How neural activity underlies visual motion perception

by

Nicolas Yvan Masse

Department of Physiology,
McGill University,
Montreal, Quebec, Canada

Thesis submitted in August 2008

A thesis submitted to the
Faculty of Graduate Studies and Research
in partial fulfillment of the degree of
Doctorate in Philosophy

Contents

List of figures

List of tables

Abstract

Over the last several decades, much progress has been made in understanding how neurons in visual cortex encode various stimuli. However, it is still unclear how their neural response underlies visual perception. The aim of this thesis is to build upon our current understanding of how visual motion perception is generated by sensory neural activity.

Many interpretations on how sensory neural activity is combined to form perceptual decisions rely upon the often-observed correlation between neural activity and behavior. Due to its importance, it is critical that we understand whether this correlation is an inherent feature in forming perceptual decisions, or whether it is the result of factors outside the brain. We found that microsaccades, which are small, involuntary eye movements, contribute a small but significant amount to this correlation. This result reinforces the view that the correlation is inherent in forming percepts, but its value has been overestimated in several studies.

Given that this correlation between neural activity and behavior is genuine, on what timescales can it operate? We found that the correlation between neural activity and behavior can exist on a timescale of about ten milliseconds. Although our results do not say anything about how visual cortex encodes motion in general, it does suggest that the brain is capable of extracting information from neural responses on fast timescales.

If neural activity can be correlated with behavior on fast timescales, can electrical microstimulation be used to demonstrate causality on this temporal regime? By comparing the temporal effects of microstimulation on behavior to an equivalent visual stimulus, we found that microstimulation has a significantly weaker and longer effect. This result raises doubt that microstimulation has the necessary temporal precision to probe causality on fast timescales, or perhaps even serve in sensory neural prosthetics.

Finally, the correlation between neural activity and behavior provides an indication of how influential a neuron is in forming perceptual decisions. We found that spatial attention modified the relationship between how informative of the stimulus a neuron was and its influence in forming decisions. This leads to several hypotheses on the role of attention in visual processing.

Résumé

Pendant les quelques dernières décades, beaucoup de progrès a été fait dans notre compréhension à savoir comment les neurones, dans les cortices visuelles, encodent différents stimulus. Toutefois, il n'est toujours pas clair comment les réponses neurales forment la base de la perception visuelle.

Le but de cette thèse est d'ajouter à notre compréhension concernant comment la perception de mouvement visuel est générée par l'activité neurale sesnsorielle.

Plusieurs interprétations, concernant comment l`activité neurale sensorielle est jointe pour former des décisions perceptuelles, comptent sur la corrélation souvent observée entre l'activité neurale et le comportement.

À cause de son importance, il est crucial de comprendre si cette corrélation est un trait inhérant dans la formation des décisions perpceptuelles ou si c'est le résultat de facteurs non reliés au cerveau. Nous avons trouvé que les microsaccades, qui sont de petits mouvements involontaires des yeux, font une petite, mais significante, contribution à cette corrélation. Ce résultat renforce l'opinion que la corrélation est inhérente dans la formation des perceptions, mais sa valeur a été surestimée dans plusieurs études.

Étant donné que cette corrélation entre l'activité neurale et le comportement est authentique, sur quelle période de temps peut-elle opérer? Nous avons trouvé que la

corrélation entre l'activité neurale et le comportement peut exister dans une période de temps d'environ 10 millisecondes.

Bien que nos résultats ne disent rien concernant commnet le mouvement visuel est encodé en général, ils suggèrent que le cerveau est capable d'extraire des informations des réponses neurales dans une période de temps rapide.

Si l'activité neurale peut être corrélée avec le comportement dans une période de temps rapide, est-ce q'une microstimulation éléctrique peut être utilisée pour démontrer une causalité sur ce régime temporel?

En comparant les effets temporels de la microstimulation sur le comportement, à un stimulus visuel équivalent, nous avons trouvé que la microstimulation a un effet beaucoup plus faible et plus long. Ce résultat soulève des doutes que la microstimulation a la précision temporelle nécessaire pour sonder la causalité dans une période de temps rapide, ou peut-être même pour servir dans la prothétique neurale sensorielle?

Finalement, la corrélation entre l'activité neurale et le comportement nous fourni une indication comment influent un neurone est dans la formation des décisions perceptuelles. Nous avons trouvé que l'attention spaciale modifie la relation entre comment informatif le stimulus d'un neurone fût et son influence dans la prise de décisions. Ceci mène à plusieurs hypothèses sur le role de l'attention dans la computation visuelle.

Acknowledgements

First and foremost, I must thank my supervisor, Erik Cook, for the tremendous amount of help he has given me over these last four years. When I first arrived at the lab, I was probably as naive as they came. Erik had the inexhaustible patience to not only teach me the necessary experimental skills, but more importantly, the qualities required to be a good scientist. I wouldn't be at this stage right now if it were not for his guidance.

Secondly, I am grateful for the colleagues I have had the pleasure of working with at the lab: Chang-an, Jackson, Navid and Pinar. I could not have asked for a nicer group of people to have around. Life at the lab was made so much easier by working with such a supportive group.

Third, a big thank you to all my friends, both inside and outside the department. Whether it was discussing science, politics, life or simply nothing at all, you all made life enjoyable.

Lastly, I cannot thank my family enough for all the support they have given me. They have always done so much to help me out, and especially so during my PhD. I am forever thankful for all you have done.

Contributions of authors

Chapter 2 was adapted from a manuscript under preparation: *Herrington,T., Masse, N.Y., Hachmeh, K., Smith, J., Assad, J.A., Cook, E.P., The effect of microsaccades on the correlation between neural activity and perception in areas MT, VIP, and LIP*. Todd Herrington and myself are listed as co-first authors. The idea for the study was developed by Todd Herrington, John Assad and Erik Cook. This manuscript was based on several experimental datasets collected by Todd Herrington, Erik Cook, Jackson Smith and myself. For my contribution of experimental data, I trained the monkeys with help from Navid Sadeghi Ghandehari with one of the monkeys, participated in the necessary surgeries, and collected the behavioral and neural data. Erik Cook wrote the necessary code to run the paradigm. I performed the data analysis and wrote the manuscript under the guidance of Erik Cook and with help from Todd Herrington.

Chapter 3 was adapted from a published manuscript: *Masse, N.Y., Cook, E.P., The Effect of Middle Temporal Spike Phase on Sensory Encoding and Correlates with Behavior during a Motion-Detection Task, Journal of Neuroscience 28(6):1343-1355, 2008*. I developed the idea for the study with help from Erik Cook. All experimental work was performed by Erik Cook while at the laboratory of John Maunsell. I performed the data analysis and wrote the manuscript under the guidance of Erik Cook.

Chapter 4 was adapted from a submitted manuscript: *Masse, N.Y., Cook, E.P., The behavioral impulse-response function of microstimulation in cortical area MT, submitted to Journal of Neurophysiology, 2008.* Erik Cook developed the idea for the study. I trained the monkeys with help from Navid Sadeghi Ghandehari with one of the monkeys, participated in the necessary surgeries, and collected the data. Erik Cook wrote the necessary code to run the paradigm. I performed the data analysis and wrote the manuscript under the guidance of Erik Cook.

Chapter 5 was adapted form a manuscript in preparation: *Masse, N.Y., Cook, E.P., The effect of spatial attention on sensory processing in cortical area VIP.* I developed the idea for the study. All experimental work was performed by Erik Cook while at the laboratory of John Maunsell. I performed the data analysis and wrote the manuscript under the guidance of Erik Cook.

Chapter 1

Introduction

Nearly all organisms, from bacteria to humans, possess the ability to sense and react to the environment. In its most basic form, sensory systems such as chemoreception allow the organism to sense the chemical nature of its environment without the need of a nervous system. However, the sensory systems have evolved tremendously, allowing some animals to recognize familiar faces by the pattern of photons falling on the retina, electric fish the ability to detect prey by disturbances in self-generated electric fields, and bats the ability to gauge distance by incoming waves of air pressures, to name a few. While much progress has been made into how neurons in these various sensory systems encode stimuli, it is still poorly understood how perception is built upon neural activity.

Current understanding of how neural activity underlies visual perception

By 1989, neuroscientists had established some of the basic tenets of neural coding in visual cortex. Seminal work by Hubel and Wiesel first showed that neurons in primary visual cortex (V1) were selective for different visual features, such as orientation (1962), and that neurons with similar preferences were grouped together in a columnar organization (1968; 1974). Additionally, it was know that different areas in visual cortex were selective for different types of visual features. In the dorsal stream, it was discovered that neurons in areas middle temporal (MT) and medial superior temporal (MST) responded preferentially to motion stimuli of certain directions, speeds and complexity (Albright, 1984; Maunsell and Van Essen, 1983; Mikami et al., 1986; Saito et al., 1986; Zeki, 1974). In the ventral stream, neurons in V4 were selective for objects of various forms (Desimone and Schein, 1987) while those in inferior temporal cortex (IT)

were also selective for forms, but of increased complexity (Gross et al., 1972; Perrett et al., 1982). Thus, while it was know that neurons in visual cortex were specialized to encode different aspects of a visual stimulus, the relationship between their activity and the perception they supposedly generated was still unknown.

It was not until studies that recorded both neural activity and behavior simultaneously did a basic picture emerge. Newsome *et. al.* were one of the first groups to perform such an experiment where they recorded neural activity from area MT of a monkey while it performed a motion discrimination task (Britten et al., 1996; Britten et al., 1992; Newsome et al., 1989; Zohary et al., 1994a). The task was for a monkey to report the direction of coherent motion that occurred in a random dot patch and lasted for two seconds. The motion could only move in one of two opposite directions and the strength of the coherent motion (the ratio of dots moving in the coherent direction to all others) was varied. During this task, they recorded from neurons in area MT that were selective for one of the two directions of motion. Both lesion studies (Baker et al., 1991; Hess et al., 1989) and microstimulation studies (Salzman et al., 1990; Salzman et al., 1992) demonstrated that neural activity in area MT is causally linked to motion perception.

These studies made surprising conclusions regarding how informative neurons were of the stimulus and their correlation with behavior. These terms appear often through out the thesis, so it is worthwhile to clarify their meaning. A neuron is informative of the stimulus if given some measure of its neural response, one can predict

the stimulus that was presented to the neuron. One method to calculate how informative a neuron was of the stimulus is to compare the neural response before and after a stimulus change. For example, the neuron in Figure 1A appears more informative than the neuron in Figure 1B since there was larger increase in neural activity after the motion stimulus started moving in a coherent fashion. This measure can be quantified by comparing the distributions of spike counts from each trial that occurred during 0% coherent motion (right panel, red bars) to the spike counts that occurred during coherent motion (right panel, blue bars). The larger the separation between the two distributions, the more informative the neuron was of the stimulus. Later on, this measure will be referred to the neurometric. Similarly, the correlation between neural activity and behavior is measured by comparing the neural response to the stimulus from trials where the subject made the correct decision to the neural response from trials where the subject made an incorrect decision. For example, the neuron in Figure 1C appears more strongly correlated with behavior than the neuron in Figure 1D since there existed a larger difference in the neural response to coherent motion between correct (blue curve) and failed (red curve) trials. As above, this measure is quantified by comparing the distribution of spike counts in response to the coherent motion between correct and failed trials. Later on, this measure will be referred to as the choice or detect probability.

The first surprising conclusion of these studies by *Newsome et. al.* was that the subject's behavioral responses were no more informative of the true motion direction than the activity of single MT neurons (1989). Given that the subject supposedly has access to the information from many MT neurons, one must have believed that the

Figure 1 Four example MT neurons. (A&B) The left column shows the average spike rate, smoothed with a Gaussian filter with standard deviation of 20 ms. The neuron is responding to 0% coherent motion until time zero, after which coherent motion in the neuron's preferred direction and speed is displayed. The right column shows the distribution of spike counts from all trial counted from 50 to 150 ms after coherent motion onset (blue bars) and from 0 to 100 ms prior to coherent motion onset (red bars). The neuron in (A), which showed a clear change in neural activity after coherent motion onset, would be considered informative of the stimulus. The neuron in (B), which showed little change in neural activity after coherent motion onset, would not be considered very informative of the stimulus. (C&D) Left column is same as above, except average spike rates are divided into correct (blue curve) and failed (red curve) trials. The right column shows the distribution in spike counts for correct (blue bars) and failed (red bars) trials, collected from 50 to 150 ms after coherent motion onset. The neuron in (C), which showed a larger separation in spike rate in response to coherent motion between correct and failed trials than the neuron in (D), was more correlated with the behavior.

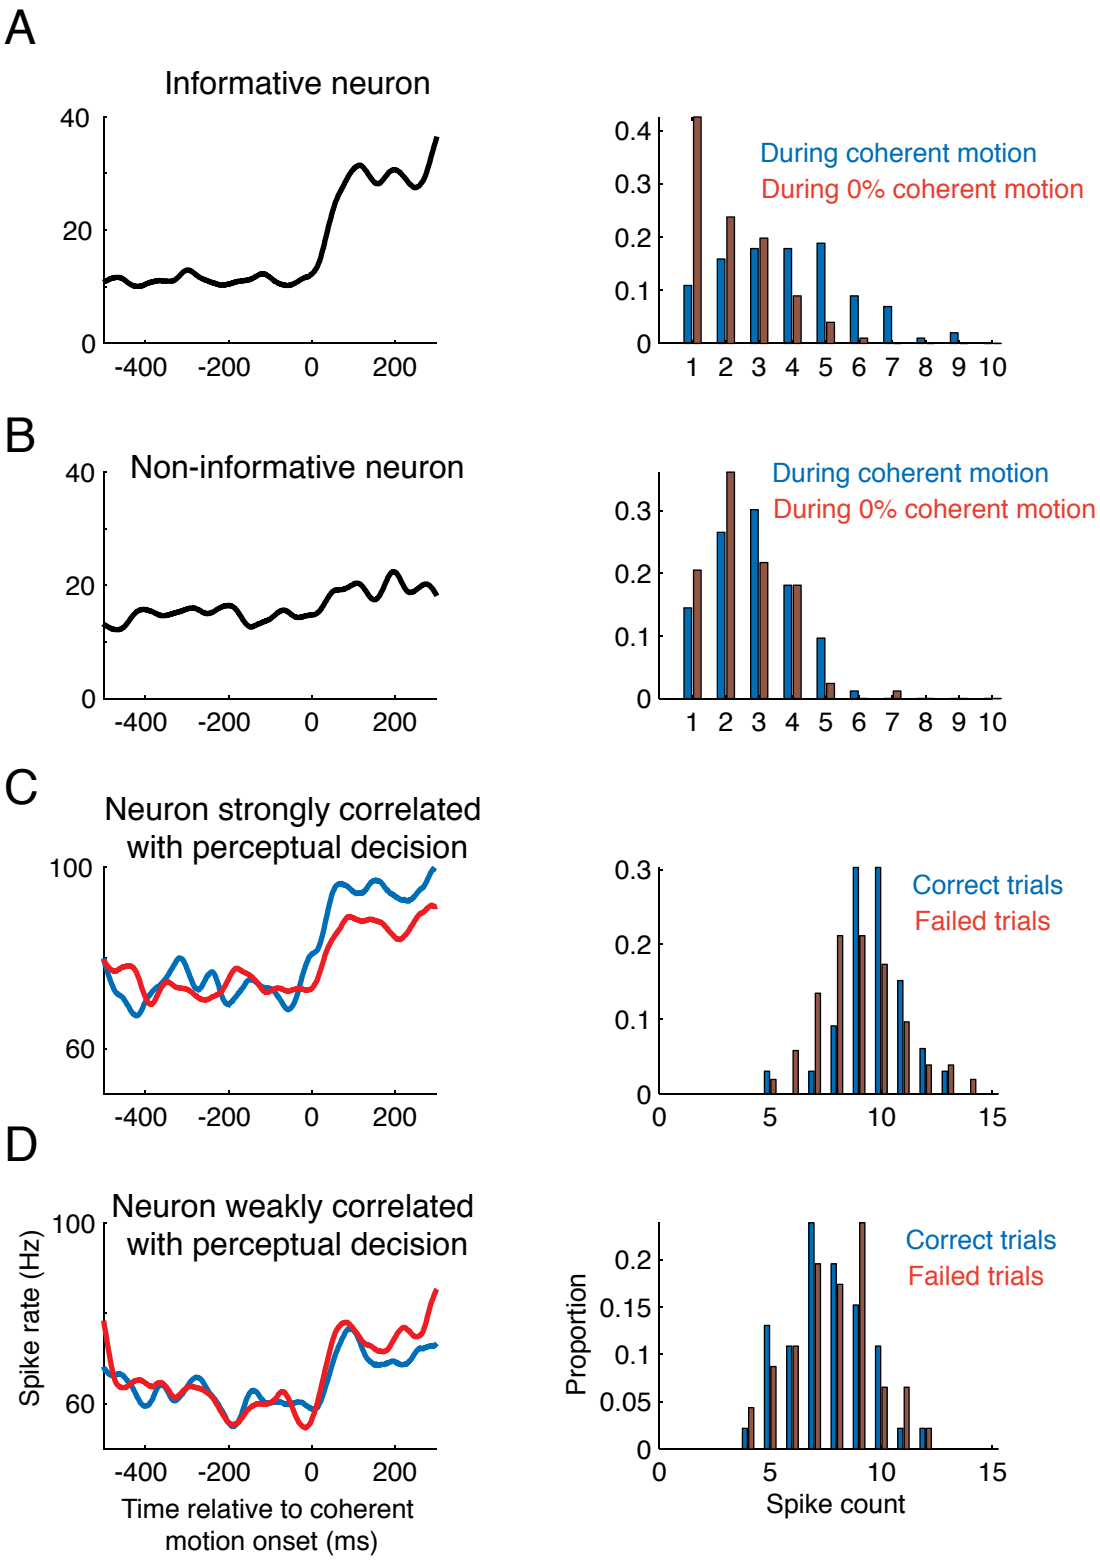
A
Informative neuron
40
20
0
-400
-200
0
200
0.4
0.3
0.2
0.1
0
During coherent motion
During 0% coherent motion
1 2 3 4 5 6 7 8 9 10
B
Non-informative neuron
40
20
0
-400
-200
0
200
0.3
0.2
0.1
0
During coherent motion
During 0% coherent motion
1 2 3 4 5 6 7 8 9 10
C
Neuron strongly correlated with perceptual decision
100
60
-400
-200
0
200
0.3
0.2
0.1
0
Correct trials
Failed trials
0
5
10
15
D
Neuron weakly correlated with perceptual decision
Spike rate (Hz)
100
60
-400
-200
0
200
Time relative to coherent motion onset (ms)
Proportion
0.2
0.15
0.1
0.05
0
Correct trials
Failed trials
0
5
10
15
Spike count

Figure 1

subject's behavioral choices would have been more informative of the stimulus than single neurons. The second surprising finding was that for trials with identical stimuli, the activity of single neurons was significantly correlated with the subject's behavioral response. At face value, significant correlations between the activity of single neurons and behavior would suggest that each neuron played a significant role in forming the decision, implying that few neurons were involved.

Another discovery around this time partly explained these two surprising observations. Recordings in the visual cortex of anesthetized cats (van Kan et al., 1985) and in the awake monkey (Gawne et al., 1996a; Gawne and Richmond, 1993) revealed that the neural response to various types of visual stimuli was correlated between neurons. Positive correlations limit the benefit of averaging the activity of many neurons when forming a decision (Zohary et al., 1994b). Additionally, correlated firing between neurons helped explain why the activity of single neurons was correlated with the behavioral outcome of the trial: if the activity of a neuron is correlated with the rest of the population response, and if the population response drives the perception, then the activity of a single neuron would be then correlated with the perception (Zohary et al., 1994b).

From these results, several groups have attempted to model how sensory information is combined to form a behavioral response (Bair et al., 2001; Schoppik et al., 2008; Shadlen et al., 1996; Zohary et al., 1994b). The conclusions of these models have not always been in accord in other known results. It was proposed that correlations

between neurons limited the benefit of pooling across a large population, reducing the capability of the subject (Bair et al., 2001; Shadlen et al., 1996; Zohary et al., 1994b). However, theoretical results have shown that correlations do not always have to be harmful (Abbott and Dayan, 1999). This is somewhat intuitive given that correlations reduce the entropy of the noise of neural responses. An example where correlations could increase the capabilities of a neural population was confirmed experimentally (Romo et al., 2003). It was also proposed that the activity of neurons that responded optimally and sub-optimally to the stimulus was combined equally when forming a decision. This was also proposed to explain why the subject performed no better than individual neurons. However, this proposal contradicted many other studies that suggested that the more informative neurons are also more influential in forming a decision (Britten et al., 1996; Dodd et al., 2001; Jazayeri and Movshon, 2007a; Jazayeri and Movshon, 2007b; Law and Gold, 2008; Parker et al., 2002; Purushothaman and Bradley, 2005; Uka and DeAngelis, 2004). In short, it is poorly understood how the activity of sensory neurons is combined when forming a behavioral response.

Perhaps one of the more important realizations in recent years has been that there is no single way that neural activity is combined when forming perceptual decisions. For example, it is well known that top-down signals can modulate how sensory neurons encode stimuli (for reviews, see Desimone and Duncan, 1995; Engel et al., 2001; Maunsell and Treue, 2006). The simplest example of such top-down modulation is the widely reported result that spatial attention increases the gain of the neurons without changing its sensitivity (e.g. Cook and Maunsell, 2004; McAdams and Maunsell, 1999b;

Treue and Martinez Trujillo, 1999). Moreover, top-down signals can change the aspect of the stimulus that is encoded by the neuron (Freedman and Assad, 2006; Lamme et al., 1998; Li et al., 2004; Pack et al., 2001), suggesting that top-down signals can modify how inputs to a neuron are combined when forming a neural response. Thus, although it is still unclear how neural activity is combined when forming a perceptual decision, it appears that the rules are not fixed.

When attempting to infer how neural activity is pooled to form perceptual decisions, an important consideration is the timescale of both the stimulus and the neural response. One potential caveat with many past studies was that constant stimuli were presented for hundreds of milliseconds and that neural activity was usually averaged over similar time periods. Unlike the stimuli used in many of these past studies, natural stimuli can be highly dynamic. Furthermore, animals have the ability to react very quickly to their environment (Rieke et al., 1999), suggesting that the brain is capable of forming decisions based on brief neural responses. In the next section, we review some of the past debate on the temporal nature of a neural response and whether perception can be correlated with neural activity on fast timescales.

The spike rate versus spike timing debate

How is information encoded in the spiking activity of neurons? This has been a highly contentious issue, with one group proposing that information is encoded in the mean firing rate of neurons while the other group favoring the notion that the precise

timing of actions potentials conveys information. These two ideas are not mutually exclusive, and there exists truth in both arguments.

The idea that information is encoded in the mean firing rate of neurons is based on the observation that cortical neurons can respond very differently to repeated presentations of similar stimuli (Britten et al., 1993; Snowden et al., 1992). It was argued that this variability was because cortical neurons can receive synaptic inputs from thousands of other neurons (Peters, 1987), out of which only a small number of excitatory inputs are required to depolarize the neuron sufficiently to emit an action potential (Mason et al., 1991; Matsumura et al., 1988; Otmakhov et al., 1993). To prevent these thousands of excitatory inputs from saturating the neuron, it was proposed that cortical neurons receive balanced excitatory and inhibitory input, preventing the neuron from being overwhelmed (Shadlen and Newsome, 1998). Although an increase in excitation would always be matched by an increase in inhibition, the firing rate would still increase because of increased variability in the number of synaptic inputs. Intuitively, if the neuron's membrane potential is normally distributed, then balanced excitation and inhibition would keep the mean constant, but increased excitation and inhibition would increase the standard deviation, leading to more threshold crossings. Thus, the consequence of balance excitation and inhibition would be inherent variability in the spiking response. It was for this reason that many proposed that variability in neural responses were too large for information to be transmitted by precise timing of action potentials (Mazurek and Shadlen, 2002; Rudolph and Destexhe, 2003; Schaette et al., 2005; Shadlen and Newsome, 1998; Zoccolan et al., 2002). By default, it was suggested

that information must be carried by the mean spiking activity of neurons, averaged over a period of tens or hundreds of milliseconds.

Recently, the claim that a state of balanced excitation and inhibition simply leads to increased variability in the neural response has been challenged. As both the level of excitation and inhibition increase, the conductance of the neuron increases as well. Several experimental studies in awake animal preparations have shown that neurons operate in a state of much higher conductance than do in anesthetized, or *in vitro* preparations (Baranyi et al., 1993; Matsumura et al., 1988; Pare et al., 1998; Steriade et al., 2001). A high-conductance state has important implication on how synaptic inputs are integrated by the neuron (for a review, see Destexhe et al., 2003). Specifically, a conductance increase results in a decrease in the membrane time constant (Azouz and Gray, 2003; Prescott et al., 2006) as well as the activation of voltage-dependent channels that hyperpolarize the neuron (Prescott et al., 2006). As a result, the neuron becomes less sensitive to slow changes to the membrane potential and becomes more sensitive to rapid fluctuations caused by synchronously arriving inputs. Because of this, the transition has referred to as the neuron going from an "integrator" to a "coincidence detector".

The claim that the responses of cortical neurons are highly variable has also been revisited. Studies that suggested that the neural responses were too variable for precise timing to play a role used slowly varying or static stimuli (Britten et al., 1993; Newsome et al., 1989; Snowden et al., 1992). As other have pointed out, natural stimuli are much more dynamic (Rieke et al., 1995; Schwartz and Simoncelli, 2001), and perhaps highly

dynamic stimuli are behaviorally more significant for the organism (Rieke et al., 1999). Thus, there was much interest in the reliability and precision of neural responses to dynamic and more natural stimuli. These studies found that the neural response can be extremely reliable with precision on the order of milliseconds (Butts et al., 2007; de Ruyter van Steveninck et al., 1997; Gur et al., 1997; Mainen and Sejnowski, 1995; Reinagel and Reid, 2000). This reliability at such a fine temporal scale allows for information to be transmitted in the precise timing of action potentials (de Ruyter van Steveninck et al., 1997). Does this mean that information is always carried by highly precise action potentials? Not necessarily; a recent study has proposed that the temporal precision of the neural response is dependent on the timescale of the stimulus (Butts et al., 2007). They put forward the hypothesis that the temporal precision is only as high as what is required to accurately encode the stimulus. This would explain past results that found highly variable neural responses to slowly changing or static stimuli; the neural responses lacked temporal precision because temporal precision was not required to describe the stimulus.

These results suggest that the nervous system is capable of encoding stimuli through a temporally precise neural response. This does not imply that the slowly varying average rate of the neural response is uninformative. On the contrary, many studies have shown the average rate to be both informative of the stimulus as well as correlated with the perception of the subject (Britten et al., 1996; Cook and Maunsell, 2002b; Dodd et al., 2001; Huk and Shadlen, 2005; Liu and Newsome, 2005; Purushothaman and Bradley, 2005). However, in an environment where highly dynamic sensory stimuli are be

prevalent as well behaviorally important (Rieke et al., 1999), it is important to know that sensory systems are capable of the temporally precise encoding necessary to represent these dynamic stimuli.

Different methods of temporal encoding and decoding

If the nervous system is capable of transmitting information in the timing of action potentials, then what are the strategies sensory systems use to temporally encode information? Alternatively, how would information contained in the timing of spikes be decoded? Strategies employing temporal encoding can be broken up into three, overlapping categories: information can be transmitted by the relative timing between spikes from different neurons, the coincident timing of spikes from different neurons, and the timing of spikes relative to a global signal (usually oscillations in the local field potential). Categories one and two appear redundant, but we shall see these two encoding schemes are decoded in separate ways.

Encoding with the relative timing between spikes

The first scheme, in which the temporal offset of spikes from different neurons conveys information, has been proposed as a method to very quickly transmit information (Thorpe et al., 2001; Van Rullen and Thorpe, 2001). Encoding information in the relative timing of spikes has been considered only applicable to transient stimuli, or to the initial part of a neural response. This is because it is difficult to determine any sense of order

between spikes during an ongoing neural response. However, in Chapter 3 we discuss a possible exception, which was only possible given the special nature of the neural response.

The basic idea behind this encoding scheme is that neurons whose preferences best match that of the stimulus will be driven to threshold first, spiking before other neurons. Thus, the time required for the neuron to first respond to the stimulus, known as the neural latency, can be informative of the stimulus. This has been shown true through out the visual system (Celebrini et al., 1993; Gawne et al., 1996b; Gollisch and Meister, 2008; Konig et al., 1995; Masse and Cook, 2008; Maunsell et al., 1999; Raiguel et al., 1999). Downstream areas will probably not have access to the neural latency, but only to the relative latency between neurons. Additionally, this relative latency can be made more precise by the fact that neural latency can be correlated between neurons (Chase and Young, 2007; Gollisch and Meister, 2008).

Several models of the neural structure required to decode relative timing have been proposed (Delorme, 2003; Thorpe et al., 2001). Both models involve upstream neurons with feed-forward connections with inhibitory neurons and a target neuron. A target neuron must receive inputs with short latencies from upstream neurons that it forms strong synaptic connections with, before feed-forward inhibition arrives. Several studies have found that a neuron's tuning is partly driven by decoding the relative timing between inputs. It was demonstrated that the neural tuning for different whisker directions in rat barrel cortex is governed by the temporal offset between excitatory and

inhibitory inputs (Wilent and Contreras, 2005). This is also true for direction selectivity in the cat primary visual cortex (Priebe and Ferster, 2005).

Encoding with synchronized spiking

Slightly different than the first scheme, another method to temporally encode information is through the simultaneous generation of action potentials. The scheme works on the principal that simultaneously arriving inputs can be much more effective at driving a target neuron than temporally spread out inputs. Such an encoding scheme is believed by many to solve the binding problem (Engel et al., 1990; Gray et al., 1990), in which different visual features must be combined to form a coherent object. However, many believe that synchronized spiking does not underlie binding of visual objects into coherent percepts (Shadlen and Movshon, 1999). If simultaneous spikes from distant neurons are to encode information, then some sort of global signal is needed to coordinate their activity. Oscillations in the local field potential, which can be coherent across different cortical areas (Destexhe et al., 1999; Donoghue et al., 1998; Engel et al., 1990), could act as the necessary timing device.

This scheme is nicely exemplified by the locust and zebrafish olfactory system. Neurons in the antennal lobe, one synapse away from the primary olfactory receptors, respond very broadly to odors. Neurons in the mushroom body, one of two targets of the antennal lobe and also an area associated with memory, respond in a very sparse fashion. Through a series of well-designed experiments (Friedrich et al., 2004; Perez-Orive et al.,

2002), it was shown that different subsets of neurons in the antennal lobe phase-lock their spiking to oscillations in the local field potentials in response to different odors. Downstream in the mushroom body, oscillations in the local field potential only provide a small window during each cycle for inputs to have an impact. Additionally, each EPSP depolarizes the target neuron very little, thus many spikes from the antennal lobe must arrive within a very short temporal window in order to sufficiently depolarize the neuron to spike. Thus, neurons in the mushroom body will only fire if they are synaptically connected to most of the subset of neurons in the antennal lobe that phase-locked their firing in response to the odor. In contrast, the anatomically similar olfactory system of *Drosophila* does not employ an oscillating local field potential to transform the broad responses of the antennal lobe in the sparse responses of the mushroom body (Turner et al., 2008).This raises the bigger question of why similar nervous structures under similar constraints have evolved to solve similar problems in completely different manners.

Encoding with the phase of spikes

Lastly, the phase of a spike relative to the oscillation in local field potential can also serve to encode information. The most prominent example are place cells in the hippocampus, where the phase of spiking encodes the location of the animal relative to the neuron's place field (O'Keefe and Recce, 1993). Such an encoding scheme has also been proposed for the olfactory system (Hopfield, 1995) and for the visual system (Fries et al., 2007; Montemurro et al., 2008), however, there is relatively little evidence for their existence. In single-unit recordings in area MT, we found no evidence to suggest that any

motion information was encoded by the phase of the spiking (personal unpublished observation; personal communication with Paul Kayhat).

Regardless of whether such an encoding scheme works outside the hippocampus, decoding the phase of the spike would not be difficult. This is because the phase of subthreshold oscillations in the membrane potential of a neuron modulates the strength of post-synaptic potentials (Lampl and Yarom, 1993; Schaefer et al., 2006). Such a scheme would require that the oscillations in the local field potential are coherent between upstream and downstream regions, which can exist (Destexhe et al., 1999; Donoghue et al., 1998; Engel et al., 1990).

Does the brain actual use temporal encoding?

Given the large amount of evidence that the temporal structure of a neural response can be informative, it is surprising that very few studies have examined whether the brain is capable of extracting temporally encoded information. None of these studies listed above have demonstrated that information in the temporal structure of the neural response was correlated to the behavior of the subject.

Only two studies that the author is aware of have attempted to demonstrate the temporally encoded information is relevant to the organism. One study was performed in the olfactory system of honeybees, which is similar to that of the locust and zebrafish. In this system, oscillations in the local field potential synchronize spiking activity across

different combinations of neurons (Stopfer et al., 1997). They chemically blocked $GABA_A$ receptors, which disrupted the odor-induced oscillations in the local field potential. The result was that the honeybees were impaired in performing fine odor discrimination tasks. The obvious concern over such a result is that it is difficult to know what other effects that picrotoxin, the $GABA_A$ antagonist, might have had on the neural response of the olfactory neurons.

In another study, the phase-locking of V4 neurons to the gamma range of the local field potential was correlated to the behavior on a trial-by-trial basis (Womelsdorf et al., 2006). They found that increased phase locking was correlated with faster reaction times in the detection of a change in a visual stimulus. Selectively attending to a target, which would also decrease reaction times, is also associated with increased phase locking to the gamma range of the local field in area V4 (Fries et al., 2001; Fries et al., 2008). Thus, it is difficult to ascertain whether increased phase locking had a direct effect on behavior or whether it was simply a correlate of greater attention. Regardless, it does suggest that attention does not simply modulate how sensory neurons integrate information, but that attention can alter how sensory activity is coordinated with the rest of the brain.

One limitation of many of these studies was that they simply correlated neural activity with behavior, and did not show that temporally encoded information was causally linked with the perceptual decision. To demonstrate such a causal link, one must turn to methods used to perturb neural activity.

Stimulation of neural tissue

The manipulation of neural activity is a well-established tool in neuroscience research, made famous by Penfield's pioneering work (Penfield, 1958; Penfield and Welch, 1949). There are two main applications for stimulating neural tissue: establishing causality between neural activity in an area of the brain and behavior and serving in sensory neural prosthetics. One classic example of the former was a set of experiments that unequivocally demonstrated that motion perception is at least partially based on the activity of neurons in area MT (Salzman et al., 1990; Salzman et al., 1992). In these experiments, a monkey discriminated between motion in one of two opposing directions, while neurons in area MT preferring one of these two directions were electrically stimulated in a subset of trials. Electrical stimulation was applied, as in many other experiments, by simply lowering a recording electrode into the desired location and passing through biphasic current at the desired current level and frequency. During trials where electrical stimulation was applied, the subjects reported observing the stimulus move in the direction matched to the preference of the stimulated neurons more often. This was one of the first examples that demonstrated that perception can be biased in a predictable fashion by stimulating sensory areas.

Several other experiments have used this same approach to establish causality in both sensory and motor areas of the brain. For sensory areas, it was used to demonstrate causality between the activity in area MT and speed perception (Liu and Newsome, 2005), MST and motion direction (Celebrini and Newsome, 1995), primary

somatosensory cortex and frequency discrimination (Romo et al., 2000), V1 and phosphene perception (Tehovnik and Slocum, 2007). For motor areas, it was used to establish causality between motor and pre-motor cortex with various body and limb movements (Graziano et al., 2002; Luppino et al., 1991; Preuss et al., 1996; Stepniewska et al., 2005), and the superior colliculus, frontal and supplementary eye fields with various types of eye movements (Gottlieb et al., 1993; Harris, 1980; Schlag and Schlag-Rey, 1987; Shibutani et al., 1984). Needless to say, the use of electrical stimulation to establish causality has proved to be an important research tool.

Secondly, there has been much interest in electrical stimulation of neural tissue for use in sensory neural prosthetics. These devices are meant to restore sensory function to patients who have lost a sense due to disease or injury. They work by electrically transmitting sensory information directly to functioning neural tissue, bypassing damaged areas. Currently, the most successful of such devices has been the cochlear implant, which restores a limited sense of hearing by directly stimulating auditory nerves in the cochlea (Middlebrooks et al., 2005; Rieke et al., 1995). Visual neural prosthetics, by comparison, are still in their infancy. Most of the focus has been placed on developing visual prosthetics that stimulate retinal cells (Eckmiller et al., 2005; Hetling and Baig-Silva, 2004; Palanker et al., 2005). However, there has been research into whether stimulating the lateral geniculate nucleus (Pezaris and Reid, 2007) or visual cortex itself (Fernandez et al., 2005) could also work.

The main drawback with electrical microstimulation is that it is neither spatially nor temporally precise. Most of these studies listed above that established causality between the neural activity and behavior have stimulated cortical areas with columnar architecture. This helped ensure that the spatial spread affected mostly neurons with similar preferences. In contrast, for an experiment involving speed perception, in which neurons in MT with similar speed tuning do not cluster, the effect of microstimulation was much less robust (Liu and Newsome, 2005). Even when performing microstimulation experiments with visual features that form columnar architecture, one must still be careful. Using too strong a current can activate multiple columns, effectively adding noise into the subject's visual perception (Murasugi et al., 1993).

Several studies have attempted to accurately measure the extent of spatial spread when microstimulating neural tissue. In a technically challenging experiment, Tolias *et. al.* (2005) microstimulated primary visual cortex of a macaque monkey while performing functional magnetic resonance imaging (fMRI). Although the blood oxygen level captured by fMRI is only an indirect level of neural activity, they reported that even low current levels of 10 micro amps applied for four seconds can perturb activity over 5 millimeters away from the electrode. In a more direct study, Butovas *et. al.* recorded the effect of a electrical stimulation on neural activity using standard multi-electrodes in both anaesthetized rat (Butovas and Schwarz, 2003) and awake mouse (Butovas et al., 2006). Even low current strength applied for less than a millisecond perturbed neural activity one to two millimeters away from the electrode tip.

These two studies by Butovas *et. al.* (2006; 2003) also measured the temporal profile that microstimulation had on the nearby neural activity. Perhaps more problematic than the large spatial spread, they found that even millisecond-long pulses of microstimulation had very long-lasting effects on neural activity. They reported that microstimulation initially induced a brief excitatory response, followed by a long a period of inhibition (over 100 ms). The observation that microstimulation leads to relatively long periods of suppressed neural activity has been previously reported is several studies (Berman et al., 1991; Chung and Ferster, 1998; Contreras et al., 1997; Shao and Burkhalter, 1996; Shao and Burkhalter, 1999). In a knockdown experiment, Butovas *et. al.* (2006) established that $GABA_B$ receptors mediated the inhibition underlying the long lasting suppression caused by microstimulation. The long-lasting effect of microstimulation on neural activity raises questions of the temporal precision of microstimulation that may be of concern for applications such as sensory neural prosthetics.

Recently, there has been much interest in a new method of stimulating neural tissue that greatly improves upon the spatial and temporal precision of electrical microstimulation. The method relies upon expressing the light gated ion-channel channelrhodopsin, usually found in green algae, in targeted neurons (Boyden et al., 2005; Li et al., 2005; Nagel et al., 2002; Nagel et al., 2003). Stimulation using light of a certain frequency opens the channel, rapidly depolarizing the neuron. This technique permits the manipulation of neural activity with temporal precision on the order of milliseconds (Boyden et al., 2005). In addition, by expressing the protein only in the desired types of

neurons, the spatial spread is known. Recent experiments have demonstrated that behavior in transgenic mice (Huber et al., 2008) and *Drosophila* (Suh et al., 2007) expressing the ion-channel can be altered in a predictable manner by photo-stimulation. Although this technique holds much promise for preparations involving lower organisms, it is currently off-limits to animals where it would be impractical to develop transgenic versions, such as primates. Thus, for the moment, electrical stimulation is and will be the method of choice to perturb neural activity in primates for the foreseeable future.

Top-down effects

To understand how neural activity underlies perception, it is insufficient to consider visual processing as a static, bottom-up computation. It is becoming clear that the visual system is highly adaptive (Brenner et al., 2000; Kourtzi and DiCarlo, 2006), and top-down signals play an important role in shaping our perception (Churchland et al., 1994; Engel et al., 2001). The most probable reason why top-down signals are required is that the visual environment simply contains too much information, and the brain is unable to efficiently process it all (Atick, 1992; Van Essen et al., 1992). This has led to a view that attention acts to filter non-attended stimuli, while allowing the attended stimulus to pass (Chun and Marois, 2002; Chun and Potter, 1995; Verghese and Pelli, 1992). As explained below, this view of top downs signals might be over simplistic. The role of the visual system is not simply to transmit stimuli through the different stages of cortex, but rather to extract useful information from the visual scene; top-down signals determine what functions are used to extract this information.

Attention and stimulus encoding

Perhaps the simplest form of top-down processing is spatial attention. Although one's eye might be directed towards one location in visual space, attention can be covertly directed to another location. Behavioral performance in a variety of visual paradigms is improved when attention is directed towards the location of the stimulus (Balz and Hock, 1997; Bashinski and Bacharach, 1980; Carrasco and Yeshurun, 1998; Cook and Maunsell, 2002a; Downing, 1988; Muller and Humphreys, 1991). There still exists a debate as to whether spatial attention changes how a neuron represents a visual stimulus. Many past studies have found that when attention is directed inside the receptive field, the neural response is scaled in multiplicative manner, without changing the sensitivity of the neuron (Cook and Maunsell, 2004; McAdams and Maunsell, 1999a; Treue and Martinez Trujillo, 1999; Treue and Maunsell, 1999). Spatial attention in this case can be though of simply turning up the volume of neurons inside the locus of attention.

Other studies have proposed that spatial attention can affect how neurons integrate incoming signals (Desimone and Duncan, 1995; Moran and Desimone, 1985; Reynolds et al., 1999). More specifically, neurons spike in response to postsynaptic potentials of different amplitudes; the proposal is that top-down signals can change these amplitudes by different amounts amongst the various input sites. This is consistent with past results

that have shown that dendrites can integrate postsynaptic potentials in different fashions depending on the context (Gasparini and Magee, 2006; Spruston, 2008).

In these attention studies mentioned above (Desimone and Duncan, 1995; Moran and Desimone, 1985; Reynolds et al., 1999), two stimuli were presented inside a neuron's receptive field, one that was effective at driving the neuron and one that was not. Attending to the effective stimulus resulted in a strong neural response, while attending to the ineffective stimulus caused a sharp decrease in the response. Thus, it was proposed that attention biases the neuron to respond to only one of the two stimuli. This implies that attention can alter which inputs a neuron will integrate. However, it was pointed out that these results could also be explained by a model in which spatial attention changed the gain of upstream neurons encoding the two stimuli (Ghose and Maunsell, 2008). Thus, it is still not clear from these studies whether attention can simply affect the gain of a neuron or whether it can alter how incoming information is integrated.

Behavioral context and stimulus encoding

Another perspective on this issue has been provided by studies that compared neural responses between awake and anesthetized subjects, in which top-down signals are possibly reduced or eliminated (Lamme et al., 1998; Pack et al., 2001). Both studies employed stimuli in which the percept does not match what a neuron would observe inside its receptive field. In a study by Lamme *et. al.* (1998), one stimulus used was a random dot field, where for 30 ms, dots inside a square patch (no actual square was

shown) would move in one direction and dots outside the square patch would move in the opposite direction. This brief movements caused the subject to perceive a square figure separated from the background, lasting several hundred milliseconds. Neurons of both awake and anesthetized subject's whose receptive field fell outside the square both responded to the transient motion. However, for awake subjects, the response for neurons inside the receptive field was much stronger and longer. For anesthetized animals, the neural response was the same for neurons inside and outside the square. Thus, anesthesia affected how neurons encoded a stimulus if it formed part of a non-local percept. A similar result was found by Pack *et. al.* (2001) in a study on the response of MT neurons to plaid and grating stimuli. A grating stimulus is simply a collection of parallel bars that, when moving, produce a percept of motion in the direction perpendicular to the bars. A plaid stimulus consists of two overlapping grating stimuli at different angles. A moving plaid stimulus can produce a motion percept in any direction regardless of the orientation of the bars. However, inside an MT neuron's receptive field, it will see motion in two directions, both perpendicular to the orientation of the two sets of bars. MT neurons in both awake and anesthetized subjects will respond in a similar fashion to a grating stimulus. However, in awake subjects, the MT response to a plaid stimulus will shift from encoding the two local directions of motion inside its receptive field to the direction of the grating as a whole. This change does not happen for neurons in an anesthetized subject. Taken together, these two studies showed that anesthesia could change which visual feature is encoded by the neuron if the stimulus formed part of a coherent percept. This cannot be explained by a change in gain, and implies that top-down signals can alter how neurons integrate incoming neural activity.

Further evidence that top-down signals can alter how neurons integrate incoming neural activity came from two studies that investigated how a sensory neural response was affected by the behavioral demands of the task. In the first, a monkey was cued to perform one of two behavioral tasks while neurons in V1 were recorded (Li et al., 2004). Neural responses to the same stimulus depended on the task cued to the monkey. Additionally, the neural responses became more informative of the behavioral relevant aspect of the stimulus. In a similar study, Freedman and Assad (2006) trained a monkey associate the direction of a motion stimulus with one of two categories while neural activity from MT and LIP was recorded. The neural response of LIP neurons was poorly modulated by the motion direction, but strongly modulated by which category the motion direction belonged to. The response of MT neurons, on the other hand, was mostly modulated by motion direction. After retraining the monkey to associate different motion direction with the two categories, the response of LIP neurons adapted to reflect the new categories. Once again, these results cannot be explain by a gain change, but can only be explained by the behavioral context affecting how neurons integrate incoming activity. Furthermore, it must be assumed that top-down signals are required to relay this behavioral context to sensory neurons.

Finally, additional evidence that top-down signals can modulate how neurons integrate information has come from studies involving simple perceptual decisions, not necessarily related to attention or top-down signaling. These studies found that the neurons most capable of encoding the stimulus given the demands of the task were the

neurons most correlated with the perceptual decision (Britten et al., 1996; Dodd et al., 2001; Law and Gold, 2008; Parker et al., 2002; Purushothaman and Bradley, 2005; Uka and DeAngelis, 2004). The popular interpretation of this result is that the neurons more informative of the stimulus given the task demands are more influential in forming decisions during the task. A modeling study has suggested that weighting the neurons differently depending how well they encode the stimulus is necessary to be optimal (Jazayeri and Movshon, 2006). Additionally, results from psychophysical showed that percept a stimulus is dependent on task demands have provided further evidence for such a weighting scheme (Jazayeri and Movshon, 2007a; Jazayeri and Movshon, 2007b). The more informative neurons given the task are not innately more heavily weighted, but are only properly weighted after extensive training (Law and Gold, 2008). Thus, as the performance of the subject improves through out training, the correlation between neural sensitivity and its influence in forming the decision increases. Since the more informative neurons are obviously dependent upon the behavioral task, top-down signals would most likely be required to mediate this weighting scheme. These results build upon the proposal of the preceding paragraph in that top-down signals can change how information is transformed between cortical areas.

This thesis, along with many of these past studies mentioned above, rely upon the correlation between neural activity and perception to infer how neural activity is pooled to form decisions. Because of the importance of this correlation, we proceed in Chapter 2 to examine whether the correlation is truly an inherent feature in forming perceptual decisions, or whether other factors play a role.

REFERENCES

Abbott, L. F., and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. Neural Comput *11*, 91-101.

Albright, T. D. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. J Neurophysiol *52*, 1106-1130.

Atick, J. J. (1992). Could information theory provide an ecological theory of sensory processing. Network *3*, 213-251.

