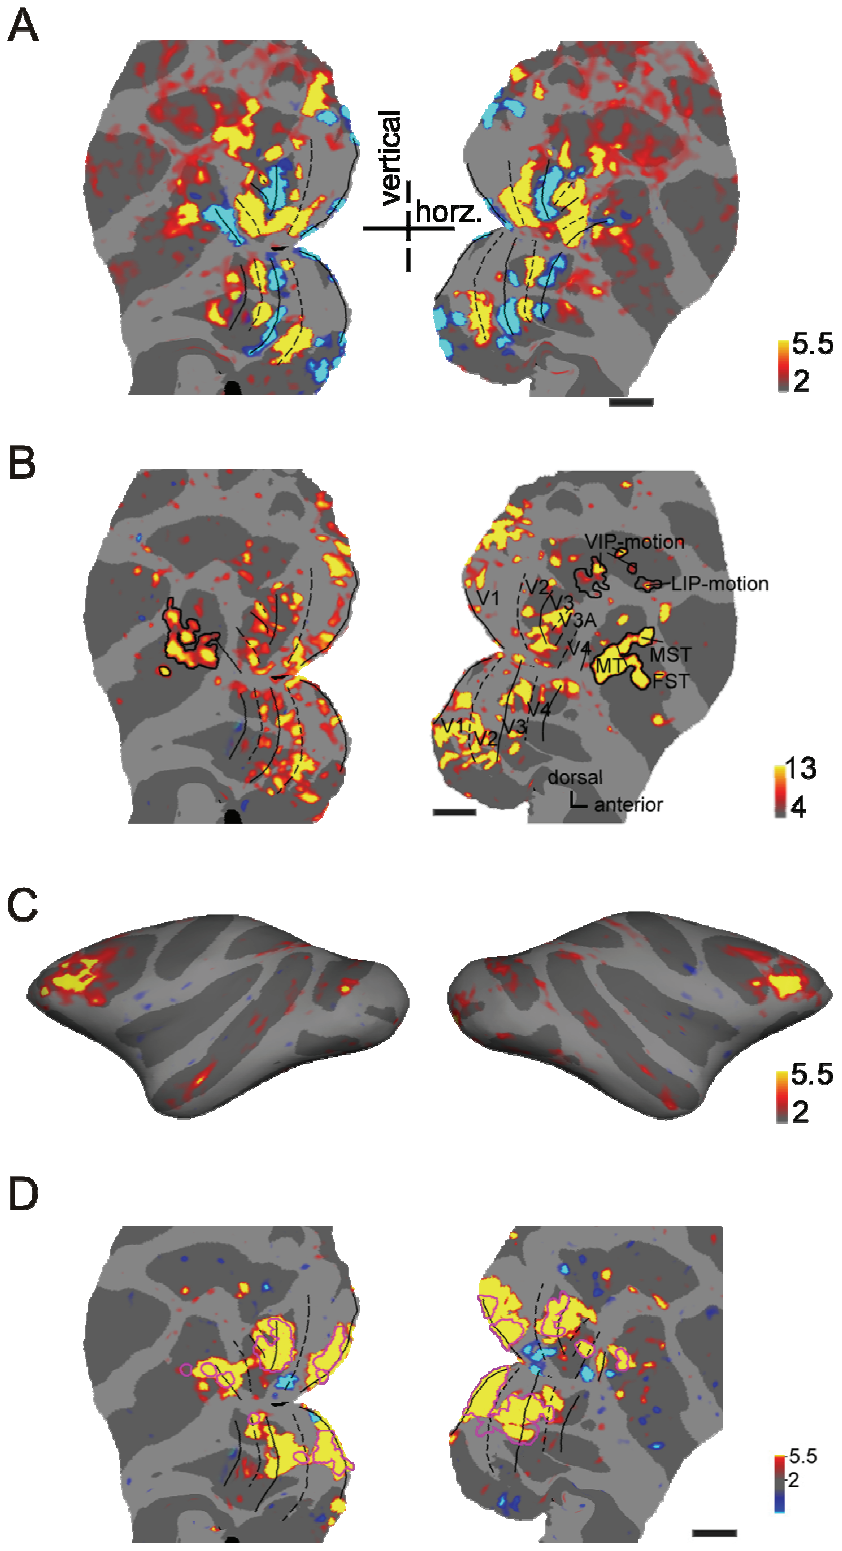

Figure 32: Monkey Q: Functional brain maps of monkey Q charted with different localization stimuli while passively looking (see Methods in 7.2). (**A**) Shows results from retinotopic mapping with a checkerboard stimulus. Yellow and blue areas depict vertical and horizontal meridian representation, respectively. (**B**) Statistical parametric map gained by contrasting moving versus static stimuli. Yellow areas mark more activity during motion. (**C**) Saccade mapping. Yellow areas show more activity during periods executing saccades. (**D**) Eccentricity localizer. Blue regions depict brain areas representing the foveal visual field while yellow regions cover peripheral parts of the visual field. Eccentricity was chosen to match the eccentricity of the random dot surfaces in the main task.

