

Contribution of Area MT to Stereoscopic Depth Perception: Choice-Related Response Modulations Reflect Task Strategy

Takanori Uka¹ and Gregory C. DeAngelis*

Department of Anatomy and Neurobiology
Washington University School of Medicine
Box 8108
660 South Euclid Avenue
St. Louis, Missouri 63110

Summary

Due to the diversity of tuning properties in sensory cortex, only a fraction of neurons are engaged in a particular task. Characterizing the tuning properties of neurons that are functionally linked to behavior is essential for understanding how activity is “read out” from sensory maps to guide decisions. We recorded from middle temporal (MT) neurons while monkeys performed a depth discrimination task, and we characterized the linkage between MT responses and behavioral choices. Trial-to-trial response fluctuations of MT neurons with odd-symmetric (“Near,” “Far”) disparity tuning were predictive of monkeys’ choices, whereas responses of neurons with even-symmetric tuning were not. This result cannot be explained by neuronal sensitivity or any other response property of MT neurons that we examined but is simply explained by the task strategy that monkeys learned during training. We suggest that this approach provides a physiological means to explore how task strategies are implemented in the brain.

Introduction

Sensory physiology has traditionally focused mainly on the information that neurons carry about external stimuli. Although stimulus-response relationships are of undeniable importance, they do not allow one to deduce whether the activity of particular neurons actually contributes to a behavioral outcome (Parker and Newsome, 1998). Experiments using reversible lesions or electrical microstimulation can provide causal links between neurons and behavior, but these manipulations affect large numbers of neurons. One would also like to be able to assess the functional linkage to behavior on a neuron-by-neuron basis after first characterizing stimulus-response relationships.

One method for relating single neurons to behavior involves analyzing trial-to-trial correlations between single-unit responses and behavioral choices. It is well known that responses of cortical neurons vary considerably from trial to trial even when presented with a fixed external stimulus (e.g., Dean, 1981; Schiller et al., 1976; Tolhurst et al., 1983). If perceptual decisions are based on sensory signals from such noisy neurons, then trial-

to-trial fluctuations in neuronal responses should be linked to trial-to-trial variations in behavioral outcomes. One can test for such a link by sorting neuronal responses according to behavioral choices and then computing the ability of an ideal observer to predict the animal’s choice from the neuronal response (“choice probability”; Britten et al., 1996; Celebrini and Newsome, 1994).

Choice probabilities have previously been used to demonstrate links between the responses of middle temporal (MT) neurons and perceptual decisions in direction discrimination (Britten et al., 1996; Shadlen et al., 1996) and structure-from-motion (Dodd et al., 2001) tasks. The implication of these studies is that neurons with large choice probabilities are functionally linked to behavior, either through bottom-up or top-down mechanisms (Parker et al., 2002). Based on this assumption, a particularly intriguing possibility is that choice probabilities, combined with conventional stimulus-response analyses, can reveal the specific (physiologically defined) cell types that are used by a subject to perform a particular task, thus providing a physiological account of task strategy.

To examine this prospect, we recorded the activity of single MT neurons during performance of a depth discrimination task in which depth was perceived from binocular disparities, small positional differences between the visual images projected onto the two retinas. Monkeys were trained to discriminate whether a circular patch of dots was presented in front of or behind the plane of fixation, and task difficulty was manipulated by adding different amounts of disparity noise (Figure 1). For this task, we have previously shown that single MT units signal depth with high fidelity (Uka and DeAngelis, 2003) and that microstimulation of disparity-tuned columns biases behavioral judgments of depth (DeAngelis et al., 1998). Consistent with these observations, we report here that many MT neurons exhibit large choice probabilities, reflecting a functional linkage between neuronal responses and behavioral choices.

Previous studies have classified disparity-tuned neurons into four basic groups: Near, Far, Tuned-Zero, and Tuned-Inhibitory (Poggio and Fischer, 1977; Poggio et al., 1988). Near and Far neurons have asymmetric tuning around zero disparity, responding selectively to stimuli presented in front of or behind the plane of fixation, respectively. Tuned-Zero and Tuned-Inhibitory neurons are symmetrically tuned around zero disparity, having a response maximum or minimum, respectively, for stimuli in the plane of fixation. We asked whether choice probabilities could reveal how these different types of disparity-tuned neurons (which are also present in area MT; DeAngelis and Uka, 2003; Maunsell and Van Essen, 1983) are monitored by the monkey to perform the depth discrimination task. In each experiment, visual stimuli were optimized to make the recorded neuron most informative to the monkey. Nevertheless, we found that neurons with odd-symmetric disparity tuning (Near, Far) usually exhibited large choice probabilities, whereas neurons with even-symmetric tuning (Tuned-Zero, Tuned-

*Correspondence: gregd@cabernet.wustl.edu

¹Present address: Department of Physiology 1, Juntendo University School of Medicine, 2-1-1 Hongo, Bunkyo, Tokyo 113-8421, Japan.

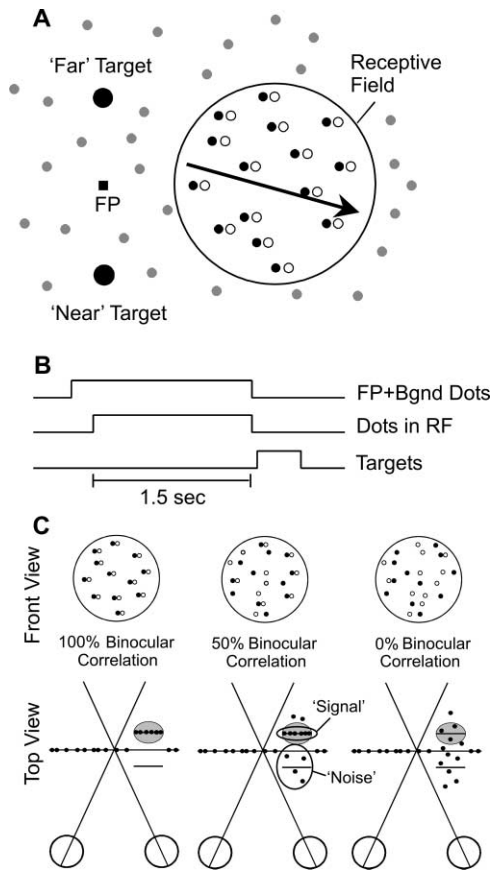


Figure 1. The Depth Discrimination Task

(A) Spatial configuration of the fixation point (FP), visual stimulus, and saccade targets. A random dot pattern was presented over the neuron's receptive field. The right-eye image (open dots) and the left-eye image (filled dots) were shifted horizontally to create a binocular disparity. The saccade targets that represent "near" and "far" choices were located below and above the fixation point, respectively.