Azouz, R., and Gray, C. M. (2003). Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. Neuron *37*, 513-523.

Bair, W., and O'Keefe, L. P. (1998). The influence of fixational eye movements on the response of neurons in area MT of the macaque. Vis Neurosci *15*, 779-786.

Bair, W., Zohary, E., and Newsome, W. T. (2001). Correlated firing in macaque visual area MT: time scales and relationship to behavior. J Neurosci *21*, 1676-1697.

Baker, C. L., Jr., Hess, R. F., and Zihl, J. (1991). Residual motion perception in a "motion-blind" patient, assessed with limited-lifetime random dot stimuli. J Neurosci *11*, 454-461.

Balz, G. W., and Hock, H. S. (1997). The effect of attentional spread on spatial resolution. Vision Res *37*, 1499-1510.

Baranyi, A., Szente, M. B., and Woody, C. D. (1993). Electrophysiological characterization of different types of neurons recorded in vivo in the motor cortex of the cat. II. Membrane parameters, action potentials, current-induced voltage responses and electrotonic structures. J Neurophysiol *69*, 1865-1879.

Bashinski, H. S., and Bacharach, V. R. (1980). Enhancement of perceptual sensitivity as the result of selectively attending to spatial locations. Percept Psychophys *28*, 241-248.

Beeler, G. W. (1967). Visual threshold changes resulting from spontaneous saccadic eye movements. Vision Res *7*, 769-775.

Berman, N. J., Douglas, R. J., Martin, K. A., and Whitteridge, D. (1991). Mechanisms of inhibition in cat visual cortex. J Physiol *440*, 697-722.

Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G., and Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. Nat Neurosci *8*, 1263-1268.

Brenner, N., Bialek, W., and de Ruyter van Steveninck, R. (2000). Adaptive rescaling maximizes information transmission. Neuron *26*, 695-702.

Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S., and Movshon, J. A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis Neurosci *13*, 87-100.

Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. J Neurosci *12*, 4745-4765.

Britten, K. H., Shadlen, M. N., Newsome, W. T., and Movshon, J. A. (1993). Responses of neurons in macaque MT to stochastic motion signals. Vis Neurosci *10*, 1157-1169.

Butovas, S., Hormuzdi, S. G., Monyer, H., and Schwarz, C. (2006). Effects of electrically coupled inhibitory networks on local neuronal responses to intracortical microstimulation. J Neurophysiol *96*, 1227-1236.

Butovas, S., and Schwarz, C. (2003). Spatiotemporal effects of microstimulation in rat neocortex: a parametric study using multielectrode recordings. J Neurophysiol *90*, 3024-3039.

Butts, D. A., Weng, C., Jin, J., Yeh, C. I., Lesica, N. A., Alonso, J. M., and Stanley, G. B. (2007). Temporal precision in the neural code and the timescales of natural vision. Nature *449*, 92-95.

Carrasco, M., and Yeshurun, Y. (1998). The contribution of covert attention to the set-size and eccentricity effects in visual search. J Exp Psychol Hum Percept Perform *24*, 673-692.

Celebrini, S., and Newsome, W. T. (1995). Microstimulation of extrastriate area MST influences performance on a direction discrimination task. J Neurophysiol *73*, 437-448.

Celebrini, S., Thorpe, S., Trotter, Y., and Imbert, M. (1993). Dynamics of orientation coding in area V1 of the awake primate. Vis Neurosci *10*, 811-825.

Chase, S. M., and Young, E. D. (2007). First-spike latency information in single neurons increases when referenced to population onset. Proc Natl Acad Sci U S A *104*, 5175-5180.

Chun, M. M., and Marois, R. (2002). The dark side of visual attention. Curr Opin Neurobiol *12*, 184-189.

Chun, M. M., and Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. J Exp Psychol Hum Percept Perform *21*, 109-127.

Chung, S., and Ferster, D. (1998). Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. Neuron *20*, 1177-1189.

Churchland, P. S., Ramachandran, V. S., and Sejnowski, T. J. (1994). A critique of pure vision. In Large-Scale Neuronal Theories of the Brain, C. Kock, and J. L. Davis, eds. (Cambridge, Massachusetts, MIT Press).

Contreras, D., Durmuller, N., and Steriade, M. (1997). Absence of a prevalent laminar distribution of IPSPs in association cortical neurons of cat. J Neurophysiol *78*, 2742-2753.

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

Cook, E. P., and Maunsell, J. H. (2002b). Dynamics of neuronal responses in macaque MT and VIP during motion detection. Nat Neurosci *5*, 985-994.

Cook, E. P., and Maunsell, J. H. (2004). Attentional modulation of motion integration of individual neurons in the middle temporal visual area. J Neurosci *24*, 7964-7977.

de Lafuente, V., and Romo, R. (2005). Neuronal correlates of subjective sensory experience. Nat Neurosci *8*, 1698-1703.

de Lafuente, V., and Romo, R. (2006). Neural correlate of subjective sensory experience gradually builds up across cortical areas. Proc Natl Acad Sci U S A *103*, 14266-14271.

de Ruyter van Steveninck, R. R., Lewen, G. D., Strong, S. P., Koberle, R., and Bialek, W. (1997). Reproducibility and variability in neural spike trains. Science *275*, 1805-1808.

Delorme, A. (2003). Early cortical orientation selectivity: how fast inhibition decodes the order of spike latencies. J Comput Neurosci *15*, 357-365.

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

Desimone, R., and Schein, S. J. (1987). Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. J Neurophysiol *57*, 835-868.

Destexhe, A., Contreras, D., and Steriade, M. (1999). Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. J Neurosci *19*, 4595-4608.

Destexhe, A., Rudolph, M., and Pare, D. (2003). The high-conductance state of neocortical neurons in vivo. Nat Rev Neurosci *4*, 739-751.

Ditchburn, R. W. (1955). Eye-movements in relation to retinal action. Opt Acta (Lond) *1*.

Dodd, J. V., Krug, K., Cumming, B. G., and Parker, A. J. (2001). Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. J Neurosci *21*, 4809-4821.

Donoghue, J. P., Sanes, J. N., Hatsopoulos, N. G., and Gaal, G. (1998). Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. J Neurophysiol *79*, 159-173.

Downing, C. J. (1988). Expectancy and visual-spatial attention: effects on perceptual quality. J Exp Psychol Hum Percept Perform *14*, 188-202.

Eckmiller, R., Neumann, D., and Baruth, O. (2005). Tunable retina encoders for retina implants: why and how. J Neural Eng *2*, S91-S104.

Engel, A. K., Fries, P., and Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. Nat Rev Neurosci *2*, 704-716.

Engel, A. K., Konig, P., Gray, C. M., and Singer, W. (1990). Stimulus-Dependent Neuronal Oscillations in Cat Visual Cortex: Inter-Columnar Interaction as Determined by Cross-Correlation Analysis. Eur J Neurosci *2*, 588-606.

Fernandez, E., Pelayo, F., Romero, S., Bongard, M., Marin, C., Alfaro, A., and Merabet, L. (2005). Development of a cortical visual neuroprosthesis for the blind: the relevance of neuroplasticity. J Neural Eng *2*, R1-12.

Freedman, D. J., and Assad, J. A. (2006). Experience-dependent representation of visual categories in parietal cortex. Nature *443*, 85-88.

Friedrich, R. W., Habermann, C. J., and Laurent, G. (2004). Multiplexing using synchrony in the zebrafish olfactory bulb. Nat Neurosci *7*, 862-871.

Fries, P., Nikolic, D., and Singer, W. (2007). The gamma cycle. Trends Neurosci *30*, 309-316.

Fries, P., Reynolds, J. H., Rorie, A. E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. Science *291*, 1560-1563.

Fries, P., Womelsdorf, T., Oostenveld, R., and Desimone, R. (2008). The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. J Neurosci *28*, 4823-4835.

Gasparini, S., and Magee, J. C. (2006). State-dependent dendritic computation in hippocampal CA1 pyramidal neurons. J Neurosci *26*, 2088-2100.

Gawne, T. J., Kjaer, T. W., Hertz, J. A., and Richmond, B. J. (1996a). Adjacent visual cortical complex cells share about 20% of their stimulus-related information. Cereb Cortex *6*, 482-489.

Gawne, T. J., Kjaer, T. W., and Richmond, B. J. (1996b). Latency: another potential code for feature binding in striate cortex. J Neurophysiol *76*, 1356-1360.

Gawne, T. J., and Richmond, B. J. (1993). How independent are the messages carried by adjacent inferior temporal cortical neurons? J Neurosci *13*, 2758-2771.

Ghose, G. M., and Maunsell, J. H. (2008). Spatial summation can explain the attentional modulation of neuronal responses to multiple stimuli in area V4. J Neurosci *28*, 5115-5126.

Gollisch, T., and Meister, M. (2008). Rapid neural coding in the retina with relative spike latencies. Science *319*, 1108-1111.

Gottlieb, J. P., Bruce, C. J., and MacAvoy, M. G. (1993). Smooth eye movements elicited by microstimulation in the primate frontal eye field. J Neurophysiol *69*, 786-799.

Gray, C. M., Engel, A. K., Konig, P., and Singer, W. (1990). Stimulus-Dependent Neuronal Oscillations in Cat Visual Cortex: Receptive Field Properties and Feature Dependence. Eur J Neurosci *2*, 607-619.

Graziano, M. S., Taylor, C. S., and Moore, T. (2002). Complex movements evoked by microstimulation of precentral cortex. Neuron *34*, 841-851.

Gross, C. G., Rocha-Miranda, C. E., and Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the Macaque. J Neurophysiol *35*, 96-111.

Grunewald, A., Bradley, D. C., and Andersen, R. A. (2002). Neural correlates of structure-from-motion perception in macaque V1 and MT. J Neurosci *22*, 6195-6207.

Gu, Y., DeAngelis, G. C., and Angelaki, D. E. (2007). A functional link between area MSTd and heading perception based on vestibular signals. Nat Neurosci *10*, 1038-1047.

Gur, M., Beylin, A., and Snodderly, D. M. (1997). Response variability of neurons in primary visual cortex (V1) of alert monkeys. J Neurosci *17*, 2914-2920.

Harris, L. R. (1980). The superior colliculus and movements of the head and eyes in cats. J Physiol *300*, 367-391.

Hegde, J. (2008). Time course of visual perception: Coarse-to-fine processing and beyond. Prog Neurobiol *84*, 405-439.

Hess, R. H., Baker, C. L., Jr., and Zihl, J. (1989). The "motion-blind" patient: low-level spatial and temporal filters. J Neurosci *9*, 1628-1640.

Hetling, J. R., and Baig-Silva, M. S. (2004). Neural prostheses for vision: designing a functional interface with retinal neurons. Neurol Res *26*, 21-34.

Hopfield, J. J. (1995). Pattern recognition computation using action potential timing for stimulus representation. Nature *376*, 33-36.

Hubel, D. H., and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J Physiol *160*, 106-154.

Hubel, D. H., and Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. J Physiol *195*, 215-243.

Hubel, D. H., and Wiesel, T. N. (1974). Sequence regularity and geometry of orientation columns in the monkey striate cortex. J Comp Neurol *158*, 267-293.

Huber, D., Petreanu, L., Ghitani, N., Ranade, S., Hromadka, T., Mainen, Z., and Svoboda, K. (2008). Sparse optical microstimulation in barrel cortex drives learned behaviour in freely moving mice. Nature *451*, 61-64.

Huk, A. C., and Shadlen, M. N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J Neurosci *25*, 10420-10436.

Jazayeri, M., and Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. Nat Neurosci *9*, 690-696.

Jazayeri, M., and Movshon, J. A. (2007a). A new perceptual illusion reveals mechanisms of sensory decoding. Nature *446*, 912-915.

Jazayeri, M., and Movshon, J. A. (2007b). Integration of sensory evidence in motion discrimination. J Vis *7*, 7 1-7.

Konig, P., Engel, A. K., Roelfsema, P. R., and Singer, W. (1995). How precise is neuronal synchronization? Neural Comput *7*, 469-485.

Kourtzi, Z., and DiCarlo, J. J. (2006). Learning and neural plasticity in visual object recognition. Curr Opin Neurobiol *16*, 152-158.

Krauskopf, J. (1967). Lack of inhibition during involuntary saccades. Am J Psychol *79*, 73-81.

Lamme, V. A., Zipser, K., and Spekreijse, H. (1998). Figure-ground activity in primary visual cortex is suppressed by anesthesia. Proc Natl Acad Sci U S A *95*, 3263-3268.

Lampl, I., and Yarom, Y. (1993). Subthreshold oscillations of the membrane potential: a functional synchronizing and timing device. J Neurophysiol *70*, 2181-2186.

Law, C. T., and Gold, J. I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat Neurosci *11*, 505-513.

Leopold, D. A., and Logothetis, N. K. (1998). Microsaccades differentially modulate neural activity in the striate and extrastriate visual cortex. Exp Brain Res *123*, 341-345.

Li, W., Piech, V., and Gilbert, C. D. (2004). Perceptual learning and top-down influences in primary visual cortex. Nat Neurosci *7*, 651-657.

Li, X., Gutierrez, D. V., Hanson, M. G., Han, J., Mark, M. D., Chiel, H., Hegemann, P., Landmesser, L. T., and Herlitze, S. (2005). Fast noninvasive activation and inhibition of neural and network activity by vertebrate rhodopsin and green algae channelrhodopsin. Proc Natl Acad Sci U S A *102*, 17816-17821.

Liu, J., and Newsome, W. T. (2005). Correlation between speed perception and neural activity in the middle temporal visual area. J Neurosci *25*, 711-722.

Luppino, G., Matelli, M., Camarda, R. M., Gallese, V., and Rizzolatti, G. (1991). Multiple representations of body movements in mesial area 6 and the adjacent cingulate cortex: an intracortical microstimulation study in the macaque monkey. J Comp Neurol *311*, 463-482.

Mainen, Z. F., and Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons. Science *268*, 1503-1506.

Martinez-Conde, S., Macknik, S. L., and Hubel, D. H. (2000). Microsaccadic eye movements and firing of single cells in the striate cortex of macaque monkeys. Nat Neurosci *3*, 251-258.

Martinez-Conde, S., Macknik, S. L., and Hubel, D. H. (2002). The function of bursts of spikes during visual fixation in the awake primate lateral geniculate nucleus and primary visual cortex. Proc Natl Acad Sci U S A *99*, 13920-13925.

Martinez-Conde, S., Macknik, S. L., Troncoso, X. G., and Dyar, T. A. (2006). Microsaccades counteract visual fading during fixation. Neuron *49*, 297-305.

Mason, A., Nicoll, A., and Stratford, K. (1991). Synaptic transmission between individual pyramidal neurons of the rat visual cortex in vitro. J Neurosci *11*, 72-84.

Masse, N. Y., and Cook, E. P. (2008). The effect of middle temporal spike phase on sensory encoding and correlates with behavior during a motion-detection task. J Neurosci *28*, 1343-1355.

Matsumura, M., Cope, T., and Fetz, E. E. (1988). Sustained excitatory synaptic input to motor cortex neurons in awake animals revealed by intracellular recording of membrane potentials. Exp Brain Res *70*, 463-469.

Maunsell, J. H., Ghose, G. M., Assad, J. A., McAdams, C. J., Boudreau, C. E., and Noerager, B. D. (1999). Visual response latencies of magnocellular and parvocellular LGN neurons in macaque monkeys. Vis Neurosci *16*, 1-14.

Maunsell, J. H., and Treue, S. (2006). Feature-based attention in visual cortex. Trends Neurosci *29*, 317-322.

Maunsell, J. H., and Van Essen, D. C. (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*, 1127-1147.

Mazurek, M. E., and Shadlen, M. N. (2002). Limits to the temporal fidelity of cortical spike rate signals. Nat Neurosci *5*, 463-471.

McAdams, C. J., and Maunsell, J. H. (1999a). Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. J Neurosci *19*, 431-441.

McAdams, C. J., and Maunsell, J. H. (1999b). Effects of attention on the reliability of individual neurons in monkey visual cortex. Neuron *23*, 765-773.

Middlebrooks, J. C., Bierer, J. A., and Snyder, R. L. (2005). Cochlear implants: the view from the brain. Curr Opin Neurobiol *15*, 488-493.

Mikami, A., Newsome, W. T., and Wurtz, R. H. (1986). Motion selectivity in macaque visual cortex. I. Mechanisms of direction and speed selectivity in extrastriate area MT. J Neurophysiol *55*, 1308-1327.

Montemurro, M. A., Rasch, M. J., Murayama, Y., Logothetis, N. K., and Panzeri, S. (2008). Phase-of-firing coding of natural visual stimuli in primary visual cortex. Curr Biol *18*, 375-380.

Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. Science *229*, 782-784.

Muller, H. J., and Humphreys, G. W. (1991). Luminance-increment detection: capacity-limited or not? J Exp Psychol Hum Percept Perform *17*, 107-124.

Murasugi, C. M., Salzman, C. D., and Newsome, W. T. (1993). Microstimulation in visual area MT: effects of varying pulse amplitude and frequency. J Neurosci *13*, 1719-1729.

Nagel, G., Ollig, D., Fuhrmann, M., Kateriya, S., Musti, A. M., Bamberg, E., and Hegemann, P. (2002). Channelrhodopsin-1: a light-gated proton channel in green algae. Science *296*, 2395-2398.

Nagel, G., Szellas, T., Huhn, W., Kateriya, S., Adeishvili, N., Berthold, P., Ollig, D., Hegemann, P., and Bamberg, E. (2003). Channelrhodopsin-2, a directly light-gated cation-selective membrane channel. Proc Natl Acad Sci U S A *100*, 13940-13945.

Newsome, W. T., Britten, K. H., and Movshon, J. A. (1989). Neuronal correlates of a perceptual decision. Nature *341*, 52-54.

O'Keefe, J., and Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus *3*, 317-330.

Otmakhov, N., Shirke, A. M., and Malinow, R. (1993). Measuring the impact of probabilistic transmission on neuronal output. Neuron *10*, 1101-1111.

Pack, C. C., Berezovskii, V. K., and Born, R. T. (2001). Dynamic properties of neurons in cortical area MT in alert and anaesthetized macaque monkeys. Nature *414*, 905-908.

Palanker, D., Vankov, A., Huie, P., and Baccus, S. (2005). Design of a high-resolution optoelectronic retinal prosthesis. J Neural Eng *2*, S105-120.

Pare, D., Shink, E., Gaudreau, H., Destexhe, A., and Lang, E. J. (1998). Impact of spontaneous synaptic activity on the resting properties of cat neocortical pyramidal neurons In vivo. J Neurophysiol *79*, 1450-1460.

Parker, A. J., Krug, K., and Cumming, B. G. (2002). Neuronal activity and its links with the perception of multi-stable figures. Philos Trans R Soc Lond B Biol Sci *357*, 1053-1062.

Parker, A. J., and Newsome, W. T. (1998). Sense and the single neuron: probing the physiology of perception. Annu Rev Neurosci *21*, 227-277.

Penfield, W. (1958). Some Mechanisms of Consciousness Discovered during Electrical Stimulation of the Brain. Proc Natl Acad Sci U S A *44*, 51-66.

Penfield, W., and Welch, K. (1949). Instability of response to stimulation of the sensorimotor cortex of man. J Physiol *109*, 358-365, illust.

Perez-Orive, J., Mazor, O., Turner, G. C., Cassenaer, S., Wilson, R. I., and Laurent, G. (2002). Oscillations and sparsening of odor representations in the mushroom body. Science *297*, 359-365.

Perrett, D. I., Rolls, E. T., and Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. Exp Brain Res *47*, 329-342.

Peters, A. (1987). Synaptic specificity in the cerebral cortex. In Synaptic function, G. Edelman, W. Gall, and W. Cowan, eds.

Pezaris, J. S., and Reid, R. C. (2007). Demonstration of artificial visual percepts generated through thalamic microstimulation. Proc Natl Acad Sci U S A *104*, 7670-7675.

Prescott, S. A., Ratte, S., De Koninck, Y., and Sejnowski, T. J. (2006). Nonlinear interaction between shunting and adaptation controls a switch between integration and coincidence detection in pyramidal neurons. J Neurosci *26*, 9084-9097.

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

Priebe, N. J., and Ferster, D. (2005). Direction selectivity of excitation and inhibition in simple cells of the cat primary visual cortex. Neuron *45*, 133-145.

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

Raiguel, S. E., Xiao, D. K., Marcar, V. L., and Orban, G. A. (1999). Response latency of macaque area MT/V5 neurons and its relationship to stimulus parameters. J Neurophysiol *82*, 1944-1956.

Reinagel, P., and Reid, R. C. (2000). Temporal coding of visual information in the thalamus. J Neurosci *20*, 5392-5400.

Reynolds, J. H., Chelazzi, L., and Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. J Neurosci *19*, 1736-1753.

Rieke, F., Bodnar, D. A., and Bialek, W. (1995). Naturalistic stimuli increase the rate and efficiency of information transmission by primary auditory afferents. Proc Biol Sci *262*, 259-265.

Rieke, F., Warland, D., van Steveninck, R., and Bialek, W. (1999). Spikes: exploring the neural code, MIT Press).

Romo, R., Hernandez, A., Zainos, A., Brody, C. D., and Lemus, L. (2000). Sensing without touching: psychophysical performance based on cortical microstimulation. Neuron *26*, 273-278.

Romo, R., Hernandez, A., Zainos, A., and Salinas, E. (2003). Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. Neuron *38*, 649-657.

Rudolph, M., and Destexhe, A. (2003). The discharge variability of neocortical neurons during high-conductance states. Neuroscience *119*, 855-873.

Saito, H., Yukie, M., Tanaka, K., Hikosaka, K., Fukada, Y., and Iwai, E. (1986). Integration of direction signals of image motion in the superior temporal sulcus of the macaque monkey. J Neurosci *6*, 145-157.

Salzman, C. D., Britten, K. H., and Newsome, W. T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. Nature *346*, 174-177.

Salzman, C. D., Murasugi, C. M., Britten, K. H., and Newsome, W. T. (1992). Microstimulation in visual area MT: effects on direction discrimination performance. J Neurosci *12*, 2331-2355.

Schaefer, A. T., Angelo, K., Spors, H., and Margrie, T. W. (2006). Neuronal oscillations enhance stimulus discrimination by ensuring action potential precision. PLoS Biol *4*, e163.

Schaette, R., Gollisch, T., and Herz, A. V. (2005). Spike-train variability of auditory neurons in vivo: dynamic responses follow predictions from constant stimuli. J Neurophysiol *93*, 3270-3281.

Schlag, J., and Schlag-Rey, M. (1987). Evidence for a supplementary eye field. J Neurophysiol *57*, 179-200.

Schoppik, D., Nagel, K. I., and Lisberger, S. G. (2008). Cortical mechanisms of smooth eye movements revealed by dynamic covariations of neural and behavioral responses. Neuron *58*, 248-260.

Schwartz, O., and Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. Nat Neurosci *4*, 819-825.

Shadlen, M. N., Britten, K. H., Newsome, W. T., and Movshon, J. A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. J Neurosci *16*, 1486-1510.

Shadlen, M. N., and Movshon, J. A. (1999). Synchrony unbound: a critical evaluation of the temporal binding hypothesis. Neuron *24*, 67-77, 111-125.

Shadlen, M. N., and Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J Neurosci *18*, 3870-3896.

Shao, Z., and Burkhalter, A. (1996). Different balance of excitation and inhibition in forward and feedback circuits of rat visual cortex. J Neurosci *16*, 7353-7365.

Shao, Z., and Burkhalter, A. (1999). Role of GABAB receptor-mediated inhibition in reciprocal interareal pathways of rat visual cortex. J Neurophysiol *81*, 1014-1024.

Shibutani, H., Sakata, H., and Hyvarinen, J. (1984). Saccade and blinking evoked by microstimulation of the posterior parietal association cortex of the monkey. Exp Brain Res *55*, 1-8.

Simoncelli, E. P., and Olshausen, B. A. (2001). Natural image statistics and neural representation. Annu Rev Neurosci *24*, 1193-1216.

Snowden, R. J., Treue, S., and Andersen, R. A. (1992). The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. Exp Brain Res *88*, 389-400.

Sperling, G. (1990). In Eye Movements and Their Role in Visual and Cognitive Processes, E. Kowler, ed. (Amsterdm, Elsevier), pp. 307-351.

Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. Nat Rev Neurosci *9*, 206-221.

Stepniewska, I., Fang, P. C., and Kaas, J. H. (2005). Microstimulation reveals specialized subregions for different complex movements in posterior parietal cortex of prosimian galagos. Proc Natl Acad Sci U S A *102*, 4878-4883.

Steriade, M., Timofeev, I., and Grenier, F. (2001). Natural waking and sleep states: a view from inside neocortical neurons. J Neurophysiol *85*, 1969-1985.

Stopfer, M., Bhagavan, S., Smith, B. H., and Laurent, G. (1997). Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. Nature *390*, 70-74.

Suh, G. S., Ben-Tabou de Leon, S., Tanimoto, H., Fiala, A., Benzer, S., and Anderson, D. J. (2007). Light activation of an innate olfactory avoidance response in Drosophila. Curr Biol *17*, 905-908.

Tehovnik, E. J., and Slocum, W. M. (2007). Phosphene induction by microstimulation of macaque V1. Brain Res Rev *53*, 337-343.

Thorpe, S., Delorme, A., and Van Rullen, R. (2001). Spike-based strategies for rapid processing. Neural Netw *14*, 715-725.

Tolias, A. S., Sultan, F., Augath, M., Oeltermann, A., Tehovnik, E. J., Schiller, P. H., and Logothetis, N. K. (2005). Mapping cortical activity elicited with electrical microstimulation using FMRI in the macaque. Neuron *48*, 901-911.

Treue, S., and Martinez Trujillo, J. C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. Nature *399*, 575-579.

Treue, S., and Maunsell, J. H. (1999). Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. J Neurosci *19*, 7591-7602.

Turner, G. C., Bazhenov, M., and Laurent, G. (2008). Olfactory representations by Drosophila mushroom body neurons. J Neurophysiol *99*, 734-746.

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

Uka, T., Tanabe, S., Watanabe, M., and Fujita, I. (2005). Neural correlates of fine depth discrimination in monkey inferior temporal cortex. J Neurosci *25*, 10796-10802.

Van Essen, D. C., Anderson, C. H., and Felleman, D. J. (1992). Information processing in the primate visual system: an integrated systems perspective. Science *255*, 419-423.

van Kan, P. L., Scobey, R. P., and Gabor, A. J. (1985). Response covariance in cat visual cortex. Exp Brain Res *60*, 559-563.

Van Rullen, R., and Thorpe, S. J. (2001). Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex. Neural Comput *13*, 1255-1283.

Verghese, P., and Pelli, D. G. (1992). The information capacity of visual attention. Vision Res *32*, 983-995.

Wilent, W. B., and Contreras, D. (2005). Dynamics of excitation and inhibition underlying stimulus selectivity in rat somatosensory cortex. Nat Neurosci *8*, 1364-1370.

Williams, Z. M., Elfar, J. C., Eskandar, E. N., Toth, L. J., and Assad, J. A. (2003). Parietal activity and the perceived direction of ambiguous apparent motion. Nat Neurosci *6*, 616-623.

Womelsdorf, T., Fries, P., Mitra, P. P., and Desimone, R. (2006). Gamma-band synchronization in visual cortex predicts speed of change detection. Nature *439*, 733-736.

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*, 549-573.

Zoccolan, D., Pinato, G., and Torre, V. (2002). Highly variable spike trains underlie reproducible sensorimotor responses in the medicinal leech. J Neurosci *22*, 10790-10800.

Zohary, E., Celebrini, S., Britten, K. H., and Newsome, W. T. (1994a). Neuronal plasticity that underlies improvement in perceptual performance. Science *263*, 1289-1292.

Chapter 2

The effect of microsaccades on the correlation between neural activity and behavior in areas MT, VIP and LIP

Abstract

It is widely reported that the activity of single neurons in visual cortex is correlated with the perceptual decision of the subject. The strength of this correlation has implications for the neuronal populations generating the percepts. Here we asked whether microsaccades, which are small, involuntary eye movements, contribute to the correlation between neural activity and behavior. We analyzed data from three different visual detection experiments, with neural recordings from the middle temporal (MT), lateral intraparietal (LIP) and ventral intraparietal (VIP) areas. All three experiments used random dot motion stimuli, with the animals required to detect a transient or sustained change in the speed or strength of motion. We found that microsaccades suppressed neural activity and inhibited detection of the motion stimulus, contributing to the correlation between neural activity and detection behavior. Microsaccades accounted for as much as 19% of the correlation for area MT, 21% for area LIP and 17% for VIP. While microsaccades only explain part of the correlation between neural activity and behavior, their effect has implications when considering the neuronal populations underlying perceptual decisions.

Introduction

Microsaccades are small, involuntary eye movements thought to counteract drift of the eyes (Cornsweet, 1956) and visual fading (Martinez-Conde et al., 2006), as well as to improve discrimination of high spatial frequencies (Bridgeman and Palca, 1980; Rucci et al., 2007). Microsaccades may also increase visual detection thresholds (Ditchburn, 1955; but see Krauskopf, 1966; Beeler, 1967; Sperling, 1990). In parallel, microsaccades have been shown to affect neural activity in the visual system, including in the lateral geniculate nucleus (Martinez-Conde et al., 2002), V1 (Leopold and Logothetis, 1998; Martinez-Conde et al., 2000; Snodderly et al., 2001), and extrastriate cortex (Bair and O'Keefe, 1998; Leopold and Logothetis, 1998). Behavioral and neuronal effects of microsaccades have not been examined in the same experiment—but if microsaccades can affect *both* perception and neuronal activity, it is possible that microsaccades could contribute to apparent correlations between the two.

Increasing evidence suggests that the activity of single neurons in visual cortex can be correlated with behavioral performance on a trial-by-trial basis. First observed in neurons located in the middle temporal (MT) area (Britten et al., 1996), this finding has been replicated in a multitude of sensory cortical areas under various experimental paradigms (for reviews, see Parker and Newsome, 1998; Romo and Salinas, 2001). The strength of the correlation between neural activity and behavior has been termed choice probability (CP) for discrimination tasks (Britten et al., 1996) or detect probability (DP) for detection tasks (Britten et al., 1996; Cook and Maunsell, 2002b). This correlation has

been used to identify which cortical areas are involved in a perceptual decision (Britten et al., 1996; Cook and Maunsell, 2002b; Grunewald et al., 2002; Williams et al., 2003; de Lafuente and Romo, 2005; Liu and Newsome, 2005; de Lafuente and Romo, 2006; Gu et al., 2007), the number of neurons involved or more importantly, correlations between these neurons (Shadlen et al., 1996; Bair et al., 2001; Dodd et al., 2001; Schoppik et al., 2008), how the activity of different neurons are weighted to form the decision (Britten et al., 1996; Purushothaman and Bradley, 2005; Uka et al., 2005; Gu et al., 2007) and how different phases of neural activity are correlated with detection (Masse and Cook, 2008).

According to this interpretation, trial-by-trial variation in a neuron's response to a repeated stimulus is assumed to be due to noise. The observation that this variability is correlated with the animal's choice suggests that the brain is "listening to" that neuron (or a population of neurons with correlated noise) in order to make the choice. However, one must be careful to rule out other sources of correlation between neural activity and behavior. For example, trial-by-trial variability in the stimulus could cause correlation between activity and behavior (Cook and Maunsell, 2002b). Because microsaccades can likewise affect both behavior and neuronal responses, they are another potential source of correlation. In this study, we addressed whether microsaccades contribute to the choice-related correlation in visual and parietal cortex during detection tasks.

We examined the simultaneous effects of microsaccades on behavior and neuronal responses, in three different visual detection paradigms, covering areas MT, LIP and VIP, and involving seven macaque monkeys. Examining the effects of microsaccades

on both behavior and neuronal responses allowed us to test the hypothesis that the correlation between neural activity and perceptual behavior is wholly or partly produced by microsaccades. We found that microsaccades inhibited the monkeys' ability to detect the visual stimuli in our three experimental paradigms. We also found that microsaccades, on average, suppressed neural activity across the cortical areas examined in our study, consistent with past results (Bair and O'Keefe, 1998; Leopold and Logothetis, 1998). The effect of microsaccades on detect probability was somewhat variable and depended on the brain area and task design. On average, the contribution of microsaccades to detect probability ranged from 7 to 19% for area MT, 21% for area LIP and 0 to 17% for area VIP depending on the time points analyzed.

Methods

To investigate the effect of microsaccades on neural activity and visual perception, we analyzed data from three different motion-detection experiments, all employing a random dot stimulus. In the first experiment, two monkeys were trained to respond to a transient increase in speed while neural activity was recorded from areas MT and LIP. In the second experiment, two monkeys were trained to detect a coherent motion step while neural activity was recorded from MT and VIP. In the third experiment, three monkeys were trained to detect a brief pulse of coherent motion in two slightly different paradigms. Neural activity was recorded from area MT in two of the monkeys. All four paradigms (the third experiment had two variations) are outline in Fig. 1.

Behavioral task for the speed-pulse experiment

Two monkeys (Macaca mulatta) were trained to perform a spatially cued speed-pulse detection task. At the beginning of a trial the stimulus consisted of a central fixation spot and two annuli, one red and one green, in opposite hemifields at equal eccentricity (Fig. 1*A*). The monkey had to maintain gaze within a fixation window throughout the trial (2˚ x 2˚ square, centered on a fixation spot). After the monkey fixated, there was a 500-ms delay before two fields of coherently moving random dots appeared within the annuli. The monkey's task was to detect a transient increase in the speed (53 ms, 4 video frames) of either dot patch and respond by releasing a touch bar within a requisite time window

(200 – 600 ms). The color of the fixation point (red or green) cued the monkey as to which patch (surrounded by red or green annulus) was more likely to contain the motion pulse (85% valid cues, 15% invalid cues). On 40% of trials the fixation point color cue switched at an unpredictable time during the trial to indicate that the likely motion pulse location had switched. Each trial had at most one cue switch. After an initial fixed delay of 400 ms, additional delays until motion pulses and cue switches, as well as between cue switches and motion pulses, were selected randomly from an exponential distribution (mean = 1 s). This main purpose of the study was to examine effects of switching spatial attention, but those results are not relevant to the questions examined here.

For the analysis of the correlation between microsaccades and behavior (Fig. 3*A*) we only used trials with valid cues. For the analysis of the neural response to coherent motion (Fig. 4*A-B*) or the correlation between neural activity and behavior (Fig. 5*A-B*) we only used trials in which the speed pulse was validly cued to occur inside the neuron's receptive field.

Visual stimulus for speed-pulse experiment

Stimuli were presented on a computer monitor positioned 57 cm in front of the animal (40˚ x 30˚, 75 Hz refresh, 1152 x 870 resolution). Background luminance was near black (0.001 cd/m^2). The fixation point was a 0.4 degree diameter red or green circle

(luminance in cd/m^2: monkey M, red: 2.7, green: 3.0; monkey B, red: 2.4, green: 5.2). Dot-patch stimuli consisted of 100% coherently moving, unlimited-lifetime, random dots. Dots were squares with 0.1-degree sides, at a density of 7 dots/$degree^2$ and moving at 12 degrees/second. Dot luminance was 0.01 cd/m^2. Annuli surrounding the moving dot patches were 0.5 degrees thick and separated from the perimeter of the dot patches by 0.5 degrees.

Where possible, dot patches were placed in the center of the receptive field of the recorded neuron. The dot-patch motion in the receptive field was set in the neuron's preferred direction as determined by a direction-mapping task that we ran before the main task for each neuron. The other dot patch was always placed at the equivalent position reflected across the fixation point and had the opposite direction of motion. The size of the dot patches was scaled with eccentricity (ranging from 4.5 to 9.4 degrees in diameter). The magnitude of the speed pulse was chosen to maintain valid correct performance in the target range (65 – 75% correct) and varied from session to session (range for monkey B: 1.6x to 2.5x, range for monkey M: 1.35x to 1.7x).

Data collection for the speed-pulse experiment

The recording chamber was placed at stereotactic coordinates P3 L10, which allowed a dorsal approach to areas MT and LIP. The chamber was outfitted with a guide-

tube/grid system (Crist Instrument). MRI was used to confirm sulcal anatomy and chamber placement. Single unit recordings were conducted using tungsten microelectrodes (Frederick Haer & Co, 75 μm diameter, 5 MΩ impedance). Single unit action potentials were isolated using a dual window discriminator (Bak Electronics) and recorded at 1-ms resolution. Horizontal and vertical eye position were monitored using a scleral search coil (Riverbend Instruments) and recorded at 200 Hz. Spike and eye-position recording, stimulus presentation and task control were handled by a Macintosh computer running custom software with a computer interface (ITC-18, Instrutech Corporation).

MT and LIP cells were identified by reference to sulcal anatomy and characteristic physiology. MT cells were characterized by highly direction-selective receptive fields with diameters roughly equal to eccentricity (Maunsell and Van Essen, 1983, 1987). LIP cells were characterized by robust, spatially-tuned responses in a memory delayed saccade task (Colby et al., 1996). Additionally, cells were considered within the target area if they were encountered between cells with characteristic properties. All such stably isolated units were recorded. In all, we recorded from 118 LIP neurons and from 67 MT neurons.

Behavioral task for the motion-step experiment

The data set analyzed for this study comes from several previous studies (Cook and Maunsell, 2002b, a, 2004; Masse and Cook, 2008). Two monkeys were trained to perform a spatially cued motion detection task (Fig. 1*B*). The monkey initiated the trial by depressing a lever and fixating upon a central point. The monkeys were trained to release a lever when coherent motion began in one of two random dot patches diametrically opposite of the fixation point. After the cue was presented, 0% coherent motion began in the two patches followed by coherent motion occurring in one of the two patches at a random time (flat hazard function) 500 to 8000 ms afterwards. The location of the coherent motion was cued to the monkey at the start of the trial with static dots and this cue was valid on 80% of the trials. The strength of coherent motion was varied between three levels (low, medium and high) with the monkey correctly detecting the coherent motion 50%, 90% and 99% of the time for the three levels. The analysis of the behavioral effect of microsaccades (Fig. 4*B*) only used trials with low-level coherent motion. Analysis of the effect of microsaccades on neural activity during 0% coherent motion used only trials where the monkey was cued to attend inside the neuron's receptive field (Fig 5*C-D*, left panel) used only used trials where the monkey was cued to attend inside the neuron's receptive field. Analysis involving the neural response to coherent motion (Fig. 5C-*D, right panel*) or the partial correlation analysis (Fig. 6*C-D*) only involved trials using low-level coherent motion that was validly cued to occur inside the neuron's receptive field.

In all experimental sessions, the direction and the speed of the coherent motion were matched to the preferred direction and speed of the neuron under study. The

coherent motion lasted 750 ms and the monkey had to release the lever from 200 to 750 ms after the onset of coherent motion to obtain a reward. Trials where the monkey failed to release the lever or released the lever too late were deemed missed trials. Only correct and missed trials were included in the analysis. Trials where the monkey released the lever too early or was unable to maintain fixation were discarded from the analysis.

Visual stimulus for the motion-step experiment

The animal sat 62 cm from a computer monitor (+/-17° × +/-13° of visual angle; 1600 × 1200 pixels; 75 Hz refresh). The stimuli consisted of two patches of white dots (each dot 0.25° diameter; 78 cd/m^2) on a dark gray background (12 cd/m^2) with a dot density of 2.1 dots/$degree^2$. Each patch of dots was updated on every other video frame (approximately every 27 ms) using the following procedure. The dots in each patch were evenly divided into two groups. On each update, one group was replaced with new, randomly positioned dots, whereas dots in the other group were displaced by a fixed distance. The dots in this latter group determined the motion coherence. For 0% coherence, all the dots in this group moved a fixed distance in a random direction. For coherent motion greater than zero, a proportion of the dots moved with a fixed distance in the same direction. This proportion determined the strength of the coherent motion. On the next update (27 ms later), the groups were switched. This arrangement insured that all the dots had a lifetime of four video frames (i.e., two stimulus updates) before they were replaced.

Data collection for the motion-step experiment

Recordings were made from well-isolated single neurons in area MT and VIP in both animals, using standard extracellular recording techniques (Gibson and Maunsell, 1997). After a neuron was isolated, the receptive field (RF) was mapped using a manually controlled bar while the animal fixated on a central spot. The preferred speed was also judged using a bar moved by hand. The animals were trained to perform the task at slow or moderate motion speeds, so neurons with a preferred speed between 4 and 12°/sec were usually selected. Once the RF location, size, preferred direction, and speed were determined, the motion detection task was then run, and the neuron was recorded from for as long as possible. The number of completed trials per coherence level for the motion detection task ranged from 15 to 175 (median, 35). The monkey's performance varied with patch location, size, and motion speed, which were determined by the response properties of the neuron under study. Consequently, different neurons were tested with different coherence levels. The animal's eye position was measured every 5 ms using a scleral search coil (Robinson, 1963; Judge et al., 1980) and the times of action potentials were recorded to the nearest millisecond. In total, 93 MT and 104 VIP neurons were recorded in two monkeys.

Behavioral task for the motion-pulse experiment

Data from two different experimental paradigms, both involving the detection of

transient motion, were combined for the purpose of this analysis. In the first (Fig. 1C), two monkeys were trained to detect a 50 ms coherent motion pulse that occurred in one or both random dot patches. Monkeys fixated on a central point and depressed a lever to initiate the trial, after which motion in two random dot patches started moving with 0% coherence. Dots could either move in the neuron's preferred or null direction, always by the same distance. Coherent motion would then appear in one or both random dot patches at a random time (flat hazard function), anytime from 500 to 10000 ms after the start trial. The coherent motion lasted 50 ms and was of a consistent strength for each random dot patch throughout the experimental session. The ratio of trials with motion in two patches to motion in one patch was varied between two to one and four to one. The monkey was rewarded with a drop by releasing the lever from 200 to 800 ms after the start of coherent motion.

In the second experiment, a single monkey was trained to detect a 33 ms coherent motion pulse in a single random patch. As above, coherent motion would begin at a random time from 500 to 10000 ms after the start of the trial, and the monkey had to release the lever from 150 to 650 ms after the motion pulse to receive a reward. This experiment was used for a microstimulation study, and some trials also contained a 33 ms probe of subthreshold microstimulation to area MT or a 33 ms probe of subthreshold motion coherence at various times relative to the main motion signal. For the purpose of this study, all trials were considered the same for this analysis.

Visual stimulus for motion-pulse experiment

Stimuli were presented on a computer monitor positioned 57 cm in front of the animal (120 Hz refresh, 1600 x 1200 resolution). Random dot patches consisted of white dots moving in either the preferred or null direction of the neuron, by always the same distance. In the first motion-pulse experiment, dots were assigned a probability of moving in the preferred direction for each frame, whereas in the second motion-pulse experiment the total number of dots moving in the preferred direction was controlled for each frame. For example, during 0 % coherent motion, each dot has a 50% probability moving the preferred direction whereas in the second experiment, exactly 50% of the dots each frame would move in the preferred direction. For both experiments, the direction that a dot moved in one frame was not related to its direction for any other frame. If a dot moved outside the random dot patch, it would be randomly replotted anywhere in the opposite hemifield of the random dot patch. Dots were circles with 0.15-degree radius, at a density of 10 dots/degree^2.

Data collection for the motion-pulse experiment

Recordings were made from well-isolated single neurons in area MT animals in the first motion-pulse experiment and from multi-unit activity in area MT in the second motion-pulse experiment. The RF was mapped using a manually controlled bar while the animal fixated on a central spot. The preferred speed and direction by recording activity in response to 100% coherent motion in various speeds and directions while the monkey

fixated. Once the RF location, size, preferred direction, and speed were determined, the motion detection task was then run, and the neuron(s) was recorded from for as long as possible. The monkey's performance varied with patch location, size, and motion speed, which were determined by the response properties of the neuron under study. Consequently, different neurons were tested with different coherence levels. The animal's eye position was measured every 5 ms using video-tracking system (ASL 6000, Applied Science Laboratories or Eyelink 1000, SR Research).

Microsaccade detection

We detected microsaccades in the eye-position records using an adaptation of a previously described technique (Martinez-Conde et al., 2000). Since eye positions were captured by different methods for each experiment, the parameters used for microsaccade detection were adjusted slightly to each one. In all experiments, horizontal and vertical eye positions, sampled at 200 Hz, were used to calculate an instantaneous eye velocity. The velocity vectors were smoothed by a 25 ms boxcar for the speed pulse and the motion-step experiments. The velocity vectors for the 33 ms motion-pulse experiment was only smoothed by a 15 ms boxcar since eye positions were filtered with a 100 Hz low-pass Butterworth filter during the course of the experiment. The velocity vectors for the 50 ms motion-pulse experiment were not smoothed since eye positions were already filtered with a 20 Hz low-pass Butterworth filter during the course of the experiment.

Instances where the eye speed crossed 8 degrees/second for the motion-step and 33 ms motion-pulse experiment and 10 degrees/second for the speed pulse and 50 ms motion-pulse experiment were considered possible microsaccades. The reason for the difference in threshold was that the eye position signal for the motion-step and 50 ms motion-pulse experiments was less noisy than the others, and thus microsaccades could be detected with greater confidence.

To be considered a microsaccade, the eye movement had to last at least 10 and no more than 300 ms, could not have started within 20 ms of a previous microsaccade, was at least 0.05˚ in length, and the eye direction could change no more than 30 degrees every 5 ms for the duration of the microsaccade. Only events that satisfied all the criteria above were deemed microsaccades, and the starting time of the microsaccade was used in all further analysis. Accuracy of the saccade algorithm was further confirmed by visual inspection of raw eye movement traces for a subset of the data.

Results

We wanted to examine the effects of microsaccades on behavioral choice and neuronal firing in visual and parietal cortex, and to determine the extent to which microsaccades contributed to the correlation between the two. Our analysis was based on three different visual detection tasks. Fig. 1 is a schematic comparing the three tasks. All three used random-dot motion as a stimulus; one involved detecting a change in speed and two involved detecting a change in motion coherence. Data were collected from three different laboratories, using a total of seven macaque monkeys as subjects.

In the first experiment (*speed-pulse* experiment, Fig. 1A) (see Methods), the monkey had to release a lever in response to a transient (53 ms) increase in speed in one of two, 100% coherently moving random dot patches. In the second experiment (*motion-step* experiment, Fig. 1B) the monkey had to release a lever in response to the onset of coherent motion in one of two, initially 0% coherent, moving random dot patches. The third experiment (*motion-pulse* experiment, Fig. 1*C-D*) was a combination of two, slightly different motion-pulse detection tasks. In the first, the monkey had to release a lever in response to a transient (50 ms) pulse of coherent motion in one of two, initially 0% coherent, moving random dot patches. In the second, the monkey detected a 33 ms pulse of coherent motion in a single random dot patch.

The correlation between neural activity and behavior has frequently been measured by employing motion-based stimuli and neural recordings from direction-

Figure 1 Experimental paradigms. (A) *Speed-pulse* experiment. The goal of this task was to detect a brief speed change. The monkey fixated on a colored fixation spot and two, diametrically opposite, colored annuli appeared on the screen. After 500 ms, random dot patches with 100% coherent motion appeared inside the annuli. The monkey's task was to release a lever in response to a transient (53 ms) increase in speed at either of the dot patches (the speed pulse). The color of the fixation point matched one of the annuli, indicating the likely location of the speed pulse. On 40% of trials the fixation point color would change mid-trial indicating that the likely speed-pulse location had changed (not shown). The monkey had to release a lever within 200-600 ms of the speed pulse to obtain a reward. (B) *Motion-step* experiment. The goal of this task was to detect the occurrence of a coherent motion step. The monkey fixated on a central point and a static random dot patch appeared indicating the likely position of the coherent motion step. Afterwards, two, diametrically opposite random dot patches began moving at 0% coherent, with one patch at the location of the static cue. 0% coherent motion was shown for a random amount of time between 500 and 8000 ms (flat hazard function), followed by a coherent motion step lasting 750ms. The monkey had to release a lever during the coherent motion step to obtain a reward. On 20% of the trials the coherent motion step occurred in the uncued patch. (C) *Motion-pulse* experiment #1. The goal was to detect a brief coherent motion pulse. After the monkey fixated, 0% coherent motion began in two random dot patches located in the same hemifield, each one matched to the preferences of two neurons simultaneously recorded. After a random amount of time between 500 and 10000 ms (flat hazard function), a 50 ms pulse was shown in one or both patches. Afterwards, 0% coherent motion would resume. The monkey had to release a lever from 200 to 800 ms after the start of the motion pulse to obtain a reward. (D) *Motion-pulse* experiment #2. The monkey initiated fixation, followed by 0% coherent motion for a random time between 500 and 10000 ms in a single random dot patch. The monkey had to release a lever from 150 to 650 ms after a 33 ms pulse of coherent motion.