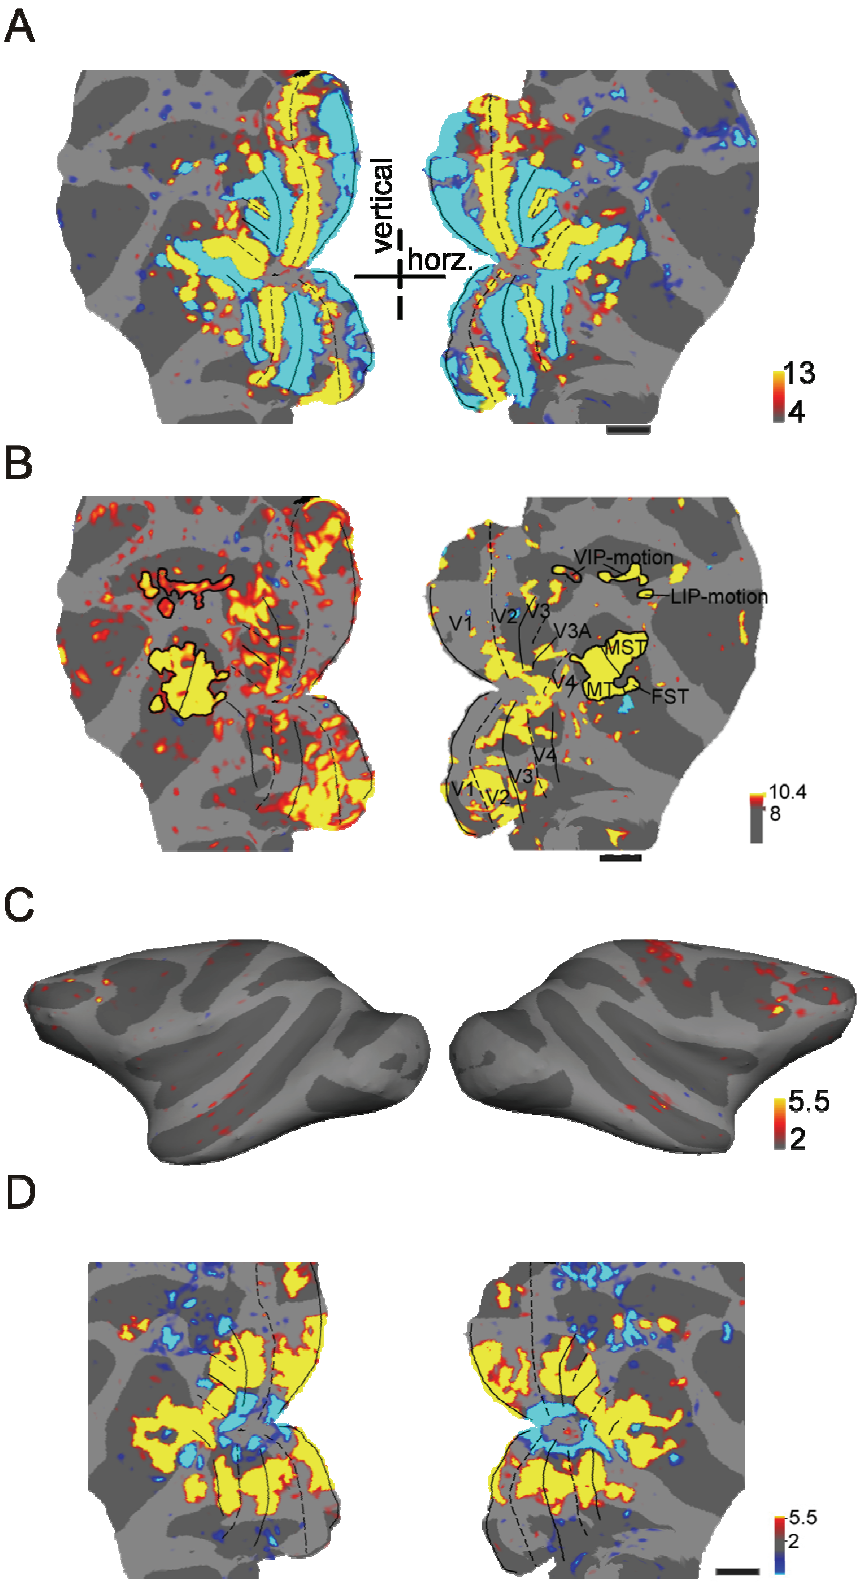

Figure 33: Monkey M: Functional brain maps for monkey M. Same convention as in **Figure 32**.

7.5.3 Reaction times

Reaction times during performance of the attentive discrimination task were, on average 285 ms for monkey Q and 296 ms for monkey M. These reaction times are measured with the beginning of the PME. Note that only 60ms after that time, was it actually beginning to be possible to differentiate the PME from the brief motion events it was embedded in.

Reaction times changed substantially with training of monkeys on the attentive motion detection task. When monkeys performed this task in alteration with the motion discrimination task, reaction times were much longer in both tasks. Monkey Q reacted, on average, 417 ms after PME start in the motion discrimination task and 415 ms after CME onset in the motion detection task. There was no statistically significant difference between the two (Mann-Whitney U-test, p=0.55). Monkey M reacted, on average, 366 ms after the PME start in the motion discrimination task and 372 ms after CME onset in the motion detection task. Again, this difference was not statistically significant (Mann-Whitney U-test, p=0.90). The differences in reaction times in the motion discrimination task *before* and *after* training (Figure on the left) were however, highly significant for both monkeys (Mann-Whitney U-test, p<0.001).

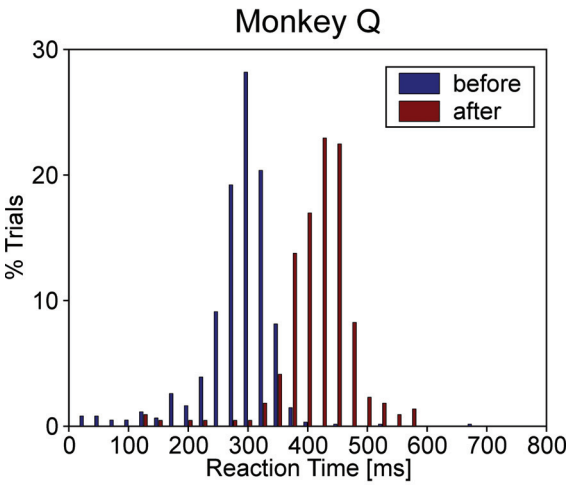

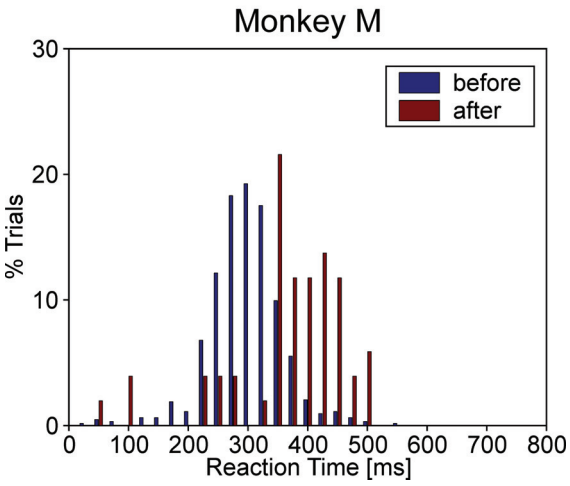

There are several possible reasons for why reaction times were rather slow after training of the motion detection task. First, performance of the motion detection task necessitated temporal integration of sensory evidence that a low coherence motion event actually occurred (Shadlen and Newsome 1996). Both monkeys could have taken all the time the CME lasted, 600ms, to do this, but rather both responded faster (Kiani, Hanks et al. 2008). Second, both tasks differed in how sensory information was to be related into a response. In the motion discrimination task, a saccade had to be withheld until the PME occurred. Then, the motion direction had to be identified in order to determine the response direction (which was not identical to the motion direction). In contrast, in the motion detection task, saccades could be planned from the cue presentation on. But when the CME

occurred, its motion direction was not to be used for determining or changing saccade direction. Thus the two tasks differed both in terms of timing of response planning, and in terms of relating motion direction into a response. Since the two tasks were performed in alternating blocks (Figure 28A), monkeys had to make sure the correct rule was applied to the currently presented stimulus material, resulting in heightened cognitive demand to control for correct task performance. Third, in case of the motion discrimination task, taking a longer time with the response also allows integrating sensory information and should allow for better detection that the PME had occurred and of the correct direction. In fact, in both monkeys performance in the motion discrimination task increased after training of the motion detection task, from 67.73% to 89.72% for Monkey Q and from 43.73% to 50.23% for Monkey M.

8 Conclusion

By applying a number of diverse theoretical and experimental methods I have investigated the neural representation of dynamic motion stimuli. The results of my experiments show that the neural responses in macaque area MT neurons are very reliable in representing the visual environment. Directional filter properties are robust and do not adapt to different stimulus statistics on short timescales. However, Kohn and Movshon (Kohn and Movshon 2003; Kohn and Movshon 2004), found a prolonged presentation of a single motion direction to induce substantial adaptation. Their experiments revealed reduction of direction-tuning bandwidth. Future research will need to identify the precise stimulus patterns to which cortical area MT neurons adapt and to which they do not.

My subsequent results show that the directional filter properties of MT neurons are subject to changes due to the employment of attention. An attended stimulus elicits a stronger response when located inside the receptive field compared to outside the receptive field with respect to the tuning curve. This effect has also been shown by other studies (Treue and Maunsell 1996; Wegener, Freiwald et al. 2004) and can be explained by gain modulation of the response. However, considering the high perceptual load put onto the visual system by our visual stimulus, I was expecting a bigger effect. Further analysis of the data needs to show whether there are additional effects in the directional filter of MT neurons. Posner and colleagues (Posner, Snyder et al. 1980) showed for example that subjects are significantly faster in responding to a cued location by an average of 50 ms. This advantage in timing should be presented somewhere in the neural response.