(B) Time course of an individual discrimination trial. Monkeys were trained to fixate within a $1.6^\circ \times 1.6^\circ$ window throughout the 1.5 s visual stimulation period. Subsequently, the fixation point and the visual stimulus were turned off, and the two choice targets appeared. The monkey made a saccade to either of the two targets to indicate their decision.

(C) Top down view of the visual stimulus. The strength of the depth signal was determined by the "binocular correlation" which corresponds to the percentage of dots that had a fixed "signal" disparity. The left panel shows the 100% correlation case where all the dots have a fixed disparity; this stimulus is seen as a single depth plane. The center panel shows the 50% correlation case where half the dots have a fixed disparity while the remaining half are assigned random disparities. This stimulus is seen as a plane of dots embedded in disparity noise. The right panel shows the 0% correlation case where all the dots are assigned a random disparity and is seen as a "cloud" of dots in 3D space. Dots outside the receptive field were presented at zero disparity and served to anchor the monkeys' vergence posture.

Inhibitory) did not. This result cannot be explained by variations in neuronal sensitivity nor by variations in any of an extensive list of response properties that we examined quantitatively. Rather, we suggest that our findings are simply explained by the task strategy adopted by the animals during training.

Our results provide new evidence that area MT plays important roles in depth perception and suggest that choice probabilities may be used as a general tool to decipher how neuronal activity is "read out" from sensory maps in the brain.

Results

We recorded from 104 neurons (52 from each of two monkeys) in area MT during performance of the depth discrimination task illustrated in Figure 1 (see Experimental Procedures for details). This is the same sample of neurons for which we previously analyzed the relationship between neuronal and psychophysical sensitivity (Uka and DeAngelis, 2003). Details of the neuronal selection criteria are described in that study and are summarized briefly here. We excluded 41/170 neurons because isolation was lost prematurely, 10/170 neurons due to lack of disparity tuning, 8/170 neurons because their disparity-tuning curves were precisely symmetric about zero disparity, and 5/170 neurons because the monkey's behavior was clearly outside the normal range of performance.

Covariation of Neuronal Responses and Behavioral Choices

How well does the response of an MT neuron predict the upcoming choice of the monkey (i.e., near versus far) in our depth discrimination task? Figure 2 illustrates the analysis that we used to address this question (see also Britten et al., 1996; Dodd et al., 2001). The disparity-tuning curve of an example neuron is shown in Figure 2A: it responded strongly to near disparities and gave little response to far disparities. The depth discrimination task was tailored to the tuning of this neuron by having the monkey discriminate between the neuron's preferred (-0.9°) and null (0.6°) disparities in the presence of noise; one of the two disparities was presented each trial. The level of noise was titrated by varying the binocular correlation in the display (the percentage of correlated or "signal" dots, see Experimental Procedures), such that the monkey worked at a range of difficulty levels around psychophysical threshold. For each different stimulus condition (2 disparities \times 7 binocular correlations), Figure 2B shows the responses of this MT neuron sorted into two groups according to the choices of the animal ("preferred" choices versus "null" choices). At low binocular correlations (e.g., 2%), the monkey makes roughly equal numbers of preferred (i.e., near) and null (i.e., far) choices for each stimulus disparity, and these proportions change as the binocular correlation increases above threshold. Importantly, the average response for preferred choices (filled bars) is larger than the average response for null choices (unfilled bars), regardless of the strength of the binocular correlation or the stimulus disparity. In other words, the coupling between behavioral choices and MT responses exists throughout the range of stimulus conditions, a point to which we will return later.

The data of Figure 2B show that we can frequently predict the choice that the monkey will make simply by counting the number of spikes that an MT neuron produces during presentation of the visual stimulus

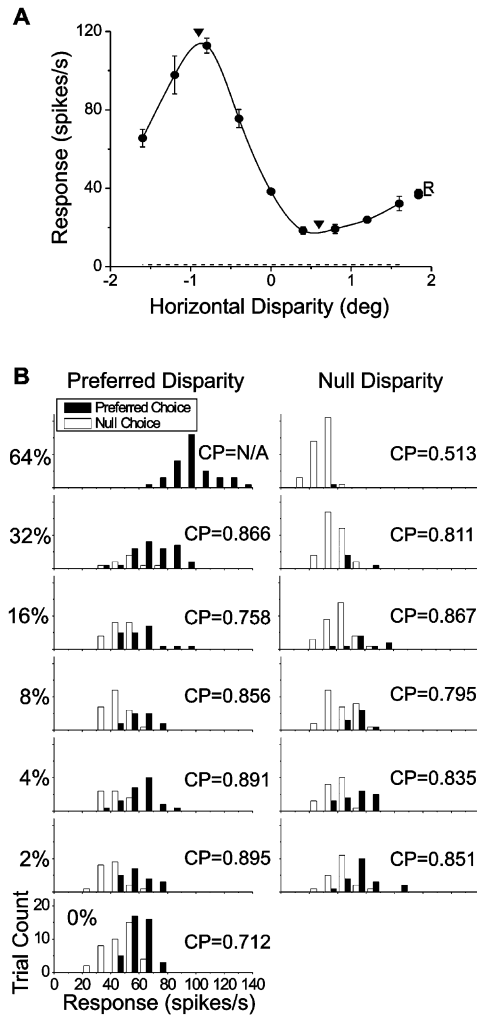


Figure 2. An Example MT Neuron Showing Choice-Related Response Modulation

(A) Disparity tuning curve. The two arrowheads indicate the preferred (-0.9°) and null (0.6°) disparities used in the discrimination task. Dotted horizontal line at 1 spike/s denotes the spontaneous activity. R and L indicate monocular responses to the right and left half images, respectively. Error bars indicate SEMs.
 (B) Frequency histograms of responses, sorted according to preferred (“near,” filled bars) and null (“far,” open bars) choices of the monkey, for each disparity and binocular correlation level. CP values indicate the choice probability for each stimulus condition.