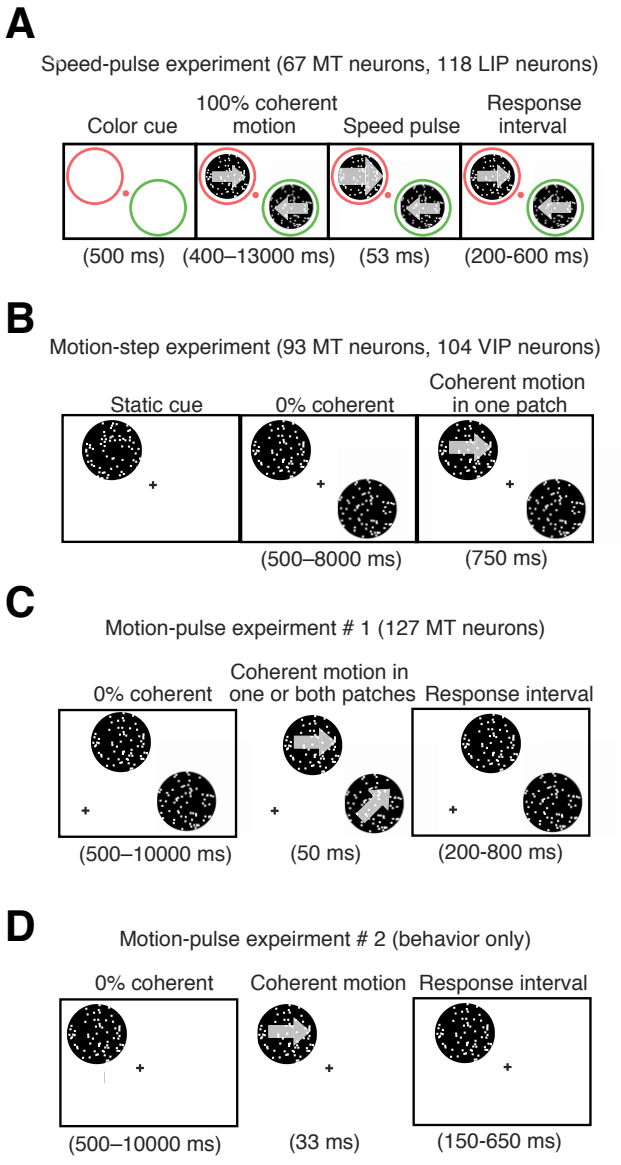
A
Speed-pulse experiment (67 MT neurons, 118 LIP neurons)
Color cue
100% coherent motion
Speed pulse
Response interval
(500 ms)
(400–13000 ms)
(53 ms)
(200-600 ms)
B
Motion-step experiment (93 MT neurons, 104 VIP neurons)
Static cue
0% coherent
Coherent motion in one patch
(500–8000 ms)
(750 ms)
C
Motion-pulse expeirment # 1 (127 MT neurons)
0% coherent
Coherent motion in one or both patches
Response interval
(500–10000 ms)
(50 ms)
(200-800 ms)
D
Motion-pulse expeirment # 2 (behavior only)
0% coherent
Coherent motion
Response interval
(500–10000 ms)
(33 ms)
(150-650 ms)

Figure 1

selective neurons (Britten et al., 1996; Bair et al., 2001; Dodd et al., 2001; Cook and Maunsell, 2002b; Grunewald et al., 2002; Krug et al., 2004; Uka and DeAngelis, 2004; Liu and Newsome, 2005; Purushothaman and Bradley, 2005). By combining data from three experiments with different motion-based stimuli along with direction-selective neural activity from three cortical areas, we wished to determine the effects of microsaccades that may apply broadly to experiments of this nature.

Microsaccade Properties

We detected microsaccades in the eye-position records using an adaptation of a previously described technique (Martinez-Conde et al., 2000) (see Methods). To confirm the accuracy of our microsaccade-detection algorithm, we compared the properties of the microsaccades for our three experiments to those of past studies on microsaccades in non-human primates (Table 1, Martinez-Conde et al., 2004). Because past studies captured the eye position and analyzed the resulting data in different ways, there is substantial variability in the reported parameters. Nonetheless, the microsaccade parameters from our experiments were consistent with those previously reported.

Previous studies have described a linear relationship between saccadic peak velocity and amplitude that is common to all saccades, including microsaccades (Zuber and Stark, 1965; Bahill et al., 1975). Microsaccades that obey this relationship are said to fall along the "main sequence". Fig. 2 is a two-dimensional histogram showing the distribution of microsaccade peak velocities and amplitudes for all of the microsaccades

Table 1 Microsaccade properties across the three experiments. Maximum and minimum values were the extreme values from references collected by Martinez-Conde *et. al.* (2004).

	Speed pulse	Motion step	Motion pulse	Past studies
Peak speed (deg/s)	18.5	16.3	25.9	9-110
Duration (ms)	16.6	16.2	15.7	20-29
Amplitude (min)	22.3	20.8	20.8	8.4-48
Frequency (Hz)	0.26	0.78	1.33	0.2-5

Figure 2 Main sequence analysis of the microsaccades collected from the three experiments. A small percentage of the microsaccades did not appear within the boundaries of the graphs. The two-dimensional histogram of the peak velocity versus amplitude for each experiment is shown, with the number of microsaccades per bin indicated by grayscale values. Bins with zero frequency are plotted in white and the bin with the greatest frequency is shown in black. The two-dimensional histogram of the peak velocity (x-axis) versus the amplitude (y-axis) of each microsaccade collected during the speed-pulse experiment (A), motion-step experiment (B), 50 ms motion-pulse experiment (C), 33 ms motion-pulse experiment (D). The speed-pulse data after a 20 Hz low-pass filter was applied to the position signal in order to mimic the data from panel B (E). The value for the maximum bin is 100, 166, 216, 241 and 667 for A, B, C, D and E, respectively.

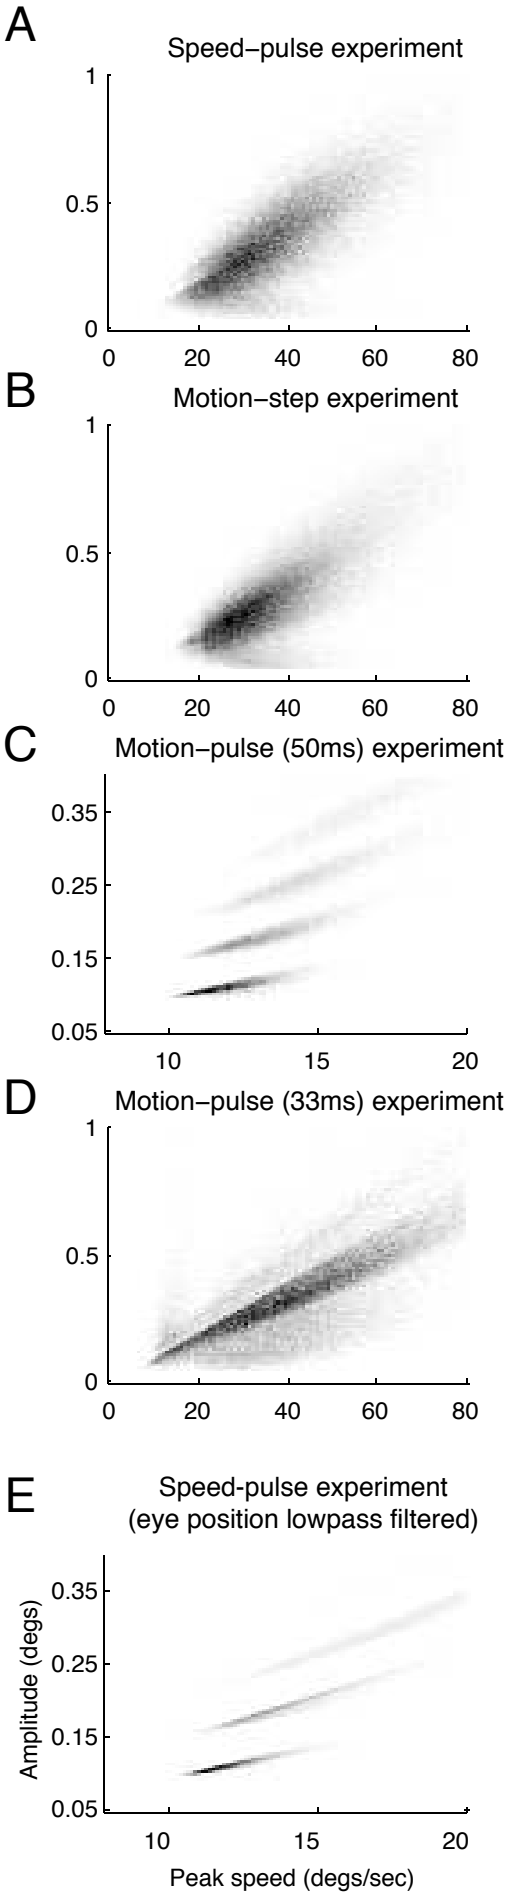
A
Speed−pulse experiment
1
0.5
0
0
20
40
60
80
B
Motion−step experiment
1
0.5
0
0
20
40
60
80
C
Motion−pulse (50ms) experiment
0.35
0.25
0.15
0.05
10
15
20
D
Motion−pulse (33ms) experiment
1
0.5
0
0
20
40
60
80
E
Speed-pulse experiment
(eye position lowpass filtered)
Amplitude (degs)
0.35
0.25
0.15
0.05
10
15
20
Peak speed (degs/sec)

Figure 2

in each of the three experiments. The number of microsaccades falling in each bin is shown by the shade of gray—a bin with a frequency of zero is white and the bin with the highest frequency for that experiment is black. Microsaccades were detected by smoothing the eye velocity as described in the Methods, but the peak velocity was calculated from unsmoothed eye velocity. All three distributions showed a strong linear relationship between peak velocity and amplitude (Spearman's rank correlation coefficients of 0.74, 0.73 and 0.77 for the speed-pulse, motion-step and 33 ms motion-pulse experiments, respectively). However, the distribution of microsaccades for the 50 ms motion-pulse experiment (Fig. 2C) showed a series of "bands". These bands were separated based on the microsaccade duration, with microsaccades lasting 10 ms in the lowest band, those lasting 15 ms in the second lowest, etc. These bands are an artifact of filtering the data with a 20 Hz low-pass filter during that experiment (see Methods). They arose because the low-pass filtering rendered the eye velocity highly positively auto-correlated at short time lags. This limited how large the velocity in any given time bin could be without the velocity in adjacent time bins crossing the 10 deg/s threshold. As the peak velocity increases, adjacent bins are dragged above the detection threshold, lengthening the saccade duration in 5 ms increments and leading to discontinuous jumps in the saccade amplitude. To confirm this we applied a 20 Hz low-pass filter to the eye position data from the speed-pulse experiment, and then detected microsaccades with no additional smoothing. The resulting distribution (Fig. 2E) closely resembled the banded distribution in Fig. 2C.

Although it was reassuring to find linear main sequences and microsaccade parameters that fell within previously observed ranges (Table 1), this does not exclude the possibility that our algorithm either missed microsaccades, misclassified non-microsaccade events as microsaccades, or both, especially for the motion-pulse experiment. The potential effect of these errors on our results is addressed in the Discussion.

Example trials

Fig. 3 provides three example trials from the *speed-pulse* task that illustrate the hypothesis that microsaccades can affect behavior and neural activity and can contribute to the correlation between the two. Fig. 3*A* shows data from a long trial with a single MT neuron's response to sustained motion in the preferred direction. The individual spike times are denoted by black dots and the instantaneous spike rate by the gray line. During this trial there were three microsaccades (asterisks) evident in the eye-speed trace (black line). For this neuron, each microsaccade was followed by a brief pause in the spike train.

Fig. 3*B* shows data from a trial in which the monkey correctly detected the speed pulse, which occurred at the vertical dashed black line. The spike rate increased after the speed pulse, possibly contributing to the monkey's correct detection. Fig. 3*C* shows data from another trial from the same experimental session. In this trial the monkey *failed* to detect the speed pulse. Notably, the monkey happened to make a microsaccade at nearly the same time as the onset of the speed pulse. Not only did the animal fail to detect the

Figure 3 Three example trials from the same speed-pulse experiment. (A) An example trial with three microsaccades (indicated by the asterisks). The black curve shows eye velocity that sharply peaks at the time of the microsaccades, while the neural activity is shown using both the raster (black dots) and the average spike rate (gray curve). Neural activity is suppressed following all three microsaccades. (B) An example correct trial that contained no microsaccades. Neural activity increased following the speed pulse (vertical dashed line). (C) An example failed trial from the same experimental session. A microsaccade was completed just after the speed pulse and resulted in a large reduction in neural activity. This example trial suggests a mechanism by which microsaccades might simultaneously affect both the neural response to a stimulus and the monkey's ability to perceive the stimulus.

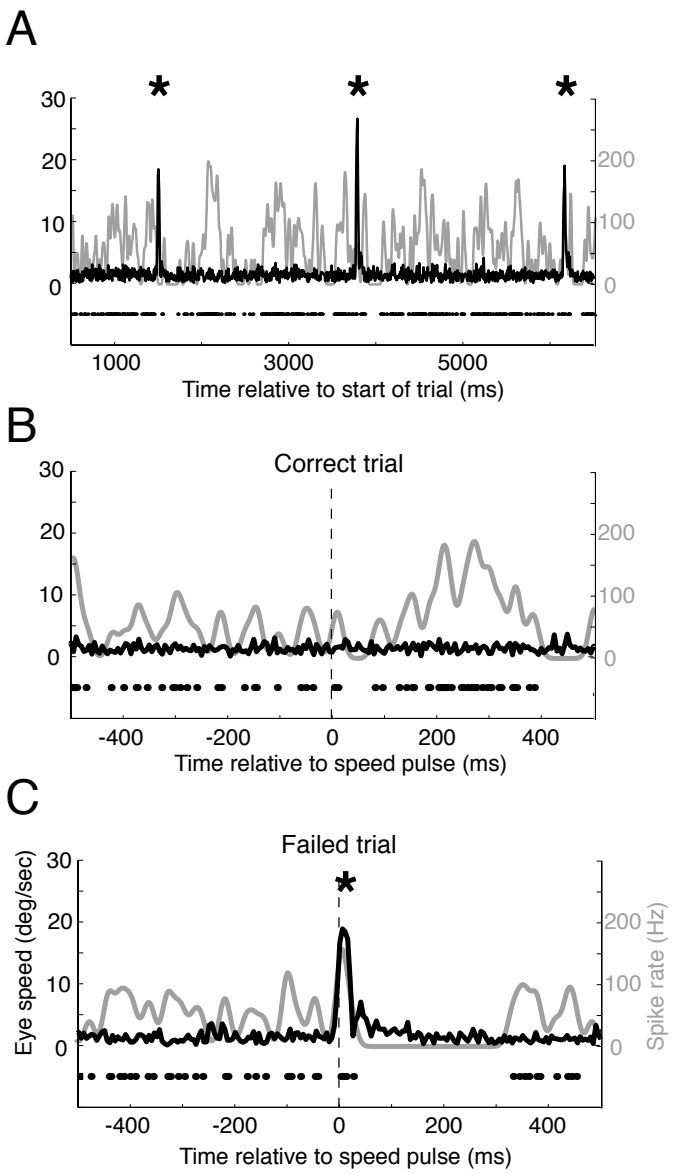
A
30
20
10
0
200
100
0
1000
3000
5000
Time relative to start of trial (ms)
B
Correct trial
-400
-200
0
200
400
Time relative to speed pulse (ms)
C
Failed trial
Eye speed (deg/sec)
Spike rate (Hz)
Time relative to speed pulse (ms)

Figure 3

speed pulse, but the neural activity was suppressed after the microsaccade, as in panel *A*. If microsaccades inhibit a subject's ability to detect visual stimuli and at the same time suppress neural activity, then it is possible that microsaccades may contribute to the detect probability. The goal of this paper is to examine this hypothesis in detail.

The effect of microsaccades on behavioral performance

We first determined what effect microsaccades had on the monkeys' ability to detect the different types of visual stimuli employed in the three tasks. Previous reports have differed on whether microsaccades modulate visual thresholds, but to our knowledge, the effect of microsaccades on detecting motion or speed signals has not been examined.

We measured the effect of microsaccades on the monkeys' detection performance in two different ways. First (Fig. 4*A-C, left panel*), the microsaccade rate is shown relative to onset of the test stimulus (the speed pulse, motion step or motion pulse) for trials in which the monkey correctly detected the stimulus (blue) and trials in which the monkey failed to detect the stimulus (red). For the experiments employing a pulse stimulus (Fig. 4*A&C, left panel*), the microsaccade rate was higher on failed than correct trials for times immediately around the onset of the test stimulus. Similarly, for the motion-step experiment (Fig. 4*B, left panel*) the microsaccade rate was greater on failed than correct trials for times after the onset of coherent motion.

Figure 4 The effect of microsaccades on the monkeys' perception. (A) In the left panel, the microsaccade rate for correct (blue curve), failed trials (red curve) and all trials (black curve) relative to stimulus onset are shown for the speed-pulse experiment. A greater number of microsaccades are completed around stimulus onset during failed trials. In the right column, the monkey's ability to detect the stimulus given a microsaccade is completed at times relative to stimulus onset is shown. The panel shows that when a microsaccade occurs near stimulus onset, the ability to correctly detect the stimulus is reduced. (B) Same as above, but for the motion-step experiment. Microsaccades after stimulus onset are also associated with a reduction in performance. (C) Same as above, but for the motion-pulse experiments.

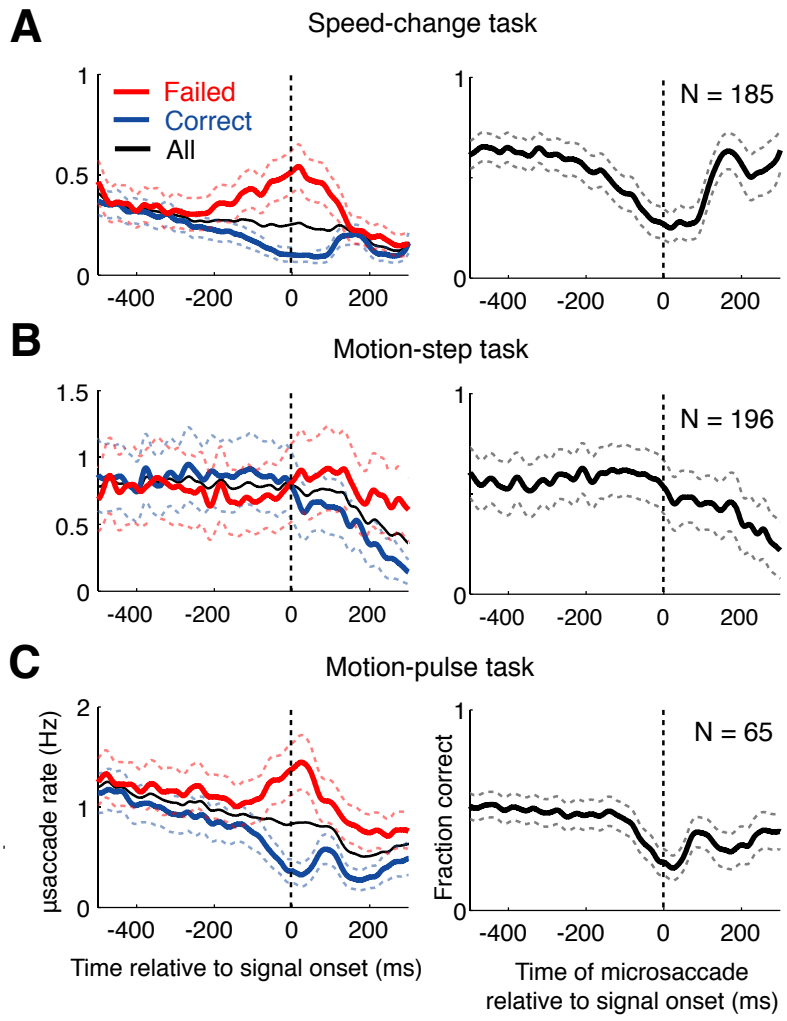
A
Speed-change task
Failed
Correct
All
N = 185
B
Motion-step task
N = 196
C
Motion-pulse task
N = 65
μsaccade rate (Hz)
Fraction correct
Time relative to signal onset (ms)
Time of microsaccade
relative to signal onset (ms)

Figure 4

Second, we calculated the monkeys' detection performance as a function of the time of microsaccades relative to the test-stimulus onset (Fig. 4*A-C, right panel*). That is, time 0 refers to trials in which a microsaccade occurred simultaneously with the onset of the test stimulus, -200 ms refers to trials in which a microsaccade occurred 200 ms prior to the onset of the test stimulus, and so on. For the two experiments employing pulse stimuli (Fig. 4*A&C, right panel*), the monkeys' exhibited dramatically reduced performance when a microsaccade occurred near the time of the test-stimulus onset. In contrast, microsaccades occurring more than 200 ms before the test-stimulus onset had no effect on performance. For the motion-step experiment (Fig. 4*B, right panel*), the monkeys' performance was reduced when a microsaccade occurred over period of hundreds of ms after stimulus onset.

These results demonstrate that microsaccades can dramatically reduce performance on a range of motion-based visual detection tasks, especially tasks involving transient motion stimuli. Importantly, there was no abrupt change in the overall rate of microsaccades around the time of test-stimulus onset (thin black curve, Fig. 4*A-C, left panel*). Rather, performance was reduced on that subset of trials in which a microsaccade happened to occur, by chance, near the time of the test-stimulus onset.

The effect of microsaccades on neural activity

The effect of microsaccades on neural activity in the visual system is variable and not completely understood. Generally, studies have found that microsaccades increase

neural activity in early stages of the visual system, including the LGN (Martinez-Conde et al., 2002) and V1 (Martinez-Conde et al., 2000; Snodderly et al., 2001), and decrease neural activity in later stages, including V2 and V4 (Leopold and Logothetis, 1998). However, clear exceptions exist (Leopold and Logothetis, 1998), and care must be taken when considering new paradigms. For example, the effect of microsaccades on neural activity may depend on the whether a stationary or moving stimulus was presented. In area MT, microsaccades have been reported to increase or decrease activity depending on the baseline firing rate and on whether the eye movement produced preferred or null direction motion (Bair and O'Keefe, 1998).

We examined the relationship between microsaccades and neural activity for the three different tasks and cortical areas. The left column of Fig. 5 shows the average neural response aligned on the time of a microsaccade. The neural response for all cells was normalized so that average activity before the microsaccade was equal to one. Because the test stimulus usually changed the firing rate of the neurons, only microsaccades that occurred at least 300 ms before the onset of the test stimulus were included in the analyses shown in the left panels in Fig. 5.

Microsaccades, on average, caused a large decrease in neural activity in the speed-pulse experiment for areas MT and LIP (Fig. 5*A-B*). In these experiments, the subjects were always viewing coherent motion in the neuron's preferred direction. One possible reason why microsaccades produced large suppression in neural activity was that any eye movement away from the null direction would produce apparent motion in a non-

Figure 5 The effect of microsaccades on neural activity. Averages were computed for each brain area and experiment type. The left column shows the average neural activity relative to each microsaccade completed before stimulus onset. Only the responses that occurred before the test-stimuli (speed or coherence change) were used. The neural response for each neuron was normalized so that its baseline rate was equal to one. The right column shows the average neural activity in response to the test stimulus separated by whether a microsaccade occurred near stimulus onset. The test stimulus was either a speed change (A and B) or a coherence change (C – E). The gray curve shows the average activity given a microsacacde occurred from 200 ms prior to 100 ms after stimulus onset for the pulse tasks, and from 100 ms prior to 200 ms after stimulus onset for the step task. The black curve shows the average activity for trials with no microsaccades completed in these windows. Only neurons with at least 10 trials with microsaccades completed in these windows and 10 trials without were included. (A) is the average response across LIP for the speed-pulse task; (B) average response for area MT for the speed-pulse task; (C) the average response for area MT for the motion-step task; (D) the average response for area VIP for the motion-step task; (E) the average response for area MT for the motion-pulse task.

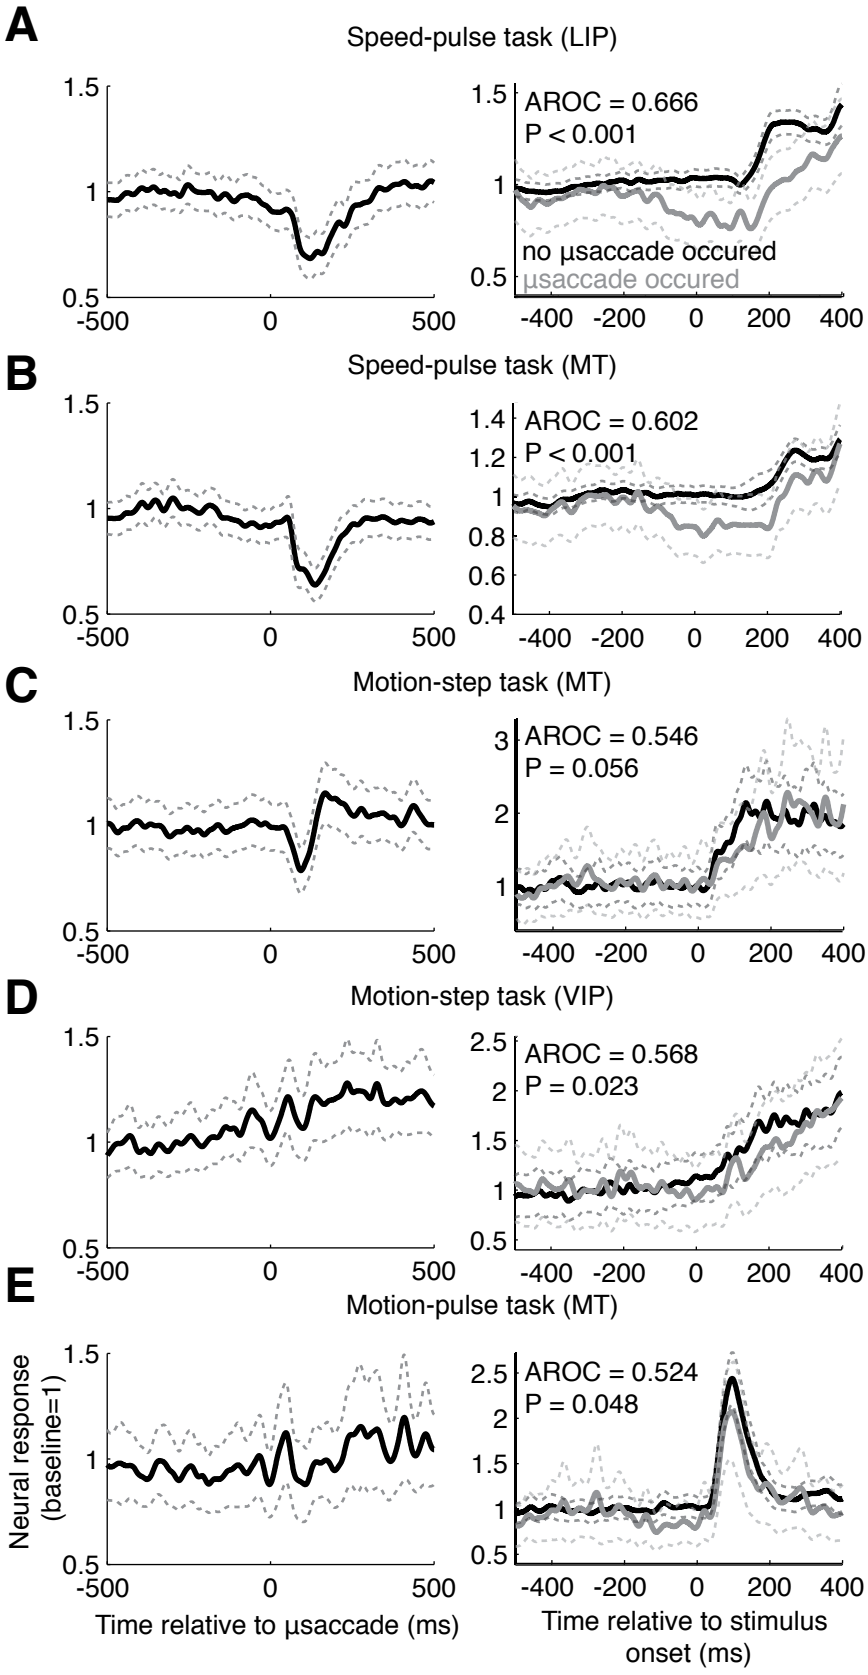
A
Speed-pulse task (LIP)
AROC = 0.666
P < 0.001
no μsaccade occured
μsaccade occured
B
Speed-pulse task (MT)
AROC = 0.602
P < 0.001
C
Motion-step task (MT)
AROC = 0.546
P = 0.056
D
Motion-step task (VIP)
AROC = 0.568
P = 0.023
E
Motion-pulse task (MT)
AROC = 0.524
P = 0.048
Neural response (baseline=1)
Time relative to μsaccade (ms)
Time relative to stimulus onset (ms)

Figure 5

optimal direction. In the other two tasks, microsaccades before the test-stimulus onset had either a smaller effect (MT in motion-pulse task, Fig. 5*C*) or no clear effect (VIP in motion-pulse task and MT in motion-step task, Fig. 5*D*-*E*). Unlike in the speed-pulse task, in these tasks the monkeys viewed 0% coherent motion that did not optimally drive the recorded neurons and therefore apparent motion induced by microsaccades had less net effect on the response. However, it is possible that microsaccades affect neural activity by other mechanisms. For example, the suppression in neural activity in LIP neurons (Fig. 5*A*) appeared to begin before the microsaccade which is not consistent with a response to a microsaccade-induced motion signal.

Although microsaccades had little or no effect on neural activity for the tasks where the monkey viewed 0% coherent motion before the test-stimulus onset (Fig. 5*C*-*E*, left column), they may have had a stronger effect on the response to the test stimulus, which is the key parameter in determining detect probability. To examine the effect of microsaccades on the response to the test stimulus we used a different approach. In the right column of Fig. 5, we show the average neural responses aligned to the onset of the test stimulus. Trials were averaged separately depending on whether there was a microsaccade (grey curve) or no microsaccade (black curve) "near" the time of the test-stimulus onset. To obtain an accurate measurement of the neural response, we only included neurons that had at least 10 trials with microsaccades near test-stimulus onset and 10 without. This condition was satisfied by 67 LIP and 47 MT neurons from the speed-pulse experiment, 28 MT and 29 VIP neurons from the motion-step experiment and 49 MT neurons from the motion-pulse experiment. For the two experiments with

transient test stimuli (the speed-pulse and motion-pulse experiments), the two trial types were separated based on whether a microsaccade occurred from 200 ms before to 100 ms after the test-stimulus onset. For the motion-step experiment the window was from 100 before to 200 ms after test stimulus onset. These windows were selected based on the results in Fig. 4, and represent the time windows for each task in which microsaccades affected detection of the test stimulus. For all three experiments and cortical areas, we found that, on average, the neurons increased their firing rate in response to the test stimulus. Moreover, the neural responses to the test stimuli were greater when no microsaccade occurred near the time of the test-stimulus onset.

To quantify the difference between the neural responses to the stimulus with and without microsaccades, we used an ROC analysis (Green and Swets, 1966). For the two experiments with transient test stimuli, we measured the spike count from 50 to 200 ms after test-stimulus onset. For the motion-step experiment, we extended this window from 50 to 300 ms after test-stimulus onset. The area under the ROC curve provides a nonparametric index of the separation between the spike-count distributions for trials with microsaccades versus trials without microsaccades. Values above 0.5 indicate that the spike count is on average greater when no microsaccades occurred near stimulus onset. For all experiments and brain areas, the average AROC (area under the ROC curve) values were above 0.5, indicating that microsaccades were associated with a reduced neural response to the test stimulus (speed-pulse experiment: LIP, AROC = 0.666, $p < 0.001$; MT, AROC = 0.602, $p < 0.001$; motion-step experiment: MT, AROC = 0.546, $p = 0.056$; VIP, AROC = 0.568, $p = 0.023$; motion-pulse experiment: MT, AROC

= 0.524, p = 0.048, all statistical tests two-sided t-tests unless otherwise specified). These results are consistent with the interpretation that at least some of the suppression in activity is the result of a less-preferred apparent motion signal caused by microsaccades (Bair and O'Keefe, 1998), but we cannot exclude the possibility that microsaccades affected neural activity through other mechanisms.

The correlation between microsaccades, neural activity and behavioral performance

We have shown that, on average, microsaccades decrease both detection performance and neural responses across three motion-based detection tasks. It follows that microsaccades may contribute some correlation between neural activity and behavior. To examine this, we calculated the partial correlation between behavioral performance, the presence of microsaccades and the neural response. Partial correlation provides a measure of the correlation between each pair of variables after removing the effect of the third.

Each of our detection tasks had two measures of behavior: detection performance (correct or failed) and reaction time. For each measure of behavior, we computed the partial correlations between neural activity, microsaccade occurrence and behavior (Fig. 6). We took the neural activity as the number of spikes that occurred in the 100 ms prior to each time point, and we took the number of microsaccades as that occurring from 100 to 200 ms prior to each time point. The partial correlations were computed at 10 ms intervals relative to stimulus onset.

Figure 6 The partial correlations between microsaccades, neural activity, and perception. The partial correlation analysis was performed at 10 ms intervals relative to stimulus onset and averaged across all neurons for each condition. Perception was either measured as the behavioral response (correct = 1, failed = 0, shown in the left column), or the reaction time for correct trials (right column). The blue curve shows the partial correlation between neural activity and perception with the effects of microsaccades removed. The red curve shows the partial correlation between neural activity and microsaccades with the effect of perception removed. The green curve shows the partial correlation between microsaccades and perception with the effect of neural activity removed. The neural activity was the number of spikes in the previous 100 ms and the number of microsaccades from 100 to 200 ms prior to each point was used in the calculation. The horizontal black and grey bars give the windows used to count spikes and microsaccades, respectively, for Fig. 7 and the ROC analyses of the effect of microsaccades on neural activity. (A) Speed pulse – LIP; (B) speed pulse – MT; (C) motion step – MT; (D) motion step – VIP; (E) motion pulse – MT.

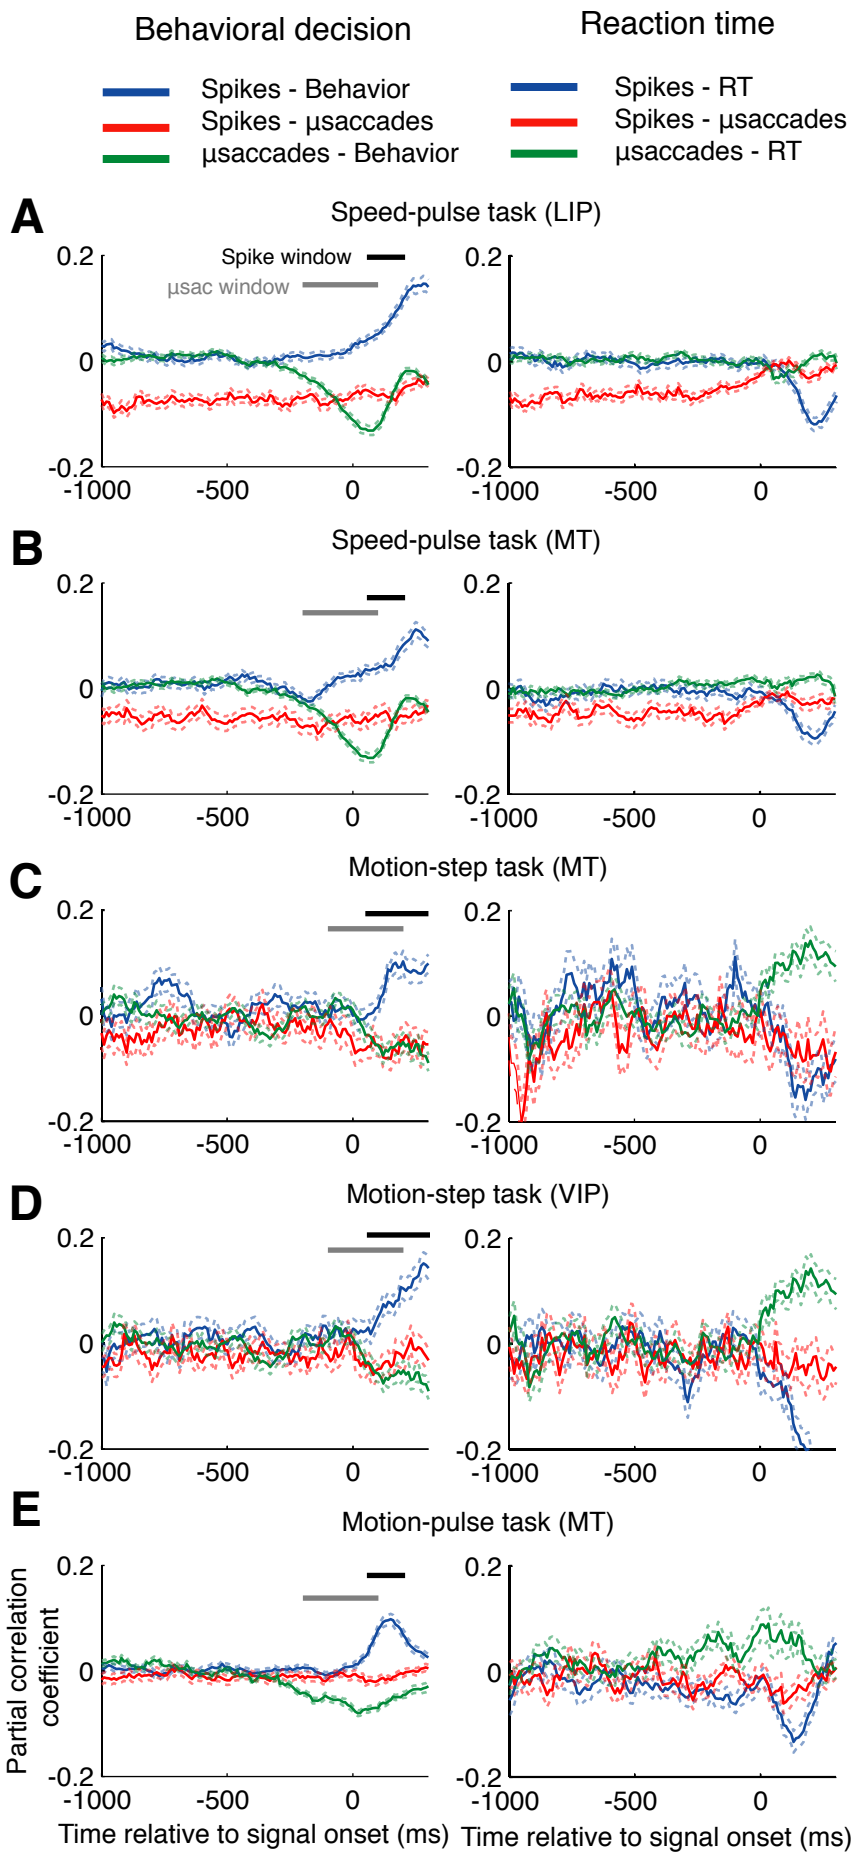
Behavioral decision
Spikes - Behavior
Spikes - μsaccades
μsaccades - Behavior
Reaction time
Spikes - RT
Spikes - μsaccades
μsaccades - RT
A
Speed-pulse task (LIP)
Spike window
μsac window
B
Speed-pulse task (MT)
C
Motion-step task (MT)
D
Motion-step task (VIP)
E
Motion-pulse task (MT)
Partial correlation coefficient
0.2
0
-0.2
-1000
-500
0
Time relative to signal onset (ms)
Time relative to signal onset (ms)

Figure 6

In the left column, the partial correlation between microsaccades and behavior (green curve) shows that the presence of microsaccades around the test-stimulus onset was negatively correlated with detection of the test stimulus, indicating that microsaccades were associated with reduced detection performance. In the right column, the presence of microsaccades around the test-stimulus onset was usually positively correlated with reaction times (Figs. 6C-E), indicating that microsaccades were associated with a delayed behavioral response. These results were consistent with the results of Fig. 4.

The partial correlation between microsaccades and neural activity is shown by the red curve in Fig. 6. For the speed-pulse task, the partial correlation between activity and microsaccades was consistently below zero, confirming that microsaccades were associated with a strong decrease in activity independent of time relative to the test-stimulus onset. For the other two experiments, this partial correlation was only marginally negative prior to stimulus onset. However, the value decreased after the onset of the test stimulus, when coherent motion was presented in the preferred direction for the neurons. This further suggests that the suppression of neural activity by microsaccades was stronger when the monkey was viewing a stimulus moving in the neuron's preferred direction. There results are consistent with the results from Fig. 5.

Lastly, we show the partial correlation between neural activity and behavior after removing the effects of microsaccades (blue curve). In all experiments, there remained a

robust positive correlation between neural activity and detection performance (left column) and a negative correlation between neural activity and reaction time (right column). Thus after removing the effects of microsaccades, greater neural activity is still associated with a correct behavioral response and shorter reaction times on a trial-by-trial basis. Although microsaccades affected both neural activity and detection, their effect cannot entirely account for the choice-related correlation between the two. This would suggest that while microsaccades might contribute to the detect probability, detect probability is not solely due to microsaccades. We confirmed this in the following analysis.

The contributions of microsaccades to the detect probability

Detect probability, which is similar to the AROC metric above, has been widely used to quantify the relationship between neural activity and behavioral outcome. Whereas the partial correlation analysis in Fig. 6 measures the strength of the linear relationship between spike rate and detection performance, the detect probability is a nonparametric measure of this relationship that corresponds to the probability that a randomly chosen spike rate from a correct trial is greater than a randomly chosen spike rate from a missed trial. We chose to compute detect probability to facilitate comparison with previous studies describing correlations between neural responses and behavior. However, using other metrics to express the link between neural activity and behavioral performance, such as d' or differences in median spike rate, produced very similar results.

To determine the effect of microsaccades on detect probability, we calculated the detect probability twice—with all trials included and with only those trials that did not have a microsaccade near the time of the test-stimulus onset (Fig. 7). We used the same windows to count spikes and microsaccades as those used in the ROC analysis of the effect of microsaccades on neural activity (see above). These windows are indicated by the black and grey horizontal bars in Fig. 6. Since the probability of a microsaccade in these windows was small on any given trial, we required a sufficient number of correct and failed trials to measure the effect of microsaccades on detect probability. Thus we only included neurons with at least 15 correct and 15 failed trials regardless of whether microsaccades occurred.

To express the change in detect probability as a percentage, we first subtracted 0.5 from both detect probability values to "zero" the detect probability. We were mindful that detect probability is limited to values between 0 and 1, and is consequently highly nonlinear for values near 1. However, over the range of detect probabilities in our data (0.55 – 0.65), the detect probability is linearly related to d' or difference in median spike rate and the percentage-change values reported here are equivalent for any of those metrics.

In all three cortical areas and detection tasks, we found that removing trials containing microsaccades reduced the detect probability (Fig. 7). For the speed-pulse experiment, eliminating trials with microsaccades reduced the detect probability (DP) by 19% for MT and by 21% for LIP (Fig. 7*A-B*, MT, all trials, mean DP = 0.558, trials

Figure 7 The detect probability with and without the effect of microsaccades. The histograms show the difference between the detect probabilities calculated using all trials and using trials with microsaccades near stimulus onset excluded. Detect probability was computed either using all trials or removing trials where a microsaccade was completed near stimulus onset. Only neurons with 15 correct and failed trials were included in this analysis. For the speed-pulse and motion-pulse experiment, trials were removed if a microsaccade occurred from 200 ms before to 100 ms after the stimulus onset. For the motion-step experiment the window was 100 ms before to 200 ms after for (C) and (D) and from 200 ms before to 50 ms after for (E). (A) Speed pulse – LIP; (B) speed pulse – MT; (C) motion step – MT; (D) motion step – VIP, late window; (E) motion step – VIP, early window ; (F) motion pulse – MT.

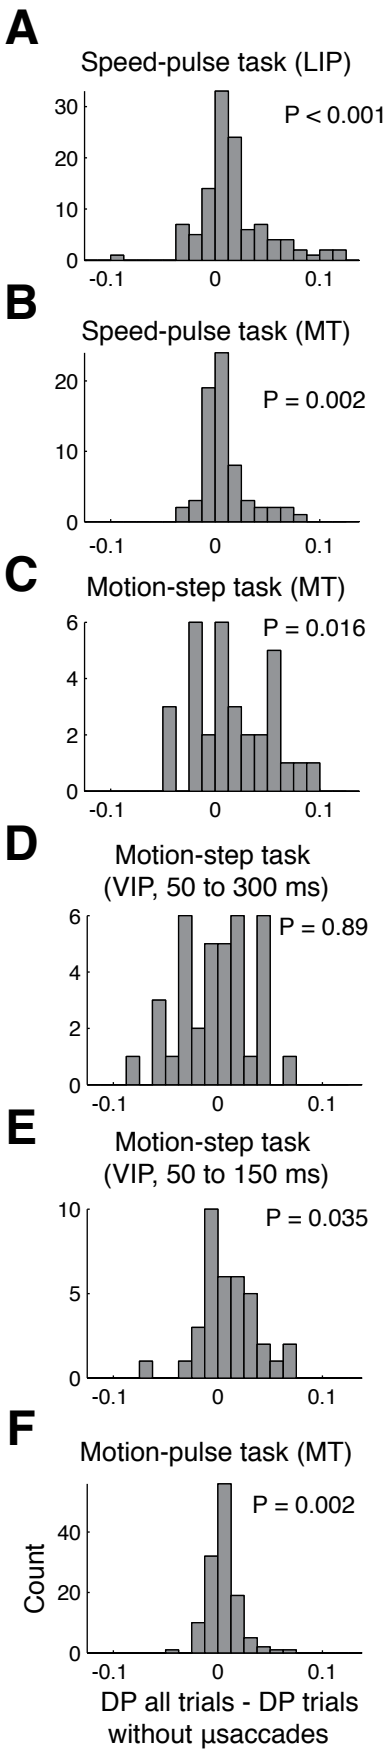
A
Speed-pulse task (LIP)
P < 0.001
30
20
10
0
-0.1
0
0.1
B
Speed-pulse task (MT)
P = 0.002
20
10
0
-0.1
0
0.1
C
Motion-step task (MT)
P = 0.016
6
4
2
0
-0.1
0
0.1
D
Motion-step task
(VIP, 50 to 300 ms)
P = 0.89
6
4
2
0
-0.1
0
0.1
E
Motion-step task
(VIP, 50 to 150 ms)
P = 0.035
10
5
0
-0.1
0
0.1
F
Motion-pulse task (MT)
P = 0.002
Count
40
20
0
-0.1
0
0.1
DP all trials - DP trials
without μsaccades

Figure 7

without microsaccades, mean DP = 0.547, p = 0.002; LIP, all trials, mean DP = 0.584, trials without microsaccades, mean DP = 0.566, p < 0.001). In the motion-step experiment, the decrease in detect probability was 19% in MT (Fig. 7*C*, all trials, mean DP = 0.588, trials without microsaccades, mean DP = 0.571, p = 0.016). However, using the windows defined above, we did not find any change in the detect probability for area VIP (Fig. 7*D*, all trials, mean DP = 0.632, trials without microsaccades, mean DP = 0.633, p = 0.89).