Using fMRI I have been able to find the areas subserving spatial attention. Attention modulated cortical activity as early as primary visual cortex. This result is in agreement with findings in human subjects (Huk and Heeger 2000). Furthermore, this finding agrees with the theoretical prediction that high perceptual load should lead to early selection (Lavie and Tsal 1994). Many areas containing motion-selective areas have been found to be modulated by attention. But motion sensitive areas MST, FST and VIP did not show consistent results across animals. This is especially surprising in the light of earlier reports on an increase in modulation effects with hierarchical rank (Lavie and Tsal 1994; Treue and Maunsell 1996; Mehta, Ulbert et al. 2000; Cook and Maunsell 2002). In the fronto-parietal network, activation was found in area LIP and FEF. Both areas have been proposed to contain salience maps (Gottlieb, Kusunoki et al. 1998; Bisley and Goldberg 2003; Thompson, Bichot et al. 2005) that can be used to direct attention. Additionally attentional modulation has been found

in a high-level cortical area in the temporal cortex. This cortical area, PITd belongs to the ventral stream has not been suspected to be involved in attentive motion processing.
It will now take electrophysiological recordings to investigate the precise nature of attention modulations within these areas. For area MT, this data is already available. The conflicting evidence for parietal areas MST, FST and VIP needs to be resolved in the same way. Furthermore, it needs to be evaluated what the neurons of ventral stream area PITd contribute to attentive motion processing. Future research also needs to reveal how areas within this network are interacting. between areas; which areas are exerting control and which areas are controlled?

Having access to a map of the network of cortical areas subserving spatial attention therefore offers a great opportunity to further investigate how the brain performs cognitive processes and how distributed cortical areas in different hierarchical position interact with each other.

9 Summary / Zusammenfassung

9.1 Summary

Several theoretical and experimental methods are employed to investigate the neural representation of dynamic motion stimuli. Starting out with an investigation of the behaviour of macaque area MT neurons, the key processing stage of visual motion, the current study presents single cell electrophysiological data that shows that the representation of dynamic motion stimuli is reliable over different stimulus statistics. Employing rapidly changing motion stimuli with different statistics of directional changes, performing extracellular recordings in area MT, reverse correlating spiketrains with the motion impulse sequence and computing direction tuning curves out of the resulting spike-triggered averages provides evidence that neural tuning curves obtained with random stimulus sequences can be directly used to predict neural tuning curves in quite different stimulus contexts (chapter 4).

Extending this view on neural reliability, a reconstruction and modelling framework is applied to predict neural responses of area MT neurons. Reconstruction is useful in estimating how much information about a physical variable is present in the activity of a neuronal population. It is shown that neurons in area MT do not adapt to the statistics of dynamical stimuli, neither in direction nor in time domain (chapter 5).

Having shown the reliability of neural coding for different stimulus statistics the modulatory influence of visual attention on spike-triggered averages and tuning curves is investigated (chapter 6). Using a visual stimulation paradigm similar to the one that has been used for the previous experiments, combining it with a spatial attention task; it is shown that attention alters the neural activity in macaque area MT in a multiplicative fashion, providing evidence in favour of a 'gain modulation' effect of attention (McAdams and Maunsell 1999; Treue and Martinez Trujillo 1999).

In a final functional imaging study (chapter 7) the functional network of cortical areas involved in representation of dynamic motion stimuli is presented. Two macaque monkeys have been trained to perform an attention task similar to the one described in chapter 6 within a 3T MR scanner. Investigation of the blood-oxygen-level-dependent signal change during performance of the attention task identified the whole set of brain areas modulated by spatial attention. In agreement with early selection accounts of visual attention (Broadbent 1958), modulated activity is found as early as V1 and continued along both parietal and temporal pathways.

9.2 Zusammenfassung

In dieser Arbeit werden theoretische und experimentelle Techniken verwendet, um die neuronale Repräsentation von dynamischen Bewegungsreizen zu untersuchen. Mit einer Untersuchung des Verhaltens einzelner Neurone im Hirnareal MT, der zentralen Verarbeitungsstufe für bewegte Reize, von Makaken wird gezeigt, dass die neuronale Repräsentation von dynamischen Bewegungsreizen über verschiedene Stimulusstatistiken hinweg sehr verlässlich ist. Unter Verwendung eines sich mit verschiedenen Statistiken schnell ändernden Bewegungsreizes, der Durchführung extrazellulärer Ableitungen im Areal MT, der Verwendung von Korrelationsmethoden und der Berechnung von Antwortprofilen (Tuningkurven) werden Hinweise präsentiert, dass die neuronalen Antwortprofile, die mit einer Zufallssequenz von Bewegungsrichtungen gewonnen wurden, direkt für die Vorhersage von neuronalem Antwortverhalten in einem sehr verschiedenartigem Stimuluszusammenhang verwendet werden können (Kapitel 4).

Durch die Verwendung eines Rekonstruktions- und Modellierungsansatzes wird das Verständnis der Verlässlichkeit von neuronalem Antwortverhalten weiter vertieft. Rekonstruktion ist ein nützliches Hilfsmittel, um die Menge der Information in der Antwort einer Gruppe von Neuronen über eine physikalische Variable bestimmen zu können. Es wird gezeigt, dass die Neurone im Areal MT sich in ihrem Antwortverhalten nicht an die Statistik der visuellen Reize anpassen, sondern sowohl in der Zeit als auch über die Richtungen hinweg verlässlich antworten.

Nachdem die Verlässlichkeit des Antwortverhaltens der Neurone gezeigt wurde, wird der modulierende Einfluss von visueller Aufmerksamkeit auf die Antwortprofile von einzelnen Zellen untersucht. Indem der oben verwendete visuelle Stimulus mit einer Aufmerksamkeitsaufgabe kombiniert wurde, wird gezeigt, dass Aufmerksamkeit das Antwortprofil von Neuronen im Areal MT auf multiplikative Weise verändert (Kapitel 6). Dieser Befund unterstützt die „gain modulation“ – Theorie von Aufmerksamkeitseffekten (McAdams and Maunsell 1999; Treue and Martinez Trujillo 1999), die besagt das die Antwortprofile von Neuronen, die ein verhaltensrelevantes Objekt und/oder Ort repräsentieren multiplikative verstärkt werden.

Im abschließenden Kapitel wird unter Verwendung von funktioneller Kernspintomographie (fMRT) das gesamte Netzwerk von Hirnarealen untersucht, das in die Verarbeitung dynamischer Bewegungsreize involviert ist. Während der Durchführung einer Aufmerksamkeitsaufgabe, wurde die Abhängigkeit der Hirnaktivität von Rhesus-Affen auf

Bewegungsreize und Aufmerksamkeit erfasst. Die Ergebnisse zeigen ein spezifisches Erregungsmuster von verschiedenen Hirnarealen, die an der Durchführung dieser Aufgabe beteiligt sind. In Übereinstimmung mit der Hypothese einer frühen Auswahl von visuellen Reizen für die weitere Verarbeitung (Broadbent 1958), wird schon in dem ersten kortikalen Areal, V1, eine Aufmerksamkeitsmodulation der Antwort gefunden. Entlang beider visueller Verarbeitungspfade, dorsal und ventral, finden sich weitere Areale, die durch die unterschiedlichen Aufmerksamkeitsanforderungen moduliert werden.

10 Acknowledgements

This work would not have been possible without support of my colleagues, friends and my family. I am deeply indebted to them.

I thank my supervisor, Prof. Dr. Andreas Kreiter, for giving me the opportunity to perform my experiment and for providing the necessary infrastructure.
Thanks to my advisor and promoter Winrich Freiwald for advice, discussions, helpful comments and support. Winrich Freiwald gave me the freedom to design and perform experiments and offered several opportunities to present the resulting data at meetings.
I would like to thank Doris Tsao. I have learned a lot from her and her permanent enthusiasm is an important source of inspiration.