(prior to any motor response). We quantify this predictability by computing a “choice probability” using ROC analysis (see Britten et al., 1996). The choice probability (which ranges from 0 to 1) gives the proportion of correct responses for an ideal observer whose task is to predict the upcoming choice of the monkey from the responses of an MT neuron during two randomly selected trials (one preferred choice and one null choice). A choice probability near 1.0 indicates that a high firing rate from the neuron predicts, with near-perfect reliability, that the monkey will choose the neuron’s preferred disparity. This is what would be expected, for example, if the monkey were to base his choices on the responses of a single MT neuron. A choice probability near 0.0

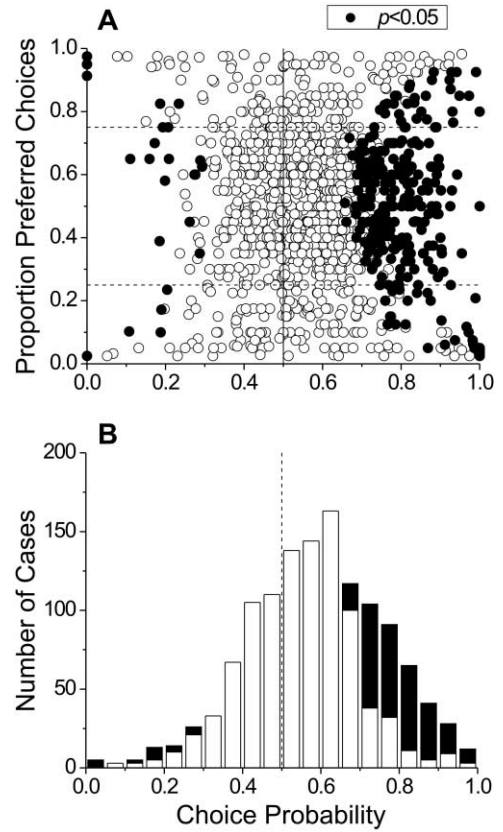


Figure 3. Distribution of Choice Probabilities Combined across All Neurons and Stimulus Conditions

(A) Choice probability is plotted against the proportion of “preferred” choices for all neurons at each disparity/binocular correlation level. Filled circles indicate choice probabilities that are significantly different from 0.5 (permutation test, $p < 0.05$). Dotted lines denote proportions of preferred choices equal to 0.25 and 0.75.
 (B) The data from panel (A) are collapsed down into a frequency histogram. Filled bars correspond to choice probabilities that are significantly different from 0.5 (permutation test, $p < 0.05$).

indicates that the monkey chooses the neuron’s preferred disparity when the neuron responds weakly; this outcome would not be expected to occur. Finally, a choice probability of 0.5 means that the ideal observer cannot predict the monkey’s choices from MT responses. For the example neuron of Figure 2B, the predictability was high for almost every stimulus condition, as evidenced by choice probabilities in the range between 0.7 and 0.9.

Figure 3A shows the relationship between choice probabilities and the proportion of preferred choices made by the monkey; this plot includes data from all neurons at all combinations of stimulus disparities and binocular correlations (one datum per unique stimulus condition). A choice probability can be calculated if there is at least one choice in each of the preferred and null groups, but the statistical significance of a choice probability will clearly depend on the relative numbers of preferred and null choices. Choice probabilities that are significantly different from 0.5 are shown as filled symbols in Figure 3A ($p < 0.05$, permutation test, see Experimental Procedures). When the proportion of preferred

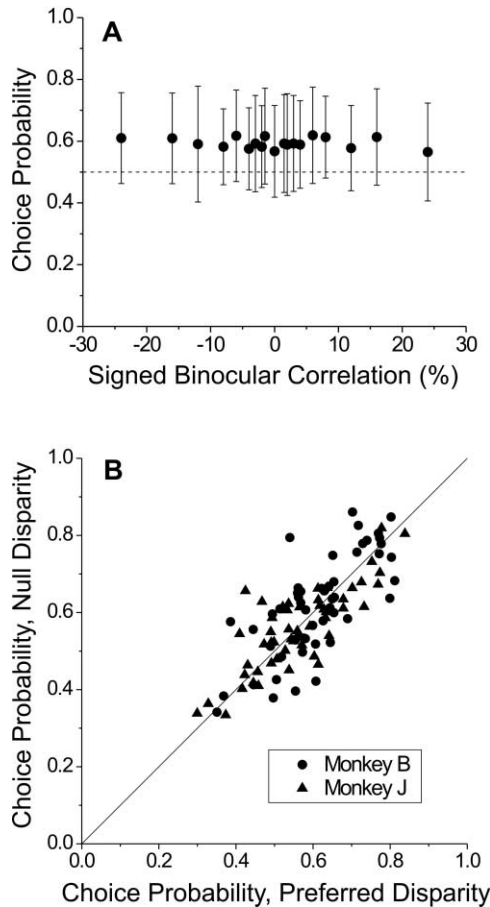


Figure 4. Choice Probabilities Do Not Depend on Visual Stimulus Parameters

(A) The average choice probability across 104 neurons is plotted as a function of signed binocular correlation (positive and negative values correspond to stimuli at the preferred and null disparities, respectively). Correlation levels at which the monkey made choices to one target more than 75% of the time were excluded from the analysis. Error bars indicate SDs.

(B) Choice probabilities calculated from responses to the preferred disparity are plotted against those for the null disparity. There is one datum for each of the 104 neurons. Responses of each neuron to each binocular correlation were normalized using z scores (subtracting the mean response and dividing by the standard deviation) and the normalized responses were then combined across binocular correlations to compute choice probabilities.

choices is close to 0 or 1, choice probabilities are widely distributed and seldom statistically significant. This is simply because there are few trials in either the preferred or null choice distribution (e.g., Figure 2B, 64% correlation). In practice, we find that most of the unreliable choice probabilities can be excluded by requiring that the proportion of preferred choices lie between 0.25 and 0.75 (dashed horizontal lines in Figure 3A). Thus, for analyses in Figure 4 and beyond, we excluded stimulus conditions for which the proportion of preferred choices fell outside these bounds.