Given that microsaccades suppressed neural activity in VIP during coherent motion, we were surprised that eliminating microsaccade-containing trials did not alter the detect probability. VIP is known to exhibit strong responses correlated with perceptual choice, which may be due to feedback (Cook and Maunsell, 2002b). One possibility is that this feedback could have overwhelmed the effect of microsaccades on neural activity. We thus repeated the detect probability calculation for VIP using a spike window of 50 to 150 ms after motion onset and a microsaccade window from 200 ms prior to 50 ms after motion onset. We reasoned that truncating the spike window at 150 ms would reduce the potential contribution of feedback on the detect probability. Using these new windows, eliminating trials with microsaccades reduced the detect probability in VIP by 17% (Fig., 7*E*, all trials, mean DP = 0.558, trials without microsaccades, mean DP = 0.548, p = 0.035).

Lastly, in the motion-pulse experiment eliminating trials with microsaccades reduced the detect probability by a relatively smaller yet still significant 7% in area MT

(Fig. 7*F*, all trials, mean DP = 0.560, trials without microsaccades, mean DP = 0.556, p = 0.002).

One possibility is that the observed decrease in detect probability is solely the result of reducing the number of trials in the analysis, and not a specific effect of removing trials with microsaccades. In theory, reducing the number of trials should not bias the outcome one way or another, although it will reduce the reliability of the measurement. To confirm this assumption, we performed a resampling analysis. For each neuron, we calculated the detect probability after eliminating a set of random trials equal to the number eliminated in the original microsaccade elimination analysis. We repeated this calculation 10,000 times for each neuron and took the mean value. As expected, eliminating the trials did not bias our population estimate of detect probability for any of the experiments (range of mean change in DP, -2×10^{-5} to 4×10^{-4}).

The contribution of microsaccades to neuronal variance

We have shown that the paired effect of microsaccades on neural activity and behavior can contribute up to 20% of the observed correlation between the two. This suggests that microsaccades are a significant source of behaviorally relevant neuronal covariance in our tasks. It does not necessarily follow, however, that microsaccades are a significant source of neuronal response variance in general. In fact, previous studies addressing this question have produced conflicting results (Gur et al., 1997; Bair and O'Keefe, 1998). To quantify the contribution of microsaccades to spike-rate variance we

calculated the variance-to-mean ratio (Fano factor) for correct trials with and without microsaccades. The time windows for detecting microsaccades and measuring the neural response were the same as those described above. For the speed-pulse task, elimination of trials with saccades produced a small but significant reduction in Fano factor for both MT (Fano factor ± SE, all trials, 1.27 ± 0.07, trials without microsaccades, 1.26 ± 0.06, paired t-test $p = 0.015$) and LIP (all trials, 1.213 ± 0.055, trials without microsaccades, 1.207 ± 0.055, $p = 0.002$). For the motion-step task there was a small, but not significant change for MT (all trials, 1.30 ± 0.09, trials without microsaccades, 1.23 ± 0.08, $p = 0.19$) and for VIP (all trials, 1.49 ± 0.09, trials without microsaccades, 1.47 ± 0.09, $p = 0.41$). The result was similar for the motion-pulse task in MT (all trials, 1.73 ± 0.15, trials without microsaccades, 1.72 ± 0.15, $p = 0.25$).

Discussion

We examined the effect of microsaccades on neural activity and behavioral performance in three motion-based detection tasks and estimated the contribution of microsaccades to the correlation between neural firing and behavior. We found that microsaccades were associated with significantly reduced detection performance for transient changes in speed or coherency or sustained changes in coherency. Microsaccades were also associated with slower response times. Additionally, in all three experiments, we found that microsaccades suppressed neural activity during stimulus presentation. All together, the dual effect of microsaccades on neurophysiology and behavior contributed 7–19% of the detect probability in area MT, 21% in area LIP, and up to 17% in area VIP, depending on the time windows used for the analysis. The observation that neural activity and behavior are correlated has profoundly shaped our view of visual system function. Our observation that microsaccades can account for up to one-fifth of this correlation between neural activity and behavior is surprising and may have implications for how these data are used to constrain neural models of the neural signals underlying sensory perception.

Accuracy of microsaccade detection

Low amplitude, brief events such as microsaccades are inherently difficult to detect and characterize. Although we believe our detection algorithm was robust, it is possible that some microsaccades went undetected or that non-microsaccade eye-

movement signals were misclassified as microsaccades. Critically, neither type of error would be expected to produce our findings. For example, when calculating the detect probability we separated trials based on whether they contained a microsaccade near stimulus onset (Fig. 6). If the algorithm missed a microsaccade, and accidentally included the trial in the no-microsaccade group, the detect probability of that group should increase, reducing our estimate of microsaccades' contribution to the detect probability. Alternatively, if the algorithm improperly classified an event as a microsaccade (akin to randomly eliminating a trial from the no-microsaccade group), this would not bias our estimate of the microsaccade-free detect probability in either direction, although it would make our estimate of the underlying value less reliable. Therefore, with respect to errors in microsaccade detection, our estimate of the contribution of microsaccades to detect probability could be considered a lower bound.

On the other hand, we do not believe that we greatly underestimated the contribution of microsaccades to the detect probability. Our velocity thresholds for detecting microsaccades (between 8 and 10 degrees/second) were chosen to be consistent with past studies (Bair and O'Keefe, 1998; Leopold and Logothetis, 1998), and the rate of microsaccades in our three experiments (0.3 - 1.3 Hz) was in the range of previous observations (Table 1). However, we were curious to know what effect a higher microsaccade rate might have on the detect probability. Shifting these thresholds more than a couple of degrees/second only reduced the contribution that microsaccades made to the detect probability (data not shown). Therefore, we do not believe that we have

grossly underestimated the contribution of microsaccades to the correlation between neural activity and behavior.

Variability of the effect of microsaccades on detect probability

The effect of microsaccades on detect probability varied among the different experiments and cortical areas: we found the strongest effects for MT and LIP in the speed-pulse experiment as well as MT for the motion-step experiment, differing effects on VIP depending on which time windows were considered and a small effect for area MT in the motion-pulse experiment (Fig. 6). However, these apparent differences should be treated cautiously. Undoubtedly, a main source of variability is that the data were collected from different animals in different laboratories using slightly different methodologies. While this experimental diversity generally underscores the robustness of our findings, it likely accounted for some of the variability in the findings. For example, the smaller effect on detect probability that we found for the motion-pulse experiment may be because we used a video-tracking system to measure eye position. The system was not optimally calibrated to detect small, rapid eye movements, so we had less confidence in our ability to detect microsaccades for that experiment. Nonetheless, despite experimental differences, it is remarkable that microsaccades contributed a relatively consistent amount to the detect probability in all three cortical areas.

For VIP during the motion-step experiment, microsaccades contributed almost nothing to the detect probability when counting spikes from 50 to 300 ms after the onset

of coherent motion, but contributed about 17% when spikes were only counted from 50 to 150 ms. One possible explanation is that decision-related "feedback", which is prominent in VIP (Cook and Maunsell, 2002b), might have overwhelmed any effect microsaccades had on the detect probability by providing a much stronger source of behaviorally correlated, neural activity. By limiting the window used to count spikes to immediately after test-stimulus onset, we likely reduced any potential role for feedback.

Microsaccades, saccadic suppression and the reduction in detection performance

Subjectively, we do not perceive the world to move during microsaccadic eye movements. This suggests that there are compensatory mechanisms to counter the motion signals that microsaccades generate (Murakami and Cavanagh, 1998, 2001). This compensatory mechanism might interfere with a subject's ability to detect a motion stimulus near the time of the microsaccade, leading to the observed reduction in performance. However, studies of the impact of microsaccades on visual detection tasks have produced contradictory results, with some reporting increased detection thresholds (Ditchburn, 1955; Beeler, 1967) and others reporting no change (Krauskopf, 1966; Sperling, 1990). Unlike our experiments, these studies examined detection thresholds for flashed stimuli. It is possible that motion detection is specifically suppressed during microsaccades. This suppression could be particularly pronounced during the highly trained and difficult motion-detection tasks that we used in our study. It would be interesting to examine the effects of microsaccades on neural responses and behavior in

other visual tasks that do not employ a motion stimulus, or which do not require a difficult detection of a threshold motion stimulus.

Another interesting question is whether our findings on microscaccades extend to saccades, which are larger, voluntary eye movements. Saccades produce full-field apparent motion of the visual scene, which is generally not perceived despite being well within the detection capability of our motion system. The perceptual suppression of saccade-induced motion, termed saccadic suppression, is thought to be secondary to a more general suppression of visual perception at the time of the saccade, though the neural mechanisms and functional significance of this phenomenon remain a matter of debate (Castet et al., 2001; Ross et al., 2001; Wurtz, 2008). Saccades are known to raise visual detection and discrimination thresholds for a variety of stimuli, with the strongest suppression for high luminance and low spatial frequency stimuli (Burr et al., 1982; Burr et al., 1994; Ross et al., 2001; Wurtz, 2008). Saccades also markedly impair motion perception (Bridgeman et al., 1975; Burr et al., 1982; Shioiri and Cavanagh, 1989; Ilg and Hoffmann, 1993; but see Castet and Masson, 2000).

Saccades and microsaccades share a fixed relationship between saccadic peak velocity and amplitude (the main sequence), suggesting a final common neural pathway (Zuber and Stark, 1965). However, microsaccades are distinct in that they are involuntary, although trained subjects have a limited capacity to suppress them (Martinez-Conde et al., 2004). In addition, studies of saccadic effects on perception have employed saccades of at least several degrees in length. Findings for microsaccades,

which are usually less than a degree in length, may be quantitatively or qualitatively different. Nonetheless, it would be interesting in the future to examine if our findings for microsaccades extend to saccades.

Extension of our findings to other paradigms

We observed that microsaccades reduced detection performance and, on average, suppressed neural firing during our tasks. The two effects were correlated on a trial-by-trial basis, so that microsaccades contributed positively to the magnitude of the correlation between activity and behavior. Minimally, our findings serve as a caution that microsaccades should be taken in to consideration when measuring correlations between neuronal activity and behavior. However, our specific results may not extend to all such experiments. For example, in our design, we used test stimuli that were chosen to match the recorded neuron's preferred direction, and we thus expected the test stimulus to trigger an increase in the neuron's response – which we found. Consider a hypothetical experiment in which motion in the preferred direction is followed by a test stimulus in a *less*-preferred direction. In such an experiment the neuron's expected response would be a decrease in firing. If, as we observed, microsaccades suppressed neural firing and decreased behavioral detection, we would instead expect microsaccades to *decrease* the magnitude of the correlation between activity and behavior.

There have been many studies with motion-based stimuli that measured choice probability in two-alternative forced-choice paradigms (e.g., Britten et al., 1996; Dodd et

al., 2001; Grunewald et al., 2002; Williams et al., 2003; Huk and Shadlen, 2005; Liu and Newsome, 2005; Purushothaman and Bradley, 2005). Could microsaccades have made a similar contribution to these results? Few previous studies have considered this issue. Dodd et al. demonstrated that the direction of microsaccades did not correlate with the animal's choice, but did not assess whether they impact the animal's behavior in other ways (by degrading accuracy, for example). As our values of detect probability were in the range of those reported for choice probability in these previous studies, even if microsaccades made a similar contribution it is unlikely that they would have significantly altered the main conclusions of these studies. In theory, microsaccades could have had a more significant impact in other studies that reported smaller choice-probability values for V2 (Nienborg and Cumming, 2006) and V1 (Palmer et al., 2007). However, these experiments did not involve motion stimuli. Moreover, one of these studies (Nienborg and Cumming, 2006) examined the possible effect of microsaccades (although in different a manner than ours) and found no contribution.

Thus it is not clear what impact microsaccades should be expected to have on choice probability. Although choice probability and detect probability are analogous quantities, there are critical distinctions in the associated paradigms. Consider a two-alternative forced-choice task in which the monkey must report whether low-coherence motion was in one of two opposed directions. First, consider the case where there is no relation between the frequency or direction of microsaccades and the choice of the subject. For example, microsaccades, regardless of direction, may not bias the subject towards one choice or the other but simply result in more incorrect choices. In this case,

microsaccades would tend to *decrease* the measured choice probability by introducing spike-rate variance that was not correlated with behavioral choice. Alternatively, consider the case where the frequency or direction of microsaccades influences the subject's choice. This would be the case if microsaccades produced a direction-specific neural response and a motion percept that biased the subject's choice depending on microsaccade direction. In this case, microsaccades would tend to *increase* the measured choice probability (assuming the neural and perceptual effects were aligned). One could easily imagine other scenarios where the effect of microsaccades on choice probability would be task and stimulus dependent.

The observation that single-trial responses from individual neurons are correlated with behavioral choice almost certainly requires that trial-by-trial variance in neural responses are correlated across many neurons (Shadlen et al., 1996), which has been demonstrated in a limited manner through paired recordings (Zohary et al., 1994). This correlation may arise from many possible sources including intrinsic properties of the neural networks, variance in the stimulus, feedback activity related to choice, fluctuations in attention or vigilance and, as we have demonstrated here, eye movements. The nature of the causal link between this correlated neural response variance and behavior is the critical question common to all choice- or detect-probability results. For example, in our experiments it is possible that microsaccades cause parallel, unrelated effects on detection behavior and neuronal firing. In this case the effect of microsaccades on detect probability might be described as "artifact". It is also possible that behavioral effects of microsaccades are mediated in part or solely through the firing of neurons such as those

recorded in this study. In this case, microsaccades might be considered a bona fide source of noise that influences behavior by modulating the firing of those neurons. While our experiment did not address this issue directly, our partial correlation analysis suggested that at least some of the effect of microsaccades on behavior was independent of the firing of *single* neurons. These questions point to the importance of better understanding the sources of variation in sensory encoding that lead to variation in behavior. Our findings suggest that microsaccades may be one such source of variation.

References

Bahill AT, Clark MR, Stark L (1975) The main sequence, a tool for studying human eye movements. Mathematical Biosciences 24:191-204.

Bair W, O'Keefe LP (1998) The influence of fixational eye movements on the response of neurons in area MT of the macaque. Vis Neurosci 15:779-786.

Bair W, Zohary E, Newsome WT (2001) Correlated firing in macaque visual area MT: time scales and relationship to behavior. J Neurosci 21:1676-1697.

Beeler GW (1967) Visual threshold changes resulting from spontaneous saccadic eye movements. Vision Res 7:769-775.

Bridgeman B, Palca J (1980) The role of microsaccades in high acuity observational tasks. Vision Res 20:813-817.

Bridgeman B, Hendry D, Stark L (1975) Failure to detect displacement of the visual world during saccadic eye movements. Vision Res 15:719-722.

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis Neurosci 13:87-100.

Burr DC, Morrone MC, Ross J (1994) Selective suppression of the magnocellular visual pathway during saccadic eye movements. Nature 371:511-513.

Burr DC, Holt J, Johnstone JR, Ross J (1982) Selective depression of motion sensitivity during saccades. J Physiol 333:1-15.

Castet E, Masson GS (2000) Motion perception during saccadic eye movements. Nat Neurosci 3:177-183.

Castet E, Jeanjean S, Masson GS (2001) 'Saccadic suppression'- no need for an active extra-retinal mechanism. Trends Neurosci 24:316-318.

Colby CL, Duhamel JR, Goldberg ME (1996) Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. J Neurophysiol 76:2841-2852.

Cook EP, Maunsell JH (2002a) Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. J Neurosci 22:1994-2004.

Cook EP, Maunsell JH (2002b) Dynamics of neuronal responses in macaque MT and VIP during motion detection. Nat Neurosci 5:985-994.

Cook EP, Maunsell JH (2004) Attentional modulation of motion integration of individual neurons in the middle temporal visual area. J Neurosci 24:7964-7977.

Cornsweet TN (1956) Determination of the stimuli for involuntary drifts and saccadic eye movements. J Opt Soc Am 46:987-993.

de Lafuente V, Romo R (2005) Neuronal correlates of subjective sensory experience. Nat Neurosci 8:1698-1703.

de Lafuente V, Romo R (2006) Neural correlate of subjective sensory experience gradually builds up across cortical areas. Proc Natl Acad Sci U S A 103:14266-14271.

Ditchburn RW (1955) Eye-movements in relation to retinal action. Opt Acta (Lond) 1.

Dodd JV, Krug K, Cumming BG, Parker AJ (2001) Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. J Neurosci 21:4809-4821.

Gibson JR, Maunsell JH (1997) Sensory modality specificity of neural activity related to memory in visual cortex. J Neurophysiol 78:1263-1275.

Green DM, Swets JA (1966) Signal Detection Theory and Psychophysics. New York: John Wiley and Sons.

Grunewald A, Bradley DC, Andersen RA (2002) Neural correlates of structure-from-motion perception in macaque V1 and MT. J Neurosci 22:6195-6207.

Gu Y, DeAngelis GC, Angelaki DE (2007) A functional link between area MSTd and heading perception based on vestibular signals. Nat Neurosci 10:1038-1047.

Gur M, Beylin A, Snodderly DM (1997) Response variability of neurons in primary visual cortex (V1) of alert monkeys. J Neurosci 17:2914-2920.

Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J Neurosci 25:10420-10436.

Ilg UJ, Hoffmann KP (1993) Motion perception during saccades. Vision Res 33:211-220.

Judge SJ, Richmond BJ, Chu FC (1980) Implantation of magnetic search coils for measurement of eye position: an improved method. Vision Res 20:535-538.

Krauskopf J (1966) Lack of inhibition during involuntary saccades. Am J Psychol 79:73-81.

Krug K, Cumming BG, Parker AJ (2004) Comparing perceptual signals of single V5/MT neurons in two binocular depth tasks. J Neurophysiol 92:1586-1596.

Leopold DA, Logothetis NK (1998) Microsaccades differentially modulate neural activity in the striate and extrastriate visual cortex. Exp Brain Res 123:341-345.

Liu J, Newsome WT (2005) Correlation between speed perception and neural activity in the middle temporal visual area. J Neurosci 25:711-722.

Martinez-Conde S, Macknik SL, Hubel DH (2000) Microsaccadic eye movements and firing of single cells in the striate cortex of macaque monkeys. Nat Neurosci 3:251-258.

Martinez-Conde S, Macknik SL, Hubel DH (2002) The function of bursts of spikes during visual fixation in the awake primate lateral geniculate nucleus and primary visual cortex. Proc Natl Acad Sci U S A 99:13920-13925.

Martinez-Conde S, Macknik SL, Hubel DH (2004) The role of fixational eye movements in visual perception. Nat Rev Neurosci 5:229-240.

Martinez-Conde S, Macknik SL, Troncoso XG, Dyar TA (2006) Microsaccades counteract visual fading during fixation. Neuron 49:297-305.

Masse NY, Cook E (2008) The effect of middle temporal spike phase on sensory encoding and correlates with behavior during a motion-detection task. J Neurosci 28:1343-1355.

Maunsell JH, Van Essen DC (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. J Neurophysiol 49:1127-1147.

Maunsell JH, Van Essen DC (1987) Topographic organization of the middle temporal visual area in the macaque monkey: representational biases and the relationship to callosal connections and myeloarchitectonic boundaries. J Comp Neurol 266:535-555.

Murakami I, Cavanagh P (1998) A jitter after-effect reveals motion-based stabilization of vision. Nature 395:798-801.

Murakami I, Cavanagh P (2001) Visual jitter: evidence for visual-motion-based compensation of retinal slip due to small eye movements. Vision Res 41:173-186.

Nienborg H, Cumming BG (2006) Macaque V2 neurons, but not V1 neurons, show choice-related activity. J Neurosci 26:9567-9578.

Palmer C, Cheng SY, Seidemann E (2007) Linking neuronal and behavioral performance in a reaction-time visual detection task. J Neurosci 27:8122-8137.

Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. Annu Rev Neurosci 21:227-277.

Purushothaman G, Bradley DC (2005) Neural population code for fine perceptual decisions in area MT. Nat Neurosci 8:99-106.

Robinson DA (1963) A Method Of Measuring Eye Movement Using A Scleral Search Coil In A Magnetic Field. IEEE Trans Biomed Eng 10:137-145.

Romo R, Salinas E (2001) Touch and go: decision-making mechanisms in somatosensation. Annu Rev Neurosci 24:107-137.

Ross J, Morrone MC, Goldberg ME, Burr DC (2001) Changes in visual perception at the time of saccades. Trends Neurosci 24:113-121.

Rucci M, Iovin R, Poletti M, Santini F (2007) Miniature eye movements enhance fine spatial detail. Nature 447:851-854.

Schoppik D, Nagel KI, Lisberger SG (2008) Cortical mechanisms of smooth eye movements revealed by dynamic covariations of neural and behavioral responses. Neuron 58:248-260.

Chapter 3

The effect of MT spike phase on sensory encoding and correlates with behavior during a motion detection task

Abstract

Past studies have shown that sensory neurons that are the most informative of the stimulus tend to be the best correlated with the subject's perceptual decision. We wanted to know if this relationship might also apply to short time segments of a neuron's response. We asked if spikes that conveyed more information about a motion stimulus were also more tightly linked to the perceptual behavior. We examined single neuron activity in area MT while monkeys performed a motion-detection task. Due to a slow stimulus update (every 27 ms), activity in many MT neurons was entrained and phase-locked to the stimulus. These stimulus-entrained neuronal oscillations allowed us to separate spikes based on phase. We observed a large amount of variability in how spikes at different phases of the oscillation encoded the stimulus, as revealed by the spike-triggered average of the motion. Spikes during certain phases of the cycle were much more informative about the presence of coherent motion than others. Importantly, we found that the phases that were the most informative about the motion stimulus were also more correlated with the behavioral performance and reaction time of the animal. Our results suggest that the relationship between a neuron's spikes, the stimulus and behavior can vary on a time scale of tens of milliseconds.

Introduction

How is the activity of neurons in the sensory areas of cortex related to our perceptual abilities? This question is fundamental to the broader understanding of how the brain underlies behavior (for reviews, see Parker and Newsome, 1998; Romo and Salinas, 2001). Past studies have shown that the activity of many cortical sensory neurons covaries with the behavior of animal subjects performing perceptual tasks (Celebrini and Newsome, 1994; Britten et al., 1996; Dodd et al., 2001; Cook and Maunsell, 2002b; Parker et al., 2002; Uka and DeAngelis, 2004; de Lafuente and Romo, 2005; Nienborg and Cumming, 2006), but see (de Lafuente and Romo, 2005). Importantly, the correlation between neuronal activity and perceptual choice tends to be strongest for neurons that are the most informative about the stimulus (Celebrini and Newsome, 1994; Britten et al., 1996; Cook and Maunsell, 2002b; Parker et al., 2002; Purushothaman and Bradley, 2005). This result suggests that a subject's perceptual choice is preferentially based on the activity of a subset of neurons that convey the most reliable task-related information.

The link between perception and neural activity has typically been studied in the context of a neuron's spike rate. However, many other possible stimulus-coding strategies have been examined. For example, different components of a neuronal response, such as the onset, sustained and offset portions, have been shown to encode different aspects of the stimulus (Lamme, 1995; Gawne et al., 1996; Friedrich and Laurent, 2001; Pack and Born, 2001; Reich et al., 2001; Bair et al., 2002; Shapley et al., 2003; Hegde and Van Essen, 2004; Roelfsema et al., 2007). Furthermore, many studies,

both theoretical (Hopfield, 1995; Thorpe et al., 2001; Mehta et al., 2002; Chacron et al., 2004; Guyonneau et al., 2005; Gutig and Sompolinsky, 2006) and experimental (Celebrini et al., 1993; O'Keefe and Recce, 1993; Konig et al., 1995; deCharms and Merzenich, 1996; Gawne et al., 1996; Victor and Purpura, 1996; de Ruyter van Steveninck et al., 1997; Mechler et al., 1998; Fries et al., 2001; Lu et al., 2001; Johansson and Birznieks, 2004; Chacron et al., 2005; Womelsdorf et al., 2006; Fries et al., 2007; Sadeghi et al., 2007) have proposed that the timing of action potentials plays a role in stimulus coding. These results raise the possibility that some spikes carry more information about a behaviorally relevant stimulus than others. Thus, if the most informative neurons exert greater influence on perceptual choice, do the most informative spikes also have a greater influence on perceptual choice?

Addressing this hypothesis would seem challenging because estimating the sensory information contained in a single spike is not readily possible. However, we were able to approach this question using data from a motion detection experiment where the visual motion stimulus was updated at a slow enough rate (approximately every 27 ms) to induce oscillations in the activity of Middle Temporal (MT) neurons. From these neuronal oscillations, we were able to group spikes based on phase and then ask if sensory and choice-related information varied as a function of phase.

Using two measures of sensory encoding (spike-triggered average and neurometric value), we found that spikes encoded a motion stimulus differently depending on the phase of the neuronal response. Additionally, spikes that occurred

during the phases that were most informative about the motion stimulus were also more correlated with both the perceptual choice and reaction time of the animal. Our analysis suggests that the link between the activity of single neurons and perceptual choice can vary on a relatively fast timescale of tens of milliseconds.

Methods

Behavioral task

The data set analyzed for this study comes from several previous studies (Cook and Maunsell, 2002a, 2004). Monkeys (Macaca mulatta) were trained to perform two variants of a motion detection task. In the first set of experiments, the animals performed a spatially cued motion detection task (Figure 1A). The trial began when the monkey depressed a lever and fixated on a central point. The goal of the task was to release the lever when coherent motion began in one of two random dot patches diametrically opposite of the fixation point. After the cue was presented, 0% coherent motion began in the two patches followed by coherent motion occurring in one of the two patches at a random time (flat hazard function) 500 to 8000 ms afterwards. The location of the coherent motion was cued to the monkey at the start of the trial with static dots and this cue was valid on 80% of the trials. Trials with invalid cues were discarded for this analysis. The strength of coherent motion was varied between three levels (low, medium and high), with the monkey correctly detecting the coherent motion 50%, 90% and 99% of the time for the three levels. The coherence levels for each experiment were adjusted depending on the eccentricity, speed and radius of the random dot patch in order to produce the target performance. 93 MT neurons were recorded in two monkeys.

The second set of experiments (Figure 1B) was similar to the first except that only a single random dot patch was presented and only two coherence levels (low and high)

were used. 25 MT neurons were recorded for this stimulus configuration in one monkey. The correlations between the neuronal response and the animal's behavioral performance was not appreciably different between the single and double patch experiments (Cook and Maunsell, 2002b).

In both sets of experiments, the direction and the speed of the coherent motion were matched to the preferred direction and speed of the neuron under study. The coherent motion lasted 750 ms and the monkey had to release the lever from 200 to 750 ms after the onset of coherent motion to obtain a juice reward. Trials where the monkey failed to release the lever or released the lever too late were deemed missed trials. Only correct and missed trials were included in the analysis. Trials where the monkey released the lever too early or was unable to maintain fixation were discarded from the analysis.

Visual stimulus

The animal sat 62 cm from a computer monitor (+/-17° × +/-13° of visual angle; 1600 × 1200 pixels; 75 Hz refresh). The stimuli consisted of two patches of white dots (each dot 0.25° diameter; 78 cd/m^2) on a dark gray background (12 cd/m^2) with a dot density of 2.1 dots/$degree^2$. Each patch of dots was updated on every other video frame (approximately every 27 ms) using the following procedure. The dots in each patch were evenly divided into two groups. On each update, one group was replaced with new, randomly positioned dots, whereas dots in the other group were displaced by a fixed distance. The dots in this latter group determined the motion coherence. For 0%

coherence, all the dots in this group moved a fixed distance in a random direction. For coherent motion greater than zero, a proportion of the dots moved with a fixed distance in the same direction. This proportion determined the strength of the coherent motion. On the next update (27 ms later), the groups were switched. This arrangement insured that all the dots had a lifetime of four video frames (i.e., two stimulus updates) before they were replaced. Because half the dots are always randomly replotted regardless of the proportion of dots moving coherently, our motion had a maximum strength of 50% coherence. For example, at 25% coherent motion, half the dots are randomly replotted, one-quarter are moving with the same fixed distance and direction, and one-quarter are moving with the same fixed distance in a random direction.

Data collection

Using standard extracellular recording techniques (Gibson and Maunsell, 1997), recordings were made from well-isolated single neurons in area MT in both animals. When a neuron was isolated, the receptive field (RF) was mapped using a manually controlled bar while the animal fixated on a central spot. The diameter of the RFs ranged from 3.9 to 10.7° (median, 7.4). RF center eccentricities ranged from 3.9 to 11.1° (median, 7.9). The preferred speed was also judged using a bar moved by hand. The animals were trained to perform the task at slow or moderate motion speeds, so neurons with a preferred speed between 4 and 12°/sec were usually selected. Once the RF location, size, preferred direction, and speed were determined, the motion detection task was then run, and the neuron was recorded from for as long as possible. For some

neurons, a memory saccade task was also run, but these data were not used in this analysis. The number of completed trials per coherence level for the motion detection task ranged from 15 to 175 (median, 35). The monkey's performance varied with patch location, size, and motion speed, which were determined by the response properties of the neuron under study. Consequently, different neurons were tested with different coherence levels. The animal's eye position was measured every 5 ms using a scleral search coil (Robinson, 1963; Judge et al., 1980) and the times of action potentials were recorded to the nearest millisecond.

Detect probability and neurometric value

Detect probability (DP) expresses the ability to predict the behavioral outcome of a trial in a detection task given the neuronal response (Cook and Maunsell, 2002b). DP is analogous to the choice probability calculation used in discrimination tasks (Celebrini and Newsome, 1994; Britten et al., 1996). The neurometric value expresses the ability to predict which one of two stimuli was presented given the neuronal response.

DP and neurometric value are formulated in the same way using traditional ROC analysis (Green, 1975). Briefly, given two random samples from different distributions, the DP and the neurometric value are the probabilities that one can correctly determine from which distribution the samples were drawn. In other words, it is a measure of how separate two distributions are from each other. Specifically, given two distributions of neuronal responses, $p(x)$ and $q(x)$, the maximum probability that we could correctly

determine from which distribution a sample was drawn from is $P(correct) = \int_0^\infty \int_x^\infty p(x)q(y)dydx$, assuming that on average, a sample drawn from $q(y)$ is greater than a sample drawn from $p(x)$.

For detect probability, $q(y)$ corresponds to the distribution of spike rates for correct trials and $p(x)$ for failed trials. For the neurometric value, $p(x)$ and $q(y)$ are the distribution of spike rates corresponding to the 0% and coherent motion, respectively. A DP of 1.0 indicates the neuronal distributions, $p(x)$ and $q(y)$, do not overlap and the behavioral outcome is completely predictable from the neuronal responses. Likewise, a neurometric value of 1.0 indicates the motion stimulus encoded (either 0% coherent or coherent) is completely predictable from the spike rate. A DP or neurometric value of 0.5 indicates a chance prediction of the behavioral outcome or the state of the coherent motion, respectively.

For the DP calculation, only neuronal responses occurring during low coherent motion trials from 50 to 150 ms after coherent motion onset were included. We only used the low coherent motion trials to compute DP because this was the only condition that produced similar numbers of correct and failed trials. For the neurometric value, we compared the neuronal response from 50 to 150 ms after coherent motion onset to the neuronal response from 100 to 0 ms before the coherent motion onset for all three coherence levels. The neurometric value was computed individually for each of the three motion coherences and then averaged.

Reaction time correlation coefficient

Reaction time (RT) correlation coefficients provide the correlation between the neuronal response and the reaction times. Given the neuronal responses and the respective RTs, the Pearson product-moment correlation coefficient is calculated between these two sets. Since neither the spike count distributions nor the reaction time distributions were normally distributed, the calculations were also performed using the Spearman's rank correlation coefficient. The correlation coefficient (γ) for each of the three coherence levels was calculated individually and then a weighted (by the number of trials) average was taken. The RT correlation coefficient was computed using the neuronal response from 50 to 150 ms after coherent motion onset.

Extracting global motion from the stimulus

We reduced our random dot stimulus to a one dimensional time series of global motion strength along the neuron's preferred/null axis (see Figure 4A). To estimate the net motion between two successive frames of random dots, we applied a previously described method (Barlow and Tripathy, 1997) that computed the correspondence between dots in two sequential frames. We calculated the motion vector between each pair of dots in successive frames. Thus, if our patch contained N dots, we computed N^2 motion vectors for each pair of successive frames, where the magnitude of the vectors represented the speed. We then scaled each vector by passing its magnitude through a speed-filter. The speed-filter was a Gaussian distribution in the log speed domain (Nover et al., 2005) centered at the preferred speed of the neuron (determined by sweeping a bar

through the receptive field) and with a standard deviation of 0.5 deg/sec. We filtered speed because speed-tuning profiles were not collected from the neurons, however, our results were insensitive to the parameters of the speed-filter. Each weighted motion vector was then projected onto the neuron's preferred/null direction of motion axis. The sum of all the projected vectors was the motion strength at that time point. This calculation was repeated for every pair of successive random dot patches to create a time series of the motion strength in the preferred/null axis of the neuron. The resultant time series was normalized to have a standard deviation of one and then smoothed with a Gaussian filter with a standard deviation of 3 ms.

Results

The goal of this study was to examine how the phase of stimulus-driven oscillatory activity in MT neurons was related to both the encoding of the motion stimulus and behavior. The results are based on the activity of single neurons recorded from two monkeys performing a motion-detection task (Figure 1 and Methods).

In this task, the monkeys released a lever when a patch of random dot motion began moving in a coherent manner. The location and size of the stimulus was matched to the classical RF of the MT neuron under observation, and the direction and speed of the coherent motion step was matched to the neuron's preferred tuning. Because we were interested in accurately correlating neuronal responses with behavioral outcome, we selected neurons for analysis based on a criterion of at least three correct and three missed trials for the low coherent motion condition. This resulted in 115 out of the original 118 MT neurons in this data set selected for analysis.

Stimulus-driven oscillations of MT neurons

Although the stimulus was presented on a monitor with a 75 Hz refresh rate, due to hardware limitations the dot locations in each patch were updated on every other refresh period. Because of this, many of our MT neurons demonstrated oscillatory firing with a frequency that matched the updates of the motion stimulus, which occurred once every 27 ms (every other frame). An example of this stimulus-driven oscillatory

Figure 1 Motion detection task. (A) The first set of experiments used two random dot patches and included 93 neurons. Once the animals fixated, a static cue was presented to indicate the most likely position of where the coherent motion would occur. Next, 0% coherent motion was presented in both the neuron's RF and in a patch diametrically opposite. Coherent motion randomly began in one of the patches from 500 to 8000 ms with a flat hazard function. The coherent motion lasted for 750 ms and the animal had to release the lever from 200 to 750 ms after the coherent motion onset to order to obtain a reward. (B) The second set of experiments used only a single random dot patch and included 25 neurons. This paradigm was identical to the one above, except that only a single random dot patch that overlapped the RF of the neuron was used.

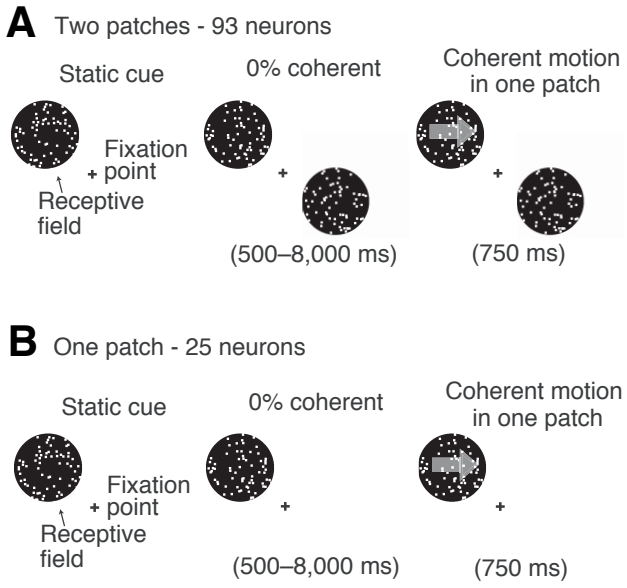
A
Two patches - 93 neurons
Static cue
0% coherent
Coherent motion in one patch
Fixation point
Receptive field
(500–8,000 ms)
(750 ms)
B
One patch - 25 neurons
Static cue
0% coherent
Coherent motion in one patch
Fixation point
Receptive field
(500–8,000 ms)
(750 ms)

Figure 1

behavior is shown for an MT neuron in Figure 2A. The activity of the neuron was entrained with a fixed phase to the stimulus updates (vertical lines). The oscillatory behavior occurred during both the 0% coherent motion and the coherent motion step. Although a 27 ms separation between motion updates may not be optimal for MT neurons (Churchland and Lisberger, 2001), it was still sufficient to robustly drive our population of cells.

To characterize the strength of the stimulus-induced oscillatory response, we collapsed the average spike rate during the 0% coherent motion to a single update, smoothed the curve using a Gaussian with a standard deviation of 1ms, normalized the area under the curve, and then centered the peak of activity at zero. For our example neuron, the neuronal response as a function of phase is illustrated in Figure 2B. The top panel of Figure 2B shows the neural activity as a function of phase aligned to the stimulus updates. As with Figure 2A, it also shows that neural activity tends to peak a few milliseconds before the next stimulus update. In the bottom panel of Figure 2B, the curve is shifted so that the peak in neural activity is aligned to the center and the phase is measured relative to this peak. All following plots of neural activity versus phase are shown in this manner. This response versus phase relation was computed for all neurons, and the oscillatory amplitude (i.e., the strength of the oscillations) was expressed as the maximum minus the minimum value.

The distribution of oscillatory amplitudes for our population of 115 neurons is shown in Figure 3A. Our example neuron in Figure 2 had prominent oscillatory firing

Figure 2 Analysis of an example neuron. (A) Raster plot and binned response of an example MT. Motion stimulus updates are thin vertical lines. The monitor's vertical refresh rate was 75 Hz or twice the stimulus update rate. The neural activity is locked to the motion updates and is greatest just prior to each motion update. Scale bar is 40 spikes per second . (B) The average spike rate during the 0% coherent motion as a function of phase for the same neuron. The data was smoothed with a 1 ms Gaussian filter and normalized so that the area under the curve was equal to one. In the top panel, the phase is relative to the stimulus updates. In the bottom panel, the curve is shifted so that the peak is aligned at the center, and the phase is given relative to this peak. The oscillatory amplitude (given in arbitrary units) is the difference between the maximum and minimum values of the curve. The black section of the curve (corresponding to 1 to 13 ms the left of the peak) is defined as the rising phase, and the gray section of the curve (corresponding to 1 to 13 ms to the right of the peak) is defined as the falling phase. (C) The STA for the rising (black) and falling (gray) phases of the example neuron. Positive amplitudes indicate motion in the neuron's preferred direction (see Figure 4A). The STA was computed from the 0% coherent motion and was much larger for the rising phase than for the falling phase. (D) The neurometric value for the rising phase (black), falling phase (gray) and the entire cycle (white) of the example neuron. The neurometric value was calculated using ROC analysis applied to the responses 100 to 0 ms before the coherent motion step and 50 to 100 ms after the coherent step. Activity on the rising phase is more informative of the coherent motion step than activity on the falling phase. (E) The average spike rate for correct (solid) and failed (dashed) trials for the low coherence motion steps only. The spike rate for correct trials is greater on the rising phase but less than the spike rate for failed trials on the falling phase. (F) Detect probability (DP) using spikes occurring on the rising phase (black), falling phase (gray) and the entire cycle (white) of the low coherence motion. DP was calculated using the response 50 to 100 ms after the onset of the coherent motion step. Activity on the rising phase is positively correlated with the animal correctly detecting the motion signal and is greater than the DP using all spikes. Activity on the falling phase is negatively correlated with the animal correctly detecting the motion stimulus.

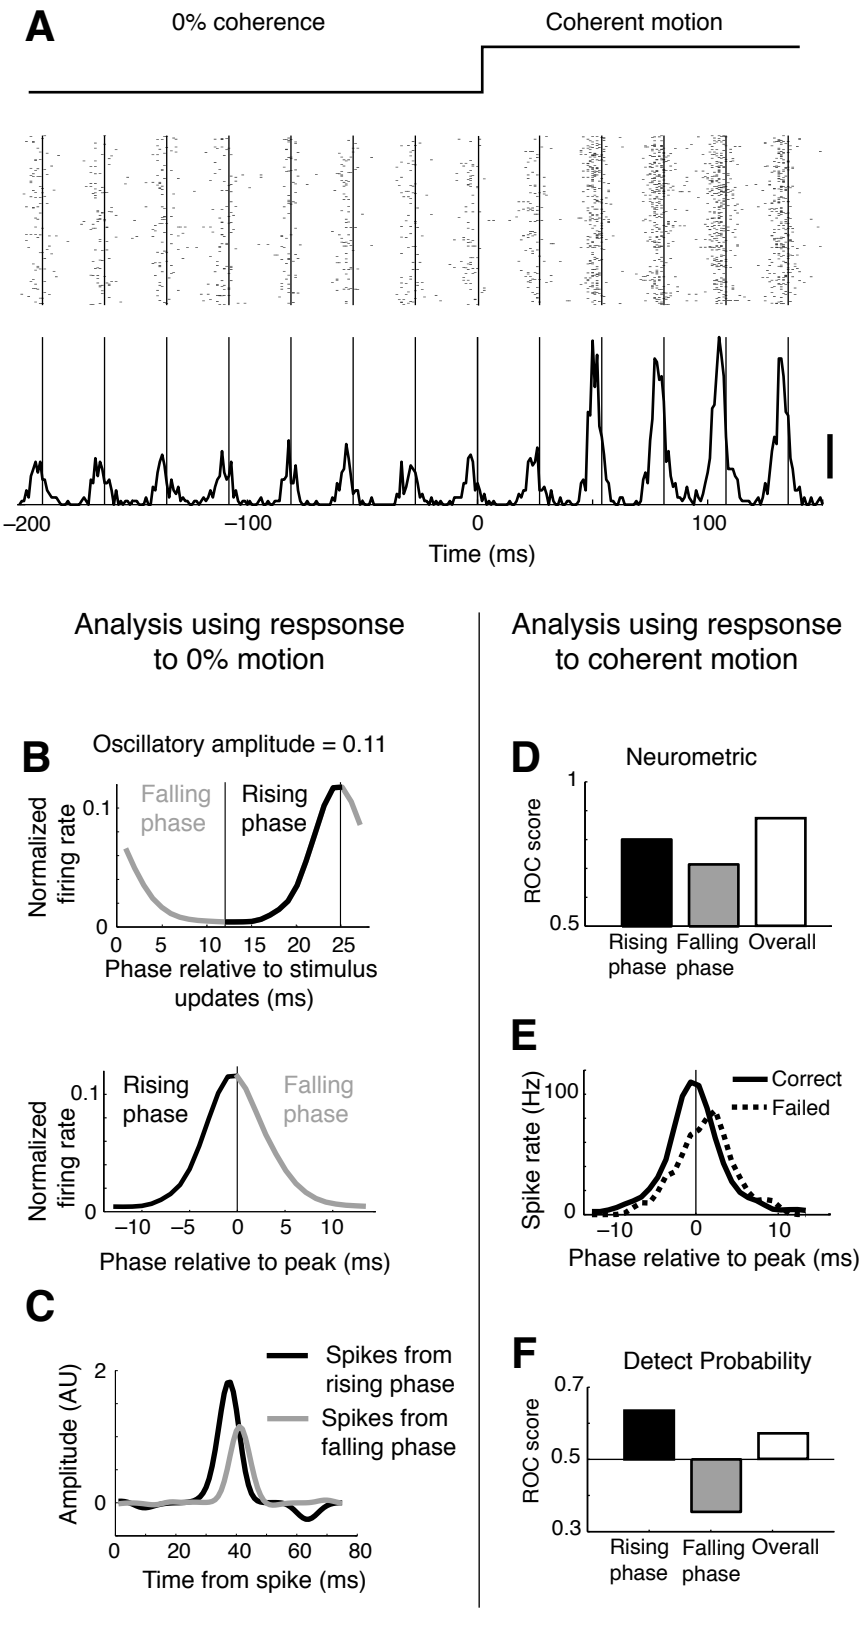
A
0% coherence
Coherent motion
–200
–100
0
100
Time (ms)
Analysis using respsonse to 0% motion
Analysis using respsonse to coherent motion
B
Oscillatory amplitude = 0.11
Normalized firing rate
0.1
0
Falling phase
Rising phase
0 5 10 15 20 25
Phase relative to stimulus updates (ms)
Rising phase
Falling phase
–10 –5 0 5 10
Phase relative to peak (ms)
C
Amplitude (AU)
2
0
Spikes from rising phase
Spikes from falling phase
0 20 40 60 80
Time from spike (ms)
D
Neurometric
ROC score
1
0.5
Rising phase
Falling phase
Overall
E
Spike rate (Hz)
100
0
Correct
Failed
–10 0 10
Phase relative to peak (ms)
F
Detect Probability
ROC score
0.7
0.5
0.3
Rising phase
Falling phase
Overall

Figure 2

and a corresponding high oscillatory amplitude of 0.11. Two other neurons with lower oscillatory scores are illustrated in Figures 3B and C. The average normalized spike rate as a function of phase for our population of neurons is shown in Figure 3D.

The dependence of stimulus encoding on neuronal phase

We first wanted to know whether action potentials at different phases of a neuron's oscillatory response differed in how they encoded the motion contained in our random-dot stimulus. To make this question more explicit, we asked whether the spike-triggered average (STA) of the motion stimulus occurring during the rising phase of the oscillatory cycle was different than the STA occurring during the falling phase of the oscillatory cycle.

An STA with a large positive lobe indicates that, on average, highly coherent motion in the preferred direction preceded each spike, which implies that the neuron was very selective for coherent motion in the preferred direction. An STA that is relatively flat indicates that the neuron was not very selective for coherent motion. We wanted to know whether the STA was dependent upon the phase of the neuronal oscillatory response. If we assume a common model of a cortical neuron as a linear filter (which is estimated by the STA) followed by a static nonlinearity, known as the LNP model (Paninski et al., 2004), then the average input to the nonlinearity is equal to the area under the STA times the average motion strength. Thus, we used the area under the STA to quantify how informative spikes are about the motion strength.

Figure 3 Oscillatory neural activity for the population of neurons. (A) Histogram of the oscillatory amplitudes of the neural oscillations. For each neuron, we calculated the spike rate at each phase, the result was smooth with a Gaussian filter and then the amplitude was taken to be the maximum minus the minimum value. (B), (C) Average spike rate as a function of phase for an example neurons display medium and low-level oscillatory behavior, respectively. (D) Average spike rate as a function of phase for all 115 MT neurons. Dashed lines are SEM.