MRI experiments took mainly place very early or very late during the day or even on the weekend. Aleksandra Nadolsky and Nicole Schweers several times bore the burden to support MRI measurements at any time. Thank you!

I thank Katrin Thoß and Ramazani Hakazimana for taking excellent care of the experimental animals.

I thank Mikhail Borisov for providing *Visiko*, electronic devices and support in software development. I am grateful to all the members of the Fachbereich 2 workshop, especially Wolfgang Fulda and Martin Nowak who supported the development of a new MRI-compatible primate chair and who provided excellent solutions to several mechanic problems.

I would like to thank my colleagues and friends Aurel Wannig and Sebastian Moeller for their assistance and encouragement. Special thanks to Aurel for motivating discussions in the wonderful alpine scenery during our hiking trips.

I would like to thank Bevil Conway and Tim Blanche for sharing exciting events while conducting their experiments next door and for fruitful discussions.

Significant contribution to this work has been provided by all the people who supported me in training of monkeys, Aleksandra Nadolsky, Israel Minguez, Tanja Munoz, and Anja Besuch,.

I would like to thank all other people of the Brain Research Institute, who have supported me in carrying out my experiments.

I want to thank my parents. Danke!

Last but not least I want to express my deep gratitude to my family for always supporting me. Thank you Elisabeth, Sophia and Katrin!

11 References

Albright, T. D. (1984). "Direction and orientation selectivity of neurons in visual area MT of the macaque." J Neurophysiol **52**(6): 1106-30.

Allman, J. M. and J. H. Kaas (1971). "A representation of the visual field in the caudal third of the middle tempral gyrus of the owl monkey (Aotus trivirgatus)." Brain Res **31**(1): 85-105.

Bair, W., C. Koch, et al. (1994). "Power spectrum analysis of bursting cells in area MT in the behaving monkey." J Neurosci **14**(5 Pt 1): 2870-92.

Bair, W. and J. A. Movshon (2004). "Adaptive temporal integration of motion in direction-selective neurons in macaque visual cortex." J Neurosci **24**(33): 7305-23.

Bair, W., E. Zohary, et al. (2001). "Correlated firing in macaque visual area MT: time scales and relationship to behavior." J Neurosci **21**(5): 1676-97.

Baizer, J. S. (1982). "Receptive field properties of V3 neurons in monkey." Invest Ophthalmol Vis Sci **23**(1): 87-95.

Baker, J. F., S. E. Petersen, et al. (1981). "Visual response properties of neurons in four extrastriate visual areas of the owl monkey (Aotus trivirgatus): a quantitative comparison of medial, dorsomedial, dorsolateral, and middle temporal areas." J Neurophysiol **45**(3): 397-416.

Barlow, H. B., R. Fitzhugh, et al. (1957). "Change of organization in the receptive fields of the cat's retina during dark adaptation." J Physiol **137**(3): 338-54.

Baylis, G. C. and J. Driver (1993). "Visual attention and objects: evidence for hierarchical coding of location." J Exp Psychol Hum Percept Perform **19**(3): 451-70.

Beauchamp, M. S., R. W. Cox, et al. (1997). "Graded effects of spatial and featural attention on human area MT and associated motion processing areas." J Neurophysiol **78**(1): 516-20.

Bisley, J. W. and M. E. Goldberg (2003). "Neuronal activity in the lateral intraparietal area and spatial attention." Science **299**(5603): 81-6.

Borghuis, B. G., J. A. Perge, et al. (2003). "The motion reverse correlation (MRC) method: a linear systems approach in the motion domain." J Neurosci Methods **123**(2): 153-66.

Boussaoud, D., L. G. Ungerleider, et al. (1990). "Pathways for motion analysis: cortical connections of the medial superior temporal and fundus of the superior temporal visual areas in the macaque." J Comp Neurol **296**(3): 462-95.

Bradley, D. C., G. C. Chang, et al. (1998). "Encoding of three-dimensional structure-from-motion by primate area MT neurons." Nature **392**(6677): 714-7.

Britten, K. H., M. N. Shadlen, et al. (1992). "The analysis of visual motion: a comparison of neuronal and psychophysical performance." J Neurosci **12**(12): 4745-65.

Broadbent, D. E. (1958). Perception and Communication. London, Pergamon Press.

Buschman, T. J. and E. K. Miller (2007). "Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices." Science **315**(5820): 1860-2.

Bushnell, M. C., M. E. Goldberg, et al. (1981). "Behavioral enhancement of visual responses in monkey cerebral cortex. I. Modulation in posterior parietal cortex related to selective visual attention." J Neurophysiol **46**(4): 755-72.

Busse, L., S. Katzner, et al. (2008). "Temporal dynamics of neuronal modulation during exogenous and endogenous shifts of visual attention in macaque area MT." Proc Natl Acad Sci U S A **105**(42): 16380-5.

Carandini, M., J. B. Demb, et al. (2005). "Do we know what the early visual system does?" J Neurosci **25**(46): 10577-97.

Chapin, J. K., K. A. Moxon, et al. (1999). "Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex." Nat Neurosci **2**(7): 664-70.

Chichilnisky, E. J. (2001). "A simple white noise analysis of neuronal light responses." Network **12**(2): 199-213.

Chun, M. M. and J. M. Wolfe (2001). Visual Attention. Blackwell Handbook of Perception. B. Goldstein. Oxford, UK, Blackwell Publishers Ltd.**:** 271-310.

Colby, C. L., J. R. Duhamel, et al. (1993). "Ventral intraparietal area of the macaque: anatomic location and visual response properties." J Neurophysiol **69**(3): 902-14.

Colby, C. L., J. R. Duhamel, et al. (1996). "Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area." J Neurophysiol **76**(5): 2841-52.

Colby, C. L. and M. E. Goldberg (1999). "Space and attention in parietal cortex." Annu Rev Neurosci **22**: 319-49.

Connor, C. E., J. L. Gallant, et al. (1996). "Responses in area V4 depend on the spatial relationship between stimulus and attention." J Neurophysiol **75**(3): 1306-8.

Connor, C. E., D. C. Preddie, et al. (1997). "Spatial attention effects in macaque area V4." J Neurosci **17**(9): 3201-14.

Conway, B. R., S. Moeller, et al. (2007). "Specialized color modules in macaque extrastriate cortex." Neuron **56**(3): 560-73.

Cook, E. P. and J. H. Maunsell (2002). "Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey." J Neurosci **22**(5): 1994-2004.

Corbetta, M., F. M. Miezin, et al. (1990). "Attentional modulation of neural processing of shape, color, and velocity in humans." Science **248**(4962): 1556-9.

Corbetta, M., F. M. Miezin, et al. (1993). "A PET study of visuospatial attention." J Neurosci **13**(3): 1202-26.

Cowey, A. and V. L. Marcar (1992). "The Effect of Removing Superior Temporal Cortical Motion Areas in the Macaque Monkey: I. Motion Discrimination Using Simple Dots." Eur J Neurosci **4**(12): 1219-1227.

Crick, F. and C. Koch (1990). "Some reflections on visual awareness." Cold Spring Harb Symp Quant Biol **55**: 953-62.

Croner, L. J. and T. D. Albright (1999). "Seeing the big picture: integration of image cues in the primate visual system." Neuron **24**(4): 777-89.

Croner, L. J. and T. D. Albright (1999). "Segmentation by color influences responses of motion-sensitive neurons in the cortical middle temporal visual area." J Neurosci **19**(10): 3935-51.

Davis, G., J. Driver, et al. (2000). "Reappraising the apparent costs of attending to two separate visual objects." Vision Res **40**(10-12): 1323-32.