In Figure 3B, the data of Figure 3A are collapsed down into a histogram. Each MT neuron contributes up to 13 data points to this distribution, one for each of the different stimulus conditions illustrated in Figure 2B. Twenty-two percent (284/1284) of the single-condition

choice probabilities are statistically significant ($p < 0.05$), and the vast majority of these (91.5%, 260/284) are larger than 0.5. The average choice probability is 0.59, which is significantly larger than 0.5 (one-sample t test, $p < 0.0001$). The responses of many MT neurons are therefore predictive of monkeys' choices.

If the data of Figures 2 and 3 reflect a true coupling between MT responses and behavioral choices, then these effects should be independent of visual stimulus properties (disparity and binocular correlation). Figure 4A shows the average choice probability (± 1 SD) as a function of signed binocular correlation, where positive and negative values correspond to signal dots presented at the preferred and null disparities, respectively. We find that choice probabilities do not depend on signed binocular correlation (ANOVA, $p > 0.5$), meaning that MT responses are larger when the monkey chooses the neuron's preferred disparity, regardless of stimulus strength and polarity. Note that the average choice probability is significantly larger than 0.5 at all binocular correlations ($p < 0.05$), including 0% correlation (mean = 0.57). The visual stimulus was ambiguous, and rewards were provided randomly in this condition; thus, the occurrence of high choice probabilities is not related to the percentage of correct responses (see also Britten et al., 1996; Dodd et al., 2001).

The independence of choice probabilities on stimulus polarity (sign of the binocular disparity) also holds true on a neuron-by-neuron basis. Figure 4B shows a scatter plot of choice probabilities for the preferred versus null disparity of each individual neuron. There is a strong correlation between these two measures ($r = 0.77$, $p < 0.001$), with no significant tendency for the data to be biased away from the unity-slope diagonal (paired t test, $p > 0.7$). These data help to rule out alternative explanations for the choice probability (see Discussion).

The finding that choice probabilities do not depend on stimulus parameters allows us to combine responses across stimulus conditions to obtain a single "grand" choice probability for each neuron. By normalizing the responses for each stimulus condition (using z scores) and combining them into a single pair of distributions (Britten et al., 1996), we remove response modulations due to the visual stimulus and isolate the effect of behavioral choice. Figure 5A shows z scored responses as a function of trial number for the same neuron as in Figure 2. It is clear that normalized responses tend to be larger for preferred choices (filled symbols) and that this effect is consistent throughout the block of 560 trials. These data are collapsed into a pair of response distributions in Figure 5B, and these distributions are used to calculate the grand choice probability, which is 0.79.

Figure 5C shows the distribution of grand choice probabilities for the 104 neurons that we studied. More than half of the neurons (62/104) have choice probabilities that are significantly different from 0.5 (filled bars, permutation test, $p < 0.05$). Of these, 88.7% (55/62) are larger than 0.5. The average choice probability across the population is 0.59 (0.61 for monkey B and 0.57 for monkey J), and this value is significantly larger than 0.5 (one-sample t test, $p < 0.0001$). Thus, the responses of a large fraction of MT neurons are positively correlated with monkeys' choices across trials, consistent with the notion that disparity signals in MT are used to perform the depth discrimination task.

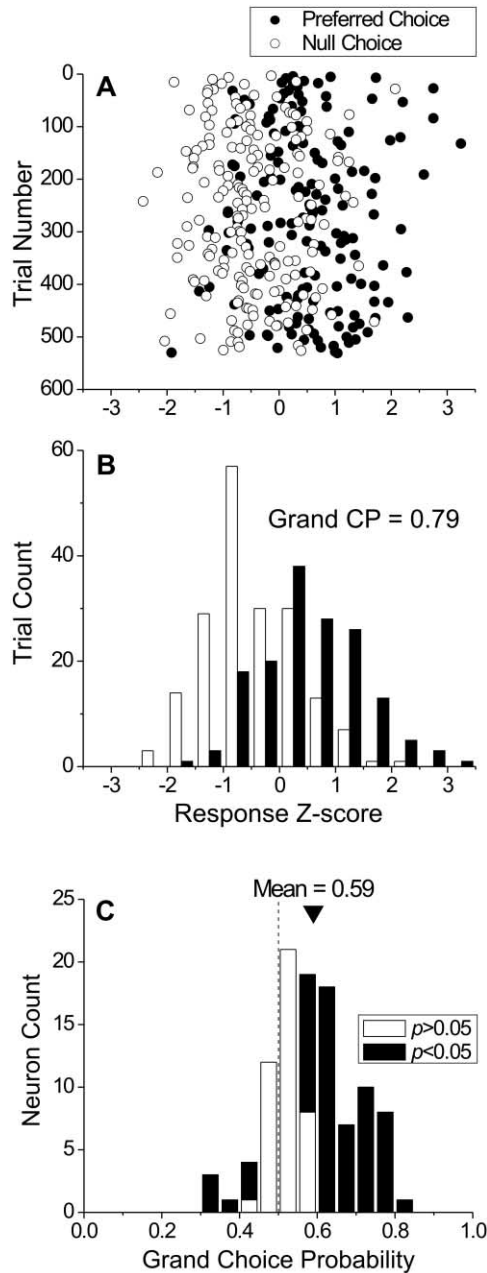


Figure 5. Calculation of “Grand” Choice Probability
 (A) Trial-by-trial z scored responses for the example neuron in Figure 2. Responses to preferred (near) choices are shown with filled circles, whereas responses to null (far) choices are shown with open circles. Trials from correlation levels where the monkey made choices preferentially to one target more than 75% of the time are excluded.
 (B) The data from panel (A) are collapsed down into a frequency histogram. Filled and open bars correspond to preferred and null choices, respectively.
 (C) Distribution of grand choice probabilities for the 104 neurons in our sample. Filled bars correspond to choice probabilities that are significantly different from 0.5 (permutation test, $p < 0.05$).

Choice Probabilities, Disparity-Tuning Symmetry, and Task Strategy

For any given behavioral task, it is likely that only a fraction of neurons in a particular cortical area provide

useful information, as dictated by their receptive field location, tuning properties, etc. If significant choice probabilities reflect a functional linkage between MT responses and decisions, then a very intriguing possibility is that choice probabilities can tell us which types of neurons the monkey monitors to perform a particular task.