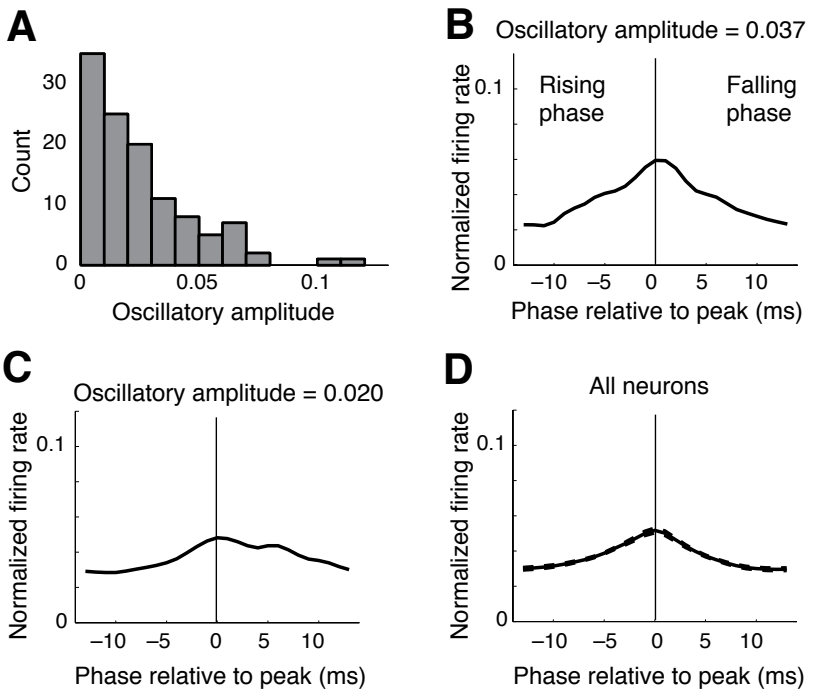
A
Count
30
20
10
0
0
0.05
0.1
Oscillatory amplitude
B
Oscillatory amplitude = 0.037
Rising phase
Falling phase
Normalized firing rate
0.1
0
–10 –5 0 5 10
Phase relative to peak (ms)
C
Oscillatory amplitude = 0.020
Normalized firing rate
0.1
0
–10 –5 0 5 10
Phase relative to peak (ms)
D
All neurons
Normalized firing rate
0.1
0
–10 –5 0 5 10
Phase relative to peak (ms)

Figure 4

To compute the STA we projected the motion in our random-dot stimulus onto the preferred/null axis as determined by the neuron's directional selectivity (see Methods). Although a previous study had described the random-dot motion using the dimensions of direction, speed and time (Cook and Maunsell, 2004), we chose to use a reduced description of the motion time series because it improved the signal-to-noise ratio of our STAs. Given the motion time series and the corresponding spike train, the STA was calculated by averaging the motion stimulus (in the preferred/null axis) that preceded each spike (Figure 4A). Only the responses to the 0% coherent motion were used to compute the STA.

The STAs for our example cell are shown in Figure 2C. We computed two STAs using either spikes occurring on the rising phase (black curve, 1 to 13 ms to the left of the peak in Figure 2B) or spikes occurring on the falling phase (gray curve, 1 to 13 ms to the right of the peak in Figure 2B). The rising and falling phases were always defined relative to the peak generated during 0% coherent motion. For this cell, the two STAs varied in size between the rising (black) and falling (gray) phases of the neuronal response. Spikes occurring on the rising phase were much more selective for coherent motion in the preferred direction than were spikes occurring during the falling phase.

Across our population, the mean area of the STA from spikes from the rising phase was 11.6% greater than that of the falling phase, although the difference was marginally significant (two-sided t-test, $p = 0.058$). This is illustrated by the average

Figure 4 Phase-dependence of the spike-triggered average. (A) Diagram of the stages used to compute the STA of a neuron. From the successive random dot patches (top panel), we computed the time series of motion strength in the preferred/null axis of the neuron, smoothed by a Gaussian filter (middle panel). Given this time series, and the corresponding spike train (bottom panel), we could compute the STA. (B) Average aligned STAs for all 115 neurons for the rising (black) and falling (grey) phases. For each neuron, the time of the peak of the overall STA (the STA calculated using all spikes) was determined, and was defined as time zero. The STAs from spikes from the rising and falling phases were then aligned to this reference time. On average, there was little difference between the two STAs. Dashed lines are SEM. (C) Histogram of the percentage difference in areas between the STA from the rising phase and falling phase for the group of oscillatory neurons (see Methods). Many neurons possessed large difference between the STAs of the rising and falling phases.

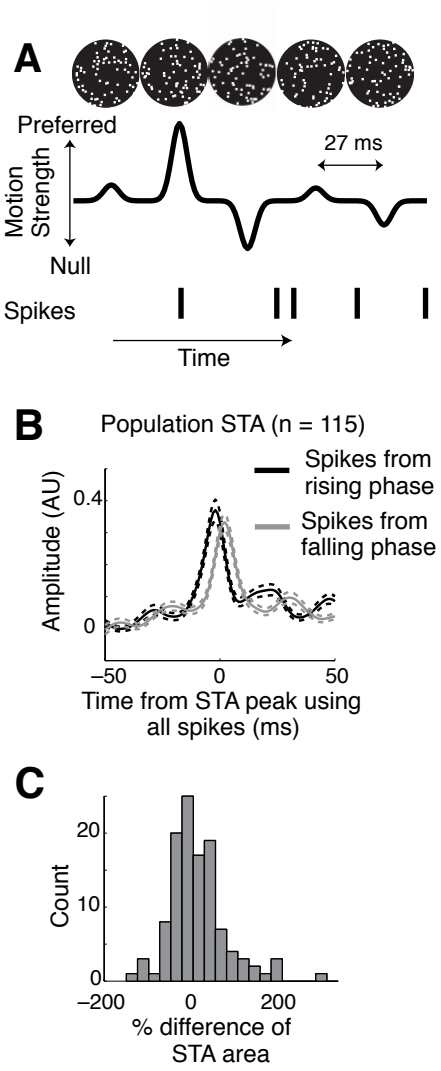
A
Preferred
27 ms
Motion Strength
Null
Spikes
Time
B
Population STA (n = 115)
Spikes from rising phase
Spikes from falling phase
Amplitude (AU)
0.4
0
–50
0
50
Time from STA peak using all spikes (ms)
C
Count
20
10
0
–200
0
200
% difference of STA area

Figure 4

STAs for our population of neurons (Figure 4B). Since neural latencies differ across neurons, we chose to align the STAs when calculating the average STAs of the population. Thus, for each neuron, we calculated the time of the peak of the overall STA (the STA calculated using all spikes), and this time was defined as time zero. The STAs from spikes from the rising and falling phases were then aligned to this reference time.

Although the mean difference in area was not appreciably different from zero, we found a surprisingly large amount of variability in the size of the STA between the two phases, as shown in Figure 4C. This histogram illustrates the distribution of the percentage difference in the areas between the STAs from the rising and falling phases for our 115 neurons. There are several possible explanations for the variability in stimulus encoding (as measured by STA area) illustrated in Figure 4C. For example, the finite number of spikes from each cell may have introduced experimental noise that produced naturally variable STAs between the two phases of the neuronal oscillations.

To test whether the variance in the difference between the two STAs was greater than expected by chance, we randomly assigned spikes with 50% probability to the rising or falling phases, recalculated the STAs for the rising and falling phases, and found the difference between the two. The population was simulated 1000 times to obtain an estimate of the amount of variance one would expect by chance. The variance in the real population was 3.1 times greater than the variance one would expect by chance (two sided z-test, $p < 10^{-7}$).

Another possibility is that some phases of the neuronal oscillation encode the coherent motion more reliably than others. This second possibility is motivated by past studies that suggest the sensory information conveyed by a neuron can evolve over time (Lamme, 1995; Gawne et al., 1996; Friedrich and Laurent, 2001; Pack and Born, 2001; Reich et al., 2001; Hegde and Van Essen, 2004; Roelfsema et al., 2007). We addressed this second hypothesis more closely by using the behavioral data that was simultaneously collected with the neuronal recordings.

Past studies have demonstrated that neurons conveying the most informative task-related information tend to show the strongest correlations with behavior (Celebrini and Newsome, 1994; Britten et al., 1996; Cook and Maunsell, 2002b; Parker et al., 2002; Purushothaman and Bradley, 2005). We wanted to know if this hypothesis extended to the phases of our oscillatory activity. In other words, was the phase of the neuronal response that most reliably encoded the coherent motion also the phase that was most strongly correlated with behavior? This is an important question because it addresses the time-scale on which the relationship between neuronal activity and sensory perception can change.

To answer this question, we first describe how we correlated the phase of the neuronal activity with behavior using standard ROC-based choice analysis. In addition to the STA described above, we introduce a second method for quantifying the phase-dependent encoding of the coherent motion using an ROC-based neurometric analysis. Finally, we combine all these measurements to show that the phase of the neuronal

oscillation that most reliably encoded the coherent motion also tended to have the strongest correlation with behavior.

Neurometric value and detect probability as a function of phase

The STA analysis above was based only on neuronal responses to the 0% coherent motion. The monkey's task, however, was to detect the onset of a coherent motion step (Figure 1). Therefore, we wanted to know if the information in the neuronal response to the coherent motion step was phase-dependent. For this, we computed a neurometric value using ROC analysis (see Methods). The neurometric value was calculated using the neuronal response occurring during the 100 ms just before the onset of the coherent motion pulse compared with the neuronal response that occurred 50 to 150 ms after the coherent motion pulse. A neurometric value of 1.0 indicates an ideal observer could discriminate the response to the coherent motion from the response to the 0% motion with perfect reliability, while a value of 0.5 indicates chance performance.

Figure 2D shows the neurometric analysis applied to our example cell. We computed the neurometric value using spikes from either the rising (black bar) or falling phases (gray bar). For this cell, the neurometric value is higher for the rising phase compared to the falling phase. As a result, spikes on the rising phase were better able to signal the presence of coherent motion than spikes on the falling phase in agreement with the phase-dependent STA for this cell shown in Figure 2C.

It is important to emphasize, however, that the neurometric value can be dependent on the number of spikes used in the analysis. Thus, Figure 2D also shows that the neurometric value using all the spikes combined from both phases provides the most informative signal for detecting the onset of the coherent motion (white bar). This suggests that if the monkey were behaving as an ideal observer, combining the spikes from both phases would be a better strategy for detecting the onset of the coherent motion. The neurometric value, however, does not provide any insight into how the monkey actually used the activity from this cell to detect the coherent motion. To address this question, we next turn to the ROC-based detect probability calculation.

For our example cell, Figure 2E compares the average response (as a function of phase) for correct and failed trials from 50 to 150 ms after the onset of low-level coherent motion only. This particular cell demonstrated a phase dependent difference in response between correct (solid line) and failed trials (dashed line). Spiking activity associated with correct trials was greater than failed trials on the rising phase and became less than failed trials on the falling phase.

To quantify how the neuronal response was correlated with the animal's behavior as a function of phase, we computed detect probability (DP, see Methods). DP measures the ability to predict the behavioral outcome of a trial given the neuronal response. We calculated DP using the neuronal responses from 50 to 150 ms after coherent motion onset and only used responses from the low coherence trials (which produced

approximately 50% correct detections). DP is an ROC calculation that is analogous to the choice probability measure commonly used to express the correlations between neuronal activity and behavioral choice in discrimination tasks (Celebrini and Newsome, 1994; Britten et al., 1996; Parker and Newsome, 1998; Uka and DeAngelis, 2004). A DP value of 1.0 indicates that an ideal observer could predict the behavioral outcome of a trial (either correct or failed) using the neuronal response. A DP of 0.5 indicates chance predictability.

For our example cell, we found DP differed between the rising and falling phase of the neuronal response (Figure 2F). The DP analysis shows that activity of the rising phase is well correlated with behavior (black bar, DP = 0.64) and is even stronger than the DP from all spikes (white bar, DP = 0.57). On the falling phase (gray bar), activity was negatively correlated with behavior (greater activity was associated with failed instead of correct trials) producing a DP of 0.36. Thus, for this particular cell, the correlation between the neuronal response and behavior was asymmetric across neuronal phase.

Population phase-dependence of motion encoding, signal detection and behavior

For our example cell, the STA, neurometric and DP calculations suggest that there was a phase asymmetry in both how the coherent motion was encoded and the correlation with the animal's behavior. We next wanted to know whether those relationships were observable in our population of neurons.

As discussed above, a large positive STA implies that the neuron is selective in its firing, and will more likely respond only to coherent motion in the preferred direction and speed. One would suspect that spikes that are more likely generated by preferred coherent motion would naturally be more predictive of whether the neuron is encoding 0% coherent or coherent motion in the preferred direction. Thus, in Figure 5A, we examined whether neurons with different STAs between the two phases would also have a corresponding difference between the neurometric values of the two phases. For our population of 115 neurons, the asymmetry between the STAs of the two phases is well correlated with asymmetry in the neurometric value ($r = 0.55$, $p<10^{-9}$). The filled symbols are neurons whose STA satisfied a minimum signal-to-noise criterion of 2.75 (discussed in the next section). Thus, the phase with the larger STA (computed from 0% coherent motion) also tended to be the phase that most reliably signaled the onset of the coherent motion.

It has generally been observed that the most informative sensory neurons also tend to show the best correlations with behavior. Thus, we wanted to know if the neuronal phase with the best neurometric value was also more correlated with behavior. In Figure 5B we show the relation between the asymmetry in the neurometric value between the two phases and the asymmetry in the DP. Consistent with this hypothesis, the phase of the cycle that is more capable of signaling the presence of coherent motion, whether it be the rising or falling phase, also tends to be more correlated with the behavior ($r = 0.32$, $p = 0.0005$). Thus, not only does the correlation between how a

Figure 5 Population phase-dependence of motion encoding, signal detection and behavior. (A) The relation between the percentage difference in the area of the STAs of the rising and falling phases (x-axis) and the difference between neurometric values for the rising and falling phases (y-axis) for our 115 neurons. The areas of the STA were computed in a 27 ms window centered at the peak of the STA from all spikes. The black dots represent neurons with STAs that satisfied the criterion of a SNR of at least 2.75 (see text for details). (B) The difference in neurometric value between the two phases (x-axis) versus the difference in the detect probability of the two phases (y-axis). (C) The relationship of the percentage difference in area between the STAs between and the two phases (x-axis) with the difference in the DP (y-axis).

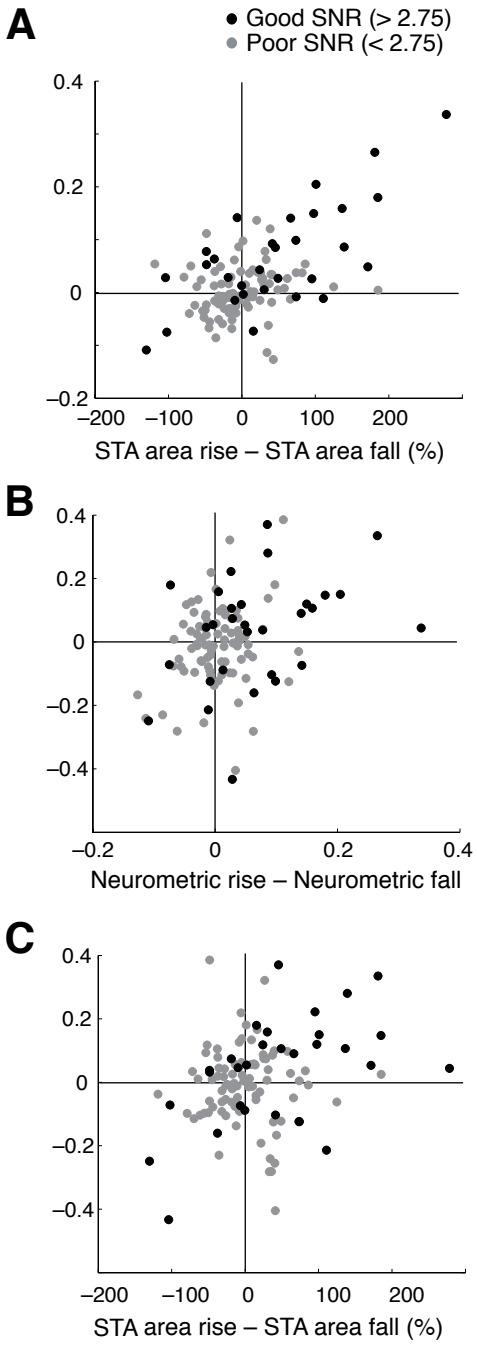
A
Good SNR (> 2.75)
Poor SNR (< 2.75)
0.4
0.2
0
–0.2
–200
–100
0
100
200
STA area rise – STA area fall (%)
B
0.4
0.2
0
–0.2
–0.4
–0.2
0
0.2
0.4
Neurometric rise – Neurometric fall
C
0.4
0.2
0
–0.2
–0.4
–200
–100
0
100
200
STA area rise – STA area fall (%)

Figure 5

stimulus is encoded (neurometric) and its relation to perceptual choice (DP) exist at the single cell level, but this relationship also seems to hold for certain spike times within a cell's response. We also examined the correlation between the phase-asymmetry in the STA with the phase-asymmetry in DP (Figure 5C). Consistent with the above results, these two asymmetries were also weakly correlated ($r = 0.25$, $p = 0.007$). Similar results were also observed for other time windows used to compute neurometric value and DP as a function of phase.

Taken together, this analysis suggests that the task-related information (as measured by STA and neurometric value) conveyed by the rising and falling phases of the neuronal response had a systematic relationship to the correlation between neuronal response and behavior (as measured by DP). A likely explanation of this result is that neuronal latencies varied depending on the strength of the pattern of moving dots in our stimulus. We examine this and other possibilities, such as possible biases induced by spike rate, below.

One limitation of the analysis in Figure 5 was that it was based on all neurons in our data set regardless of their tendency to oscillate. In addition, the STA estimates were inherently noisy for many neurons. The goal of our next analysis was to overcome these limitations by focusing on a subset of neurons with the most robust STAs and oscillations.

The strong- and weak-encoding phases

If one phase of the neuronal oscillation was more informative about the stimulus than the other, how much more correlated is that phase with behavior? To answer this, we grouped activity based on the size of the STA. We refer to the phase with the larger STA as the *strong-encoding* phase and the phase with the smaller STA the *weak-encoding* phase. We then examined how the strong- and weak-encoding phases encoded the motion stimulus and were related to the formation of a perceptual decision.

Since this calculation relies on an accurate measurement of the STA, we only included neurons whose STA satisfied a minimum signal-to-noise ratio (SNR). We based our SNR on a similar definition as used in a past analysis of this data (Cook and Maunsell, 2004). Our SNR was defined as the variance amongst the values of the STA in a 27 ms window centered at the peak (signal) divided by the variance amongst the values of the STA in the two 27 ms windows lying adjacent to this central window (noise). The idea was that the two adjacent windows represent the noise in our STA by capturing the chance correlations between the stimulus and neuronal response. We arbitrarily set the minimum SNR to 2.75, resulting in 30 neurons that satisfied the criteria. The STA of our example cell shown in Figure 2C had a robust SNR of 189.9. Figure 5 (filled symbols) illustrates that our 30 best cells with the most reliable STA spanned the range of our population with respect to asymmetries in the STA, neurometric and DP between the rising and falling phases.

In Figure 6A, we show the average STA from our 30 best neurons for the strong-encoding and weak-encoding phases. Because the strong-encoding phase was defined as the phase with the larger STA, the strong-encoding phase was naturally more selective for motion in the preferred direction compared to the weak-encoding phase. As there was a correlation between the SNR of the STA and oscillatory amplitude ($r = 0.55$, $p < 10^{-9}$), our best neurons also demonstrated stronger oscillations in response to our 0% coherent stimulus compared to the overall population (Figure 6B).

The prediction from Figure 6A is that the strong-encoding phase (i.e., phase with the largest STA) should also be the most informative phase about the onset of the coherent motion. Thus, we compared the neurometric values for the strong- and weak-encoding phases of our subset of neurons with high SNR. It is important to note that the strong /weak classification was based on the response to the 0% coherent motion, while the neurometric value was based on both the 0% coherent and coherent motion. The neurometric value for the strong-encoding phase (mean: 0.65 +/- 0.02 SEM, black bar) was significantly greater than the average value for the weak-encoding phase (mean: 0.59 +/- 0.01, gray bar) as shown in Figure 6C (pair wise difference two-sided t-test, $p = 0.009$). The histogram on the right shows a shift in the distribution of the pair-wise difference in the neurometric value between the strong- and weak-encoding phases. Thus, the strong-encoding phase tended to be more informative about the onset of the coherent motion than the weak-encoding phase. Since activity on both phases (usually) increased during coherent motion, both provided information about the onset of the

Figure 6 The strong-encoding and weak-encoding phases. (A) The average STA for the strong-encoding (black) and weak-encoding (gray) phase from the group of 30 neurons whose STA had a SNR above 2.75 (filled symbols in Figure 5). The phase with the larger STA by area was designated the encoding phase. (B) The average spike rate as a function of phase for the 30 neurons. The oscillatory amplitude of this group was larger than the oscillatory amplitude of the entire population (Figure 3D). (C) The average neurometric of these 30 neurons using spikes from the strong-encoding phase (black), weak-encoding phase (gray) and the entire cycle (white). The associated histogram to the right shows the pair-wise differences between the strong-encoding and weak-encoding values for each neuron. The vertical arrow indicates the median value. (C) Same as (B), but showing the average DP. (D) Same as (B), but showing the average RT Pearson's correlation coefficient.

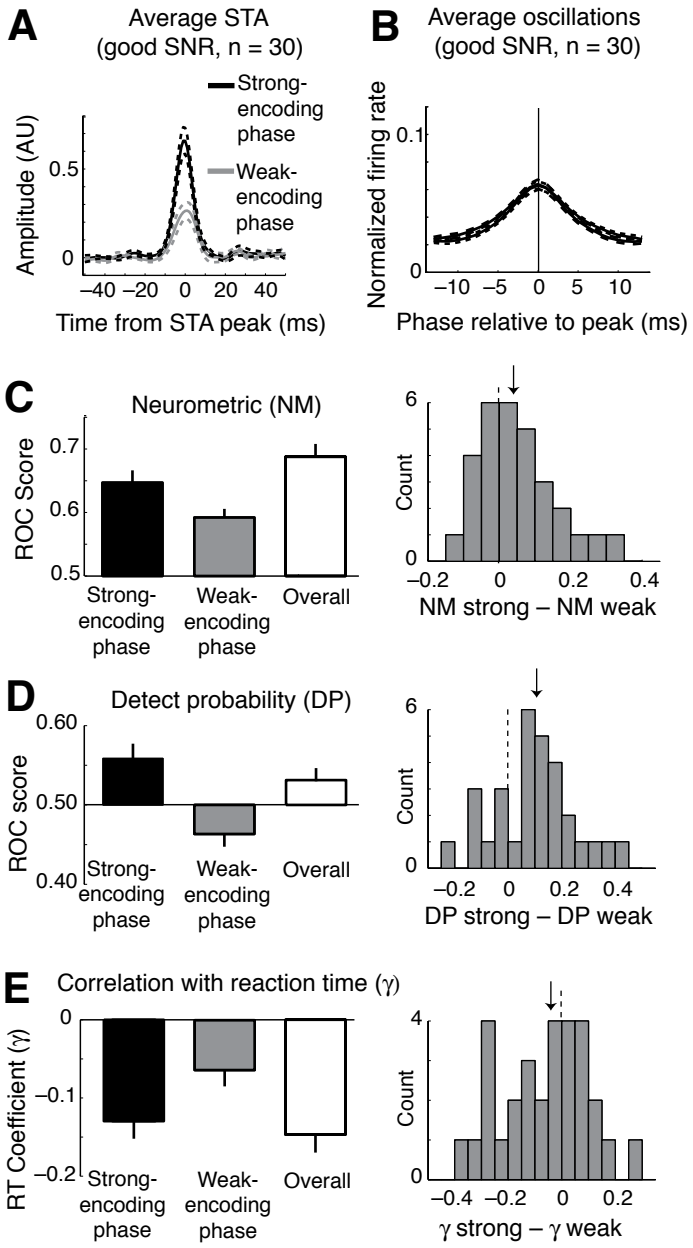
A
Average STA
(good SNR, n = 30)
Strong-encoding phase
Weak-encoding phase
Amplitude (AU)
0.5
0
–40 –20 0 20 40
Time from STA peak (ms)
B
Average oscillations
(good SNR, n = 30)
Normalized firing rate
0.1
0
–10 –5 0 5 10
Phase relative to peak (ms)
C
Neurometric (NM)
ROC Score
0.7
0.6
0.5
Strong-encoding phase
Weak-encoding phase
Overall
Count
6
0
–0.2 0 0.2 0.4
NM strong – NM weak
D
Detect probability (DP)
ROC score
0.60
0.50
0.40
Strong-encoding phase
Weak-encoding phase
Overall
Count
6
0
–0.2 0 0.2 0.4
DP strong – DP weak
E
Correlation with reaction time (γ)
RT Coefficient (γ)
0
–0.1
–0.2
Strong-encoding phase
Weak-encoding phase
Overall
Count
4
0
–0.4 –0.2 0 0.2
γ strong – γ weak

Figure 6

coherent motion. Thus, using all the spikes provided a slight increase in the neurometric value (white bar, Figure 6C).

We next examined how neural activity was correlated with perceptual choice for the strong- and weak-encoding phases. Figure 6D shows that the detect probability (DP) for the strong-encoding phase (0.56 +/- 0.02, black bar) was marginally greater than the DP from all spikes (0.53 +/- 0.02, white bar, pair wise difference, two-sided t-test, $p = 0.082$). On the other hand, the DP for the weak-encoding phase was significantly below chance at 0.46 +/- 0.02. The associated histogram reveals that neural activity on the strong-encoding phase was more positively correlated with perceptual choice than the weak-encoding phase in 22 of the 30 neurons, and the mean difference between the two DPs was significantly greater than zero (pair wise difference two-sided t-test, $p = 0.0018$). Thus, spikes occurring during the weak-encoding phase for our subset of 30 neurons were significantly and positively correlated with the animal failing to detect the motion stimulus. Spikes in the strong-encoding phase, by comparison, were positively correlated with the animal correctly detecting the motion. This raises the possibility that the brain could combine activity from both phases by weighting them with positive and negative weights, respectively. However, since activity is correlated between subsequent phases (see below), this might not be advantageous.

To further confirm that the strong-encoding and weak-encoding phases had different relationships with behavior, we also examined the (Pearson's) correlation between the neuronal response and reaction time (γ, see Methods). In this analysis, a

negative value indicates that greater activity is correlated with faster reaction times. In Figure 6E, we show that activity on the strong-encoding phase had an average correlation coefficient (–0.130 +/- 0.023) that nearly matched the correlation coefficient using all spikes (–0.147 +/- 0.023). However, activity on the weak-encoding phase was less correlated with reaction time (Womelsdorf et al., 2006)(–0.064 +/- 0.021). The associated histogram illustrates that the RT coefficient for the strong-encoding phase is greater for a majority of the neurons and the mean difference of –0.066 was significantly less than zero (pair wise difference, two-sided t-test, $p = 0.033$). Since neither the spike count distributions nor the reaction time distributions were not normally distributed, we repeated our calculations using the Spearman's rank correlation coefficient. The pair wise difference between the correlation coefficient of the strong-encoding and weak-encoding phases was –0.078, also significantly less than zero (pair wise difference, two-sided t-test, $p = 0.024$).

Since it has been previously reported that gamma band synchronization between spiking and activity and the LFP prior to a stimulus change in a detection task is correlated with reaction times (Womelsdorf et al., 2006), we also calculated the Pearson's correlation coefficient between neural activity and reaction time from 100 to 0 ms prior to coherent motion onset. Although there was a small difference between the correlation on the strong-encoding and weak-encoding phases (strong: –0.022 +/- 0.015; weak: 0.020 +/- 0.022; overall: –0.008 +/- 0.021), the differences were not significant.

By analyzing our 30 best neurons based on the quality of their STA, we found a phase-dependent relationship between the encoding of the motion stimulus and how the neuronal activity was correlated with behavior. We next examine several potential mechanisms that may account for the observed phase-dependent encoding of motion information.

Potential mechanisms

One possibility is that the mean or variance in the number of spikes varied between phases, biasing our results. The correlation between the asymmetry of the neurometric and the asymmetry in the STA (Figure 5A) is because there was a greater increase in firing rate on the strong-encoding phase. For low-coherent motion (used for the DP analysis), the average spike count was 13% greater (pair wise mean) and the standard deviation was 10% greater for the strong-encoding phase than the weak-encoding phase. To control for the difference in spike rate between phases, we calculated the difference in spike rate between the strong-encoding and weak-encoding phase during low coherent motion for each neurons. Based on this difference, we randomly removed spikes from the phase with greater activity in order to equalize the mean activity between the two phases. Each trial was simulated 100 times, and then we recalculated the correlations between the changes DP and the STA and NM. These new values are consistent with our previous results (change NM versus change DP: $r = 0.30$, $p = 0.001$; change STA versus change DP: $r = 0.24$, $p = 0.010$). We also performed the same calculation equalizing total spike count between the two phases, instead of spike rate.

Again, the results were left almost unchanged (change NM versus change DP: $r = 0.32$, $p = 0.0006$; change STA versus change DP: $r = 0.24$, $p = 0.009$). Since equalizing the spike rate and the spike count between the two phases when calculating the DP does not alter its significant correlation with the NM and the STA, we do not believe that the number of spikes was a major factor in the asymmetric relationship between the phase of stimulus encoding and behavior.

Another possible reason for the observed asymmetry in the DP between phases could be due to stimulus variability. More precisely, if there was more coherent motion for correct trials than for failed trials, this could have produced more spikes during the strong-encoding phase as well increased the chance the animal detects the coherent motion. If this was the case, then the percentage difference between the amount of coherent motion in the preferred direction between correct and failed trials should be correlated with the difference in the DP between the strong-encoding and weak-encoding phases. However, we found no such correlation for the entire population of 115 neurons ($r = –0.10$, $p = 0.30$) or for the population of 30 neurons with high SNR STAs ($r = 0.05$, $p = 0.80$). Thus, we do not believe that stimulus variability was the reason for the difference in the DP between the strong-encoding and weak-encoding phases.

A third possible explanation for our finding was that it was the result of correlations between the neural activity in each rising phase and the subsequent falling phase. For example, suppose activity on the rising and falling phases were negatively correlated, so that greater activity on the rising phase tended to reduce the amount of

activity on the subsequent falling phase. If the onset of coherent motion first increased neural activity on the rising phase, then activity on the falling phase would be suppressed, shifting the average phase of activity to the left. This possibility is somewhat similar to previous studies on how short-term plasticity can induce phase shifts in neural activity (Fortune and Rose, 2002; Fortune, 2006).

To explore this possibility, we computed the correlation coefficient between the numbers of spikes in each rising phase with the subsequent falling phase (similar results were obtained using the spike count in each falling phase and the subsequent rising phase) during 0% coherent motion. Although the mean correlation coefficient for all 115 neurons was quite weak ($r = 0.039$), most neurons possessed a significant ($p < 0.05$) correlation coefficient, with 39 having a negative correlation coefficient (32 significant) and 76 having a positive correlation coefficient (66 significant).

This correlation between the activity in each rising and subsequent falling phases was associated with a phase shift in the neural activity towards the encoding phase when coherent motion occurred. This was revealed by the significant correlation ($r = –0.20$, $p = 0.030$) between the correlation coefficient between activity on the rising and falling phases and the change in the neurometric value between the strong-encoding and weak-encoding phases. This is fairly intuitive since the encoding phase was generally the phase with the shorter latency (see Discussion), and likely to be first affected by the onset of coherent motion. For a neuron with negative correlation between activity on the rising and falling phases, coherent motion would increase activity disproportionably more on

the encoding phase, resulting in a greater neurometric value on the strong-encoding phase compared to the weak-encoding phase. Although the correlation coefficient between rising and falling phases was correlated to the change in neurometric values between the strong-encoding and weak-encoding phases, no significant correlation existed with the change in DP between the two phases ($r = –0.02$, $p = 0.82$). Thus, while the correlation between the activity on the rising and falling phases might shift the neural activity, it does not appear to be the source of the asymmetry in the correlation with behavior. However, we cannot say what role it might play, if any, in encoding the motion stimulus.

Although we cannot rule out the existence of other possible explanations for our observations (see Discussion), our result suggests that the spikes that convey more information about the presence of coherent motion are weighted more heavily in forming a perceptual decision.

Discussion

In this study, the activity of many MT neurons in response to a random moving dot stimulus oscillated with the same frequency as the stimulus updates. This allowed us to investigate how neural activity at different phases was correlated with both the stimulus and the behavior. We found that neuronal phases that were more informative about the stimulus, both in terms of the STA and the neurometric value, were also more correlated with the animal's detection performance and reaction time. Importantly, the asymmetry between stimulus encoding and behavior occurred on a relatively short time scale of a few tens of milliseconds.

Encoding of the motion stimulus

Oscillations seem to serve two broad, and possibly nonexclusive roles, in stimulus encoding. In the first, oscillations can synchronize ensembles of neurons to fire together (Gray et al., 1989; Singer, 1999; Usrey and Reid, 1999; Perez-Orive et al., 2002; Friedrich et al., 2004). Active conductances in dendrites can then sum these synchronously arriving spikes in a superlinear a manner, a term dubbed "coincidence detection" (Softky, 1994; Wang et al., 2000; Stuart and Hausser, 2001; Schaefer et al., 2003). These oscillations in the local field potential thus act as a timing device, and since these oscillations can remain coherent across large distances (Engel et al., 1990; Donoghue et al., 1998; Destexhe et al., 1999), they are capable of synchronizing neural activity across different areas in the brain (Engel et al., 1991; Destexhe et al., 1999). This

mechanism is widely observed through out the brain, and thought to play a role in a wide range of systems (for reviews, see (Engel et al., 2001; Averbeck and Lee, 2004; Sejnowski and Paulsen, 2006).

In the second role, the stimulus is encoded by the phase at which an action potential is generated. This is the encoding scheme used by place cells in the hippocampus, where the phase of the spike encodes the relative location of the animal inside the neuron's receptive field (O'Keefe and Recce, 1993). There are two important properties for such a encoding scheme. First, in order for the phase of the spike to be meaningful downstream, oscillations must be coherent between the upstream and downstream areas. Second, incoming EPSPs must have different effects on an oscillating target cell depending on the phase of its arrival (Lampl and Yarom, 1993; Schaefer et al., 2006).

Unlike the encoding by synchronization scheme mentioned above, a phase encoding scheme as only been observed in the hippocampus. However, it has been proposed that a phase encoding scheme might be more prevalent, and serve as a more general encoding scheme through out the cortex (Hopfield, 1995; Fries et al., 2007). Although the oscillations we observed in our data were specific to the stimulus used, the results presented here suggest that a cortical phase encoding scheme may be possible.

What are the mechanisms behind our observed phase encoding? One possibility is that spikes at phases that encoded stronger coherent motion were the same spikes

generated with a short latency. It has been shown that latencies are shortest in response to preferred stimuli (Bolz et al., 1982; Celebrini et al., 1993; Lagae et al., 1994; Gawne et al., 1996; Lisberger and Movshon, 1999; Maunsell et al., 1999; Raiguel et al., 1999; Johansson and Birznieks, 2004). If this is the case, then there should be a temporal offset between the STAs of the two different phases, with the peak of the larger STA occurring before the peak of the smaller STA. To verify this, we took the difference between the time of the peak of the STA from the rising phase minus the time of the peak from the falling phase. We then correlated the sign of this temporal difference with the sign of the difference of the area of the rising STA minus the area of the falling STA. There was significant negative correlation between the two sets ($r = -0.254$, $p = 0.0061$), indicating that the STA with a peak that occurred first was, on average, the STA with the larger area. This suggests that the variability in encoding between the phases was due, in part, to spikes produced with different latencies.

It has been proposed that the visual cortex may encode aspects of the visual stimulus through its response latency. However, these past studies have usually considered latency as the first spikes produced by a step-change in a visual stimulus. Thus, a latency code usually only refers to the initial response of the neuron to a stimulus. Oscillations, by comparison, provide a reference point by which latency of a spike can be determined from the phase in the presence of ongoing activity. The idea that oscillations could provide a reference allowing for a continuous latency encoding scheme has also been recently proposed by Fries et al. (2007), and it bears many similarities to other studies proposing that oscillations may serve to convert signal strength into a phase code

(Hopfield, 1995; Thorpe et al., 2001; Mehta et al., 2002). Although the phase encoding observed in our study is the result of the long delay between motion updates, it is of interest whether oscillations of the local field potential in visual cortex would allow for this encoding scheme to work under more general conditions, as observed in the hippocampus (O'Keefe and Recce, 1993). Alternatively, the relative phase difference in firing between neurons with different stimulus preferences might allow for this phase encoding scheme to function without the need for coherent oscillations between MT and downstream areas (Konig et al., 1995; Thorpe et al., 2001). The idea is that during coherent motion, MT neurons tuned to the direction of the coherent motion will fire in advance of neurons tuned to other directions, with the biggest phase difference for neurons tuned to the opposite direction. If downstream neurons could detect the temporal differences between incoming EPSPs arriving from differently tuned neurons (as in the auditory cortx ((Carr and Konishi, 1990), then this might eliminate the need for coherent oscillations to exist between MT and downstream.

Relation to the behavioral response

It is well documented that the activity of a single sensory neuron can covary with perceptual report (Celebrini and Newsome, 1994; Britten et al., 1996; Dodd et al., 2001; Cook and Maunsell, 2002b; Parker et al., 2002; Uka and DeAngelis, 2004; de Lafuente and Romo, 2005; Nienborg and Cumming, 2006). Additionally, it has also been shown in several visual detection/discrimination tasks that the neurons that most reliably encode the stimulus are also the neurons most correlated with the behavior (Celebrini and

Newsome, 1994; Britten et al., 1996; Cook and Maunsell, 2002b; Parker et al., 2002; Purushothaman and Bradley, 2005). This had led to the hypothesis that perceptual decisions depend more on a small subset of highly influential neurons that carry the most reliable sensory information.

In our study, certain spike times encoded the coherent motion more reliably than others. By comparing the difference in encoding between phases with the difference in the correlation with behavior between phases, we determined that the phase most informative about the motion stimulus was also more (positively) correlated with the perceptual behavior of the animal. This relation was present as early as 50 to 150 ms after the coherent motion onset and is unlikely that top-down attentional modulation of MT contributed to this result.

There does exist evidence linking synchronous neural activity, as described in the first part of the Discussion, to behavioral performance in a discrimination task (Stopfer et al., 1997; Womelsdorf et al., 2006). However, to our knowledge, no one has shown phase-dependent neural activity to be correlated with behavior.

Implications for decision models

When we restricted our analysis to the 30 neurons with the best-defined STAs, we found that the difference between how well the two phases are correlated with behavior is large. In fact, our DP analysis showed that the activity on the phase with the smaller STA

is actually correlated with the animal's failure to detect the coherent motion. This result is interesting because current models on how simple decisions are formed in the brain propose that evidence is accumulated towards one decision or another by integrating the spiking activity from sensory areas (Kim and Shadlen, 1999; Mazurek et al., 2003; Huk and Shadlen, 2005). In these models, spike times are irrelevant and the spike rate is the determining factor in decision-making. This implies that activity at any phase of the cycle should be positively correlated with the animal's detection performance. It is unclear how our observation that activity during half the cycle of the oscillation can be correlated with the animal failing the trial is consistent with these integrator models. Further studies will be required to determine how oscillating sensory activity is temporally processed to form decisions.

REFERENCES

Averbeck BB, Lee D (2004) Coding and transmission of information by neural ensembles. Trends Neurosci 27:225-230.

Bair W, Cavanaugh JR, Smith MA, Movshon JA (2002) The timing of response onset and offset in macaque visual neurons. J Neurosci 22:3189-3205.

Barlow H, Tripathy SP (1997) Correspondence noise and signal pooling in the detection of coherent visual motion. J Neurosci 17:7954-7966.

Bolz J, Rosner G, Wassle H (1982) Response latency of brisk-sustained (X) and brisk-transient (Y) cells in the cat retina. J Physiol 328:171-190.

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis Neurosci 13:87-100.

Carr CE, Konishi M (1990) A circuit for detection of interaural time differences in the brain stem of the barn owl. J Neurosci 10:3227-3246.

Celebrini S, Newsome WT (1994) Neuronal and psychophysical sensitivity to motion signals in extrastriate area MST of the macaque monkey. J Neurosci 14:4109-4124.

Celebrini S, Thorpe S, Trotter Y, Imbert M (1993) Dynamics of orientation coding in area V1 of the awake primate. Vis Neurosci 10:811-825.

Chacron MJ, Lindner B, Longtin A (2004) Noise shaping by interval correlations increases information transfer. Phys Rev Lett 92:080601.

Chacron MJ, Maler L, Bastian J (2005) Electroreceptor neuron dynamics shape information transmission. Nat Neurosci 8:673-678.

Churchland MM, Lisberger SG (2001) Shifts in the population response in the middle temporal visual area parallel perceptual and motor illusions produced by apparent motion. J Neurosci 21:9387-9402.

Cook EP, Maunsell JH (2002a) Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. J Neurosci 22:1994-2004.

Cook EP, Maunsell JH (2002b) Dynamics of neuronal responses in macaque MT and VIP during motion detection. Nat Neurosci 5:985-994.

Cook EP, Maunsell JH (2004) Attentional modulation of motion integration of individual neurons in the middle temporal visual area. J Neurosci 24:7964-7977.

de Lafuente V, Romo R (2005) Neuronal correlates of subjective sensory experience. Nat Neurosci 8:1698-1703.

de Ruyter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W (1997) Reproducibility and variability in neural spike trains. Science 275:1805-1808.

deCharms RC, Merzenich MM (1996) Primary cortical representation of sounds by the coordination of action-potential timing. Nature 381:610-613.

Destexhe A, Contreras D, Steriade M (1999) Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. J Neurosci 19:4595-4608.

Dodd JV, Krug K, Cumming BG, Parker AJ (2001) Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. J Neurosci 21:4809-4821.

Donoghue JP, Sanes JN, Hatsopoulos NG, Gaal G (1998) Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. J Neurophysiol 79:159-173.

Engel AK, Fries P, Singer W (2001) Dynamic predictions: oscillations and synchrony in top-down processing. Nat Rev Neurosci 2:704-716.

Engel AK, Konig P, Gray CM, Singer W (1990) Stimulus-Dependent Neuronal Oscillations in Cat Visual Cortex: Inter-Columnar Interaction as Determined by Cross-Correlation Analysis. Eur J Neurosci 2:588-606.

Engel AK, Kreiter AK, Konig P, Singer W (1991) Synchronization of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. Proc Natl Acad Sci U S A 88:6048-6052.

Fortune ES (2006) The decoding of electrosensory systems. Curr Opin Neurobiol 16:474-480.

Fortune ES, Rose GJ (2002) Roles for short-term synaptic plasticity in behavior. J Physiol Paris 96:539-545.

Friedrich RW, Laurent G (2001) Dynamic optimization of odor representations by slow temporal patterning of mitral cell activity. Science 291:889-894.

Friedrich RW, Habermann CJ, Laurent G (2004) Multiplexing using synchrony in the zebrafish olfactory bulb. Nat Neurosci 7:862-871.

Fries P, Nikolic D, Singer W (2007) The gamma cycle. Trends Neurosci 30:309-316.

Fries P, Neuenschwander S, Engel AK, Goebel R, Singer W (2001) Rapid feature selective neuronal synchronization through correlated latency shifting. Nat Neurosci 4:194-200.

Gawne TJ, Kjaer TW, Richmond BJ (1996) Latency: another potential code for feature binding in striate cortex. J Neurophysiol 76:1356-1360.

Gibson JR, Maunsell JH (1997) Sensory modality specificity of neural activity related to memory in visual cortex. J Neurophysiol 78:1263-1275.

Gray CM, Konig P, Engel AK, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. Nature 338:334-337.

Green DM, Swets, J.A. (1975) Signal Detection Theory and Psychophysics. Huntington, New York: Robert E. Krieger Publishing Co.

Gutig R, Sompolinsky H (2006) The tempotron: a neuron that learns spike timing-based decisions. Nat Neurosci 9:420-428.

Guyonneau R, VanRullen R, Thorpe SJ (2005) Neurons tune to the earliest spikes through STDP. Neural Comput 17:859-879.

Hegde J, Van Essen DC (2004) Temporal dynamics of shape analysis in macaque visual area V2. J Neurophysiol 92:3030-3042.

Hopfield JJ (1995) Pattern recognition computation using action potential timing for stimulus representation. Nature 376:33-36.

Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J Neurosci 25:10420-10436.

Johansson RS, Birznieks I (2004) First spikes in ensembles of human tactile afferents code complex spatial fingertip events. Nat Neurosci 7:170-177.

Judge SJ, Richmond BJ, Chu FC (1980) Implantation of magnetic search coils for measurement of eye position: an improved method. Vision Res 20:535-538.

Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat Neurosci 2:176-185.

Konig P, Engel AK, Roelfsema PR, Singer W (1995) How precise is neuronal synchronization? Neural Comput 7:469-485.

Lagae L, Maes H, Raiguel S, Xiao DK, Orban GA (1994) Responses of macaque STS neurons to optic flow components: a comparison of areas MT and MST. J Neurophysiol 71:1597-1626.

Lamme VA (1995) The neurophysiology of figure-ground segregation in primary visual cortex. J Neurosci 15:1605-1615.

Lampl I, Yarom Y (1993) Subthreshold oscillations of the membrane potential: a functional synchronizing and timing device. J Neurophysiol 70:2181-2186.

Lisberger SG, Movshon JA (1999) Visual motion analysis for pursuit eye movements in area MT of macaque monkeys. J Neurosci 19:2224-2246.

Lu T, Liang L, Wang X (2001) Temporal and rate representations of time-varying signals in the auditory cortex of awake primates. Nat Neurosci 4:1131-1138.

Maunsell JH, Ghose GM, Assad JA, McAdams CJ, Boudreau CE, Noerager BD (1999) Visual response latencies of magnocellular and parvocellular LGN neurons in macaque monkeys. Vis Neurosci 16:1-14.

Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. Cereb Cortex 13:1257-1269.

Mechler F, Victor JD, Purpura KP, Shapley R (1998) Robust temporal coding of contrast by V1 neurons for transient but not for steady-state stimuli. J Neurosci 18:6583-6598.

Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. Nature 417:741-746.

Nienborg H, Cumming BG (2006) Macaque V2 neurons, but not V1 neurons, show choice-related activity. J Neurosci 26:9567-9578.

Nover H, Anderson CH, DeAngelis GC (2005) A logarithmic, scale-invariant representation of speed in macaque middle temporal area accounts for speed discrimination performance. J Neurosci 25:10049-10060.

O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus 3:317-330.

Pack CC, Born RT (2001) Temporal dynamics of a neural solution to the aperture problem in visual area MT of macaque brain. Nature 409:1040-1042.

Paninski L, Pillow JW, Simoncelli EP (2004) Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model. Neural Comput 16:2533-2561.

Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. Annu Rev Neurosci 21:227-277.

Parker AJ, Krug K, Cumming BG (2002) Neuronal activity and its links with the perception of multi-stable figures. Philos Trans R Soc Lond B Biol Sci 357:1053-1062.

Perez-Orive J, Mazor O, Turner GC, Cassenaer S, Wilson RI, Laurent G (2002) Oscillations and sparsening of odor representations in the mushroom body. Science 297:359-365.

Purushothaman G, Bradley DC (2005) Neural population code for fine perceptual decisions in area MT. Nat Neurosci 8:99-106.

Raiguel SE, Xiao DK, Marcar VL, Orban GA (1999) Response latency of macaque area MT/V5 neurons and its relationship to stimulus parameters. J Neurophysiol 82:1944-1956.

Reich DS, Mechler F, Victor JD (2001) Temporal coding of contrast in primary visual cortex: when, what, and why. J Neurophysiol 85:1039-1050.

Robinson DA (1963) A Method Of Measuring Eye Movement Using A Scleral Search Coil In A Magnetic Field. IEEE Trans Biomed Eng 10:137-145.

Roelfsema PR, Tolboom M, Khayat PS (2007) Different Processing Phases for Features, Figures, and Selective Attention in the Primary Visual Cortex. Neuron 56:785-796.

Romo R, Salinas E (2001) Touch and go: decision-making mechanisms in somatosensation. In: Annu Rev Neurosci, pp 107-137.