Dayan, P. and L. F. Abbott (2001). Theoretical Neuroscience: Computational and mathematical modelling if neural systems. Boston, MIT Press.

de Boer, R. and P. Kuyper (1968). "Triggered correlation." IEEE Trans Biomed Eng **15**(3): 169-79.

De Weerd, P., M. R. Peralta, 3rd, et al. (1999). "Loss of attentional stimulus selection after extrastriate cortical lesions in macaques." Nat Neurosci **2**(8): 753-8.

DeAngelis, G. C. and W. T. Newsome (1999). "Organization of disparity-selective neurons in macaque area MT." J Neurosci **19**(4): 1398-415.

Desimone, R. and J. Duncan (1995). "Neural mechanisms of selective visual attention." Annu Rev Neurosci **18**: 193-222.

Desimone, R. and L. G. Ungerleider (1986). "Multiple visual areas in the caudal superior temporal sulcus of the macaque." J Comp Neurol **248**(2): 164-89.

Deutsch, J. A. and D. Deutsch (1963). "Some theoretical considerations." Psychol Rev **70**: 80-90.

Dodd, J. V., K. Krug, et al. (2001). "Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT." J Neurosci **21**(13): 4809-21.

Downing, P., J. Liu, et al. (2001). "Testing cognitive models of visual attention with fMRI and MEG." Neuropsychologia **39**(12): 1329-42.

Dragoi, V., J. Sharma, et al. (2002). "Dynamics of neuronal sensitivity in visual cortex and local feature discrimination." Nat Neurosci **5**(9): 883-91.

Dubner, R. and S. M. Zeki (1971). "Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey." Brain Res **35**(2): 528-32.

Duncan, J. (1984). "Selective attention and the organization of visual information." J Exp Psychol Gen **113**(4): 501-17.

Duncan, J., G. Humphreys, et al. (1997). "Competitive brain activity in visual attention." Curr Opin Neurobiol **7**(2): 255-61.

Enroth-Cugell, C. and J. G. Robson (1966). "The contrast sensitivity of retinal ganglion cells of the cat." J Physiol **187**(3): 517-52.

Eriksen, B. A. and C. W. Eriksen (1974). "Effects of noise letters upon the identification of a target letter in a non-search task." Percept Psychophys **16**: 143-149.

Ernst, M. O. and M. S. Banks (2002). "Humans integrate visual and haptic information in a statistically optimal fashion." Nature **415**(6870): 429-33.

Eurich, C. W. and S. D. Wilke (2000). "Multidimensional encoding strategy of spiking neurons." Neural Comput **12**(7): 1519-29.

Felleman, D. J., A. Burkhalter, et al. (1997). "Cortical connections of areas V3 and VP of macaque monkey extrastriate visual cortex." J Comp Neurol **379**(1): 21-47.

Felleman, D. J. and J. H. Kaas (1984). "Receptive-field properties of neurons in middle temporal visual area (MT) of owl monkeys." J Neurophysiol **52**(3): 488-513.

Felleman, D. J. and D. C. Van Essen (1991). "Distributed hierarchical processing in the primate cerebral cortex." Cereb Cortex **1**(1): 1-47.

Ferrera, V. P., K. K. Rudolph, et al. (1994). "Responses of neurons in the parietal and temporal visual pathways during a motion task." J Neurosci **14**(10): 6171-86.

Fischer, B. and R. Boch (1985). "Peripheral attention versus central fixation: modulation of the visual activity of prelunate cortical cells of the rhesus monkey." Brain Res **345**(1): 111-23.

Fries, P., J. H. Reynolds, et al. (2001). "Modulation of oscillatory neuronal synchronization by selective visual attention." Science **291**(5508): 1560-3.

Friston, K. J., A. P. Holmes, et al. (1994). "Statistical parametric maps in functional imaging: A general linear approach." Human Brain Mapping **2**(4): 189-210.

Galletti, C., P. P. Battaglini, et al. (1990). "'Real-motion' cells in area V3A of macaque visual cortex." Exp Brain Res **82**(1): 67-76.

Gattass, R. and C. G. Gross (1981). "Visual topography of striate projection zone (MT) in posterior superior temporal sulcus of the macaque." J Neurophysiol **46**(3): 621-38.

Georgopoulos, A. P., J. T. Lurito, et al. (1989). "Mental rotation of the neuronal population vector." Science **243**(4888): 234-6.

Georgopoulos, A. P., A. B. Schwartz, et al. (1986). "Neuronal population coding of movement direction." Science **233**(4771): 1416-9.

Goodale, M. A. and A. D. Milner (1992). "Separate visual pathways for perception and action." Trends Neurosci **15**(1): 20-5.

Gottlieb, J. P., M. Kusunoki, et al. (1998). "The representation of visual salience in monkey parietal cortex." Nature **391**(6666): 481-4.

Gruber, T., M. M. Muller, et al. (1999). "Selective visual-spatial attention alters induced gamma band responses in the human EEG." Clin Neurophysiol **110**(12): 2074-85.

Haenny, P. E., J. H. Maunsell, et al. (1988). "State dependent activity in monkey visual cortex. II. Retinal and extraretinal factors in V4." Exp Brain Res **69**(2): 245-59.

Haenny, P. E. and P. H. Schiller (1988). "State dependent activity in monkey visual cortex. I. Single cell activity in V1 and V4 on visual tasks." Exp Brain Res **69**(2): 225-44.

Hartline, H. K. (1938). "The response of single optic fibres of the vertebrate eye to illumination of the retina." American Journal of Physiology **121**: 400-415.

Hastie, T., R. Tibshirani, et al. (2001). The elements of statistical learning theory. New York, Springer.

Heeger, D. J., S. P. Gandhi, et al. (2001). Neuronal correlates of attention in human visual cortex. Visual attention and cortical circuits. J. Braun, C. Koch and J. L. Davis. Cambridge, Massachusetts, MIT Press**:** 25-47.

Hikosaka, K. (1997). "Responsiveness of neurons in the posterior inferotemporal cortex to visual patterns in the macaque monkey." Behav Brain Res **89**(1-2): 275-83.

Hopf, J. M., S. J. Luck, et al. (2006). "The neural site of attention matches the spatial scale of perception." J Neurosci **26**(13): 3532-40.

Hosoya, T., S. A. Baccus, et al. (2005). "Dynamic predictive coding by the retina." Nature **436**(7047): 71-7.

Hubel, D. H. and T. N. Wiesel (1959). "Receptive fields of single neurones in the cat's striate cortex." J Physiol **148**: 574-91.

Hubel, D. H. and T. N. Wiesel (1962). "Receptive fields, binocular interaction and functional architecture in the cat's visual cortex." J Physiol **160**: 106-54.

Hubel, D. H. and T. N. Wiesel (1968). "Receptive fields and functional architecture of monkey striate cortex." J Physiol **195**(1): 215-43.

Hubel, D. H. and T. N. Wiesel (1972). "Laminar and columnar distribution of geniculo-cortical fibers in the macaque monkey." J Comp Neurol **146**(4): 421-50.

Huk, A. C. and D. J. Heeger (2000). "Task-related modulation of visual cortex." J Neurophysiol **83**(6): 3525-36.

James, W. (1890). The principles of Psychology. New York, Dover.

Juan, C. H., S. M. Shorter-Jacobi, et al. (2004). "Dissociation of spatial attention and saccade preparation." Proc Natl Acad Sci U S A **101**(43): 15541-4.

Kanwisher, N. and E. Wojciulik (2000). "Visual attention: insights from brain imaging." Nat Rev Neurosci **1**(2): 91-100.

Kastner, S. and L. G. Ungerleider (2000). "Mechanisms of visual attention in the human cortex." Annu Rev Neurosci **23**: 315-41.

Kay, S. M. (1993). Fundamentals of Statistical Processing: estimation theory. New Jersey, Prentice Hall, Inc.