We were able to address this issue by taking advantage of the manner in which monkeys were trained to perform our depth discrimination task. Initially, each monkey was trained to discriminate between two coarse disparities (one near, one far) that were symmetrically arranged around the plane of fixation (a typical example would be -0.4° versus $+0.4^\circ$). An effective strategy for solving this task would be to compare the activity level of a Near-tuned pool of neurons with the activity of a Far-tuned pool of neurons. In MT, such Near and Far cells are most prevalent and usually have odd-symmetric tuning around zero disparity (DeAngelis and Uka, 2003). In contrast, neurons that are even-symmetrically tuned around zero disparity (Tuned-Zero or Tuned-Inhibitory neurons) would provide no useful information for solving the task, and the monkey would be wise to exclude them from the decision process. By this reasoning, we predict that choice probabilities should be greater than 0.5 for Near/Far neurons but not for Tuned-Zero/Tuned-Inhibitory neurons.

In the latter stages of training, the two stimulus disparities were sometimes placed asymmetrically around zero disparity (e.g., -0.2° versus $+0.6^\circ$) in order to accommodate variations in the peak and trough locations of MT disparity-tuning curves. However, one disparity was always negative and the other positive. Hence, the monkey’s strategy of comparing Near versus Far pools of neurons would not need to change, and the above prediction would still be expected to hold.

We measured the symmetry of disparity-tuning curves around zero disparity by fitting each curve with a Gabor function having a Gaussian envelope that was centered at zero disparity (Equation 1). The phase parameter, ϕ , can then be used to gauge the symmetry of the tuning curve around zero disparity. Figure 6A shows the relationship between grand choice probability and tuning symmetry for our sample of 104 MT neurons. Phases near 0° correspond to Tuned-Zero or Tuned-Inhibitory neurons, whereas phases near 90° correspond to Near or Far neurons. The latter outnumber the former, consistent with our previous analyses of tuning shape in MT (DeAngelis and Uka, 2003). The striking result in Figure 6A is that choice probabilities tend to be larger than 0.5 for Near/Far neurons but not for Tuned-Zero/Tuned-Inhibitory neurons, consistent with our prediction based on task strategy. The median choice probability for neurons with $\phi > 60$ is 0.62, whereas the median for neurons with $\phi < 30$ is 0.52. This difference is highly significant (Mann-Whitney U Test, $p < 0.002$). In addition, the variance of choice probabilities for $\phi > 60$ is significantly larger than the variance for $\phi < 30$ (Levene’s test, $p < 0.01$). Overall, there is a significant positive correlation between grand choice probability and disparity tuning symmetry as well (Spearman $r = 0.24$, $p < 0.02$). Note the lack of data points in the upper-left corner of Figure 6A, indicating that no Tuned-Zero or Tuned-Inhibitory neuron had a large choice probability.

It is important to emphasize that the stimulus dispar-

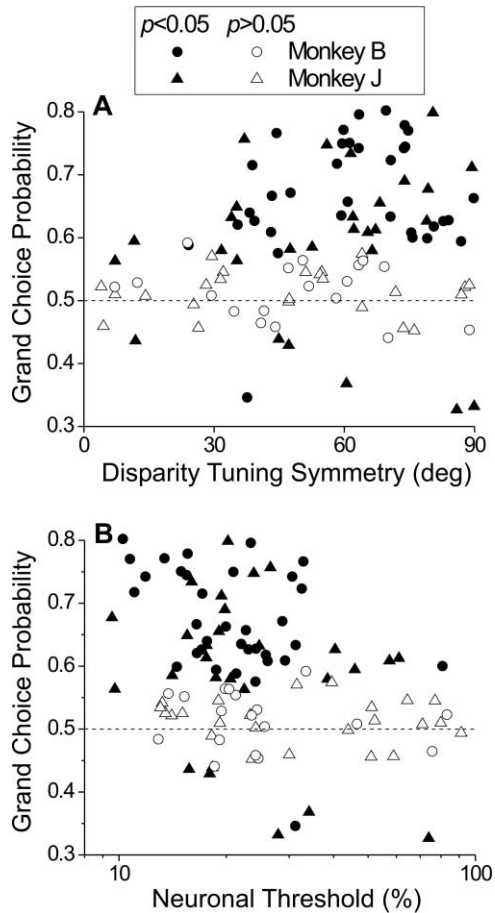


Figure 6. Dependence of Choice Probabilities on Neuronal Tuning Properties

(A) Choice probability is plotted against the symmetry of the disparity-tuning curve. Disparity-tuning curves were fit with a Gabor function with the center of the Gaussian envelope fixed at zero disparity. The phase parameter was then wrapped into the range from 0° to 90° . A phase near 0° indicates even-symmetric tuning around zero disparity, whereas a phase near 90° indicates odd-symmetric tuning. (B) Choice probabilities are plotted against neuronal thresholds for our sample of neurons. The neuronal threshold was calculated using ROC analysis, and low thresholds correspond to neurons that are highly sensitive to binocular correlation (see Uka and DeAngelis, 2003).

ities were tailored to the peak and trough of the disparity tuning curve for each neuron in Figure 6A. In the case of Tuned-Zero/Tuned-Inhibitory neurons, this means that one disparity was close to zero and the other was farther away with the opposite sign (e.g., -0.1° versus $+0.5^\circ$). Thus, we find weaker choice probabilities for these neurons despite the fact that the stimuli were tailored to make them most informative to the monkey. One possible explanation for this result is that Tuned-Zero/Tuned-Inhibitory neurons are inherently less sensitive to depth at low binocular correlations. Indeed, previous studies (Britten et al., 1996; Parker et al., 2002) have found that choice probabilities are correlated with neuronal sensitivity, suggesting that the most sensitive neurons have the strongest links to behavioral performance. Figure 6B shows that we find a similar result in our task, with

a significant negative correlation between grand choice probability and neuronal threshold ($r = -0.29$, $p < 0.005$).