Sadeghi SG, Chacron MJ, Taylor MC, Cullen KE (2007) Neural variability, detection thresholds, and information transmission in the vestibular system. J Neurosci 27:771-781.

Schaefer AT, Larkum ME, Sakmann B, Roth A (2003) Coincidence detection in pyramidal neurons is tuned by their dendritic branching pattern. J Neurophysiol 89:3143-3154.

Schaefer AT, Angelo K, Spors H, Margrie TW (2006) Neuronal oscillations enhance stimulus discrimination by ensuring action potential precision. PLoS Biol 4:e163.

Sejnowski TJ, Paulsen O (2006) Network oscillations: emerging computational principles. J Neurosci 26:1673-1676.

Shapley R, Hawken M, Ringach DL (2003) Dynamics of orientation selectivity in the primary visual cortex and the importance of cortical inhibition. Neuron 38:689-699.

Singer W (1999) Neuronal synchrony: a versatile code for the definition of relations? Neuron 24:49-65, 111-125.

Softky W (1994) Sub-millisecond coincidence detection in active dendritic trees. Neuroscience 58:13-41.

Stopfer M, Bhagavan S, Smith BH, Laurent G (1997) Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. Nature 390:70-74.

Stuart GJ, Hausser M (2001) Dendritic coincidence detection of EPSPs and action potentials. Nat Neurosci 4:63-71.

Thorpe S, Delorme A, Van Rullen R (2001) Spike-based strategies for rapid processing. Neural Netw 14:715-725.

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

Usrey WM, Reid RC (1999) Synchronous activity in the visual system. Annu Rev Physiol 61:435-456.

Victor JD, Purpura KP (1996) Nature and precision of temporal coding in visual cortex: a metric-space analysis. J Neurophysiol 76:1310-1326.

Wang SS, Denk W, Hausser M (2000) Coincidence detection in single dendritic spines mediated by calcium release. Nat Neurosci 3:1266-1273.

Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) Gamma-band synchronization in visual cortex predicts speed of change detection. Nature 439:733-736.

Chapter 4

The behavioral impulse-response function of microstimulation in cortical area MT

Abstract

Electrical stimulation of the brain is a valuable research tool and has shown therapeutic promise in the development of new sensory neural prosthetics. In spite of its widespread usage, we still do not fully understand how current passed through a microelectrode interacts with functioning neural circuits. Past behavioral studies have suggested that electrical stimulation of sensory areas of cortex produces percepts that are similar to those generated by normal sensory stimuli. In contrast, electrophysiological studies using in vitro or anesthetized preparations have shown that neural activity produced by brief electrical stimulation is radically different than normal responses. To help reconcile these two aspects of electrical stimulation, we examined the temporal properties that electrical stimulation has on visual perception. We found that brief application of subthreshold electrical stimulation in the Middle Temporal (MT) area of visual cortex produced smaller and longer lasting effects on motion perception when compared to an equivalent visual stimulus.

Introduction

Electrical microstimulation of the brain is an important research tool for establishing causality between neural activity and behavior (for reviews, see Cohen and Newsome 2004; Romo and Salinas 1999) and serves as the basis for supplying sensory inputs in neural prosthetics (Bradley et al. 2005; Fernandez et al. 2005; Girvin 1988; McIntyre and Grill 2000; Middlebrooks et al. 2005; Normann et al. 1999; Tehovnik and Slocum 2007; Troyk et al. 2003). Both of these applications rely on the assumption that microstimulation can generate percepts that are reasonably similar to those produced by naturally occurring stimuli.

Many past behavioral studies suggest that microstimulation of cortical sensory areas is equivalent to natural inputs in its ability to influence sensory perception (Bisley et al. 2001; Carey et al. 2005; Celebrini and Newsome 1995; de Lafuente and Romo 2005; DeAngelis and Newsome 2004; Ditterich et al. 2003; Hanks et al. 2006; Liu and Newsome 2005; Murasugi et al. 1993; Nichols and Newsome 2002; Romo et al. 2000; Romo et al. 1998; Salzman et al. 1990; Salzman et al. 1992; Uka and DeAngelis 2006). In these experiments, however, microstimulation was usually applied for hundreds of milliseconds to seconds, which in many ways differs from the high temporal variability of natural sensory stimuli.

In contrast to these behavioral studies, neural activity measured in response to microstimulation in both in vitro and in vivo preparations has usually been characterized

as a short excitatory response followed by a long period of inhibition (Berman et al. 1991; Butovas et al. 2006; Butovas and Schwarz 2003; Chung and Ferster 1998; Contreras et al. 1997; Shao and Burkhalter 1996; 1999). Thus, microstimulation can produce a "temporal spread" of neural activity that is fundamentally different from normal neurophysiological responses. We do not know the extent to which the temporal spread of neural activity produced by microstimulation influences perception. If microstimulation is to be used to probe the causality between sensory activity and perception or supply sensory input for neural prosthetics, then it is important to understand how its temporal properties affect perception. In this study, we specifically ask whether microstimulation and visual inputs generate different percepts.

To explore these differences, we measured the time course (i.e., the impulse-response function) that brief microstimulation in area MT has on the perception of a motion stimulus. Importantly, we designed our experiment so that we could compare the temporal effect of microstimulation on behavior with that of a visual input of equivalent duration. We found the ability of brief microstimulation to influence the detection of the motion stimulus was weaker and decayed significantly more slowly compared to a control visual input. These results suggest that the temporal spread of microstimulation could place design challenges on its use in a cortical visual prosthesis and for probing the relationship between cortical activity and behavior on fast timescales.

Methods

Behavioral task

Two male monkeys (Macaca mulatta) were trained to detect coherent motion in a random dot patch. The trial began when the monkey depressed a lever and fixated on a central point. Eye movements were recorded with an eye camera system and the trial ended if the monkey's eye position deviated by more than 1.5 degrees from the fixation point. After the lever was depressed, 0% coherent motion began in a random dot patch. A 33ms coherent motion signal was presented to the monkey at a random time from 500 to 10000 ms after the start of the trial (flat hazard function). For some trials, 33 ms of subthreshold coherent motion (visual probe) or 33 ms of electrical microstimulation (microstimulation probe) was also presented at various times relative to the motion signal (Figure 1A&C). The monkey was rewarded if he released the lever from 150 to 650 ms after the end of the motion signal or after the end of the probe. The strength of the coherent motion signal was adjusted for each experimental session so that the monkey detected the motion signal about 50% of the time and the strength of both the visual and microstimulation probes were set so that they were detected about 5% of the time. The strength of the signal and the probes were adjusted before data collection began by running probe only and signal only trials. The possible time intervals between the visual probe and the visual motion signal (given by the start of the visual motion signal minus the end of the probe) were 0, 25, 50, 75, 100, 150, 250, 350, or 450 ms. For the

microstimulation probe, the possible intervals were -100, –75, –50, –25, 0, 50, 150, 250 or 350 ms. Negative values indicate the probe was presented after the signal. The length of the time between the visual motion signal and probe was denoted by τ (Figure 1C). Other trials consisted of only a visual motion signal or only a visual motion or microstimulation probe. In all, there were 21 possible trial types: 9 different microstimulation probe plus signal trials, 9 different visual motion probe plus signals trials, a microstimulation probe only trial, a visual probe only trial and a visual motion signal only trial. These 21 different trial types were presented to the subject in a random order with approximately equal proportions.

Visual stimulus

The monkey was positioned 62 cm from a computer monitor (34° × 26° of visual angle; 1600×1200 pixels; 120 Hz refresh). The stimuli consisted of a random dot patch of white dots (each dot 0.25° diameter) on a dark gray background with a dot density of 2.1 dots/degree^2. The location, radius, speed and direction of the random dot patch was matched to the preferences to the site in area MT that we were recording and stimulating for that experiment (Figure 1B). The preferred location and radius of the MT site were determined by manual mapping while the speed and direction were determined by recording multi-unit activity in response to 100% coherent motion in different directions and speeds. On every frame update, dots were assigned to either move in the preferred or null direction. The ratio of dots moving in the preferred and null direction determined the coherence level of the motion. Thus, for 0% coherent motion, an equal number of dots

moved in the preferred and null directions. All dots moved the same distance (preferred speed) each frame update. If a dot moved outside the receptive field, it was re-plotted in a random position in the opposite side of the receptive field. This stimulus constrained motion to the preferred/null axis. When 0% coherent motion was viewed, there was no net or transparent motion perceived.

Electrical microstimulation

All animal care and procedures followed guidelines set forth by McGill University's Animal Care Committee and the Canadian Council for Animal Care. The monkeys received head-posts and recording chambers under aseptic surgeries following standard techniques (Cook and Maunsell 2002). Structural MRI brain scans (1.5 T) and the neurophysiological properties of the recording sites confirmed our electrode placement in area MT.

Electrical microstimulation was delivered to an area of MT with low-impedance (250 to 1000KΩ at 1 KHz) tungsten microelectrodes using a constant-current biphasic stimulator (Bak Electronics, Mount Airy, MD). The microstimulation probe consisted of 8 biphasic pulses (200 Hz with 40 μs pulse width) over a time of 33ms. The current amplitude was set at the beginning of an experiment so that the animal detected the microstimulation probe alone approximately 5% of the time. The current values used ranged from 3 to 50 μA with a median of 12 μA.

Determining the impulse-response functions

The impulse-response functions were a measure of how well the monkey detected the signal plus the probe together compared to how well the monkey detected the signal as a function of τ. The impulse-response function is the best linear estimate of the effect of the probes on motion detection. Although we do not claim that the neural mechanisms of motion detection is a linear process, if the linear estimates of the effects of visual and microstimulation probes differ, than the underlying non-linear effects must also differ.

During each experimental session, the monkey was presented with trials consisting of either a visual motion signal only, a visual probe only, a microstimulation probe only, a visual probe plus a visual motion signal or a microstimulation probe plus a visual motion signal (Figure 1C). The probability of a correct detection in these five conditions was denoted as:

p_{signal}, p^{V}_{probe}, p^{M}_{probe}, p^{V}_{τ} and p^{M}_{τ}, respectively.

For any correct trial, there was always a chance that the monkey simply guessed correctly instead of actually perceiving coherent motion. To account for this, we calculated the rate at which the monkey false alarmed (released the lever before the coherent motion signal occurred) for each experimental session to first modify the probability of a correct detection:

$$p_{signal} \rightarrow p_{signal} - FA,$$
$$p^{V}_{probe} \rightarrow p^{V}_{probe} - FA,$$
$$p^{M}_{probe} \rightarrow p^{M}_{probe} - FA,$$
$$p^{V}_{\tau} \rightarrow p^{V}_{\tau} - FA(\tau),$$
$$p^{M}_{\tau} \rightarrow p^{M}_{\tau} - FA(\tau),$$

where FA is the probability that the monkey would false alarm in a 500ms interval, which was the length of the response window. For a trial with a probe plus visual motion signal, the monkey was rewarded if he released the lever after the probe or after the signal. Thus the response window was adjusted according to τ. No value was adjusted below zero, and the mean FA was 2.8%.

This modification was based on the assumption that a false alarm and correct signal detection are mutually exclusive events. Although this assumption is only an approximation, the modifications were small and were applied equally to both visual and microstimulation results, and the overall results were left unchanged if no adjustments for false alarm rates were performed.

Next, we defined two values that were the probabilities of correctly detecting a signal or a probe assuming the two were independent:

$$p^{V}_{signal \cup probe} = p_{signal} + p^{V}_{probe} - p_{signal} p^{V}_{probe}$$
$$p^{M}_{signal \cup probe} = p_{signal} + p^{M}_{probe} - p_{signal} p^{M}_{probe}.$$

Although detecting the probe and detecting the signal are most likely not independent events, the differences approach zero as the probability of a probe detection becomes small and the same modifications are performed to both visual and microstimulation trials.

Since a detection task results in a binomial distribution (correct/failed), we can calculate the means,

$$m^V_{signal|probe} = n_{signal} p^V_{signal|probe}$$
$$m^M_{signal|probe} = n_{signal} p^M_{signal|probe}$$
$$m^V_\tau = n^V_\tau p^V_\tau$$
$$m^M_\tau = n^M_\tau p^M_\tau ,$$

and the variances,

$$s^V_{signal|probe} = n_{signal} p^V_{signal|probe}(1 - p^V_{signal|probe})$$
$$s^M_{signal|probe} = n_{signal} p^M_{signal|probe}(1 - p^M_{signal|probe})$$
$$s^V_\tau = n^V_\tau p^V_\tau (1 - p^V_\tau)$$
$$s^M_\tau = n^M_\tau p^M_\tau (1 - p^M_\tau).$$

Here, the variable n refers to the total number of correct plus missed trials of each type.

Given the mean and variances, the visual and microstimulation impulse-response functions were defined using the t-statistic based on the Student's t-test:

$$IRF^{V}(\tau) = \frac{m_{\tau}^{V} - m_{signal|probe}^{V}}{\sqrt{\frac{(n_{\tau}^{V}-1)s_{\tau}^{V} + (n_{signal}-1)s_{signal|probe}^{V}}{n_{\tau}^{V} + n_{signal} - 2}(\frac{1}{n_{\tau}^{V}} + \frac{1}{n_{signal}})}},$$

$$IRF^{M}(\tau) = \frac{m_{\tau}^{M} - m_{signal|probe}^{M}}{\sqrt{\frac{(n_{\tau}^{M}-1)s_{\tau}^{M} + (n_{signal}^{M}-1)s_{signal|probe}^{M}}{n_{\tau}^{M} + n_{signal} - 2}(\frac{1}{n_{\tau}^{E}} + \frac{1}{n_{signal}})}},$$

where $IRF^{V}(\tau)$ and $IRF^{M}(\tau)$ denote the visual and microstimulation impulse-response functions, respectively.

Since there is a delay before visual motion signals reach area MT, we shifted the microstimulation impulse-response function to compensate for this delay. To compute the neural latency, we computed the average neural activity from multi-unit activity in response to only the visual motion signal. The neural response was smoothed with a Gaussian with a standard deviation of 4ms, and the latency was defined as the point in time where the average neural activity crossed 3 standard deviations above baseline activity after coherent motion onset. Since the neural response was smoothed with a non-causal filter, this slightly lowered our estimate of the latency.

Overall, the average latency was 53.6 ms with a standard deviation of 18.7 ms. To compensate for this latency, we shifted our microstimulation temporal effect 50 ms. Thus the original interval time between the microstimulation probe and the signal of –100, –75, –50, –25, 0, 50, 150, 250 and 350 ms were shifted to –50, –25, 0, 25, 50, 100, 200, 300

and 400 ms. We do not show the effect of the microstimulation probe at negative intervals, which corresponds to the visual motion signal reaching MT on average before the microstimulation probe. It should be noted that the effect of the microstimulation probe at negative intervals was excitatory but weaker than the effect of the probe at 0 ms.

Significance tests for comparing the two probes

The significance test for the strength of the probe for $\tau \leq 50$ was calculated using real, non-fitted data. For each experimental session, we averaged the effect of the probes at $\tau = 0$, 25 and 50. We then performed a two-sided t-test between these averages for the microstimulation and visual probes.

Bootstrap analysis was used to compute significance when comparing how quickly the effects of the probes decayed. It was performed on both fitted and un-fitted data. Given *n* experimental sessions in our population, we sampled at random *n* experimental sessions with replacement. We then calculated the decay rate in two different ways. In the first, we divided the effect of the probe averaged between $\tau = 50$ and 100 ms by the effect of probe averaged between $\tau = 0$ and 25 ms. In the second method, we then fit the average effect of the probes at all time points with a single exponential function and determined its time constant. If one or both of these fits were poor, then both fits were discarded and new experimental sessions were resampled. These calculations were performed simultaneously for both electrical and visual probes so that both were calculated using the same *n* experimental sessions. These steps were then

repeated 20000 times. The average values reported in the text refer to the average of the 20000 values calculated during the bootstrap. The p-value is the proportion of times that the microstimulation decay value was greater than the visual decay value calculated form the same sample set of *n* experimental sessions.

Principal component analysis

To demonstrate that the differences between the visual and microstimulation impulse-response functions were consistent on an experiment-by-experiment basis, we used principal component analysis. This was because individual experiments had considerable noise and thus could not be reliably fit by our single exponential function. Principal component analysis was used to reduce the noise by lowering the dimensionality of the data set.

Since the time intervals (τ) used for the visual and microstimulation temporal effects did not perfectly match, we first interpolated missing quantities by averaging adjacent values to obtain an identical set of time intervals for both sets of temporal effects. The time intervals used were 0, 25, 50, 75, 100, 150, 200, 250, and 300 ms. From the population averages, there did not appear to be much difference between visual and microstimulation probes after 300 ms. Next, we computed the covariance matrix of all microstimulation and visual temporal effects combined together, and found the two eigenvectors with the two largest associated eigenvalues. We reconstructed the temporal

effects (from 2 to 9 dimensions) by multiplying the two-dimensional (row) vector with the 2 by 9 matrix of the two principal eigenvector

Results

We first trained two monkeys (Rhesus macaque) to perform a simple motion detection task (Figure 1A). The visual stimulus was 0% coherent random dot motion and the monkey had to quickly release a lever after the occurrence of a brief (33 ms) coherent motion pulse. The selection of such a brief stimuli was important because it allowed us to easily determine how the motion information was integrated in time. The coherent motion (referred to as the *signal*) occurred at a random time and its strength was set to produce threshold performance of ~50% correct.

We applied microstimulation to area MT because it has been shown to be highly selective for coherent motion (Albright et al. 1984; Maunsell and Van Essen 1983; Van Essen et al. 1981; Zeki 1974) and its activity is linked to motion perception (Salzman et al. 1990). The stimulus location, motion speed and motion direction were matched to that preferred by the location of MT containing the microelectrode (Figure 1B). In this way, the microstimulated region of MT would contribute to the monkey's detection of the coherent signal (Bisley et al. 2001; Ditterich et al. 2003; Murasugi et al. 1993; Nichols and Newsome 2002; Salzman et al. 1992). To measure the temporal effect of microstimulation on behavior, we briefly microstimulated on some random trials for 33 ms (referred to as the *probe*) at various times relative to the main visual motion signal (Figure 1C, top). We denote the time between the probe and the signal as τ. This experimental design is similar to the "two-pulse interaction" approach developed to psychophysically measure temporal integration of visual stimuli in humans (Rashbass

Figure 1 Experimental design and example result. (A) Two monkeys were trained to detect a brief (33 ms) coherent motion signal in a moving random dot patch. The coherent motion occurred at a random time and the strength of the motion was set to produce threshold detection performance. The animals had from 150 to 650 ms to release a lever once the coherent motion occurred. (B) Microstimulation was applied to a site in area MT whose preferred direction, speed, size and location matched the motion stimulus. It was assumed that neural activity in area MT was linearly integrated to support the animal's detection of the coherent motion signal. (C) Some trials contained a probe that preceded the coherent motion signal. The probe was either 33 ms of subthreshold microstimulation or coherent motion. The separation between the signal and the probe is denoted by τ. Microstimulation consisted of eight 200 Hz biphasic 40 μs pulses with amplitudes that ranged from 3 to 50 μA. (D) Results from an example experiment. The animal's probability of detecting the motion signal as a function of τ is shown on the vertical axis. Visual probe (blue) and microstimulation probe (red). Dashed black line is the monkey's probability of detecting the coherent motion signal alone, and the dashed blue and red lines represent the probability of detecting the visual and microstimulation probes alone, respectively.

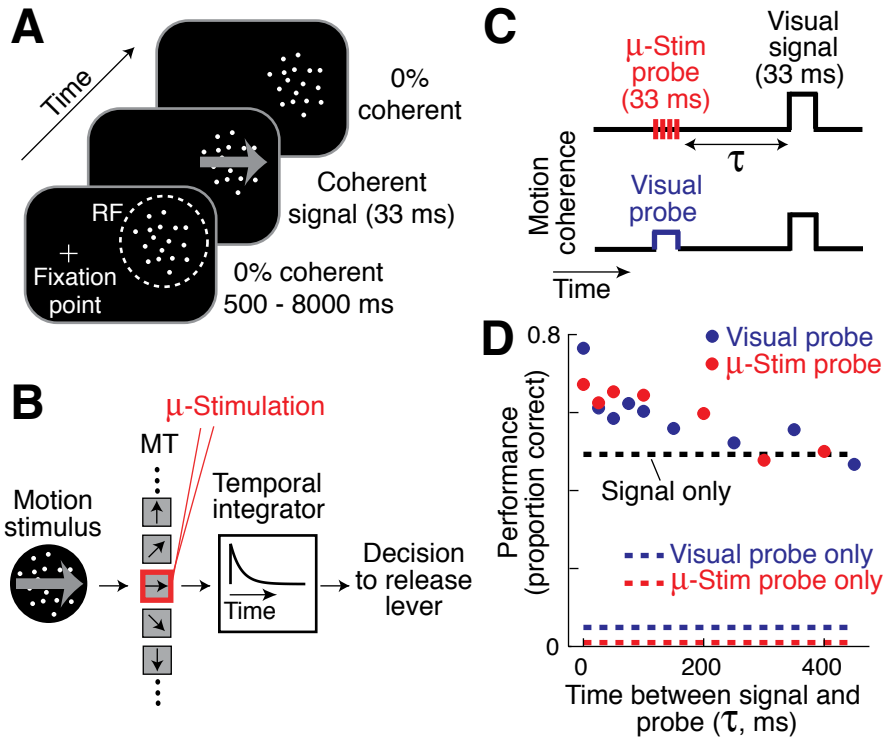
A
Time
0%
coherent
Coherent
signal (33 ms)
RF
Fixation
point
0% coherent
500 - 8000 ms
B
μ-Stimulation
MT
Motion
stimulus
Temporal
integrator
Time
Decision
to release
lever
C
μ-Stim
probe
(33 ms)
Visual
signal
(33 ms)
τ
Motion
coherence
Visual
probe
Time
D
0.8
Visual probe
μ-Stim probe
Performance
(proportion correct)
Signal only
Visual probe only
μ-Stim probe only
0
0
200
400
Time between signal and
probe (τ, ms)

Figure 1

1970; Simpson 1994). It is important to point out that the animals were rewarded for responding to either the visual motion signal or probe.

To determine how the effect of microstimulation differs from that of a normal visual input, on some trials we used a 33 ms probe of subthreshold visual coherent motion (Figure 1C, bottom). Comparing the animal's ability to detect the coherent motion signal as a function of τ for both types of probes allowed us to reveal the temporal differences between microstimulation and visual motion.

Trials containing visual and microstimulation probes were randomly interleaved and τ was shifted by 50 ms for the microstimulation probe trials to account for neural latency of MT. Although the neural latency varied between sites, our results were unaffected when we controlled for any differences in latency (see Supplementary Figures 1&2). By comparing the probability that the monkey detected the visual motion signal and probe for each τ to the probability the monkey detected the visual motion signal only, we could estimate the underlying time course of the effect that each type of probe had on motion perception.

Figure 1D shows the proportion of trials that the animal correctly detected the visual motion motion as a function of the separation between the probe and signal during a single experimental session. For this example, both microstimulation and visual probes presented just before the signal (small τ) increased the monkey's ability to detect the coherent motion signal compared to the trials where we presented only the signal. Probes further in time from the signal (large τ) had little or no effect on behavioral performance.

It was important that the strength of both the visual and microstimulation probe was the same. At the beginning of each experiment we adjusted both to be approximately equivalent by presenting trials consisting of only the probes. During the experiment we continued to monitor their strength by occasionally presenting probe only trials. Figure 2A shows the pairwise difference between the probability of detecting the visual and microstimulation probes. These probabilities have been modified to account for the monkey detecting the probe by chance (see Methods). Although there was an appreciable amount of variability, on average, the animals detected the visual and microstimulation probes with similar probability (visual: 5.2 ± 0.6%; microstimulation: 4.5 ± 0.7%; $p = 0.37$, all p-values are pairwise differences, two-sided t-tests). We do not believe that this variability was a factor in our results (see Supplementary Methods). The average reaction times (RT) for detecting the visual and electrical probe only trials were also similar after accounting for the 50 ms difference in latency (Figure 2B, visual: 390.3 ± 7.1 ms; microstimulation + 50 ms: 397.5 ± 10.0 ms; $p = 0.55$). Thus, the visual and microstimulation probes presented alone were weak (i.e., subthreshold) and produced relatively equivalent effects on behavior.

In order to combine results across experimental sessions, we normalized the effect of the visual and microstimulation probes (see Methods). We first accounted for the probability of detecting the probes (red and blue dashed lines in Figure 1D) using the probe only trials. We then converted this probe-detect adjusted data to a t-score which indicated how effective a probe was at influencing the detection performance compared

Figure 2 Visual and microstimulation probes were of equivalent strength. (A) The histogram of the pairwise difference between the animal's ability to detect the visual motion probe alone and the microstimulation probe alone. (B) The histogram of the pairwise difference between the animal's reaction time to respond to the visual motion probe alone and the microstimulation probe alone. The difference is shifted by 50 ms to account for the neural latency of MT.

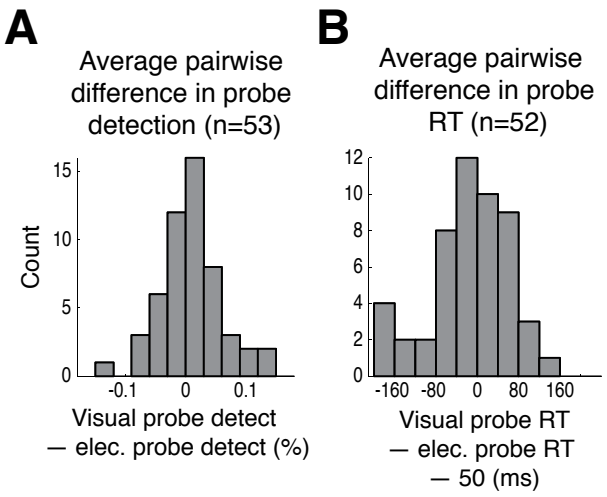
A
Average pairwise difference in probe detection (n=53)
Count
15
10
5
0
-0.1
0
0.1
Visual probe detect — elec. probe detect (%)
B
Average pairwise difference in probe RT (n=52)
12
10
8
6
4
2
0
-160
-80
0
80
160
Visual probe RT — elec. probe RT — 50 (ms)

Figure 2

to the visual motion signal only condition (black dashed line in Figure 1D). Figure 3A shows the normalized data for our example experiment and illustrates that both the visual and electrical probes initially enhanced the detection of the visual motion signal for small τ but then had no or a slight inhibitory effect on behavior for longer τ.

To compare the effects between visual and microstimulation probes, we averaged the normalized t-scores of the probes at each τ across all experiments (Figure 3B). We then averaged the effects of each type of probe to obtain its impulse-response function (IRF). The IRF is a linear approximation of the effect of the probe on motion perception across time. The IRF is by definition the output of a linear system given that the input is an impulse (a pulse of vanishingly small width in time). In our experiment, we approximated the impulse with a 33 ms probe, and the output of the system was the change in detecting the visual motion signal. We do not claim that motion perception is a linear system, but this method can still be used to obtain the best linear estimates of the temporal effects of the probes.

Although somewhat similar, the IRFs of the microstimulation and visual probes on motion perception were significantly different in two ways. Firstly, the effect of the microstimulation probe was on average weaker than the effect of the visual probe for short $\tau \leq 50$ ms ($p < 0.001$, two-sided t-test, see Methods). Secondly, the decay rate of the microstimulation IRF was significantly longer than the visual probe. We compared the decay rate of the IRFs using bootstrap analysis on both fitted and un-fitted data. When fitting the IRFs with single exponentials, the time constant for the microstimulation probe

Figure 3 The behavioral impulse-response function of MT microstimulation. (A) Example experiment from Figure 1D. The effect of the probes on the animal's ability to detect the motion signal was normalized to a t-score (see Methods). Positive/negative values indicated that the probe enhanced/suppressed the monkey's ability to detect the motion signal. (B) The average temporal effect of the probes on normalized performance (n = 53). Averages were fit with a single exponential. (C) Same as in (B), but using the probes effect on reaction time instead.

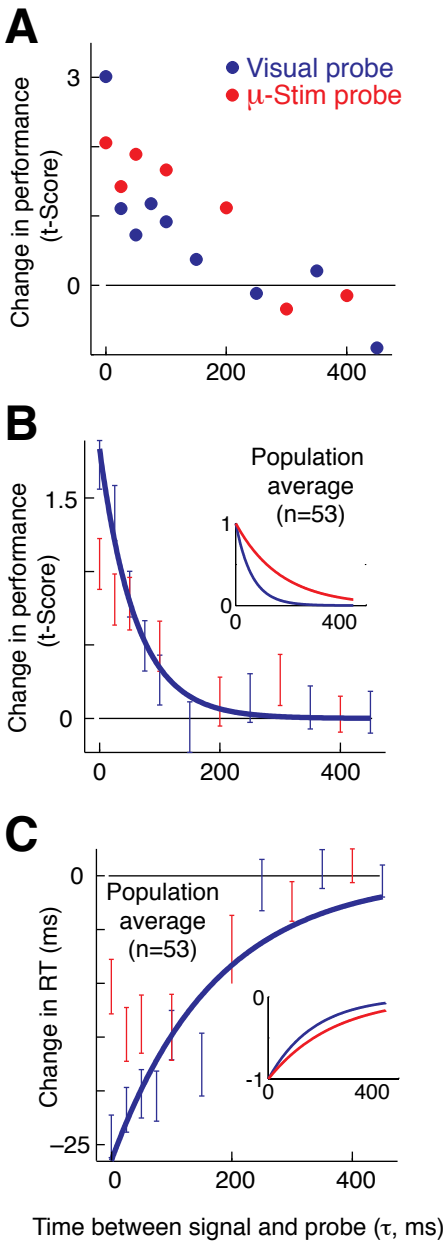
A
Visual probe
μ-Stim probe
Change in performance (t-Score)
3
0
0
200
400
B
Population average (n=53)
Change in performance (t-Score)
1.5
0
1
0
0
400
0
200
400
C
Population average (n=53)
Change in RT (ms)
0
–25
0
-1
0
400
0
200
400
Time between signal and probe (τ, ms)

Figure 3

was much longer (visual probe time constant: 61.5 ± 12.9 ms, microstimulation probe time constant: 170.9 ± 91.5 ms, p = 0.002, bootstrap, see Methods). The inset in Figure 3B shows the fitted IRFs of the two probes normalized so that the value of the fit is set to unity at time zero. This clearly shows the difference in decay rates between the two IRFs. We also calculated the decay rate of the two types of IRFs from the non-fitted day. We compared the effect of the probe averaged between τ = 50 and 100 ms to the effect averaged between τ = 0 and 25 (see Methods). For the microstimulation probe, the effect averaged between τ =50 and 100 ms decayed to 68% from the effect averaged between τ = 0 and 25 ms. On the other hand, the effect of the visual probe for these same time points decayed to 35% (p = 0.002, bootstrap).

The average effect of the probes on the animal's RT for correct trials showed a similar trend (Figure 3C). Microstimulation probes were weaker for $\tau \leq 50$ ms ($p < 10^{-5}$, two-sided t-test). The time constant for the microstimulation probe was marginally greater than the time constant for the visual probe (visual probe time constant: 174.1 ± 23.2 ms, microstimulation probe time constant: 250.9 ± 64.2 ms, p = 0.089). The inset of the normalized fitted IRFs highlights this difference. Furthermore, the effect of the microstimulation probe averaged between τ =50 and 100 ms increased to 112% of the average effect from τ = 0 and 25 ms while the effect of the visual probe for these same time points decayed to 76% (p = 0.005, bootstrap). These differences between the two probes were consistent for both monkeys (see Supplementary Methods).

To further demonstrate that microstimulation had a weaker and longer lasting effect on behavior, we compared the IRFs of microstimulation to that of visual motion on a site-by-site basis using principal component analysis (see Methods). Principal component analysis is a widely used technique to extract the most salient features of the data. We chose this analysis because there was considerable variability within many of our experiments, and we were not able to reliably fit exponentials to each experiment individually. The IRF of either probe is measured at nine time points; we used principal component analysis to project these nine values down to two data points. In Figure 4A, we show the projections of the two IRFs from the example experiment shown in Figure 1D and 3A. The coordinates of each projection are the dot product between the IRF and the two vectors (the eigenvectors) that capture the most variance out of all microstimulation and visual IRFs combined from all experimental sessions. We shifted both projections by equal amounts so that the projection of the visual IRF is at the origin. The microstimulation projection is above the visual projection, but the meaning of this difference is not evident from this two-dimensional representation. To reveal is meaning, we reconstructed IRFs from these two-dimensional projections (Figure 4B). The reconstructed IRFs are linear combinations of the two vectors (eigenvectors) that capture the most variance. Thus, the reconstructed IRFs will in general not be equal to the original IRFs. By comparing the constructed IRFs in Figure 4B to the originals in Figure 3A, we observe that that the reconstructed IRFs are effectively smoothed versions of the originals.

Figure 4 Principal component analysis of an example experimental session. (A) The IRF of the microstimulation and visual probes shown in Figure 1D and 3A are projected upon the two eigenvectors that capture the most variance across all experiments. The projection of the microstimulation IRF is given by the red circle while the visual projection is given by the blue square. Both projections have been shifted by equal amounts so that the visual projection is at the origin. The method used to calculate the angle between the two projections used in Figure 5B&E is shown. (B) IRFs are reconstructed from the two projections. The reconstructed microstimulation IRF is shown in red while the visual reconstruction is shown in blue. The reconstruction is the sum of the eigenvectors weighted by the coordinates of the projection.

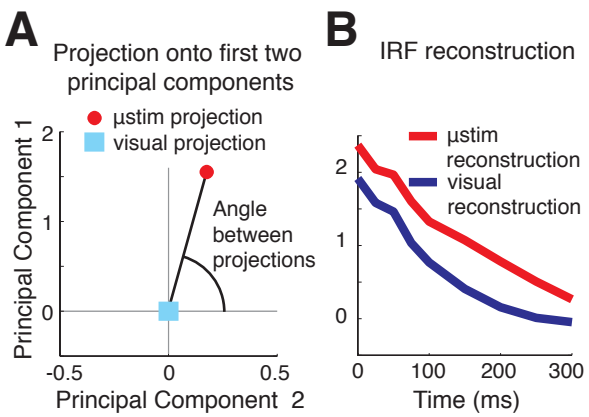
A
Projection onto first two principal components
μstim projection
visual projection
Angle between projections
Principal Component 1
2
1
0
-0.5
0
0.5
Principal Component 2
B
IRF reconstruction
μstim reconstruction
visual reconstruction
2
1
0
0
100
200
300
Time (ms)

Figure 4

In Figure 5A, we project the microstimulation (red dots) and visual (light blue square) IRFs of each experimental session and then shift the values so that the projection of the visual IRF is at the origin. Hence, all light blue squares overlap at the origin. For most experimental sessions, the microstimulation projections were to the right of the visual projections. This is quantified by the histogram of angles between the visual and microstimulation projections (Figure 5B), where the angle is calculated according to Figure 4A. The blue-to-red line gives the average direction between the microstimulation and visual projections. The distribution of angles between the visual and microstimulation projections (Figure 5B) was significantly different from uniform (Rayleigh test for non-uniformity, $p = 0.0032$). This implies that there was a consistent difference between the visual and microstimulation impulse-response functions, although this difference is not clear in this two-dimensional representation.

To understand the meaning of this difference between microstimulation and visual projections, we reconstructed IRFs as done in Figure 4B. We reconstructed IRFs (shown in Figure 5C) from points along the blue-to-red line in Figure 5A that points in the average direction between the two types of projections. Points along the blue-to-red line in Figure 5A are reconstructed to the IRF of matching color in Figure 5C. As one moves from blue to red along the line in Figure 5A, the reconstructed IRFs in Figure 5C become weaker and decay more slowly. This shows that moving rightward in the two-dimensional representation in Figure 5A corresponds to IRFs becoming weaker at short τ and decaying more slowly. Since most microstimulation projections are in the rightward

Figure 5 Principal component analysis of all experimental sessions. (A) Microstimulation (red dots) and visual temporal effects (light blue square) for each experiment were projected onto the first two eigenvectors. Each pair of microstimulation and visual projections from the same experimental session was shifted so that the visual projection was always centered at the origin. The line changing from blue to red line indicates the average direction between the visual and microstimulation projections. (B) Distribution of the angles between the projections of visual and microstimulation impulse-response functions from the same experimental session. (C) Reconstruction of the impulse-response functions from the projections. The dark red and blue lines are the reconstruction (see Methods) of the projections of the average microstimulation and visual impulse-response functions, respectively. The light blue-to-red lines are the reconstructions of the projections of points along the blue-to-red line in (A) of matching color. Thus, as one moves along the average direction going from visual to microstimulation projections, the reconstructed temporal effects becomes weaker at short time intervals and decays more slowly. (D-F) Same as above, except that the probes' effects on reaction times are used instead of the effect on detection performance.

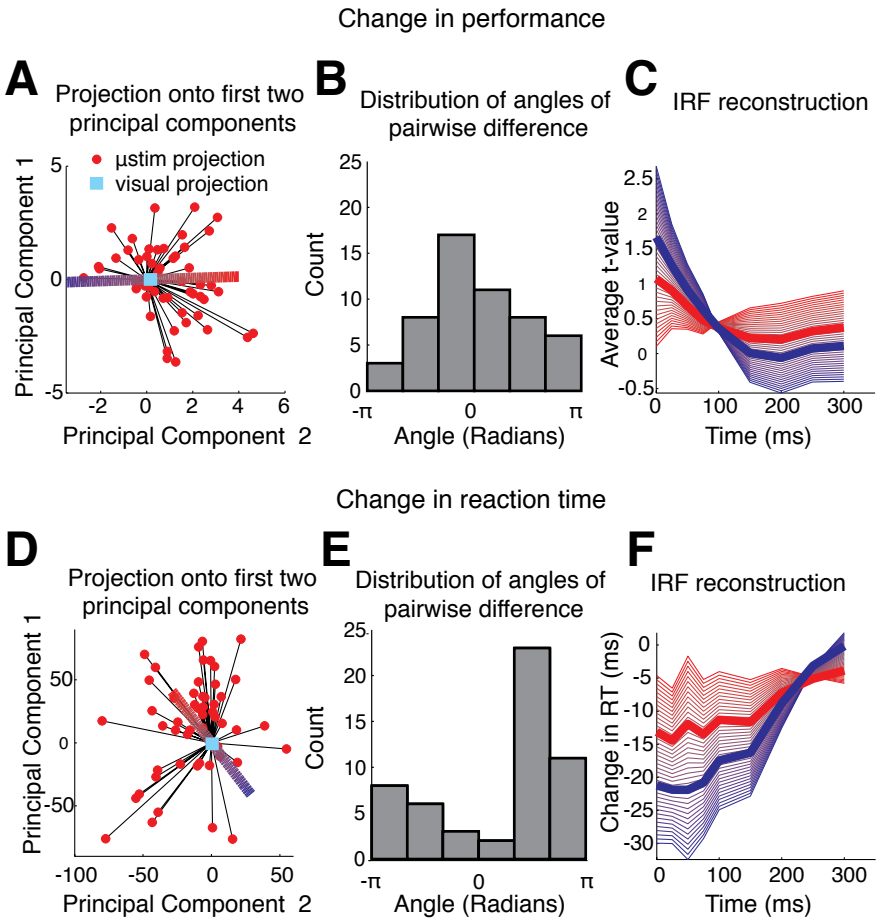
Change in performance
A
Projection onto first two principal components
μstim projection
visual projection
Principal Component 1
Principal Component 2
B
Distribution of angles of pairwise difference
Count
Angle (Radians)
C
IRF reconstruction
Average t-value
Time (ms)
Change in reaction time
D
Projection onto first two principal components
Principal Component 1
Principal Component 2
E
Distribution of angles of pairwise difference
Count
Angle (Radians)
F
IRF reconstruction
Change in RT (ms)
Time (ms)

Figure 5

direction (Figure 5B), microstimulation IRFs are consistently weaker at short τ and decay more slowly.

We repeated this calculation for the effect of the probes of the RT and found similar results (Figure 4D-F). The distribution of directions between the microstimulation and visual projections is shown in Figure 4E and was also significantly different from uniform ($p = 0.0003$). All together, the analysis in Figure 4 confirms the main finding in Figure 3 that the effect of microstimulation on behavior was weaker and lased longer compared to an equivalent visual stimulus.

Our analysis so far suggests that the visual and microstimulation probes had different effects on the behavioral detection of the motion signal. We next wanted to examine the mechanism of how the probe enhanced detection performance. One possibility is that the temporal effects of the probe was due to the integration of MT activity as illustrated in Figure 1B and proposed in many neural models of decision-making (Cook and Maunsell 2002; Huk and Shadlen 2005). An alternative possibility is that there was no integration of MT activity and the temporal effects of the probes was simply because the probes altered how area MT responds to the motion signal. To differentiate between these possibilities, we recorded multi-unit activity in area MT in response to the visual probes and signal in 47 out of 53 experiments. Because we could not record and stimulate simultaneously, this was done before we began the main microstimulation component of the experiment.

In Figure 6, we show the average normalized neural response to the motion signal when preceded by the different visual probes. Responses were normalized to the neural activity produced by the motion signal alone and expressed as the mean from 40 to 90 ms after the onset of the coherent motion signal (the peak neural response occurred 65 ms after stimulus onset). When the visual probe was adjacent the signal ($\tau = 0$ ms), the neural response to the motion signal was significantly greater ($p = 0.01$). However, the rest of the probe times did not consistently change the neural response in area MT to the motion signal (all other probe times, $p > 0.3$). The time constant of decay as a function of the separation between the probe and signal in Figure 6 was 26.4 ms and much shorter than the behavioral time constants in Figure 3B and C. Thus, the temporal effects of the probes on behavior cannot be solely explained by their effect on the neural responses to the motion signal. We conclude that the behavioral effects of both types of probes are at least partly due to the integration of MT activity, which then forms the basis of the decision to release the lever (illustrated in Figure 1B).

Figure 6 The time course of neural activity. In a subset of experiments (n = 47), multi-unit activity in area MT was recorded in response to the motion signal and visual probes. The neural response was defined as the average spike count within 25 ms of the peak neural response to visual motion (which had an average latency of 65 ms). The spike count to the motion signal only is assigned a value of 1 and all other values are relative to this baseline.

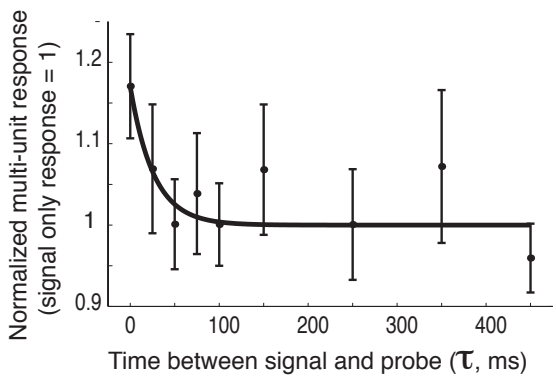
Normalized multi-unit response
(signal only response = 1)
1.2
1.1
1
0.9
0
100
200
300
400
Time between signal and probe (τ, ms)

Figure 6

Discussion

We measured the behavioral impulse-response function of brief microstimulation in area MT. An important aspect of our experimental design was matching the strength of visual and microstimulation probes, which allowed us to compare the impulse-response function of the microstimulation to that of an equivalent visual stimulus. We found that the effect of the microstimulation probe on motion perception was weaker and decayed more slowly compared to the effect of the equivalent visual probe. Given that future applications of microstimulation may require high temporal precision, this study raises the possibility that cortical microstimulation is ill suited to generate the required temporally precise percepts.

Understanding the temporal effects of microstimulation on behavior is important because the sensory environment, along with the perception it generates, is highly dynamic (Hegde 2008; Simoncelli and Olshausen 2001). Animals can perceive (Thorpe et al. 1996) and react to changes in their environment on very short timescales (Rieke et al. 1999) , and sensory perception itself evolves over time to even static stimuli (Hegde 2008). Although highly informative, previous studies probing the causality between neural activity and perception have usually examined the effects of microstimulating for much longer periods of hundreds or thousands of milliseconds. Recent studies have begun to reveal how sensory areas encode natural stimuli (Butts et al. 2007; David et al. 2004; Felsen et al. 2005; Lesica and Stanley 2004; Nemenman et al. 2008; Sharpee et al. 2004; Simoncelli and Olshausen 2001) and it is useful to know the extent to which

microstimulation can be used to reveal the causality between sensory activity and perception on these fast timescales. Furthermore, microstimulation will be the basis for future advanced sensory neural prosthetics (Bradley et al. 2005; Fernandez et al. 2005; Girvin 1988; McIntyre and Grill 2000; Middlebrooks et al. 2005; Normann et al. 1999; Tehovnik and Slocum 2007; Troyk et al. 2003). Thus, it is imperative we understand the effects of microstimulation in the same temporal regime as natural sensory inputs. In addition to neural prosthetics, deep brain stimulation has shown therapeutic promise for a range of neurological conditions (Perlmutter and Mink 2006). Both of these applications will benefit from a precise description of how current injected from microelectrodes is translated into behavior and perception.

Several past studies have previously examined the temporal spread of microstimulation on cortical activity, both *in vitro* (Shao and Burkhalter 1996; 1999) and in anaesthetized, *in vivo* preparations (Berman et al. 1991; Butovas et al. 2006; Butovas and Schwarz 2003; Chung and Ferster 1998; Contreras et al. 1997). These studies revealed that microstimulation induced long-lasting suppression in neural activity, mediated by $GABA_B$-receptor synaptic inhibition (Butovas et al. 2006). It is likely that in our detection task, a microstimulation-induced reduction in the activity of neurons encoding the stimulus could have been used as evidence for the presence of the coherent motion signal. Since we do not fully understand the rules by which neural activity is transformed into behavior, we cannot claim that long lasting effects on neural activity translates into long lasting effects on behavior. However, these past electrophysiology

studies combined with our behavioral study provide a consistent picture that microstimulation affects perception for several hundred milliseconds after it stops.

There has been a great deal of effort to understand the spatial spread of cortical activity in response to microstimulation (Ranck 1975; Salzman et al. 1992; Stoney et al. 1968; Tehovnik et al. 2005; Tehovnik et al. 2003; Tolias et al. 2005). Outside the electrophysiological studies above, few behavioral studies have investigated the temporal properties of microstimulation. In an elegant study, Brecht and colleagues demonstrated that eye movements produced by two stimulating electrodes in the superior colliculus were sensitive to temporal offsets of less than 10 ms (Brecht et al. 2004). In addition, *in vivo* results have demonstrated that the precise timing of stimulation pulses can result in measurable changes in neural plasticity (Baranyi and Feher 1981; Jackson et al. 2006). In agreement with our results, Tehovnik and colleagues showed that the effect of microstimulation on the latency of visually guided saccades decayed over a time course of about 200 ms (Tehovnik et al. 2004). However, their experiment could not determine if the long interaction time between microstimulation and saccade latency was due to the natural temporal integration of the cortex or due to the temporal properties of the microstimulation. This long integration time is perhaps best exemplified in a study that observed that weak 100 ms pulses of motion perturbed both motion perception as well as neural activity in area LIP, believed to integrate activity from MT, over a period of 800 ms (Huk and Shadlen, 2005). By using a visual probe as a control, we are in a better position to show that the long interaction time is a result of the temporal spread of

microstimulation and cannot be fully explained by the natural temporal integration of cortex.

We believe that the microstimulation used in our study only activated a single cortical column involved in encoding the motion stimulus. Although we could not test this assumption directly, past results clearly indicate that our microstimulation parameters (low amplitude and very short duration) likely produced activity that was tightly localized around the tip of our electrode (Murasugi et al. 1993; Salzman et al. 1992). The low detection rate (~5%) of the probes also suggests our microstimulation parameters produced subthreshold behavioral effects. Our amplitudes were similar to recently reported detection thresholds in MT (Murphey and Maunsell 2007), however, this other study used stimulation durations that were eight times longer.