Kiani, R., T. D. Hanks, et al. (2008). "Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment." J Neurosci **28**(12): 3017-29.

Kobatake, E. and K. Tanaka (1994). "Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex." J Neurophysiol **71**(3): 856-67.

Kohn, A. and J. A. Movshon (2003). "Neuronal adaptation to visual motion in area MT of the macaque." Neuron **39**(4): 681-91.

Kohn, A. and J. A. Movshon (2004). "Adaptation changes the direction tuning of macaque MT neurons." Nat Neurosci **7**(7): 764-72.

Kording, K. P. and D. M. Wolpert (2004). "Bayesian integration in sensorimotor learning." Nature **427**(6971): 244-7.

Kuffler, S. W. (1953). "Discharge patterns and functional organization of mammalian retina." J Neurophysiol **16**(1): 37-68.

Kumada, T. (2001). "Feature-based control of attention: evidence for two forms of dimension weighting." Percept Psychophys **63**(4): 698-708.

Lavie, N. and Y. Tsal (1994). "Perceptual load as a major determinant of the locus of selection in visual attention." Percept Psychophys **56**(2): 183-97.

Lehky, S. R. and T. J. Sejnowski (1990). "Neural model of stereoacuity and depth interpolation based on a distributed representation of stereo disparity." J Neurosci **10**(7): 2281-99.

Liu, J. and W. T. Newsome (2003). "Functional organization of speed tuned neurons in visual area MT." J Neurophysiol **89**(1): 246-56.

Livingstone, M. and D. Hubel (1988). "Segregation of form, color, movement, and depth: anatomy, physiology, and perception." Science **240**(4853): 740-9.

Livingstone, M. S. and D. H. Hubel (1987). "Connections between layer 4B of area 17 and the thick cytochrome oxidase stripes of area 18 in the squirrel monkey." J Neurosci **7**(11): 3371-7.

Logothetis, N. K., H. Guggenberger, et al. (1999). "Functional imaging of the monkey brain." Nat Neurosci **2**(6): 555-62.

Logothetis, N. K., J. Pauls, et al. (2001). "Neurophysiological investigation of the basis of the fMRI signal." Nature **412**(6843): 150-7.

Luck, S. J., L. Chelazzi, et al. (1997). "Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex." J Neurophysiol **77**(1): 24-42.

Luck, S. J. and S. A. Hillyard (2000). The operation of selective attention at multiple stages of processing: evidence from human and monkey electrophysiology. The new cognitive neurosciences. M. S. Gazzaniga. Cambridge, Massachusetts, MIT Press**:** 687-700.

Machens, C. K., M. S. Wehr, et al. (2004). "Linearity of cortical receptive fields measured with natural sounds." J Neurosci **24**(5): 1089-100.

Malonek, D., R. B. Tootell, et al. (1994). "Optical imaging reveals the functional architecture of neurons processing shape and motion in owl monkey area MT." Proc Biol Sci **258**(1352): 109-19.

Martinez-Trujillo, J. C. and S. Treue (2004). "Feature-based attention increases the selectivity of population responses in primate visual cortex." Curr Biol **14**(9): 744-51.

Mather, G., F. Verstraten, et al. (1998). The Motion Aftereffect: A Modern Perspective. Cambridge, Massachusetts, MIT Press.

Maunsell, J. H., G. Sclar, et al. (1991). "Extraretinal representations in area V4 in the macaque monkey." Vis Neurosci **7**(6): 561-73.

Maunsell, J. H. and D. C. van Essen (1983). "The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey." J Neurosci **3**(12): 2563-86.

Maunsell, J. H. and D. C. Van Essen (1983). "Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation." J Neurophysiol **49**(5): 1127-47.

Maunsell, J. H. and D. C. Van Essen (1983). "Functional properties of neurons in middle temporal visual area of the macaque monkey. II. Binocular interactions and sensitivity to binocular disparity." J Neurophysiol **49**(5): 1148-67.

Mazer, J. A., W. E. Vinje, et al. (2002). "Spatial frequency and orientation tuning dynamics in area V1." Proc Natl Acad Sci U S A **99**(3): 1645-50.

McAdams, C. J. and J. H. Maunsell (1999). "Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4." J Neurosci **19**(1): 431-41.

McAdams, C. J. and J. H. Maunsell (2000). "Attention to both space and feature modulates neuronal responses in macaque area V4." J Neurophysiol **83**(3): 1751-5.

Mehta, A. D., I. Ulbert, et al. (2000). "Intermodal selective attention in monkeys. I: distribution and timing of effects across visual areas." Cereb Cortex **10**(4): 343-58.

Miller, E. K., P. M. Gochin, et al. (1991). "Habituation-like decrease in the responses of neurons in inferior temporal cortex of the macaque." Vis Neurosci **7**(4): 357-62.

Mirabella, G., G. Bertini, et al. (2007). "Neurons in area V4 of the macaque translate attended visual features into behaviorally relevant categories." Neuron **54**(2): 303-18.

Mitchell, J. F., G. R. Stoner, et al. (2003). "Attentional selection of superimposed surfaces cannot be explained by modulation of the gain of color channels." Vision Res **43**(12): 1323-8.

Moore, T. and K. M. Armstrong (2003). "Selective gating of visual signals by microstimulation of frontal cortex." Nature **421**(6921): 370-3.

Moran, J. and R. Desimone (1985). "Selective attention gates visual processing in the extrastriate cortex." Science **229**(4715): 782-4.

Motter, B. C. (1993). "Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli." J Neurophysiol **70**(3): 909-19.

Motter, B. C. (1994). "Neural correlates of attentive selection for color or luminance in extrastriate area V4." J Neurosci **14**(4): 2178-89.

Motter, B. C. (1994). "Neural correlates of feature selective memory and pop-out in extrastriate area V4." J Neurosci **14**(4): 2190-9.

Movshon, J. A., I. D. Thompson, et al. (1978). "Receptive field organization of complex cells in the cat's striate cortex." J Physiol **283**: 79-99.

Movshon, J. A., I. D. Thompson, et al. (1978). "Spatial summation in the receptive fields of simple cells in the cat's striate cortex." J Physiol **283**: 53-77.

Muhlenen, A., M. I. Rempel, et al. (2005). "Unique temporal change is the key to attentional capture." Psychol Sci **16**(12): 979-86.

Muller, M. M., T. Gruber, et al. (2000). "Modulation of induced gamma band activity in the human EEG by attention and visual information processing." Int J Psychophysiol **38**(3): 283-99.

Nakamura, K. and C. L. Colby (2000). "Visual, saccade-related, and cognitive activation of single neurons in monkey extrastriate area V3A." J Neurophysiol **84**(2): 677-92.

Nelissen, K., W. Vanduffel, et al. (2006). "Charting the lower superior temporal region, a new motion-sensitive region in monkey superior temporal sulcus." J Neurosci **26**(22): 5929-47.

Newsome, W. T., K. H. Britten, et al. (1989). "Neuronal correlates of a perceptual decision." Nature **341**(6237): 52-4.

Newsome, W. T. and E. B. Pare (1988). "A selective impairment of motion perception following lesions of the middle temporal visual area (MT)." J Neurosci **8**(6): 2201-11.

Niebur, E. and C. Koch (1994). "A model for the neuronal implementation of selective visual attention based on temporal correlation among neurons." J Comput Neurosci **1**(1-2): 141-58.

Niebur, E., C. Koch, et al. (1993). "An oscillation-based model for the neuronal basis of attention." Vision Res **33**(18): 2789-802.

Niessing, J., B. Ebisch, et al. (2005). "Hemodynamic signals correlate tightly with synchronized gamma oscillations." Science **309**(5736): 948-51.