Could this relationship between choice probability and sensitivity account for the data of Figure 6A? We find that it cannot. There is no significant correlation between neuronal thresholds and tuning symmetry ($r = -0.1$, $p > 0.15$, data not shown). Moreover, when the trend between choice probability and neuronal threshold is removed, we still find that choice probabilities are significantly larger for $\phi > 60$ than for $\phi < 30$ (Mann-Whitney U Test on residuals, $p < 0.02$). We also performed a multiple regression analysis to examine the dependence of grand choice probabilities on the following tuning properties of MT neurons: disparity tuning symmetry (ϕ), disparity discrimination index, disparity tuning width (size of Gaussian envelope), peak firing rate, preferred direction, direction discrimination index, direction tuning width, preferred speed, speed discrimination index, eccentricity, ocular dominance, optimal stimulus size, and percentage of surround inhibition (see DeAngelis and Uka, 2003, for the quantitative methods used to extract these parameters). As expected from Figure 6A, we found a significant partial correlation ($r = 0.28$, $p < 0.01$) between choice probability and tuning symmetry, but we found no significant partial correlations with any of the other independent variables ($p > 0.1$ in all cases). The same result was found in a forward stepwise regression. Thus, disparity symmetry (ϕ) was the only tuning property that accounted for significant variance in the distribution of choice probabilities. We can also reject the possibility that the result of Figure 6A reflects variation in behavioral performance due to asymmetry of the preferred and null disparities, as disparity tuning symmetry had no significant effect on the psychophysical thresholds of the monkeys ($r = -0.15$, $p = 0.14$). Together, these analyses confirm that choice probabilities are larger for Near and Far neurons than for Tuned-Zero and Tuned-Inhibitory neurons, consistent with the prediction outlined above (see also Discussion).

Choice Probabilities and Performance Fluctuations

It is a common observation that psychophysical performance (sensitivity) fluctuates gradually over time within an experimental session. Are these slow fluctuations in behavior due to variations in the responses of low-level neurons that provide the sensory information for the task, or do they reflect some higher-level cognitive process (e.g., motivation)? If large choice probabilities reflect a functional linkage between MT neurons and perceptual decisions and if behavioral fluctuations are due to slow variations in the responses of sensory neurons, then we might expect neuronal sensitivity to wax and wane along with psychophysical sensitivity. Moreover, the coupling between neuronal and psychophysical sensitivity should be stronger for neurons with larger choice probabilities, as the response of these neurons is more strongly coupled to the monkey's reports.

To examine this issue, we calculated neuronal and psychophysical thresholds (as described in Uka and DeAngelis, 2003) within a moving window having a length of five stimulus repetitions (70 trials). This window

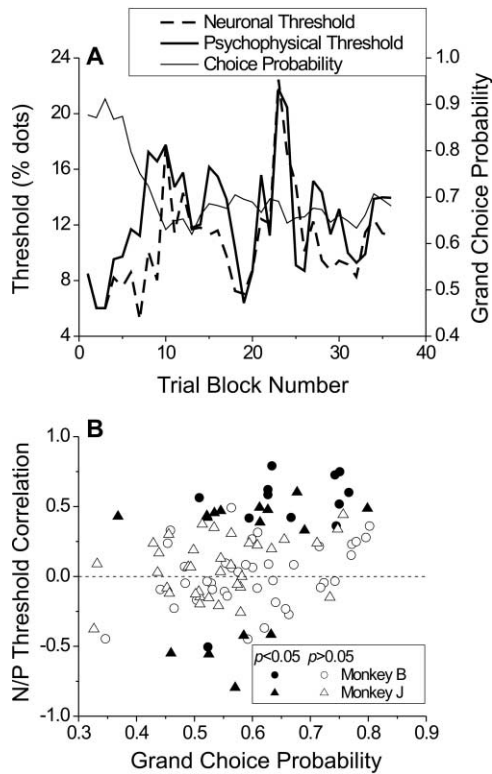


Figure 7. Analysis of Neuronal/Psychophysical Threshold Correlations and Their Dependence on Choice Probabilities

(A) Neuronal and psychophysical thresholds (see Uka and DeAngelis, 2003) were calculated within a window of five stimulus repetitions (70 trials) that was slid across the data in increments of one repetition. The thick solid curve shows the time series of psychophysical thresholds, whereas the thick dashed curve shows the time series of neuronal thresholds for an example neuron. The correlation coefficient between the two time series is 0.73. The thin solid curve shows the time course of the choice probability for this neuron.

(B) The neuronal/psychophysical (N/P) threshold correlation is plotted as a function of grand choice probability for our population of 104 MT neurons.

was slid across the data in increments of one repetition. An example of this analysis is shown in Figure 7A. The monkey performed 40 repetitions (560 trials) of the depth discrimination task during an ~50 min recording session. Our sliding window analysis thus yielded a time series of 36 threshold values. The thick solid curve in Figure 7A shows the time series of psychophysical thresholds, and it can be seen that sensitivity fluctuates considerably within the session. Note that these variations in sensitivity are on a *much* slower time scale than the trial-to-trial response variations that give rise to choice probabilities (Figure 5A). Interestingly, these slow fluctuations in psychophysical performance are mirrored quite closely by variations in the neuronal threshold (thick dashed curve), such that the correlation coefficient between the two time series is 0.73 ($p < 0.001$).

The thin solid curve in Figure 7A shows the time course of the grand choice probability based on a similar sliding window analysis. The grand choice probability across

all trials was 0.74, but there was some variation in this value across blocks of trials in the experiment. There is a modest negative correlation ($r = -0.49$, $p < 0.003$) between time series for the choice probability and the neuronal threshold, mainly due to anticorrelation during the first several blocks of trials. Overall, however, the time course of the choice probability is largely uncoupled from that of the neuronal and psychophysical thresholds.

Across the population of 104 neurons, neuronal and psychophysical thresholds tend to be positively correlated (average $r = 0.09$, one-sample t test $p < 0.006$, see also Zohary et al., 1994). If these slow covariations in threshold are predictable from the fast trial-to-trial fluctuations in response that underlie choice probabilities, then larger correlations between neuronal and psychophysical thresholds should be seen for neurons with larger choice probabilities. Indeed, Figure 7B reveals a significant positive correlation between these variables ($r = 0.30$, $p < 0.003$). That this trend is rather weak is not surprising, given that neurometric and psychometric functions computed from only 70 trials tend to be quite noisy, leading to a large scatter in N/P threshold correlations. Despite this scatter, it is clear that the coupling between neuronal sensitivity and behavioral performance is strongest for neurons with large choice probabilities.