Although there is good reason to think our microstimulation activated a local cortical column in MT, the notion that we were always activating neurons representing "coherent motion" may be over simplistic. DeAngelis and Newsome suggest that stimulating areas in MT that are also strongly tuned to disparity can reduce the effect of microstimulation in a motion discrimination task (DeAngelis and Newsome 2004). Born and colleagues also showed that the effectiveness of microstimulation may depend on the center surround properties of the column surrounding the electrode (Born et al. 2000). That MT is not a homogeneous motion encoding area provides a likely explanation of why many studies have found behavioral thresholds for the amplitude of microstimulation vary from site to site. However, it seems unlikely that the particular

sensory representation of a cortical column in MT would influence how the effect of microstimulation spreads across time.

This study is different than many previous studies that measure the effect of microstimulation on behavior in that we employ a detection paradigm, as opposed to a two-alternative forced choice or discrimination paradigm (Bisley et al. 2001; Celebrini and Newsome 1995; de Lafuente and Romo 2005; Ditterich et al. 2003; Hanks et al. 2006; Liu and Newsome 2005; Romo et al. 1998; Salzman et al. 1990). Our motivation for this task design was to ensure the visual and microstimulation probes were equally, yet rarely detected because the probes could occur at any time during a 10 second period. Our detection task required the animals to adopt a low guess rate and allowed us to quickly measure the effects of the both types of probes (because the variance of a binomial distribution is smallest at the extremes). These constraints on our paradigm meant that we could not consider several variants of our experiment that might otherwise have been interesting. For example, it has previously been shown microstimulation in an area of MT that is not behaviorally relevant does not affect the perceptual decision of the subject (Salzman et al. 1992). However, since we rewarded the monkey for detecting either probe, the monkey could always use the microstimulation to help detect the visual motion, no matter its origin within area MT.

Was the difference between the behavioral effects of the two probes the result of a difference in how they were processed by the brain? For example, the visual probe might have provided sensory evidence towards a coherent motion percept, while the

microstimulation probe might have acted to cue the subject to the upcoming motion stimulus. Past studies have shown that longer pulses of microstimulation to MT provide sensory evidence towards coherent motion perception (Ditterich et al. 2003; Nichols and Newsome 2002; Salzman et al. 1990). However, if microstimulation acted to cue the subject to the upcoming motion stimulus, then one must also believe that the visual motion probe must have done the same; both probes were subthreshold stimuli that generated neural responses in the same cortical area. Even if the visual system did not process the two probes in the same fashion, then this would only serve to strengthen the result that brief microstimulation is not equivalent to an equivalent sensory stimulus.

Supplementary Methods

The variability in neural latency

One possible explanation for the difference between the microstimulation and visual impulse-response functions could be the variability in neural latency of sites in MT we stimulated. Although our shift of 50 ms was very close to the average neural latency of 53.6ms, the variability in the latency could have smoothed our estimate of microstimulation impulse-response function, producing the weaker effect that we observed at short time intervals and increasing the decay time.

To test this hypothesis, we repeated our calculation using only experimental sessions whose neural latency was between 43 and 57 ms inclusive (arbitrarily set to give us 30 experimental sessions). Overall, the fundamental differences between the visual and microstimulation temporal effects remained the same (Supplementary Figure 1A). On average, microstimulation probes had a weaker effect on performance for short τ ($\tau \leq 50$: $p = 0.060$; $\tau \leq 25$: $p = 0.005$) and decayed more slowly (visual time constant: 54.2 ms, microstimulation time constant: 178.3 ms, $p = 0.006$, bootstrap). We also calculated the decay rate of the two types of probes from the non-fitted data. We compared the effect of the probe averaged between $\tau = 50$ and 100 ms to the effect averaged between $\tau = 0$ and 25 (see Methods). For the microstimulation probe, the effect averaged between $\tau = 50$ and 100 ms decayed to 67% from the effect averaged between $\tau = 0$ and 25 ms. On the other

Supplementary Figure 1 Population averages for sites with similar neural latency. (A) The average temporal effects of the visual (red dots) and microstimulation (blue dots) probes for 30 sites with neural latency between 43 and 57 ms, inclusive. As in Figure 3B, all results are first normalized to a t-score before being averaged. Both averages were fit with an exponential of the form $a\exp(-b\tau)$. (B) Same as above, but using the probes effect on reaction time instead.

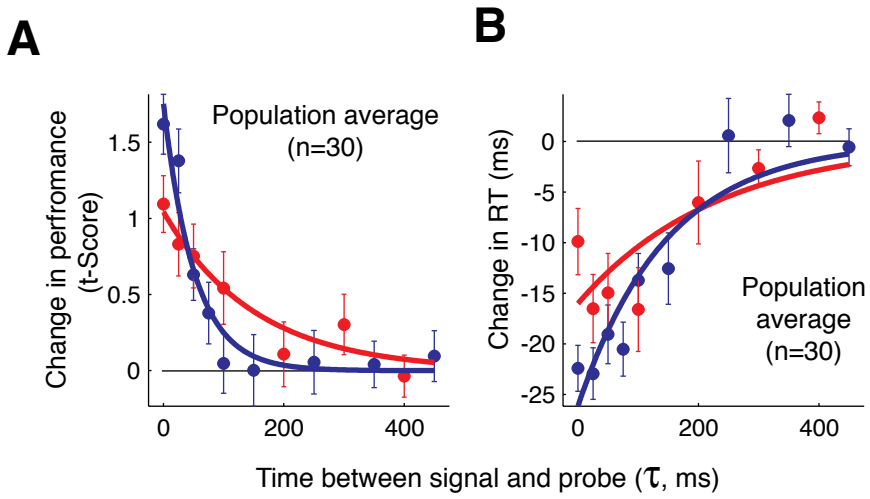

Supplementary Figure 1

hand, the effect of the visual probe for these same time points decayed to 22% ($p = 0.001$, bootstrap).

Similarly, the differences between the microstimulation and visual probes' effect on reaction times did not change (Supplementary Figure 1B). Microstimulation probes at short τ had less effect than visual probes ($\tau \leq 50$: $p = 0.006$; $\tau \leq 25$: $p < 10^{-3}$) and decayed more slowly (visual time constant: 150.4 ms, microstimulation time constant: 235.0 ms, $p = 0.034$, bootstrap). Using the actual data (not fitted by an exponential), for the microstimulation probe, the effect averaged between τ =50 and 100 ms increased to 120% from the effect averaged between $\tau = 0$ and 25 ms. On the other hand, the effect of the visual probe for these same time points decayed to 72% ($p = 0.004$, bootstrap).

To confirm that these results were also consistent on site-by-site basis for these 30 experiments, we repeated the principal component analysis on this reduced data set (Supplementary Figure 2). Both distributions of the directions were significantly different from uniform (Supplementary Figure 2B&E, change in performance: $p = 0.011$; change in reaction time: $p = 0.043$), indicating that the differences between the two impulse-response functions were consistent across experimental sessions. Additionally, the reconstruction of the impulse-response functions (Supplementary Figure 2C&F) once again showed that the effect of microstimulation on performance and RT was initially weaker and decayed more slowly. All together, these results suggest that variability in

Supplementary Figure 2 Principal component analysis for sites with similar neural latency. Same as Figure 4, except only 30 sites with neural latencies between 43 and 57 ms were used for the analysis.

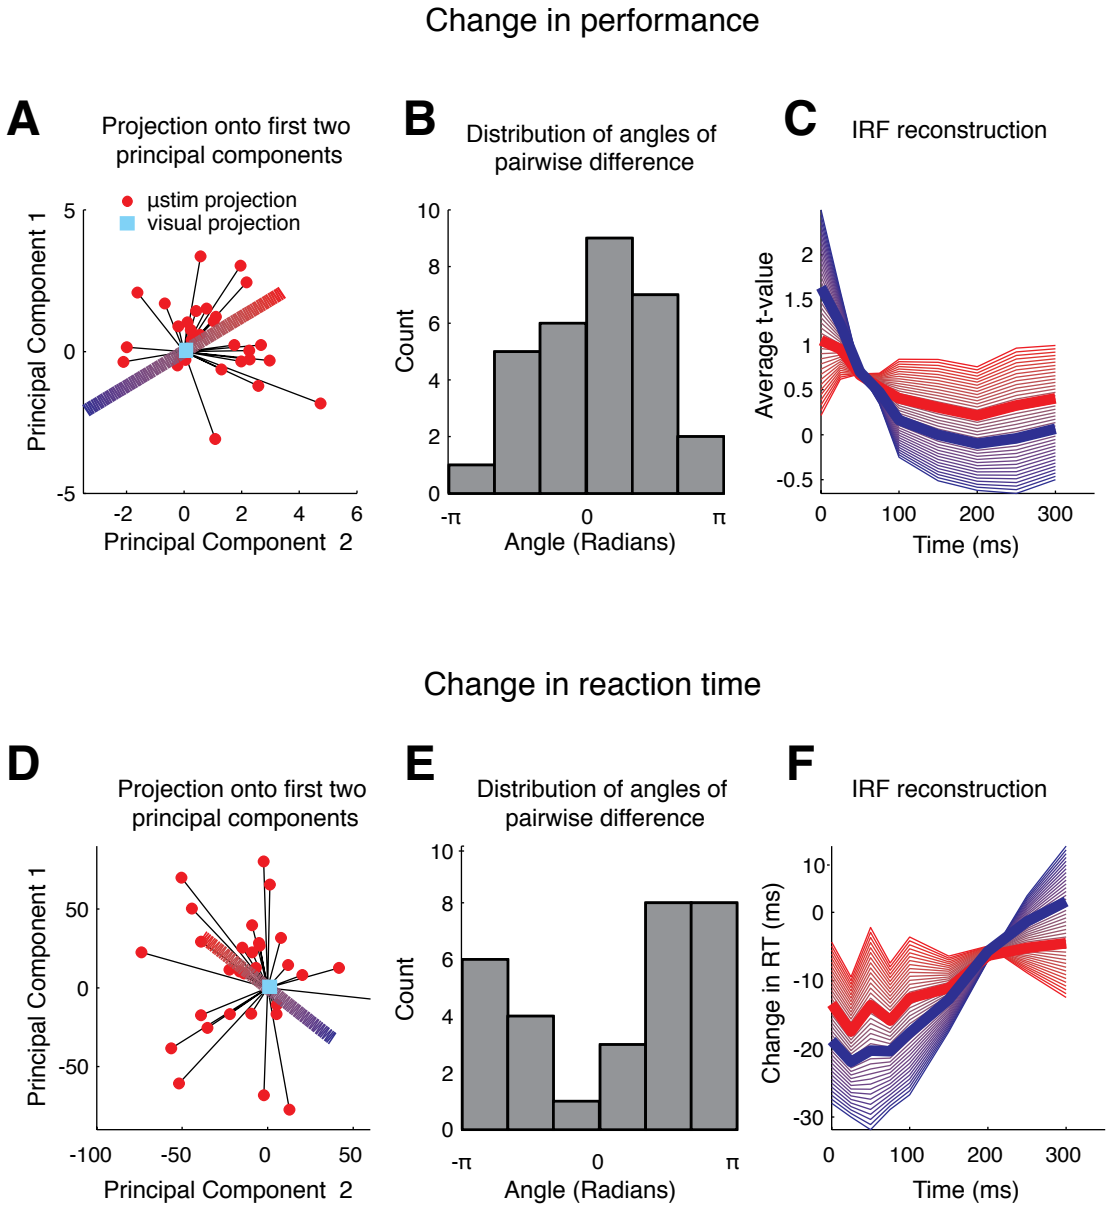

Supplementary Figure 2

neural latency was not the reason for the difference between microstimulation and visual impulse-response functions.

The variability in probe strength

In order to make a fair comparison between the two probes, it was important to set the strength of the microstimulation probe so that it was approximately equal to that of an equivalent visual probe. To ensure this equality, we measured how often the animal could detect microstimulation and visual probes alone (Figure 2A&B). On average, the two types of probes were detected with low and almost equal probability. Although there was appreciable variability between the strength of the visual and microstimulation probes, we do not believe that this could have produced the observed differences between the temporal effects of the two types of probes. To check this, we recalculated the temporal effects of the probes on performance and RT for the case where visual probe only trials were detected more often than microstimulation probe only trials. The effect of the microstimulation probe was weaker than the effect of the visual probe at small $\tau \leq 50$ ms (changed in performance: $p = 0.033$; change in RT: $p < 10^{-4}$, two-sided t-test, n=29) and time constant was longer for the microstimulation probe (change in performance: 67.3 ms for visual and 295.7 ms for microstimulation, $p < 0.001$; time constant for change RT: 228.2 ms for visual and 363.2 ms for microstimulation, $p = 0.27$, bootstrap, n = 29). Using the actual data (not fitted by an exponential), for the microstimulation probe, the effect on the change in performance averaged between τ =50 and 100 ms decayed to 78% from the effect averaged between τ = 0 and 25 ms while the effect of the visual

probe for these same time points decayed to 37% ($p = 0.005$, bootstrap). For the change in RT, the effect of the micrsotimulation probe averaged between τ =50 and 100 ms increased to 106% from the effect averaged between $\tau = 0$ and 25 ms while the effect of the visual probe for these same time points decayed to 82% ($p = 0.133$, bootstrap).

For the case when the microstimulation probe only trials were detected more often than visual probe only trials, the results were similar. The effect of the microstimulation probe was weaker the visual probe at small $\tau \leq 50$ ms (changed in performance: $p = 0.02$; change in RT: $p = 0.039$, two-sided t-test, $n = 22$) and time constant was longer for the microstimulation probe (change in performance: 62.5 ms for visual and 100.2 ms for microstimulation, $p = 0.28$; time constant for change RT: 175.8 ms for visual and 405.1 ms for microstimulation, $p = 0.004$, bootstrap, $n = 22$). For the microstimulation probe, the effect on the change in performance averaged between $\tau = 50$ and 100 ms decayed to 51% from the effect averaged between $\tau = 0$ and 25 ms while the effect of the visual probe for these same time points decayed to 28% ($p = 0.111$, bootstrap). For the change in RT, the effect of the microstimulation probe averaged between τ =50 and 100 ms increased to 117% from the effect averaged between $\tau = 0$ and 25 ms while the effect of the visual probe for these same time points decayed to 78% ($p = 0.022$, bootstrap).

Regardless of the strength of the probe, microstimulation probes had significantly weaker effect on visual perception at $\tau \leq 50$ than visual probes. Additionally, the decay time from both the actual and fitted data was longer for the microstimulation probes.

Although not all of the differences in decay time were significant, this was to be expected given that the population was split into two smaller subgroups. On the whole, the strength of the probe was not a factor in the differences between the visual and microstimulation probes.

Individual analysis of both monkeys

Both monkeys showed similar difference between the effects of the mcirostimulation and visual probes. For monkey W, who performed 30 experimental sessions, the microstimulation probe was weaker at $\tau <= 50$ ms (change in performance: $p = 0.034$; change in RT: $p = 0.009$, two-sided t-test). The time constant from the fitted data was longer for the microstimulation probe (change in performance: microstimulation time constant = 45.5 ms, visual time constant = 202.0 ms, $p = 0.077$; change in RT: microstimulation time constant = 430.2 ms, visual time constant = 152.6 ms, $p = 0.004$, bootstrap). Additionally, there was less of a decay going from $\tau = 0$ and 25 ms to $\tau = 50$ and 100 ms (change in performance: $p = 0.015$; change in RT: $p = 0.009$, bootstrap).

For monkey G, who performed 23 experimental sessions, the same differences existed. The microstimulation probe was weaker at $\tau <= 50$ ms (change in performance: $p = 0.034$; change in RT: $p = 0.009$, two-sided t-test). For the change in performance, the time constant from the fitted data was longer for the microstimulation probe while there was no significant difference for the change in RT (change in performance: microstimulation time constant = 86.5 ms, visual time constant = 167.2 ms, $p = 0.008$;

change in RT: microstimulation time constant = 147.3 ms, visual time constant = 210.4 ms, $p = 0.088$, bootstrap). It should be noted that this calculation of the change in RT is noisy given it consists of only the correct trials from 23 exprimental sessions. Finally, the effect of the microstimulation probe decayed from $\tau = 0$ and 25 ms to $\tau = 50$ and 100 ms more slowly, although the difference was not as great as it was for monkey W (change in performance: $p = 0.199$; change in RT: $p = 0.67$, bootstrap).

REFERENCES

Albright TD, Desimone R, and Gross CG. Columnar organization of directionally selective cells in visual area MT of the macaque. J Neurophysiol 51: 16-31, 1984.

Baranyi A, and Feher O. Synaptic facilitation requires paired activation of convergent pathways in the neocortex. Nature 290: 413-415, 1981.

Berman NJ, Douglas RJ, Martin KA, and Whitteridge D. Mechanisms of inhibition in cat visual cortex. J Physiol 440: 697-722, 1991.

Bisley JW, Zaksas D, and Pasternak T. Microstimulation of cortical area MT affects performance on a visual working memory task. J Neurophysiol 85: 187-196, 2001.

Born RT, Groh JM, Zhao R, and Lukasewycz SJ. Segregation of object and background motion in visual area MT: effects of microstimulation on eye movements. Neuron 26: 725-734, 2000.

Bradley DC, Troyk PR, Berg JA, Bak M, Cogan S, Erickson R, Kufta C, Mascaro M, McCreery D, Schmidt EM, Towle VL, and Xu H. Visuotopic mapping through a multichannel stimulating implant in primate V1. J Neurophysiol 93: 1659-1670, 2005.

Brecht M, Singer W, and Engel AK. Amplitude and direction of saccadic eye movements depend on the synchronicity of collicular population activity. J Neurophysiol 92: 424-432, 2004.

Butovas S, Hormuzdi SG, Monyer H, and Schwarz C. Effects of electrically coupled inhibitory networks on local neuronal responses to intracortical microstimulation. J Neurophysiol 96: 1227-1236, 2006.

Butovas S, and Schwarz C. Spatiotemporal effects of microstimulation in rat neocortex: a parametric study using multielectrode recordings. J Neurophysiol 90: 3024-3039, 2003.

Butts DA, Weng C, Jin J, Yeh CI, Lesica NA, Alonso JM, and Stanley GB. Temporal precision in the neural code and the timescales of natural vision. Nature 449: 92-95, 2007.

Carey MR, Medina JF, and Lisberger SG. Instructive signals for motor learning from visual cortical area MT. Nat Neurosci 8: 813-819, 2005.

Celebrini S, and Newsome WT. Microstimulation of extrastriate area MST influences performance on a direction discrimination task. J Neurophysiol 73: 437-448, 1995.

Chung S, and Ferster D. Strength and orientation tuning of the thalamic input to simple cells revealed by electrically evoked cortical suppression. Neuron 20: 1177-1189, 1998.

Cohen MR, and Newsome WT. What electrical microstimulation has revealed about the neural basis of cognition. Curr Opin Neurobiol 14: 169-177, 2004.

Contreras D, Durmuller N, and Steriade M. Absence of a prevalent laminar distribution of IPSPs in association cortical neurons of cat. J Neurophysiol 78: 2742-2753, 1997.

Cook EP, and Maunsell JH. Dynamics of neuronal responses in macaque MT and VIP during motion detection. Nat Neurosci 5: 985-994, 2002.

David SV, Vinje WE, and Gallant JL. Natural stimulus statistics alter the receptive field structure of v1 neurons. J Neurosci 24: 6991-7006, 2004.

de Lafuente V, and Romo R. Neuronal correlates of subjective sensory experience. Nat Neurosci 8: 1698-1703, 2005.

DeAngelis GC, and Newsome WT. Perceptual "read-out" of conjoined direction and disparity maps in extrastriate area MT. PLoS Biol 2: E77, 2004.

Ditterich J, Mazurek ME, and Shadlen MN. Microstimulation of visual cortex affects the speed of perceptual decisions. Nat Neurosci 6: 891-898, 2003.

Felsen G, Touryan J, Han F, and Dan Y. Cortical sensitivity to visual features in natural scenes. PLoS Biol 3: e342, 2005.

Fernandez E, Pelayo F, Romero S, Bongard M, Marin C, Alfaro A, and Merabet L. Development of a cortical visual neuroprosthesis for the blind: the relevance of neuroplasticity. J Neural Eng 2: R1-12, 2005.

Girvin JP. Current status of artificial vision by electrocortical stimulation. Can J Neurol Sci 15: 58-62, 1988.

Hanks TD, Ditterich J, and Shadlen MN. Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. Nat Neurosci 9: 682-689, 2006.

Hegde J. Time course of visual perception: Coarse-to-fine processing and beyond. Prog Neurobiol 84: 405-439, 2008.

Huk AC, and Shadlen MN. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J Neurosci 25: 10420-10436, 2005.

Jackson A, Mavoori J, and Fetz EE. Long-term motor cortex plasticity induced by an electronic neural implant. Nature 444: 56-60, 2006.

Lesica NA, and Stanley GB. Encoding of natural scene movies by tonic and burst spikes in the lateral geniculate nucleus. J Neurosci 24: 10731-10740, 2004.

Liu J, and Newsome WT. Correlation between speed perception and neural activity in the middle temporal visual area. J Neurosci 25: 711-722, 2005.

Maunsell JH, and Van Essen DC. Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. J Neurophysiol 49: 1127-1147, 1983.

McIntyre CC, and Grill WM. Selective microstimulation of central nervous system neurons. Ann Biomed Eng 28: 219-233, 2000.

Middlebrooks JC, Bierer JA, and Snyder RL. Cochlear implants: the view from the brain. Curr Opin Neurobiol 15: 488-493, 2005.

Murasugi CM, Salzman CD, and Newsome WT. Microstimulation in visual area MT: effects of varying pulse amplitude and frequency. J Neurosci 13: 1719-1729, 1993.

Murphey DK, and Maunsell JH. Behavioral detection of electrical microstimulation in different cortical visual areas. Curr Biol 17: 862-867, 2007.

Nemenman I, Lewen GD, Bialek W, and de Ruyter van Steveninck RR. Neural coding of natural stimuli: information at sub-millisecond resolution. PLoS Comput Biol 4: e1000025, 2008.

Nichols MJ, and Newsome WT. Middle temporal visual area microstimulation influences veridical judgments of motion direction. J Neurosci 22: 9530-9540, 2002.

Normann RA, Maynard EM, Rousche PJ, and Warren DJ. A neural interface for a cortical vision prosthesis. Vision Res 39: 2577-2587, 1999.

Perlmutter JS, and Mink JW. Deep Brain Stimulation. Annu Rev Neurosci 2006.

Ranck JB, Jr. Which elements are excited in electrical stimulation of mammalian central nervous system: a review. Brain Res 98: 417-440, 1975.

Rashbass C. The visibility of transient changes of luminance. J Physiol 210: 165-186, 1970.

Rieke F, Warland D, van Stevenick RR, and Bialek W. Spikes: exploring the neural code. MIT Press, 1999.

Romo R, Hernandez A, Zainos A, Brody CD, and Lemus L. Sensing without touching: psychophysical performance based on cortical microstimulation. Neuron 26: 273-278, 2000.

Romo R, Hernandez A, Zainos A, and Salinas E. Somatosensory discrimination based on cortical microstimulation. Nature 392: 387-390, 1998.

Romo R, and Salinas E. Sensing and deciding in the somatosensory system. Curr Opin Neurobiol 9: 487-493, 1999.

Salzman CD, Britten KH, and Newsome WT. Cortical microstimulation influences perceptual judgements of motion direction. Nature 346: 174-177, 1990.

Salzman CD, Murasugi CM, Britten KH, and Newsome WT. Microstimulation in visual area MT: effects on direction discrimination performance. J Neurosci 12: 2331-2355, 1992.

Shao Z, and Burkhalter A. Different balance of excitation and inhibition in forward and feedback circuits of rat visual cortex. J Neurosci 16: 7353-7365, 1996.

Shao Z, and Burkhalter A. Role of GABAB receptor-mediated inhibition in reciprocal interareal pathways of rat visual cortex. J Neurophysiol 81: 1014-1024, 1999.

Sharpee T, Rust NC, and Bialek W. Analyzing neural responses to natural signals: maximally informative dimensions. Neural Comput 16: 223-250, 2004.

Simoncelli EP, and Olshausen BA. Natural image statistics and neural representation. Annu Rev Neurosci 24: 1193-1216, 2001.

Simpson WA. Temporal summation of visual motion. Vision Res 34: 2547-2559, 1994.

Stoney SD, Jr., Thompson WD, and Asanuma H. Excitation of pyramidal tract cells by intracortical microstimulation: effective extent of stimulating current. J Neurophysiol 31: 659-669, 1968.

Tehovnik EJ. Electrical stimulation of neural tissue to evoke behavioral responses. J Neurosci Methods 65: 1-17, 1996.

Tehovnik EJ, and Slocum WM. Phosphene induction by microstimulation of macaque V1. Brain Res Rev 53: 337-343, 2007.

Tehovnik EJ, Slocum WM, Carvey CE, and Schiller PH. Phosphene induction and the generation of saccadic eye movements by striate cortex. J Neurophysiol 93: 1-19, 2005.

Tehovnik EJ, Slocum WM, and Schiller PH. Microstimulation of V1 delays the execution of visually guided saccades. Eur J Neurosci 20: 264-272, 2004.

Tehovnik EJ, Slocum WM, and Schiller PH. Saccadic eye movements evoked by microstimulation of striate cortex. Eur J Neurosci 17: 870-878, 2003.

Thorpe S, Fize D, and Marlot C. Speed of processing in the human visual system. Nature 381: 520-522, 1996.

Tolias AS, Sultan F, Augath M, Oeltermann A, Tehovnik EJ, Schiller PH, and Logothetis NK. Mapping cortical activity elicited with electrical microstimulation using FMRI in the macaque. Neuron 48: 901-911, 2005.

Troyk P, Bak M, Berg J, Bradley D, Cogan S, Erickson R, Kufta C, McCreery D, Schmidt E, and Towle V. A model for intracortical visual prosthesis research. Artif Organs 27: 1005-1015, 2003.

Uka T, and DeAngelis GC. Linking neural representation to function in stereoscopic depth perception: roles of the middle temporal area in coarse versus fine disparity discrimination. J Neurosci 26: 6791-6802, 2006.

Van Essen DC, Maunsell JH, and Bixby JL. The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization. J Comp Neurol 199: 293-326, 1981.

Zeki SM. Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. J Physiol 236: 549-573, 1974.

Shadlen MN, Britten KH, Newsome WT, Movshon JA (1996) A computational analysis of the relationship between neuronal and behavioral responses to visual motion. J Neurosci 16:1486-1510.

Shioiri S, Cavanagh P (1989) Saccadic suppression of low-level motion. Vision Res 29:915-928.

Snodderly DM, Kagan I, Gur M (2001) Selective activation of visual cortex neurons by fixational eye movements: implications for neural coding. Vis Neurosci 18:259-277.

Sperling G (1990) In: Eye Movements and Their Role in Visual and Cognitive Processes (Kowler E, ed), pp 307-351. Amsterdm: Elsevier.

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

Uka T, Tanabe S, Watanabe M, Fujita I (2005) Neural correlates of fine depth discrimination in monkey inferior temporal cortex. J Neurosci 25:10796-10802.

Williams ZM, Elfar JC, Eskandar EN, Toth LJ, Assad JA (2003) Parietal activity and the perceived direction of ambiguous apparent motion. Nat Neurosci 6:616-623.

Wurtz RH (2008) Neuronal mechanisms of visual stability. Vision Res.

Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. Nature 370:140-143.

Zuber BL, Stark L (1965) Microsaccades and the velocity-amplitude relationship for saccadic eye movements. Science 150:1459-1460.

Zohary, E., Shadlen, M. N., and Newsome, W. T. (1994b). Correlated neuronal discharge rate and its implications for psychophysical performance. Nature *370*, 140-143.

Chapter 5

The effect of spatial attention on sensory processing in area Ventral Intraparietal

Abstract

Visual attention has been shown to increase behavioral performance in a variety of tasks. However, the neural mechanisms that underlie this improvement are poorly understood. We hypothesize that the increase in behavioral performance is not simply because attention modulates how visual neurons encode stimuli, but because attention affects how neural activity is transformed between cortical areas. We tested this hypothesis by recording direction-selective neurons from ventral intraparietal area (VIP) during a motion detection task. We found that the neurons more sensitive to the stimulus were also more influential in forming the perceptual decision only when spatial attention was directed inside their receptive fields. This suggests that attention increases visual performance by properly weighing sensory activity and that properly weighing activity represents a bottleneck in visual processing. Unrelated to this finding, we also observed that the top-down signal of attention does not simply modulate how bottom-up signals are integrated, but can strongly interact with other top-down signals.

Introduction

Our environment contains an overwhelming amount of visual information (Jacobson, 1951); attention allows one to suppress irrelevant stimuli and focus on those that are behaviorally important. One form of attention is spatial attention, in which attention is directed towards a specific location in one's visual field. It is well known that spatial attention lowers detection and discrimination thresholds for visual stimuli at the attended location (Bashinski and Bacharach, 1980; Downing, 1988; Muller and Humphreys, 1991; Balz and Hock, 1997; Carrasco and Yeshurun, 1998; Cook and Maunsell, 2002a). However, the neural mechanisms of attention that lead to this increase in performance are still unknown.

At least two, non-exclusive mechanisms have been proposed to explain the increase in performance from attention. The first is the "noise reduction" or change in bias model, in which signals from attended locations are weighted more heavily than signals from non-attended locations (Dosher and Lu, 1998; Lu and Dosher, 1998; Morgan et al., 1998; Baldassi and Burr, 2000). This is supported by studies that have reported that neural activity is increased when the attended location matches the receptive field, and decreased when attention is outside (Treue and Maunsell, 1996; McAdams and Maunsell, 1999; Treue and Maunsell, 1999; Cook and Maunsell, 2002a). However, psychophysical experiments performed under "noiseless" conditions have observed increased visual sensitivity at the attended location, suggesting that attention can also be modeled as a

"signal enhancement" (Bashinski and Bacharach, 1980; Yeshurun and Carrasco, 1999; Carrasco et al., 2000). This model proposes that the visual system can encode stimuli at the attended location with better sensitivity, leading to increased behavioral performance. However, most studies have found no increase in the signal-to-noise ratio of the neural response with attention (McAdams and Maunsell, 1999; Treue and Martinez Trujillo, 1999), placing doubt on this hypothesis. Although experimental data has not ruled out the possibility that these two possible explanations serve a role in attention, it does suggest that our understanding of the neural mechanisms of attention is incomplete.

Other studies, not necessarily related to attention, have asked how information from sensory neurons is combined when forming perceptual decisions. These studies suggest that the activity of sensory neurons is not treated equally, but is weighted depending on how informative the neural activity is of the stimulus given the task (Britten et al., 1996; Dodd et al., 2001; Parker et al., 2002; Uka and DeAngelis, 2004; Purushothaman and Bradley, 2005; Jazayeri and Movshon, 2006, 2007b, a; Law and Gold, 2008). The proper weighting of sensory activity is not innate, but is the result of training (Law and Gold, 2008). Since the neurons most informative of the stimulus are dependent on the task, it would suggest that top-down feedback would be required to properly weight the activity of sensory neurons.

In this study, we ask whether the top-down signal of spatial attention properly weighs the activity of sensory neurons when forming a decision. We find that when spatial attention was directed inside VIP neurons' receptive fields, the neurons more

informative of the stimulus were more correlated with the behavioral outcome of the trial, as previously observed in the studies mentioned above. However, this correlation is eliminated when attention was directed outside the neurons' receptive fields. This change is at least partly because spatial attention increased the sensitivity of the neurons' more correlated with the behavior. We hypothesize that spatial attention increases visual performance by properly weighing sensory activity.

In an observation unrelated to the findings above, we found that many VIP neurons received strong feedback just prior to a correct response, even when no motion was present inside the neuron's receptive field. However, when attention was directed outside the neuron's receptive field, this feedback was significantly reduced. This suggests that the top-down signal of attention does not simply modulate a neuron's response to incoming sensory stimuli, but can gait other top-down feedback, complicating our view of how attention functions.

Methods

Behavioral task

The data set analyzed for this study comes from several previous studies (Cook and Maunsell, 2002a, 2004). Two monkeys (Macaca mulatta) were trained to perform a spatially cued motion detection task (Figure 1). The trial began when the monkey depressed a lever and fixated on a central point. The goal of the task was to release the lever when coherent motion began in one of two random dot patches diametrically opposite of the fixation point. After the cue was presented, 0% coherent motion began in the two patches followed by coherent motion occurring in one of the two patches at a random time (flat hazard function) 500 to 8000 ms afterwards. The location of the coherent motion was cued to the monkey at the start of the trial with static dots and this cue was valid on 80% of the trials. Trials where the spatial cue was inside the neuron's receptive field were termed attend in (Figure 1A) trials while those where the spatial cue was outside were termed attend out (Figure 1B). The strength of coherent motion was varied between three levels (low, medium and high), with the monkey correctly detecting the coherent motion 50%, 90% and 99% of the time for the three levels. Valid trials comprised of all three coherence levels, while only medium-level coherence was used for invalid trials. The coherence levels for each experiment were adjusted depending on the eccentricity, speed and radius of the random dot patch in order to produce the target performance. 93 MT neurons and 104 VIP neurons were recorded in two monkeys.

In all experimental sessions, the direction and the speed of the coherent motion were matched to the preferred direction and speed of the neuron under study. The coherent motion lasted 750 ms and the monkey had to release the lever from 200 to 750 ms after the onset of coherent motion to obtain a juice reward. Trials where the monkey failed to release the lever or released the lever too late were deemed missed trials. Only correct and missed trials were included in the analysis. Trials where the monkey released the lever too early or was unable to maintain fixation were discarded from the analysis.

Visual stimulus

The animal sat 62 cm from a computer monitor (+/-17° ´ +/-13° of visual angle; 1600 ´ 1200 pixels; 75 Hz refresh). The stimuli consisted of two patches of white dots (each dot 0.25° diameter; 78 cd/m2) on a dark gray background (12 cd/m2) with a dot density of 2.1 dots/degree2. Each patch of dots was updated on every other video frame (approximately every 27 ms) using the following procedure. The dots in each patch were evenly divided into two groups. On each update, one group was replaced with new, randomly positioned dots, whereas dots in the other group were displaced by a fixed distance. The dots in this latter group determined the motion coherence. For 0% coherence, all the dots in this group moved a fixed distance in a random direction. For coherent motion greater than zero, a proportion of the dots moved with a fixed distance in the same direction. This proportion determined the strength of the coherent motion. On the next update (27 ms later), the groups were switched. This arrangement insured that all the dots had a lifetime of four video frames (i.e., two stimulus updates) before they were

replaced. Because half the dots are always randomly replotted regardless of the proportion of dots moving coherently, our motion had a maximum strength of 50% coherence. For example, at 25% coherent motion, half the dots are randomly replotted, one-quarter are moving with the same fixed distance and direction, and one-quarter are moving with the same fixed distance in a random direction.

Data collection

Using standard extracellular recording techniques (Gibson and Maunsell, 1997), recordings were made from well-isolated single neurons in area MT in both animals. When a neuron was isolated, the receptive field (RF) was mapped using a manually controlled bar while the animal fixated on a central spot. The diameter of the RFs ranged from 3.9 to 10.7° (median, 7.4). RF center eccentricities ranged from 3.9 to 11.1° (median, 7.9). The preferred speed was also judged using a bar moved by hand. The animals were trained to perform the task at slow or moderate motion speeds, so neurons with a preferred speed between 4 and 12°/sec were usually selected. Once the RF location, size, preferred direction, and speed were determined, the motion detection task was then run, and the neuron was recorded from for as long as possible. For some neurons, a memory saccade task was also run, but these data were not used in this analysis. The monkey's performance varied with patch location, size, and motion speed, which were determined by the response properties of the neuron under study. Consequently, different neurons were tested with different coherence levels. The animal's

eye position was measured every 5 ms using a scleral search and the times of action potentials were recorded to the nearest millisecond.

Detect probability, neurometric and spike ROC values

Detect probability (DP) expresses the ability to predict the behavioral outcome of a trial in a detection task given the neuronal response (Cook and Maunsell, 2002b). DP is analogous to the choice probability calculation used in discrimination tasks (Celebrini and Newsome, 1994; Britten et al., 1996). The neurometric value expresses the ability to predict which one of two stimuli was presented given the neuronal response. The spike ROC value expresses the ability to predict whether an action potential was generated given the preceding motion stimulus and the STA.

All three values are formulated in the same way using traditional ROC analysis (Green, 1975). Briefly, given two random samples from different distributions, the ROC value are the probabilities that one can correctly determine from which distribution the samples were drawn. In other words, it is a measure of how separate two distributions are from each other. Specifically, given two distributions of neuronal responses, $p(x)$ and $q(x)$, the maximum probability that we could correctly determine from which distribution a sample was drawn from is $P(correct) = \int_0^\infty \int_x^\infty p(x)q(y)dydx$, assuming that on average, a sample drawn from $q(y)$ is greater than a sample drawn from $p(x)$.

For detect probability, $q(y)$ corresponds to the distribution of spike rates for correct trials and $p(x)$ for failed trials. For the neurometric value, $p(x)$ and $q(y)$ are the distribution of spike rates corresponding to the 0% and coherent motion, respectively. For the spike ROC, $q(y)$ corresponds to the distribution of dot-product between the STA and motion stimulus at times that spikes occurred and $p(x)$ is the distribution of the same dot-product, but taken over all time points.

For the DP calculation, only neuronal responses occurring during low coherent motion trials from 50 to 150 ms after coherent motion onset were included for the attend in condition (except for Figure 3A which used a sliding window for the DP). We only used the low coherent motion trials to compute DP because this was the only condition that produced similar numbers of correct and failed trials. For the attend out trials only consisted of medium-level coherent motion. The medium-level coherent motion trials during the attend out condition also produced roughly similar amounts of correct and failed trials.

For the neurometric value, we compared the neuronal response from 50 to 150 ms after coherent motion onset to the neuronal response from 100 to 0 ms before coherent motion onset. Except for the value for the example neuron in Figure 2, the neurometric values for attend in conditions was calculated using low-level coherent motion trials. Attend out trials, by default, used medium-level coherent motion trials.

The spike ROC values used the segment from 200 ms after the start of the 0% coherent random motion up until coherent motion onset for either attend in or attend out

trials. The motion stimulus was reconstructed from the dot positions (see below) which allowed for the STA to be calculated for both the attend in and attend out conditions. Two distributions were then computed for both the attend in and out conditions: the dot-product between the STA and motion stimulus for all time points during 0% coherent motion, and the dot-product between the STA and motion stimulus only at times an action potential was generated during 0% coherent motion. The spike ROC was then calculated from these two distributions.

Neuron selection criteria for Figure 3

For Figure 3A, we required at least three correct and three failed trials for the attend in condition to accurately measure the detect probability. Additionally, we excluded neurons with baseline firing rates bellow 1 Hz for either attentional condition to accurately measure the spike ROC. For Figure 3B&C, where we compared correlations involving detect probability between attentional conditions, we required at least three correct and failed trials for both attentional conditions. As for Figure 3A, Figure 3C excluded neurons with baseline firing rates below 1 Hz for either attentional condition. There was no such restriction for Figure 3B, which measure the neurometric instead of the spike ROC.

Extracting global motion from the stimulus and the linear-nonlinear cascade

We reduced our random dot stimulus to a one dimensional time series of global motion strength along the neuron's preferred/null axis. To estimate the net motion between two successive frames of random dots, we applied a previously described method (Barlow and Tripathy, 1997) that computed the correspondence between dots in two sequential frames. We calculated the motion vector between each pair of dots in successive frames. Thus, if our patch contained N dots, we computed N^2 motion vectors for each pair of successive frames, where the magnitude of the vectors represented the speed. We then scaled each vector by passing its magnitude through a speed-filter. The speed-filter was a Gaussian distribution in the log speed domain (Nover et al., 2005) centered at the preferred speed of the neuron (determined by sweeping a bar through the receptive field) and with a standard deviation of either 0.1, 0.2 or 0.5 deg/sec. The speed-filter was used since speed-tuning profiles were not collected from the neurons. Spike ROCs (see above) were calculated for the three standard deviation for both the attend in and out conditions. The standard deviation that led to the largest spike ROC, for either attend in or out, was then used for all further calculations. Each weighted vector was then projected onto the neuron's preferred/null direction of motion axis. The sum of all the projected vectors was the motion strength at that time point. This calculation was repeated for every pair of successive random dot patches to create a time series of the motion strength in the preferred/null axis of the neuron. All time point between every pair of successive dot patches was assigned the same motion strength.

Results

The results of this analysis are based on a previously analyzed motion detection experiment (Cook and Maunsell, 2002b, a, 2004; Masse and Cook, 2008). Two monkeys were trained to detect a 750 ms coherent motion step in one of two diametrically opposite random dot patches (Figure 1). The likely location of the coherent motion step was cued to the monkey at the start of the trial. The cue was valid 75% of the time. One of these random dot patches was located in the receptive field of the neuron under study, and the direction and speed of the coherent motion were matched to the preferences of the neuron. Trials where the monkey was cued to the patch inside the neuron's receptive field are termed "attend in" trials while trials where the monkey was cued to the other patch are termed "attend out" trials.

The activity of 104 direction-selective neurons from ventral intraparietal area (VIP) and 93 neurons from middle temporal area (MT) was recorded. This study is primarily focused on the neural activity from VIP, however the results from MT will be discussed in the Discussion. VIP is a multimodal area located in the parietal cortex and is situated adjacent to the lateral intraparietal area (LIP). It receives strong inputs from area MT, and although it has not been causally linked to motion perception, its activity is strongly correlated to the perceptual decision in a motion detection task (Cook and Maunsell, 2002b) .

Figure 1 Motion detection task. The experiment used two moving random dot patches. Once the animals fixated, a static cue was presented to indicate the most likely position of where the coherent motion would occur. Next, 0% coherent motion was presented in both the neuron's RF and in a patch diametrically opposite. Coherent motion randomly began in one of the patches from 500 to 8000 ms with a flat hazard function. The coherent motion lasted for 750 ms and the animal had to release the lever from 200 to 750 ms after the coherent motion onset in order to obtain a reward. In 80% of the trials, the coherent motion would occur in the location spatially cued to the animal (A). In the other 20% of the trials, the cue would be invalid and the coherent motion would occur in the location not cued to the animal (B).

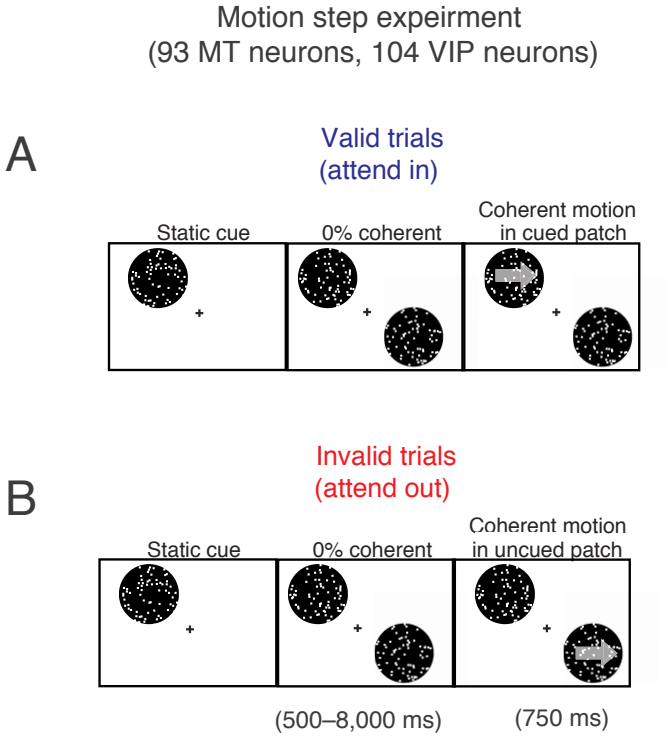
Motion step expeirment
(93 MT neurons, 104 VIP neurons)
A
Valid trials
(attend in)
Static cue
0% coherent
Coherent motion
in cued patch
B
Invalid trials
(attend out)
Static cue
0% coherent
Coherent motion
in uncued patch
(500–8,000 ms)
(750 ms)

Figure 1

Previous analysis of this data revealed that attending inside VIP neurons' receptive fields increased baseline neural activity by about 35%, significantly more than the approximate 15% increase in MT (Cook and Maunsell, 2002a). Additionally, by using ROC analysis to compare the neural response evoked by coherent motion to the neural response prior to coherent motion, it was found that VIP neurons were more sensitive to coherent motion when attention was directed inside the receptive field (Maunsell and Cook, 2002). The activity of an example VIP neuron, during both the attend in (Figure 2A) and attend out (Figure 2B) conditions, highlights these effects of attention. The black dots in the spike raster correspond to spikes during correct trials while the red dots are spikes from failed trials. The black curve just below gives the average response for correct and failed trials combined. Both figures show the neural response for trials where the same-strength coherent motion step was shown inside the neuron's receptive field.

Baseline activity for this neuron more than doubled during the attend in condition, which was not atypical for VIP neurons. Additionally, the neuron responded vigorously within 70 ms to the coherent motion step during the attend in condition, but barely responded within the first 150 ms after coherent motion onset during the attend out condition. We quantified the responsiveness of the neuron to the coherent motion step using the neurometric (NM) (see Methods). The neurometric is an ROC comparison between the spike count from 50 to 150 ms after coherent motion onset (light blue stripe, Figure 2A&B) to the spike count from 0 to 100 ms prior to the motion step (light gray stripe). Values above 0.5 indicate the neural response after motion onset was usually

Figure 2 Example VIP neuron. The neural response to medium-level coherent motion is shown for both the attend in condition (A), and the attend out condition (B). The black dots represent the time of action potentials that occurred during correct trials, while red dots are those from failed trials. The panels below the spike raster show the average neural response relative to coherent motion onset average over all correct and failed trials. The light blue (50 to 150 ms after coherent motion onset) and gray (0 to 100 ms prior) bars are the windows used to calculate the neurometric value. The light blue bar is also the window used to calculate the detect probability (except for Figure 3A which employed a sliding window for the detect probability). (C) The steps involved in calculating the linear-nonlinear cascade that was used to measure the neuron's sensitivity to the motion stimulus. The calculation was performed separately for attend in and attend out trials. From the motion stimulus (see Methods) and the spike train (Box 1), the spike triggered average was calculated (Box 2). From there, two distribution were formed. The first was the dot product of the STA and motion stimulus performed at all time points (black curve, Box 3). The other was formed by taking the dot product between the STA and the motion stimulus only at times action potentials were generated (gray curve, Box 3). The spike ROC was an ROC comparison between these two distributions (see Methods). Finally, dividing the two distributions gives the nonlinearity function, which maps the dot product between the STA and motion stimulus into the predicted spike rate (Box 4). (D) The nonlinearity for the both the attend in (blue curve) and attend out (red curve) conditions.