O'Craven, K. M., P. E. Downing, et al. (1999). "fMRI evidence for objects as the units of attentional selection." Nature **401**(6753): 584-7.

O'Craven, K. M., B. R. Rosen, et al. (1997). "Voluntary attention modulates fMRI activity in human MT-MST." Neuron **18**(4): 591-8.

Ogawa, S., T. M. Lee, et al. (1990). "Brain magnetic resonance imaging with contrast dependent on blood oxygenation." Proc Natl Acad Sci U S A **87**(24): 9868-72.

Ogawa, S., T. M. Lee, et al. (1990). "Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields." Magn Reson Med **14**(1): 68-78.

Olshausen, B. A. and D. J. Field (1996). "Emergence of simple-cell receptive field properties by learning a sparse code for natural images." Nature **381**(6583): 607-9.

Olshausen, B. A. and D. J. Field (2005). "How close are we to understanding v1?" Neural Comput **17**(8): 1665-99.

Orban, G. A., H. Kennedy, et al. (1986). "Velocity sensitivity and direction selectivity of neurons in areas V1 and V2 of the monkey: influence of eccentricity." J Neurophysiol **56**(2): 462-80.

Osborne, L. C., W. Bialek, et al. (2004). "Time course of information about motion direction in visual area MT of macaque monkeys." J Neurosci **24**(13): 3210-22.

Paninski, L., J. W. Pillow, et al. (2004). "Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model." Neural Comput **16**(12): 2533-61.

Paradiso, M. A. (1988). "A theory for the use of visual orientation information which exploits the columnar structure of striate cortex." Biol Cybern **58**(1): 35-49.

Parker, A. J. and W. T. Newsome (1998). "Sense and the single neuron: probing the physiology of perception." Annu Rev Neurosci **21**: 227-77.

Pasternak, T. and W. H. Merigan (1994). "Motion perception following lesions of the superior temporal sulcus in the monkey." Cereb Cortex **4**(3): 247-59.

Perge, J. A., B. G. Borghuis, et al. (2005). "Dynamics of directional selectivity in MT receptive field centre and surround." Eur J Neurosci **22**(8): 2049-58.

Perge, J. A., B. G. Borghuis, et al. (2005). "Temporal dynamics of direction tuning in motion-sensitive macaque area MT." J Neurophysiol **93**(4): 2104-16.

Posner, M. I. (1980). "Orienting of attention." Q J Exp Psychol **32**(1): 3-25.

Posner, M. I., C. R. Snyder, et al. (1980). "Attention and the detection of signals." J Exp Psychol **109**(2): 160-74.

Press, W. H., S. A. Teukolsky, et al. (1992). Numerical Recipes in C. Cambridge, Cambridge University Press.

Purushothaman, G. and D. C. Bradley (2005). "Neural population code for fine perceptual decisions in area MT." Nat Neurosci **8**(1): 99-106.

Raiguel, S. E., D. K. Xiao, et al. (1999). "Response latency of macaque area MT/V5 neurons and its relationship to stimulus parameters." J Neurophysiol **82**(4): 1944-56.

Rao, R. P. (2004). "Bayesian computation in recurrent neural circuits." Neural Comput **16**(1): 1-38.

Recanzone, G. H. (2000). "Spatial processing in the auditory cortex of the macaque monkey." Proc Natl Acad Sci U S A **97**(22): 11829-35.

Reynolds, J. H. and L. Chelazzi (2004). "Attentional modulation of visual processing." Annu Rev Neurosci **27**: 611-47.

Richmond, B. J. and T. Sato (1987). "Enhancement of inferior temporal neurons during visual discrimination." J Neurophysiol **58**(6): 1292-306.

Richmond, B. J., R. H. Wurtz, et al. (1983). "Visual responses of inferior temporal neurons in awake rhesus monkey." J Neurophysiol **50**(6): 1415-32.

Riecansky, I., A. Thiele, et al. (2005). "Chromatic sensitivity of neurones in area MT of the anaesthetised macaque monkey compared to human motion perception." Exp Brain Res **167**(4): 504-25.

Rieke, F., D. Warland, et al. (1999). Spikes: Exploring the Neural Code. Cambridge, Massachusetts, MIT Press.

Ringach, D. L., M. J. Hawken, et al. (1997). "Dynamics of orientation tuning in macaque primary visual cortex." Nature **387**(6630): 281-4.

Robinson, D. A. (1963). "A Method Of Measuring Eye Movement Using A Scleral Search Coil In A Magnetic Field." IEEE Trans Biomed Eng **10**: 137-45.

Rodman, H. R., C. G. Gross, et al. (1989). "Afferent basis of visual response properties in area MT of the macaque. I. Effects of striate cortex removal." J Neurosci **9**(6): 2033-50.

Rodman, H. R., C. G. Gross, et al. (1990). "Afferent basis of visual response properties in area MT of the macaque. II. Effects of superior colliculus removal." J Neurosci **10**(4): 1154-64.

Roelfsema, P. R., V. A. Lamme, et al. (1998). "Object-based attention in the primary visual cortex of the macaque monkey." Nature **395**(6700): 376-81.

Rossi, A. F. and M. A. Paradiso (1995). "Feature-specific effects of selective visual attention." Vision Res **35**(5): 621-34.

Saalmann, Y. B., I. N. Pigarev, et al. (2007). "Neural mechanisms of visual attention: how top-down feedback highlights relevant locations." Science **316**(5831): 1612-5.

Saenz, M., G. T. Buracas, et al. (2002). "Global effects of feature-based attention in human visual cortex." Nat Neurosci **5**(7): 631-2.

Saenz, M., G. T. Buracas, et al. (2003). "Global feature-based attention for motion and color." Vision Res **43**(6): 629-37.

Salzman, C. D., K. H. Britten, et al. (1990). "Cortical microstimulation influences perceptual judgements of motion direction." Nature **346**(6280): 174-7.

Schwartz, A. B. (1994). "Direct cortical representation of drawing." Science **265**(5171): 540-2.

Seidemann, E. and W. T. Newsome (1999). "Effect of spatial attention on the responses of area MT neurons." J Neurophysiol **81**(4): 1783-94.

Sereno, M. I., A. M. Dale, et al. (1995). "Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging." Science **268**(5212): 889-93.

Seung, H. S. and H. Sompolinsky (1993). "Simple models for reading neuronal population codes." Proc Natl Acad Sci U S A **90**(22): 10749-53.

Shadlen, M. N., K. H. Britten, et al. (1996). "A computational analysis of the relationship between neuronal and behavioral responses to visual motion." J Neurosci **16**(4): 1486-510.

Shadlen, M. N. and W. T. Newsome (1994). "Noise, neural codes and cortical organization." Curr Opin Neurobiol **4**(4): 569-79.

Shadlen, M. N. and W. T. Newsome (1996). "Motion perception: seeing and deciding." Proc Natl Acad Sci U S A **93**(2): 628-33.

Shamir, M. and H. Sompolinsky (2004). "Nonlinear population codes." Neural Comput **16**(6): 1105-36.

Shao, J. and D. Tu (1995). The Jackknife and Bootstrap. New York, Springer.

Shapley, R. M. and C. Enroth-Cugell (1974). "Letter: The sensitivity distribution, not the receptive field size, is a fixed property of retinal ganglion cells." Vision Res **14**(1): 153-6.

Shipp, S. and S. Zeki (1985). "Segregation of pathways leading from area V2 to areas V4 and V5 of macaque monkey visual cortex." Nature **315**(6017): 322-5.

Shipp, S. and S. Zeki (1989). "The Organization of Connections between Areas V5 and V1 in Macaque Monkey Visual Cortex." Eur J Neurosci **1**(4): 309-332.