One possible explanation for choice probabilities is that a “top-down” choice-dependent signal is fed back to MT from higher-level areas and modulates the response of MT neurons to the visual stimulus once a decision is made (see Discussion). Slow variations in the strength or quality of such a decision signal are likely to be correlated with slow variations in psychophysical thresholds, as both could be driven by variations in attention or motivation, for example. If the top-down signal modulates the responses of MT neurons (producing choice probabilities), then slow temporal variations in the top-down signal could also drive slow variations in neuronal sensitivity (as seen in Figure 7A). In this scheme, the choice probability itself should covary slowly with neuronal and psychophysical thresholds, and this covariation should be strongest for neurons with the largest choice probabilities. We therefore computed correlation coefficients between the time course of the choice probability (thin solid curve in Figure 7A) and the time courses of neuronal and psychophysical thresholds (thick curves in Figure 7A). We find a weak negative correlation, on average, between the time series of choice probabilities and neuronal thresholds (average $r = -0.15$, one-sample t test $p < 0.01$), but no correlation between time series of choice probabilities and psychophysical thresholds (average $r = -0.013$, one-sample t test $p = 0.68$). More importantly, the strength of these correlations does not depend on the value of the grand choice probability across our population of MT neurons ($r = -0.058$, $p = 0.56$ for neuronal thresholds; $r = -0.16$, $p = 0.11$ for psychophysical thresholds, data not shown). Thus, whereas neuronal and psychophysical thresholds covary slowly for neurons with large choice probabilities (Figure 7B), these slow variations are not explained by the time course of the choice probability itself. This finding is inconsistent

with the simple top-down explanation for choice probabilities given above.

Together, our findings suggest that trial-to-trial covariations between neuronal and behavioral responses transduce slow modulations of neuronal response properties into slow variations in behavioral sensitivity. We examine these issues further in the Discussion.

Effects of Vergence Posture

Uncontrolled fluctuations in vergence posture (the difference between the left and right eye positions) are a potential source of artifacts in our analysis since vergence errors will directly affect the absolute binocular disparities on the retina. If the monkey were to systematically vary his vergence posture in accord with his upcoming choice, and if MT responses are modulated by these vergence changes, then vergence fluctuations could give rise to artifactual choice probabilities.

To address this issue, we determined whether the responses of each MT neuron (after z scoring) depended on vergence angle. Indeed, a substantial fraction of MT neurons (50/104) showed a significant correlation between vergence angle and neuronal response (linear regression, $p < 0.05$). To determine whether these correlations contribute to choice probabilities, we removed the linear trend between response and vergence angle and computed a grand choice probability from the residuals. Figure 8A shows that the effect of choice survives removal of vergence effects. Original and vergence-corrected choice probabilities are strongly correlated ($r = 0.99$, $p < 0.0001$), indicating that systematic vergence fluctuations had little effect on our results. In fact, there was no significant change in the mean choice probability (paired t test, $p > 0.3$) following correction for vergence.

Effects of Variability in the Visual Stimulus

Another potential difficulty with our choice probability calculations concerns trial-to-trial variations in the visual stimulus (see also Britten et al., 1996; Uka and DeAngelis, 2003). Although the disparities of noise dots were chosen randomly from a uniform distribution having zero mean, the exact distribution of noise disparities in a stimulus presentation varied from trial to trial. If both the monkey and the MT neurons were sensitive to these small trial-to-trial variations in the distribution of noise disparities, this could produce artifactual correlations between MT responses and behavioral choices.

To address this possibility, we tested 61/104 MT neurons with random-dot patterns that were identical across repeated trials. For these neurons, one-half of the 0% correlation trials had a fixed distribution of dot positions and noise disparities (NOVAR condition), whereas the other one-half of trials contained the normal trial-to-trial variation in dot locations (VAR condition). As we have described previously (Figure 5A of Uka and DeAngelis, 2003), mean firing rates were not significantly different between the NOVAR and VAR conditions, but the variance of spike counts was 1.51 times larger for the VAR condition than the NOVAR condition. Despite this change in response variance, we find no significant difference between mean choice probabilities for the VAR and NOVAR conditions (paired t test, $p = 0.83$). Moreover, there is a significant correlation between

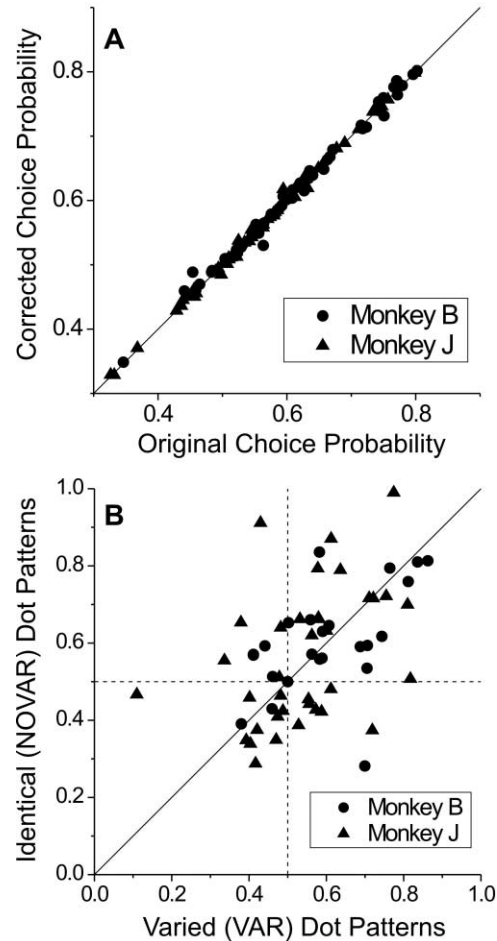


Figure 8. Dependence of Choice Probabilities on Vergence Angle and Stimulus Variations

(A) Effect of vergence angle on choice probabilities. The linear trend between z scored responses and vergence angle was removed, and a vergence-corrected grand choice probability was computed from the residuals. The original (uncorrected) choice probability is plotted against the vergence-corrected choice probability for each neuron from the two monkeys.