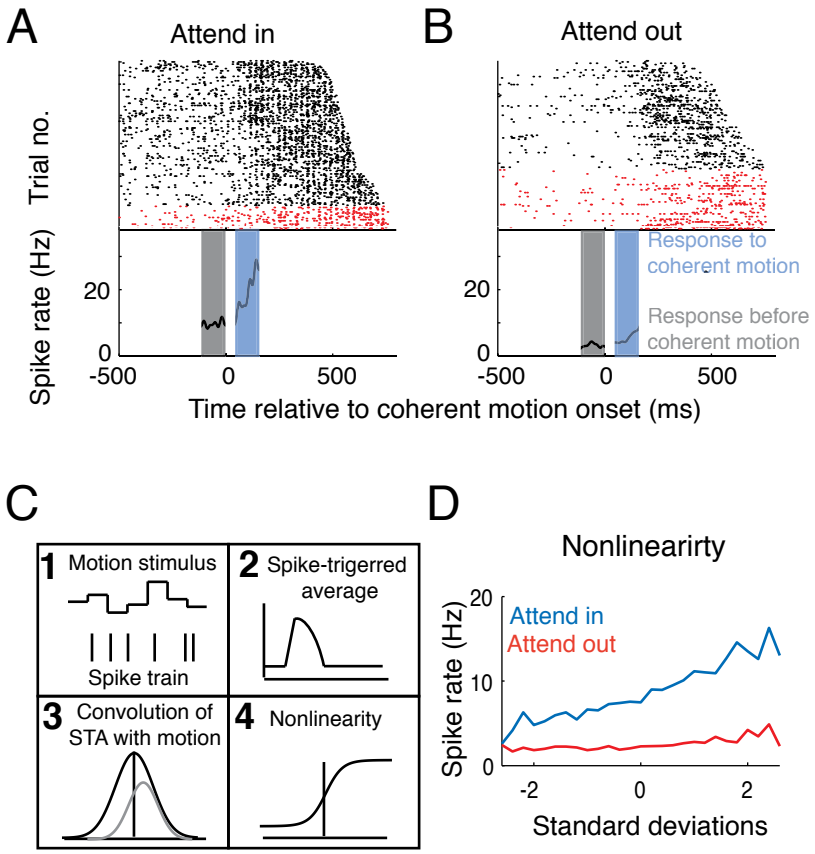
A
Attend in
B
Attend out
Trial no.
Spike rate (Hz)
20
0
-500
0
500
Response to coherent motion
Response before coherent motion
Time relative to coherent motion onset (ms)
C
1 Motion stimulus
Spike train
2 Spike-trigerred average
3 Convolution of STA with motion
4 Nonlinearity
D
Nonlinearirty
Attend in
Attend out
Spike rate (Hz)
20
10
0
-2
0
2
Standard deviations

Figure 2

greater the spike count before onset. For this example neuron, the neurometric was 0.678 for the attend in condition but only 0.537 for attend out.

We wanted to know whether this increase in responsiveness was at least partly due to a change in how the neuron encoded the stimulus, rather than a change in feedback, arousal, etc. To determine whether this was the case, we modeled the relation between the motion stimulus and the neural activity as a linear-nonlinear cascade (Paninski et al., 2004; Pillow et al., 2005; Rust et al., 2006). The cartoon in Figure 2C gives the steps involved in the calculation. From the motion stimulus during 0% coherent motion and the spike train (Box 1), we calculate the average motion preceding each spike, termed the spike triggered-average (STA) (Box 2). From the STA and the motion stimulus, we can form two distributions (Box 3). One is formed by taking the dot product between the STA and motion stimulus at all time points (black curve, Box 3), and the other from the dot product between the STA and the motion stimulus only at times an action potential was recorded (gray curve, Box 3). The separation between the two distributions gives a measure of how sensitive the neuron was for motion that matched its STA. We calculated the sensitivity of the neuron by performing ROC analysis between these two distributions, which we termed the spike ROC. Finally, by dividing these two distributions, one obtains the nonlinearity, also known as the gain function (Box 4). The nonlinearity is a function that maps the dot product of the STA and motion stimulus (x-axis) with the predicted spike rate (y-axis). Graphically, a greater slope indicates greater sensitivity.

These calculations were performed separately for the 0% coherent motion segment of the attend in and attend out trials. The nonlinearity of this example neuron is shown in Figure 2D for both the attend in condition (blue curve) and attend out condition (red curve). The slope of the attend in nonlinearity was much greater than the attend out nonlinearity, implying that sensitivity increased with spatial attention. In agreement, the spike ROC for the attend in condition was 0.572, compared to 0.539 for the attend out condition. This suggests that the greater neural response to coherent motion when attention was directed inside the receptive field was partly due to an increase in the sensitivity of the neuron.

The result that some VIP neurons increased sensitivity to the motion stimulus during the attend in condition was not very surprising. VIP is a multimodal area, and if attention increased the strength of motion signals arriving from area MT, this could possibly explain the observed effect. We really wanted to know whether the attentional change in sensitivity was related to how the neuron was weighted in forming the perceptual decision. Unfortunately, there exists no way of directly measuring how the activity of a neuron was weighted. Instead, we can exploit the intuitive notion that the more influential the activity of a neuron is in forming a decision, the more correlated it is with the behavioral outcome of the trial. This notion that the most heavily weighted neurons are most correlated with behavioral outcome has been supported by theoretical studies as well as by neural and psychophysical data (Britten et al., 1996; Dodd et al., 2001; Parker et al., 2002; Uka and DeAngelis, 2004; Purushothaman and Bradley, 2005; Jazayeri and Movshon, 2006, 2007b, a; Law and Gold, 2008).

In Figure 3A, we looked at whether the attentional change in sensitivity of the neuron was related to its correlation with the perceptual decision. For each time point relative to the coherent motion onset, we measured the spike count in the preceding 100 ms, and computed the ROC measure between correct trials and failed trials. This value has been termed the detect probability (Cook and Maunsell, 2002b), analogous to the choice probability used in discrimination tasks (Britten et al., 1996). If detect probability indicates how the activity of the neuron was weighted in forming the decision, then this calculation would show whether the attentional change in sensitivity was linked to how the neurons were weighted in forming the decision. The change in sensitivity was the spike ROC value for the attend in condition minus spike ROC of the attend out condition. Neurons had to have a minimum number of correct and failed trials to accurately measure the detect probability, as well as a minimum spike rate to measure the spike ROC (see Methods for the criteria used).

If the detect probability was only correlated to the sensitivity during the attend in state, and not the difference between attentional conditions, this might result in a false positive. Thus, we took the partial correlation between the detect probability and the change in sensitivity between attentional states after removing the contribution of the neural sensitivity during the attend in state. We found that from around 100 to 150 ms after coherent motion onset, the detect probability was significantly correlated with a positive change in sensitivity from attend out to attend in (Figure 3A). Thus, the neurons

Figure 3 The relation between neural sensitivity and detect probability. (A) In the left column, the Spearman's rank coefficient of the partial correlation between the detect probability and the difference in the spike ROC between attentional conditions given the spike ROC of the attend in condition. The spike ROCs were held fixed, while the detect probability compared total spike counts in the preceding 100 ms for each time point relative to coherent motion onset. (B) The associated P-value of the Spearman rank partial correlation coefficient. (B) The scatter plot of the neurometric value and the detect probability for each neuron for the attend in condition (left panel) and attend out conditions (right panel). The neurometric and detect probability values for the attend in condition used lo-level coherent motion trials while the attend out condition used medium-level coherent motion trials. The time windows used for these calculations were those given by the vertical bars in Figure 2A: 50 to 150 ms after coherent motion onset for the detect probability while the neurometric compared values from 50 to 150 ms after coherent motion onset to 0 to 100 ms prior. (C) The scatter plot of the spike ROC and detect probability values for the attend in (left panel) and attend out (right panel) conditions. The detect probability was calculated using the same windows and trials as stated above. The spike ROCs were calculated from the 0% coherent motion segment of all attend in or attend out trials.

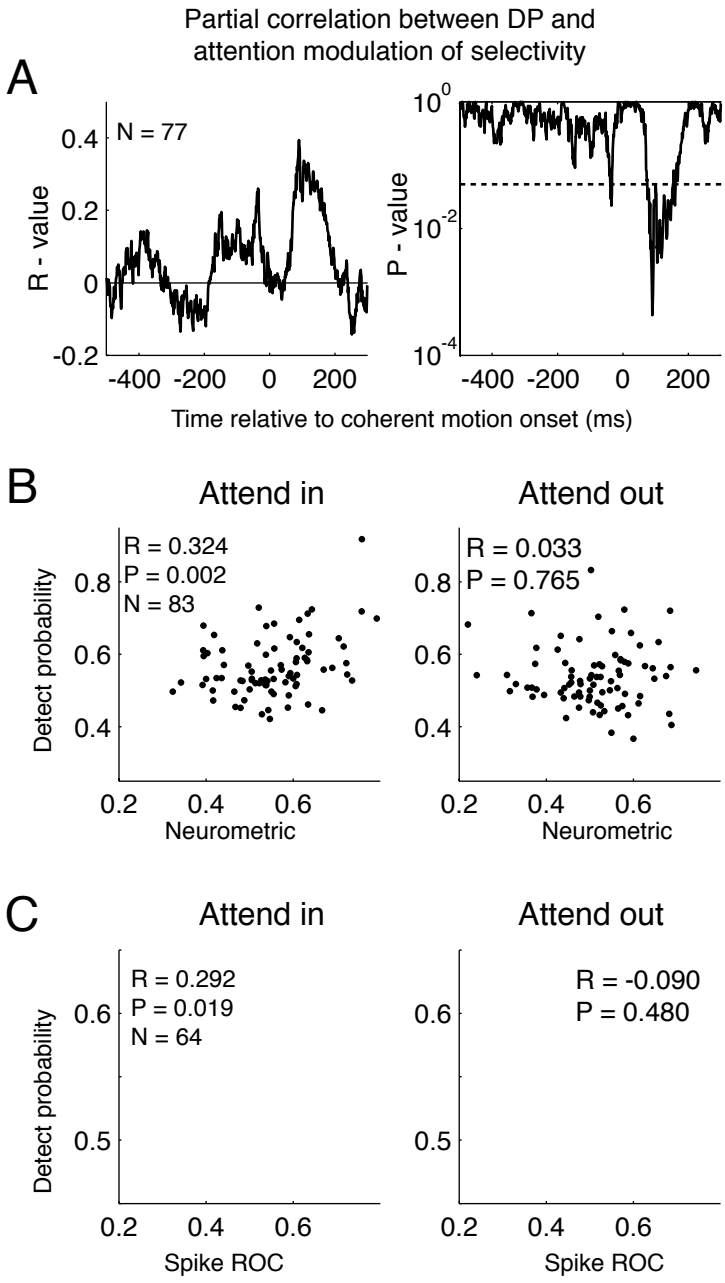
A
Partial correlation between DP and
attention modulation of selectivity
N = 77
R - value
P - value
Time relative to coherent motion onset (ms)
B
Attend in
Attend out
R = 0.324
P = 0.002
N = 83
R = 0.033
P = 0.765
Detect probability
Neurometric
C
Attend in
Attend out
R = 0.292
P = 0.019
N = 64
R = -0.090
P = 0.480
Detect probability
Spike ROC

Figure 3

more influential in forming the decision increased sensitivity when attention was directed inside their receptive fields.

We next wondered how this change in sensitivity affected the relation between how informative the neuron was and how its activity was weighted in forming a decision. We used two metrics to quantify how informative the neuron was: the neurometric that measured the neural response to the coherent motion (Figure 3B), and the spike ROC, that measured the sensitivity during 0% coherent motion (Figure 3C). We compared both metrics to the detect probability calculated using spikes from 50 to 150 ms after coherent motion onset. Both metrics were significantly correlated with the detect probability during the attend in condition, but were not correlated during the attend out condition. Thus, only when attention was directed inside a neuron's receptive field were the more informative neurons weighted more heavily.

When calculating both the neurometric and the detect probability, we only used the neural response up to 150 ms after coherent motion onset. This was because many VIP neurons showed a strong increase in neural activity just prior to a correct response, and we wanted to minimize its effect on our measurements. This increase can be clearly seen in the example neuron in Figure 4A, which showed a large increase in activity before a correct response even when no coherent motion was presented inside the neuron's receptive field. We quantified the strength of this increase by using ROC analysis to compare the spike count in a 100 ms window before lever release to the spike count from 100 to 200 ms after coherent motion onset. Although we cannot be totally

Figure 4 The increase in neural activity prior to a correct response. (A) An example VIP neuron showing all correct trials where coherent motion occurred outside the neuron's receptive field. The neurons showed a large increase in activity just prior to the lever release. (B) This increase was quantified by an ROC comparison between the spike count in the 100 ms prior to lever release to the spike count from 100 to 200 ms after coherent motion onset. The scatter plot shows a significant negative correlation between the ROC value and the neurometric calculated for the attend condition. (C) An example VIP neuron where spatial attention did not affect this increase in activity between the attend in (left panel) and attend out (right panel) conditions. (D) and (E) are two more example neurons where the increase in activity prior to a correct response was largely eliminated when spatial attention was outside the neuron's receptive field. (F) Histogram of the difference in this increase in activity prior to a correct response for medium-level coherence trials between the attend in and attend out conditions. The ROC value was calculated using the same windows described above. (G) The scatter plot of this increase in activity prior to a correct response when coherent motion was outside the neuron's receptive field and the attention modulation of baseline activity. The modulation in baseline activity was calculated as (AI-AO)/(AI+AO) where AI and AO represent the average spike rate during 0% coherent motion during attend and attend out trials, respectively.

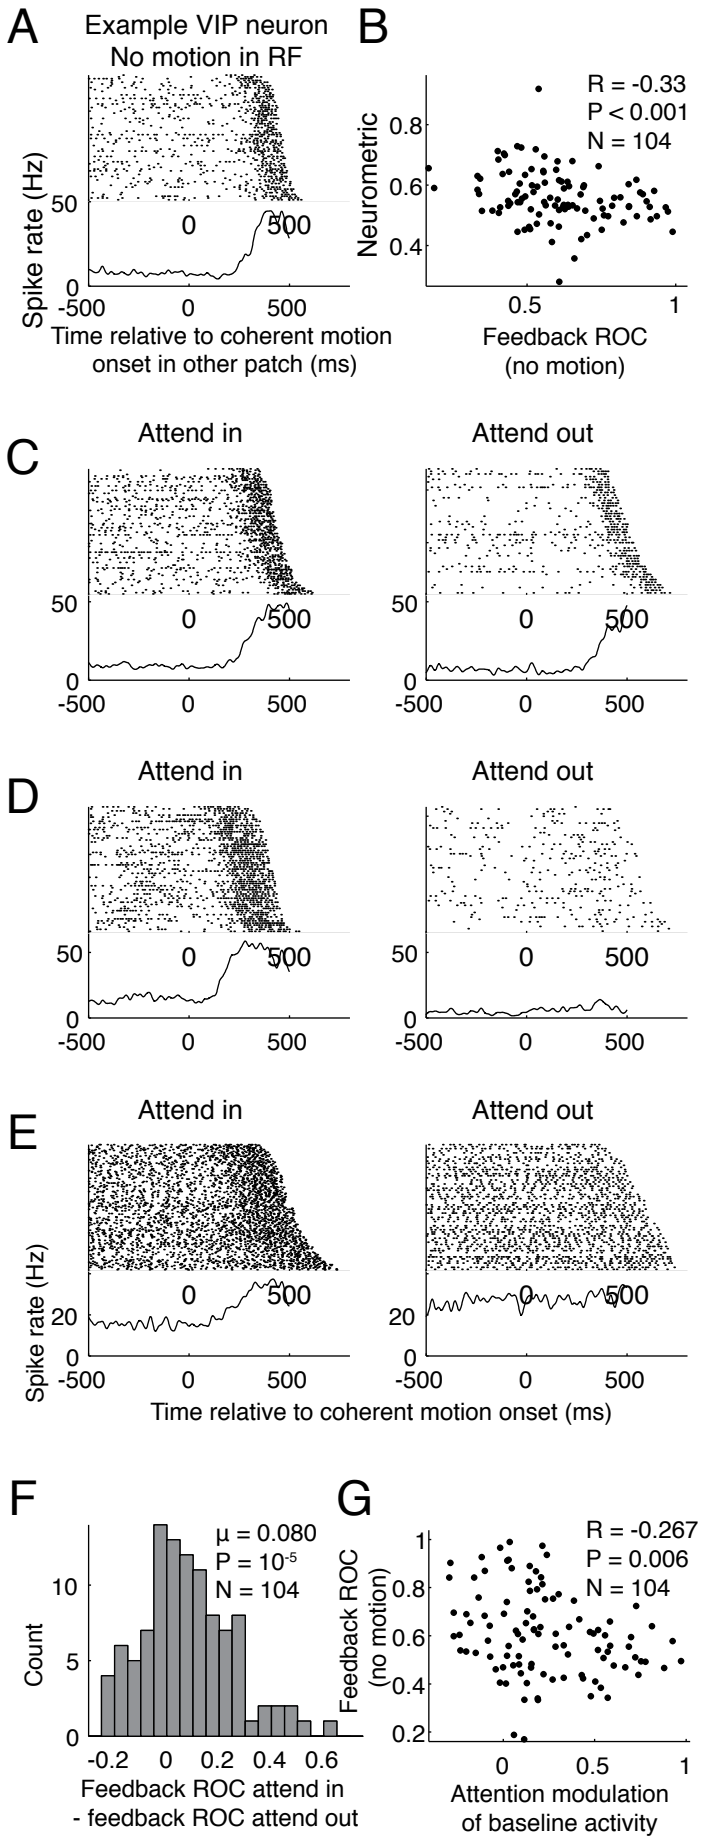
A
Example VIP neuron
No motion in RF
Spike rate (Hz)
Time relative to coherent motion onset in other patch (ms)
B
R = -0.33
P < 0.001
N = 104
Neurometric
Feedback ROC (no motion)
C
Attend in
Attend out
D
Attend in
Attend out
E
Attend in
Attend out
Spike rate (Hz)
Time relative to coherent motion onset (ms)
F
μ = 0.080
P = 10-5
N = 104
Count
Feedback ROC attend in - feedback ROC attend out
G
R = -0.267
P = 0.006
N = 104
Feedback ROC (no motion)
Attention modulation of baseline activity

Figure 4

certain of the source of this increase (see Discussion), we will simply refer to it as feedback. Interestingly, the strength of this feedback was negatively correlated with the neurometric (Figure 4B). Thus, the neurons least informative of the stimulus showed the largest increase in activity prior to a correct response.

Surprisingly, this feedback was significantly affected by the attentional state. Although attention did not affect this feedback in some neurons (Figure 4C), the feedback was all but eliminated for many neurons when attention was outside the receptive field (Figures 4D&E). Across the entire population, the feedback was significantly reduced during the attend out condition (Figure 4F). Additionally, the reduction was correlated with attentional modulation of the baseline spike rate (Figure 4G). Thus, the greater the increase in baseline neural activity during attend in, the smaller the feedback prior to a correct response during the attend out trials.

Discussion

In this study, we examined the neural mechanisms that possibly underlie how attention increases behavioral performance. We found that when attention was directed inside VIP neurons' receptive fields, the detect probability of the neuron was correlated with how informative the neurons was of the stimulus. This correlation was eliminated when attention was shifted outside a neuron's receptive field. This suggests that spatial attention acts to properly weigh sensory neural activity when forming a perceptual decision. We propose this as a possible explanation why spatial attention increases behavioral performance in visual tasks. Unrelated to this conclusion, we also found that attention can gait other top-down signal, implying that attention does not simply modify how bottom-up signals are integrated.

Results from area MT

In this experiment, we also recorded neurons from area MT. The analysis used above was difficult to apply to MT because of the nature of the neural response generated by the motion stimulus. The dots in the random dot patch were updated at a fairly slow rate (about once every 27 ms), which led to oscillations in the neural response of many MT neurons. Furthermore, we found in a previous study that neural responses at different phases of the oscillation varied in how informative they were of the stimulus (Masse and Cook, 2008). Additionally, the correlation between the perceptual decision and neural activity was also highly dependent on the phase, with the more informative phases more

correlated with the behavioral outcome. Importantly, the more informative phases were more correlated with the perceptual decision during the attend in, but not during the attend out condition (data not shown).

This variability in stimulus encoding and correlation with behavior across different phases made it difficult to determine how the activity was combined between phases. Thus, we were unsure whether our neurometric and detect probability calculations that count spikes over a period of 100 ms accurately reflects how activity is processed in the brain. Although the same argument can be made with the neural response from VIP, these motion-induced oscillations in the neural response were much less pronounced. Additionally, our results were invariant to small changes to our window or to whether spikes were convolved with a box filter or exponential filter.

The various top-down signals

Previous studies have found that, on average, spatial attention seemed to act in a multiplicative manner on the neural response, without affecting neural sensitivity (McAdams and Maunsell, 1999; Treue and Martinez Trujillo, 1999). What was surprising was that we could not find any correlation between the attentional modulation of neural activity in response to 0% coherent motion and any change in detect probability or neuronal sensitivity. We did however find a significant correlation between an increase in neural activity and a strong decrease in this feedback prior to a correct response a neuron receives during attend out trials. Thus, while the strength of this attentional modulation

on the firing rate affects other forms of top-down feedback, we do not know what role in plays, if any, in changing neuronal sensitivity or how activity is weighted downstream.

In Figure 4, we showed that many VIP neurons increased neural activity just prior to a correct response, and that this increase was attenuated during the attend out condition. Since the occurrence of this increase in activity was time-locked to the lever release, and that the least informative neurons showed the greatest increase (Figure 4B), this suggested that this increase in activity was encoding a motor response. However, it would be hard to believe the motor response would be so affected by spatial attention. This suggests that the increase in activity prior to a correct response was the result of top-down feedback. We do not know what functional role this feedback might serve. If feedback was the source of the increase, then either attending outside the neuron's receptive field prevented this feedback from being sent to VIP, or the lack of attention gaited this feedback at the level of VIP. Either way, it might not suffice to think of attention as simply a top-down signal that modulates how sensory information is processed.

Top-down signals, neural encoding and decoding

Many have suggested that top-down signals can modify not just gain, but how a neuron encodes different visual features. Anesthesia, which likely reduces top-down signals to sensory areas, inhibits neurons from encoding global features of a stimulus (Lamme et al., 1998; Pack et al., 2001). The behavioral demands of a task, most likely

mediated through top-down signaling, modifies neural responses in V1 (Crist et al., 2001; Li et al., 2004). Additionally, attending to different stimuli within a neuron's receptive field biases the neuron to selectively encode the attended stimulus (Moran and Desimone, 1985; Reynolds et al., 1999).

If top-down signals alter the relationship between the stimulus and the neural response, the one must assume that this change is accompanied by a change in how the neural response is decoded. Many studies have proposed that neural activity is weighted differently depending on how well the neuron encodes the stimulus given the demands of the task (Britten et al., 1996; Dodd et al., 2001; Parker et al., 2002; Uka and DeAngelis, 2004; Purushothaman and Bradley, 2005; Jazayeri and Movshon, 2006, 2007b, a; Law and Gold, 2008). Since the most informative neurons are dependent on the demands of the task, these results imply that top-down signals are necessary to change how neural responses are decoded. Thus, it is not very surprising that attention can change both how sensory information is encoded and how the subsequent neural response is processed. Although we cannot be certain that the increase in behavioral performance during the attend in trials was the result of properly weighting sensory neurons, it has been shown that during training in a perceptual task, as performance improves, so does the correlation between the neuron's sensitivity and its choice probability (Law and Gold, 2008). All together, the results of this study highlight the need to not just consider how attention modifies neural response, but how attention modifies how neural responses are transformed between cortical areas.

Attention and the bottleneck in visual processing

It has been widely reported that subjects can only attend to so many objects before performance is compromised (Desimone and Duncan, 1995; Kastner and Ungerleider, 2000). What is unknown is the nature of the bottleneck that prevents subject from processing information from multiple objects or locations. In our study, subjects were trained to respond to coherent motion whether it was validly or invalidly cued. Thus, one would believe it to be advantageous for the most informative neurons to be weighted accordingly regardless of where spatial attention was directed. Since this was not the case, one wonders whether this represented a bottleneck in visual processing. More specifically, perhaps the brain is not capable of properly weighing the neural activity representing information from multiple stimuli, especially if this neural activity all converges to a common downstream target.

Possible mechanisms

What might be the neural mechanism that allows spatial attention to properly match the neural sensitivity to its appropriate weight? Since it would necessarily involve coordinating activity across cortical regions, a mechanism involving oscillations in the local field potential, which can be coherent across separate cortical areas (Engel et al., 1990; Donoghue et al., 1998; Destexhe et al., 1999), could serve as a possible mechanism. If incoming activity to a neuron that was highly informative of the stimulus was temporally offset from incoming activity that was less informative, oscillations in the

local field potential could serve to selectively increase the impact (or weight) of the most informative inputs. As a result, the spiking activity of that neuron would become more sensitive to the stimulus. Attention has already been shown to synchronize spiking activity to the gamma frequency range of the local field potential (Fries et al., 2001), and increase synchrony between parietal and frontal areas (Buschman and Miller, 2007). Thus, it would be interesting to know whether the degree of synchrony was related to how informative the neuron was of the stimulus given the task.

REFERENCES

Baldassi S, Burr DC (2000) Feature-based integration of orientation signals in visual search. Vision Res 40:1293-1300.

Balz GW, Hock HS (1997) The effect of attentional spread on spatial resolution. Vision Res 37:1499-1510.

Bashinski HS, Bacharach VR (1980) Enhancement of perceptual sensitivity as the result of selectively attending to spatial locations. Percept Psychophys 28:241-248.

Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis Neurosci 13:87-100.

Buschman TJ, Miller EK (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. Science 315:1860-1862.

Carrasco M, Yeshurun Y (1998) The contribution of covert attention to the set-size and eccentricity effects in visual search. J Exp Psychol Hum Percept Perform 24:673-692.

Carrasco M, Penpeci-Talgar C, Eckstein M (2000) Spatial covert attention increases contrast sensitivity across the CSF: support for signal enhancement. Vision Res 40:1203-1215.

Cook EP, Maunsell JH (2002a) Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. J Neurosci 22:1994-2004.

Cook EP, Maunsell JH (2002b) Dynamics of neuronal responses in macaque MT and VIP during motion detection. Nat Neurosci 5:985-994.

Cook EP, Maunsell JH (2004) Attentional modulation of motion integration of individual neurons in the middle temporal visual area. J Neurosci 24:7964-7977.

Crist RE, Li W, Gilbert CD (2001) Learning to see: experience and attention in primary visual cortex. Nat Neurosci 4:519-525.

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

Destexhe A, Contreras D, Steriade M (1999) Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. J Neurosci 19:4595-4608.

Dodd JV, Krug K, Cumming BG, Parker AJ (2001) Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. J Neurosci 21:4809-4821.

Donoghue JP, Sanes JN, Hatsopoulos NG, Gaal G (1998) Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. J Neurophysiol 79:159-173.

Dosher BA, Lu ZL (1998) Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. Proc Natl Acad Sci U S A 95:13988-13993.

Downing CJ (1988) Expectancy and visual-spatial attention: effects on perceptual quality. J Exp Psychol Hum Percept Perform 14:188-202.

Engel AK, Konig P, Gray CM, Singer W (1990) Stimulus-Dependent Neuronal Oscillations in Cat Visual Cortex: Inter-Columnar Interaction as Determined by Cross-Correlation Analysis. Eur J Neurosci 2:588-606.

Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. Science 291:1560-1563.

Jacobson H (1951) The informational capacity of the human eye. Science 113:292-293.

Jazayeri M, Movshon JA (2006) Optimal representation of sensory information by neural populations. Nat Neurosci 9:690-696.

Jazayeri M, Movshon JA (2007a) A new perceptual illusion reveals mechanisms of sensory decoding. Nature 446:912-915.

Jazayeri M, Movshon JA (2007b) Integration of sensory evidence in motion discrimination. J Vis 7:7 1-7.

Kastner S, Ungerleider LG (2000) Mechanisms of visual attention in the human cortex. Annu Rev Neurosci 23:315-341.

Lamme VA, Zipser K, Spekreijse H (1998) Figure-ground activity in primary visual cortex is suppressed by anesthesia. Proc Natl Acad Sci U S A 95:3263-3268.

Law CT, Gold JI (2008) Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat Neurosci 11:505-513.

Li W, Piech V, Gilbert CD (2004) Perceptual learning and top-down influences in primary visual cortex. Nat Neurosci 7:651-657.

Lu ZL, Dosher BA (1998) External noise distinguishes attention mechanisms. Vision Res 38:1183-1198.

Masse NY, Cook EP (2008) The effect of middle temporal spike phase on sensory encoding and correlates with behavior during a motion-detection task. J Neurosci 28:1343-1355.

Maunsell JH, Cook EP (2002) The role of attention in visual processing. Philos Trans R Soc Lond B Biol Sci 357:1063-1072.

McAdams CJ, Maunsell JH (1999) Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. J Neurosci 19:431-441.

Moran J, Desimone R (1985) Selective attention gates visual processing in the extrastriate cortex. Science 229:782-784.

Morgan MJ, Ward RM, Castet E (1998) Visual search for a tilted target: tests of spatial uncertainty models. Q J Exp Psychol A 51:347-370.

Muller HJ, Humphreys GW (1991) Luminance-increment detection: capacity-limited or not? J Exp Psychol Hum Percept Perform 17:107-124.

Pack CC, Berezovskii VK, Born RT (2001) Dynamic properties of neurons in cortical area MT in alert and anaesthetized macaque monkeys. Nature 414:905-908.

Paninski L, Pillow JW, Simoncelli EP (2004) Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model. Neural Comput 16:2533-2561.

Parker AJ, Krug K, Cumming BG (2002) Neuronal activity and its links with the perception of multi-stable figures. Philos Trans R Soc Lond B Biol Sci 357:1053-1062.

Pillow JW, Paninski L, Uzzell VJ, Simoncelli EP, Chichilnisky EJ (2005) Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. J Neurosci 25:11003-11013.

Purushothaman G, Bradley DC (2005) Neural population code for fine perceptual decisions in area MT. Nat Neurosci 8:99-106.

Reynolds JH, Chelazzi L, Desimone R (1999) Competitive mechanisms subserve attention in macaque areas V2 and V4. J Neurosci 19:1736-1753.

Rust NC, Mante V, Simoncelli EP, Movshon JA (2006) How MT cells analyze the motion of visual patterns. Nat Neurosci 9:1421-1431.

Treue S, Maunsell JH (1996) Attentional modulation of visual motion processing in cortical areas MT and MST. Nature 382:539-541.

Treue S, Maunsell JH (1999) Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. J Neurosci 19:7591-7602.

Treue S, Martinez Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. Nature 399:575-579.

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

Yeshurun Y, Carrasco M (1999) Spatial attention improves performance in spatial resolution tasks. Vision Res 39:293-306.

Chapter 6

Summary and future directions

Understanding how the brain generates conscience perception has been arguably one of the most important, and hardest, problems in neuroscience. One approach that has yielded much success has been to record neural activity from sensory areas in behaving animals while they performed simple perceptual tasks. These tasks, which usually involve the animal trying to detect a stimulus or discriminate between two simple alternatives, provides a quantitative measure of the perception of the animal through its behavioral response. Importantly, both neural activity and the behavioral response are recorded simultaneously, allowing for a direct comparison between the two. These types of experiments led to the discovery that the activity of a single sensory neuron can be correlated to the perceptual decision of the animal (Britten et al., 1996).

The finding that neural activity from sensory neurons can be correlated with perception is important for two reasons. Firstly, it places constraints on the neural population responsible for generating the perception (Bair et al., 2001; Zohary et al., 1994): either the population size has to be relatively small or the response of neurons within the population must be correlated. Current evidence suggests that the later prediction is correct, but does not exclude the former. Secondly, this correlation between activity and perception provides a window on how the activity from a neuron is used by the brain in forming a perceptual decision. If the activity of a neuron is highly correlated with the decision, then it is assumed that the activity was highly influential, or heavily weighted, in forming the decision.

One of the main findings that resulted from this notion is that the neurons that are more informative of the signal are also more influential in forming a decision. This finding has been reported in several experimental studies (Britten, 1992; Celebrini and Newsome, 1994; Cook and Maunsell, 2002b; Purushothaman and Bradley, 2005; Uka and DeAngelis, 2004), and is supported by theoretical (Jazayeri and Movshon, 2006) and psychophysical work (Jazayeri and Movshon, 2007). At one level, this result is almost a truism. For example, it is known that neurons from auditory cortex would not influence visual perceptual decisions. However, this relationship holds amongst groups of neurons whose preferences match the properties of the stimulus. This suggests that the ability to properly weigh sensory neural activity can function on a fine level.

It is in the author's opinion that the ability to measure how neural activity from different neurons is weighted when forming a decision will play a key role in determining how neural activity underlies perception. In this thesis, the author has presented two examples (Chapters 3 and 5) of how this measure can help us understand the rules by which sensory neural activity is transformed into perception. Additionally, it also raises interesting questions on the mechanisms involved in properly weighing sensory activity. However, this all relies upon that the correlation between neural activity and perceptual decisions is an inherent feature in decision-making, and not simply an artifact. For example, measuring the correlation between activity and behavior requires using an equivalent stimulus across trials. However, since many studies employ stochastic stimuli, it is possible that the observed correlation between neural activity and perception is simply due to fluctuations in the stimulus. This was ruled out in one study (Cook and

Maunsell, 2002b), but other possibilities remain. In Chapter 2, it was examined whether small, involuntary eye movements termed microsaccades could explain the correlation between activity and perception. Consistent with past studies (Bair and O'Keefe, 1998; Martinez-Conde et al., 2000, 2004), it was found that microsaccades affect both visual motion perception and neural activity in three motion sensitive areas: MT, VIP and LIP. Consequently, microsaccades added a small, yet significant, amount to the correlation between neural activity and perception. This study revealed the importance of considering microsaccades when an accurate measure of the correlation between neural activity and perception is required. However, since microsaccades added a relatively small amount to this correlation, up to 20%, it strengthens the belief that the correlation is truly the result of how the brain transforms neural activity into perception. Thus, existing results that used the correlation between neural activity and perception are still valid.

Given this, Chapter 3 is the first of two examples where the author uses the correlation between activity and behavior to further understand the neural basis of perception. Neural activity is not simply integrated across neurons when forming a decision, but is also integrated over time. Many past studies have shown that the information contained in a sensory neural response can change over time (Chacron et al., 2005; de Ruyter van Steveninck et al., 1997; Osborne, 2004; Pack and Born, 2001; Sadeghi et al., 2007), however it is poorly understood whether the brain can extract this information contained in the temporal properties of the response. The study in Chapter 3 examined data from a motion detection experiment in which the motion stimulus generated oscillations (one cycle = 27 ms) in the response of MT neurons. The study

showed that different phases of the oscillation differed in how reliably they encoded the motion signal. Crucially, the phases in which the stimulus was reliably encoded were the same phases that were more correlated with the perceptual decision. This suggests that the brain is capable of recognizing a highly dynamic neural response, and weighing the response at different times accordingly.

The conclusion of this study was supported by two recent studies (Engineer et al., 2008; Gu et al., 2007). Given that it widely known that decision-making can occur on fast time scales (Thorpe et al., 2001; VanRullen et al., 2005), and that information can be contained in the temporal structure of the neural response, it is comforting to know that the brain is capable of properly weighing a dynamic response. The obvious question that this raises is what are the mechanisms that allow the brain to properly weigh information at different times. The role of oscillations in the local field potential, important in temporally coordinating activity across different brain areas in many systems (Donoghue et al., 1998; Fries et al., 2001; Gray et al., 1992; Perez-Orive et al., 2002), is poorly understood in area MT. Further experiments are obviously required to understand the mechanisms behind the results of this chapter and that of Chapter 5 (see below).

One weakness of this study was that it simply measured correlation, and not causality, between information at different times (or phases in this case) in the response and the perceptual decision. To demonstrate causality, one requires tools such as electrical microstimulation to perturb the information contained in different times in the

response (Salzman et al., 1990). However, it is unknown whether electrical microstimulation has the necessary temporal precisions to probe causality on these timescales. In Chapter 4, we thus examined the time course of the effect of microstimulation on visual motion perception compared to an equivalent motion stimulus. We found that microstimulation in area MT has a much longer effect on motion perception compared to an equivalent visual stimulus, raising doubts as to whether microstimulation has the necessary temporal resolution to probe these timescales or to serve in cortical sensory prosthetics. However, new methods exists that could supplant microstimulation as the method of choice to perturb activity. Photostimulating neurons expressing the channelrhodopsin protein can achieve temporal and spatial precision significantly greater than microstimulating neural tissue (Nagel et al., 2002; Zhang et al., 2006). However, its use is currently limited to those animals for which it is feasible to generate transgenic versions. Thus, it may become advantageous to study the neural mechanisms of decision-making in animals such as mice and *Drosophila* (see below).

Finally, Chapter 5 explored the limits of how well the brain can properly weigh sensory neural activity. Given that the neurons that provide the most reliable signals can change from task to task, the brain must be capable of dynamically switching how it weighs sensory activity. This process may not be trivial, and the brain may only be capable of properly weighing a certain number of neurons or from only a few locations in the brain. We examined data from a motion detection experiment where the spatial attention of the animal was directed towards the location likely to contain the motion

stimulus. In the past, studies on the neural correlates of attention have focused on the effect on the gain of the neural response (Cook and Maunsell, 2002a; McAdams and Maunsell, 2000; Treue and Martinez Trujillo, 1999). Instead, we examined whether the animal was able to properly weigh sensory activity from neurons focused on the attended and unattended location. Even though motion could occur in either the attended or the unattended location, we found that the brain was able to properly weigh VIP activity only when attention was directed inside the neuron's receptive field. This suggests that the brain is limited in how it can weigh sensory activity, and may represent a bottleneck in visual processing. Furthermore, the ability to properly weigh sensory activity could explain the increased behavioral performance with attention.

As with Chapter 3, the results of this chapter naturally lead to questions about the mechanisms involved. Obviously, the mechanism could occur downstream of sensory areas where sensory activity would be integrated Alternatively, sensory areas could encode sensory information differently so that the activity of the most informative neurons is maximally effective downstream. For example, it was discussed above and in the Introduction that the correlation between neural activity and perception is likely because neural activity amongst different sensory neurons is correlated. This may not be an artifact; correlations between different sensory neurons may serve to increase or decrease the impact of their activity in downstream targets. Synchronous firing between neurons is already known as one such a mechanism that accomplishes this (Azouz, 2003; Perez-Orive et al., 2002); however further experimental and theoretical studies would be

required to determine whether correlations between sensory neurons on longer timescales could serve the same purpose.

Future directions

How do we proceed to further understand how neural activity generates conscience perception? There are many reasons why primate electrophysiology will continue be an important part of the answer. As our closest relative, studies involving primates are our best hope of shedding light on human neural function. Secondly, it is possible to train primates in a larger variety of more complex behavioral tasks than would be possible lower species. This permits the simultaneous recording of neural activity and perceptual decisions under a much larger variety of behavioral states. Finally, new recording technology is allowing scientists to measure the activity of tens to hundreds of neurons simultaneously, allowing observation of how populations of neurons contribute to perception (Brown et al., 2004).

However, certain question exists that may be difficult to answer using primate models. Experiments on primates are currently limited to one of two types: correlating cortical activity with perceptual choice, or perturbing cortical activity through electrical stimulation and measuring the effect on behavior. Although these studies have been very influential in shaping current thinking about decision making in the brain, they say relatively little about how this is implemented in neural hardware. The highly complicated nature of the primate cortex is undoubtedly a major stumbling block along

with the inability to selectively record or manipulate specific neurons within these circuits.

To actually determine the neural mechanisms that lead to perception and behavior, it may be advantageous to turn to simpler models. We will discuss one model system that is relatively simple compared to primate models and allows for a wide range of experimental techniques: the olfactory system of *Drosophila*. This preparation allows for patch-clamp recording of targeted neurons (Wilson et al., 2004) as well as high-resolution neural anatomy (Jefferis et al., 2007). These two methods are currently impossible to carry out in primates and may prove vital in understanding the neural computations that lead to perception.

In the olfactory system of *Drosophila*, which bears remarkable anatomical similarities with both other insects and mammals, odors first activate olfactory receptor neurons, of which there are about 50 different classes in *Drosophila*. These send axons into the antennal lobe, where second-order neurons transmit the information to two different higher order centers: the mushroom body and the lateral horn. Several ablation studies have shown that the mushroom body is necessary for olfactory learning (Dubnau et al., 2001) while, by exclusion, the lateral horn appears to be required for more "innate" olfactory-generated behavior (Heimbeck et al., 2001).

A perceptual decision is the product of two ingredients: bottom-up sensory information interacting with top-down signals of the behavioral state of the animal. While

much progress has been made understanding the bottom-up odor encoding (Turner et al., 2008; Wilson et al., 2004), little is known about how these neural signals react with top-down information. The ablation study above suggests that the lateral horn might be the first stage where bottom-up sensory information is converted into a behavioral response.

To uncover how perceptual decisions are formed, one approach would be to keep the stimulus constant and modify the behavioral state (much like attention studies). For example, the male pheromone cVA could be presented to unmated (sexually receptive) and mated (unreceptive) females. cVA is produced by male flies and inhibits courtship by other males while increasing female sexual arousal (Kurtovic et al., 2007). This pheromone produces different behavioral responses in the two sets of females. In such a setup, it may be possible to isolate the top-down signals that transform sensory neural activity into a behavioral response. Although this is but one example, it highlights the possibility of using simple systems such as the olfactory system of *Drosophila* to uncover basic mechanisms involved in forming perceptual decisions. Whether the mechanisms uncovered in these organisms would be applicable to higher-order species would be another question. Even if the rules governing how top-down signals convert sensory information into a behavioral response are completely different amongst different species, then it is still important that we understand why. Regardless, it is highly likely that a multitude of different approaches will needed to understand how neural activity underlies perception.

REFERENCES

Azouz, R., Gray, C. G. (2003). Adaptive Coincidence Detection and Dynamic Gain Control in Visual Cortical Neurons In Vivo. Neuron *37*, 513-523.

Bair, W., and O'Keefe, L.P. (1998). The influence of fixational eye movements on the response of neurons in area MT of the macaque. Vis Neurosci *15*, 779-786.

Bair, W., Zohary, E., and Newsome, W.T. (2001). Correlated firing in macaque visual area MT: time scales and relationship to behavior. J Neurosci *21*, 1676-1697.

Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis Neurosci *13*, 87-100.

Britten, K.H., Shadlen, M. N., Newsome, W. T., Movshon, J. A. (1992). The analysis of visual motion: a comparaison of neuronal and psychophysical performance. Journal of Neuroscience *12*, 4745-4765.

Brown, E.N., Kass, R.E., and Mitra, P.P. (2004). Multiple neural spike train data analysis: state-of-the-art and future challenges. Nat Neurosci *7*, 456-461.

Celebrini, S., and Newsome, W.T. (1994). Neuronal and psychophysical sensitivity to motion signals in extrastriate area MST of the macaque monkey. J Neurosci *14*, 4109-4124.

Chacron, M.J., Maler, L., and Bastian, J. (2005). Electroreceptor neuron dynamics shape information transmission. Nat Neurosci *8*, 673-678.

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

Cook, E.P., and Maunsell, J.H. (2002b). Dynamics of neuronal responses in macaque MT and VIP during motion detection. Nat Neurosci *5*, 985-994.

de Ruyter van Steveninck, R.R., Lewen, G.D., Strong, S.P., Koberle, R., and Bialek, W. (1997). Reproducibility and variability in neural spike trains. Science *275*, 1805-1808.

Donoghue, J.P., Sanes, J.N., Hatsopoulos, N.G., and Gaal, G. (1998). Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. J Neurophysiol *79*, 159-173.

Dubnau, J., Grady, L., Kitamoto, T., and Tully, T. (2001). Disruption of neurotransmission in Drosophila mushroom body blocks retrieval but not acquisition of memory. Nature *411*, 476-480.

Engineer, C.T., Perez, C.A., Chen, Y.H., Carraway, R.S., Reed, A.C., Shetake, J.A., Jakkamsetti, V., Chang, K.Q., and Kilgard, M.P. (2008). Cortical activity patterns predict speech discrimination ability. Nat Neurosci *11*, 603-608.

Fries, P., Neuenschwander, S., Engel, A.K., Goebel, R., and Singer, W. (2001). Rapid feature selective neuronal synchronization through correlated latency shifting. Nat Neurosci *4*, 194-200.

Gray, C.M., Engel, A.K., Konig, P., and Singer, W. (1992). Synchronization of oscillatory neuronal responses in cat striate cortex: temporal properties. Vis Neurosci *8*, 337-347.

Gu, Y., DeAngelis, G.C., and Angelaki, D.E. (2007). A functional link between area MSTd and heading perception based on vestibular signals. Nat Neurosci *10*, 1038-1047.

Heimbeck, G., Bugnon, V., Gendre, N., Keller, A., and Stocker, R.F. (2001). A central neural circuit for experience-independent olfactory and courtship behavior in Drosophila melanogaster. Proc Natl Acad Sci U S A *98*, 15336-15341.

Jazayeri, M., and Movshon, J.A. (2006). Optimal representation of sensory information by neural populations. Nat Neurosci *9*, 690-696.

Jazayeri, M., and Movshon, J.A. (2007). A new perceptual illusion reveals mechanisms of sensory decoding. Nature *446*, 912-915.

Jefferis, G.S., Potter, C.J., Chan, A.M., Marin, E.C., Rohlfing, T., Maurer, C.R., Jr., and Luo, L. (2007). Comprehensive maps of Drosophila higher olfactory centers: spatially segregated fruit and pheromone representation. Cell *128*, 1187-1203.

Kurtovic, A., Widmer, A., and Dickson, B.J. (2007). A single class of olfactory neurons mediates behavioural responses to a Drosophila sex pheromone. Nature *446*, 542-546.

Martinez-Conde, S., Macknik, S.L., and Hubel, D.H. (2000). Microsaccadic eye movements and firing of single cells in the striate cortex of macaque monkeys. Nat Neurosci *3*, 251-258.

Martinez-Conde, S., Macknik, S.L., and Hubel, D.H. (2004). The role of fixational eye movements in visual perception. Nat Rev Neurosci *5*, 229-240.

McAdams, C.J., and Maunsell, J.H. (2000). Attention to both space and feature modulates neuronal responses in macaque area V4. J Neurophysiol *83*, 1751-1755.

Nagel, G., Ollig, D., Fuhrmann, M., Kateriya, S., Musti, A.M., Bamberg, E., and Hegemann, P. (2002). Channelrhodopsin-1: a light-gated proton channel in green algae. Science *296*, 2395-2398.

Osborne, L.C., Bialek, W., Lisberger, S. G. (2004). Time Course of Information about Motion Direction in Visual Area MT of Macaque Monkeys. Journal of Neuroscience *24*, 3210-3222.

Pack, C.C., and Born, R.T. (2001). Temporal dynamics of a neural solution to the aperture problem in visual area MT of macaque brain. Nature *409*, 1040-1042.

Perez-Orive, J., Mazor, O., Turner, G.C., Cassenaer, S., Wilson, R.I., and Laurent, G. (2002). Oscillations and sparsening of odor representations in the mushroom body. Science *297*, 359-365.

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

Sadeghi, S.G., Chacron, M.J., Taylor, M.C., and Cullen, K.E. (2007). Neural variability, detection thresholds, and information transmission in the vestibular system. J Neurosci *27*, 771-781.

Salzman, C.D., Britten, K.H., and Newsome, W.T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. Nature *346*, 174-177.

Thorpe, S., Delorme, A., and Van Rullen, R. (2001). Spike-based strategies for rapid processing. Neural Netw *14*, 715-725.

Treue, S., and Martinez Trujillo, J.C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. Nature *399*, 575-579.

Turner, G.C., Bazhenov, M., and Laurent, G. (2008). Olfactory representations by Drosophila mushroom body neurons. J Neurophysiol *99*, 734-746.

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

VanRullen, R., Guyonneau, R., and Thorpe, S.J. (2005). Spike times make sense. Trends Neurosci *28*, 1-4.

Wilson, R.I., Turner, G.C., and Laurent, G. (2004). Transformation of olfactory representations in the Drosophila antennal lobe. Science *303*, 366-370.

Zhang, F., Wang, L.P., Boyden, E.S., and Deisseroth, K. (2006). Channelrhodopsin-2 and optical control of excitable cells. Nat Methods *3*, 785-792.

Zohary, E., Shadlen, M.N., and Newsome, W.T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. Nature *370*, 140-143.