Simoncelli, E. P. and B. A. Olshausen (2001). "Natural image statistics and neural representation." Annu Rev Neurosci **24**: 1193-216.

Simoncelli, E. P., L. Paninski, et al. (2004). Characterization of neural responses wih stochastic stimuli. The new Cognitive Neurosciences. M. S. Gazzaniga. Boston, MIT Press.

Sincich, L. C. and J. C. Horton (2003). "Independent projection streams from macaque striate cortex to the second visual area and middle temporal area." J Neurosci **23**(13): 5684-92.

Singer, W. and C. M. Gray (1995). "Visual feature integration and the temporal correlation hypothesis." Annu Rev Neurosci **18**: 555-86.

Smirnakis, S. M., M. J. Berry, et al. (1997). "Adaptation of retinal processing to image contrast and spatial scale." Nature **386**(6620): 69-73.

Smith, M. A., N. J. Majaj, et al. (2005). "Dynamics of motion signaling by neurons in macaque area MT." Nat Neurosci **8**(2): 220-8.

Spitzer, H., R. Desimone, et al. (1988). "Increased attention enhances both behavioral and neuronal performance." Science **240**(4850): 338-40.

Spitzer, H. and B. J. Richmond (1991). "Task difficulty: ignoring, attending to, and discriminating a visual stimulus yield progressively more activity in inferior temporal neurons." Exp Brain Res **83**(2): 340-8.

Standage, G. P. and L. A. Benevento (1983). "The organization of connections between the pulvinar and visual area MT in the macaque monkey." Brain Res **262**(2): 288-94.

Steinmetz, P. N., A. Roy, et al. (2000). "Attention modulates synchronized neuronal firing in primate somatosensory cortex." Nature **404**(6774): 187-90.

Stemmann, H. and W. A. Freiwald (2005). Effects of Attention During Rapid Serial Visual Presentation on Macque Area MT Neurons. Society for Neuroscience. Washington, DC, 2005 Abstract Viewer/Itinerary Planner.

Stemmann, H., W. A. Freiwald, et al. (2005). "Encoding of dynamic visual stimuli by primate area MT neurons." Neurocomputing **65-66**: 135.

Swindale, N. V. (1998). "Orientation tuning curves: empirical description and estimation of parameters." Biol Cybern **78**(1): 45-56.

Taylor, K., S. Mandon, et al. (2005). "Coherent oscillatory activity in monkey area v4 predicts successful allocation of attention." Cereb Cortex **15**(9): 1424-37.

Theunissen, F. E., S. V. David, et al. (2001). "Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli." Network **12**(3): 289-316.

Thiele, A., K. R. Dobkins, et al. (2000). "Neural correlates of contrast detection at threshold." Neuron **26**(3): 715-24.

Thompson, K. G., N. P. Bichot, et al. (2005). "Frontal eye field activity before visual search errors reveals the integration of bottom-up and top-down salience." J Neurophysiol **93**(1): 337-51.

Touryan, J., B. Lau, et al. (2002). "Isolation of relevant visual features from random stimuli for cortical complex cells." J Neurosci **22**(24): 10811-8.

Treue, S. (2001). "Neural correlates of attention in primate visual cortex." Trends Neurosci **24**(5): 295-300.

Treue, S. and J. C. Martinez Trujillo (1999). "Feature-based attention influences motion processing gain in macaque visual cortex." Nature **399**(6736): 575-9.

Treue, S. and J. H. Maunsell (1996). "Attentional modulation of visual motion processing in cortical areas MT and MST." Nature **382**(6591): 539-41.

Tsao, D. Y., W. A. Freiwald, et al. (2003). "Faces and objects in macaque cerebral cortex." Nat Neurosci **6**(9): 989-95.

Tsao, D. Y., W. A. Freiwald, et al. (2006). "A cortical region consisting entirely of face-selective cells." Science **311**(5761): 670-4.

Tsao, D. Y., N. Schweers, et al. (2008). "Patches of face-selective cortex in the macaque frontal lobe." Nat Neurosci **11**(8): 877-9.

Tsao, D. Y., W. Vanduffel, et al. (2003). "Stereopsis activates V3A and caudal intraparietal areas in macaques and humans." Neuron **39**(3): 555-68.

Tsotsos, J. K., S. M. Culhane, et al. (2001). From Foundational Principles to a Hierarchical Selection Circuit for Attention. Visual Attention and Cortical Circuits. J. Braun, C. Koch and J. L. Davis. Cambridge, MIT Press**:** 285-306.

Uka, T. and G. C. DeAngelis (2004). "Contribution of area MT to stereoscopic depth perception: choice-related response modulations reflect task strategy." Neuron **42**(2): 297-310.

Ungerleider, L. G. and R. Desimone (1986). "Cortical connections of visual area MT in the macaque." J Comp Neurol **248**(2): 190-222.

Ungerleider, L. G. and M. Mishkin (1982). Two Cortical visual systems. The Analysis of Visual Beavior. D. J. Ingle, R. J. W. Mansfield and M. S. Goodale. Cambridge, MA, MIT Press**:** 549-86.

Valdes-Sosa, M., M. A. Bobes, et al. (1998). "Switching attention without shifting the spotlight object-based attentional modulation of brain potentials." J Cogn Neurosci **10**(1): 137-51.

Van Essen, D. C., J. H. Maunsell, et al. (1981). "The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization." J Comp Neurol **199**(3): 293-326.

Vanduffel, W., D. Fize, et al. (2001). "Visual motion processing investigated using contrast agent-enhanced fMRI in awake behaving monkeys." Neuron **32**(4): 565-77.

Vanduffel, W., D. Fize, et al. (2002). "Extracting 3D from motion: differences in human and monkey intraparietal cortex." Science **298**(5592): 413-5.

Vanduffel, W., R. B. Tootell, et al. (2000). "Attention-dependent suppression of metabolic activity in the early stages of the macaque visual system." Cereb Cortex **10**(2): 109-26.

Vogels, R. and G. A. Orban (1994). "Activity of inferior temporal neurons during orientation discrimination with successively presented gratings." J Neurophysiol **71**(4): 1428-51.

Wannig, A., V. Rodriguez, et al. (2007). "Attention to surfaces modulates motion processing in extrastriate area MT." Neuron **54**(4): 639-51.

Wegener, D., W. A. Freiwald, et al. (2004). "The influence of sustained selective attention on stimulus selectivity in macaque visual area MT." J Neurosci **24**(27): 6106-14.

Wilke, S. D. and C. W. Eurich (2002). "Representational accuracy of stochastic neural populations." Neural Comput **14**(1): 155-89.

Wilson, M. A. and B. L. McNaughton (1993). "Dynamics of the hippocampal ensemble code for space." Science **261**(5124): 1055-8.

Wojciulik, E., N. Kanwisher, et al. (1998). "Covert visual attention modulates face-specific activity in the human fusiform gyrus: fMRI study." J Neurophysiol **79**(3): 1574-8.

Zeki, S. (1980). "The response properties of cells in the middle temporal area (area MT) of owl monkey visual cortex." Proc R Soc Lond B Biol Sci **207**(1167): 239-48.

Zeki, S. M. (1974). "Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey." J Physiol **236**(3): 549-73.

Zeki, S. M. (1978). "Uniformity and diversity of structure and function in rhesus monkey prestriate visual cortex." J Physiol **277**: 273-90.

Zhang, K., I. Ginzburg, et al. (1998). "Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells." J Neurophysiol **79**(2): 1017-44.

Zhang, K. and T. J. Sejnowski (1999). "Neuronal tuning: To sharpen or broaden?" Neural Comput **11**(1): 75-84.