(B) Effect of stimulus variations on choice probabilities. For 61 neurons, choice probabilities were calculated separately from 0% correlation trials with dot patterns that changed from trial-to-trial (VAR condition) and from 0% correlation trials with dot patterns that were identical from trial-to-trial (NOVAR condition). Choice probabilities for varying (VAR) dot patterns are plotted against those for fixed (NOVAR) dot patterns.

these two measures ($r = 0.42$, $p < 0.001$), as shown in Figure 8B. The slope of the best-fitting regression line was not significantly different from unity ($p > 0.05$). These observations indicate that large choice probabilities were not driven by trial-to-trial variations in the fine structure of the visual stimulus.

Time Course of Choice Probabilities

There are two basic explanations, which are not mutually exclusive, for the finding of large choice probabilities. The “bottom-up” explanation is that trial-to-trial fluctuations in MT responses drive variations in behavioral choices (for a fixed visual stimulus). The “top-down”




Figure 9. Time Course of Choice-Related Modulations in Response (A) Average (\pm SEM) choice probability is plotted as a function of time after stimulus onset. For each time point, grand choice probabilities were calculated within a 100 ms window of time, and these values were then averaged across all 104 neurons. (B) Average (\pm SEM) normalized response is plotted as a function of time for preferred and null choices separately. Spikes were counted in 20 ms bins and normalized to the peak response at each binocular correlation level, then averaged across correlation levels for each neuron. Correlation levels where the monkey made choices preferentially to one target more than 75% of the time were excluded. Normalized responses were then averaged across all 104 neurons.

explanation is that decision-related signals from higher-level areas are fed back to MT and modulate responses to visual stimuli. For a top-down mechanism, the choice-related change in MT response would occur after the decision is formed, and the timing of this change would vary somewhat from trial to trial. Averaged across stimulus repetitions, one might expect choice probabilities to rise gradually during the trial epoch, becoming maximal after the latest time that a decision is made and before execution of the saccade.

To examine the time course of choice-related modulations, grand choice probabilities were calculated within a 100 ms sliding window that was shifted across the data in increments of 20 ms. Figure 9A shows the time course of the average choice probability for our sample of 104 neurons. Each time point indicates the center of the 100 ms time window; e.g., the value at 50 ms corresponds to the choice probability calculated in a 0–100 ms time window. The average choice probability is close to 0.5 at the beginning of the trial and gradually increases over time, reaching a plateau \sim 400–500 ms after stimulus onset. It is also worth noting that we see a clear and significant rise in choice probability only 50–60 ms after stimulus onset.

To compare the time course of choice modulation with the dynamics of stimulus-driven excitation, we also calculated the time course of MT population responses. Responses to stimuli that evoked preferred and null choices were normalized to the peak response at each correlation level and then averaged across different correlation levels (in 20 ms bins). Figure 9B shows the average data for all 104 neurons. Responses to stimuli eliciting preferred and null choices started \sim 60 ms after stimulus onset. The choice-related difference in response (dot-dashed line, right axis) started to appear immediately after the onset of the responses and followed a time course very similar to that of the choice probability (Figure 9A). Note that there is no predictive activity in the period before the visual response, suggesting that predictive activity is not due to response bias or expectation (see Discussion).

Discussion

We measured the degree of correlation between MT responses and monkeys' behavioral choices in a depth discrimination task. The average choice probability was substantially larger than 0.5, and many MT neurons reliably predicted the monkey's choices. Vergence eye movements or trial-to-trial stimulus variation cannot explain these effects. Previous work has shown that single MT neurons have sensitivity comparable to the monkey in this task (Uka and DeAngelis, 2003) and that electrical microstimulation of MT biases depth judgments in this task (DeAngelis et al., 1998). In total, this body of data leaves little doubt that MT transmits sensory signals used by the monkey in our depth discrimination task.

The more general implication of our findings is that choice probabilities can be used to characterize which types of (physiologically defined) neurons contribute to the performance of a given task. Our monkeys were initially trained to discriminate between two disparities placed symmetrically around zero disparity (the fixation plane). In this case, Near/Far neurons would be informative, whereas Tuned-Zero/Tuned-Inhibitory neurons would not. If monkeys learned to selectively monitor Near/Far neurons during training, this would readily explain the pattern of choice probabilities that we observe. Hence, we suggest that choice probabilities can allow one to assess the strategy used by the animal to read out sensory signals from a population of neurons.

There is an important distinction to make between our findings (Figure 6A) and those of previous studies. Britten et al. (1996) showed that choice probabilities are abolished when the receptive field of the recorded neuron does not overlap the stimulus. Similarly, both Britten et al. (1996) and Parker et al. (2002) have shown that choice probabilities tend to be stronger for more sensitive neurons (see also our Figure 6B). These results show that choice probabilities depend on the quality of the signals that neurons carry about the stimulus, and it is not surprising that the functional linkage of neurons to behavior should depend on signal fidelity (see DeAngelis et al., 1998). However, our finding that choice probabilities depend on the shape of the disparity-tuning curve (Figure 6A) cannot be explained in terms of signal fidelity nor by any of the tuning properties of

MT neurons that we examined. This shows that signal fidelity alone cannot be used to infer the involvement of particular sensory neurons in a task.

We conclude that choice probabilities are low for neurons with even-symmetric (Tuned-Zero or Tuned-Inhibitory) disparity tuning solely because the monkeys learned not to monitor signals from these neurons during training on our task. Thus, although our experiments were not specifically designed to probe task strategy, our results provide a strong indication that choice probabilities can be used to explore how task strategies are implemented in the read out of population codes. Future experiments can now be directed squarely at the goal of manipulating task strategy while monitoring the functional linkage of neurons to behavior using choice probabilities.

Source of Choice-Related Response Modulations

What gives rise to the predictive activity that underlies large choice probabilities in MT? Before we can conclude that choice probabilities reflect a functional linkage between MT responses and behavioral choices, it is important to consider alternative explanations. The analyses of Figure 8 exclude the uninteresting possibilities that choice probabilities are driven by eye movements or stimulus variations. We now consider the potential roles of spatial attention and choice bias.

Spatial attention is known to modulate the responses of MT neurons during behavior (Cook and Maunsell, 2002; Seidemann and Newsome, 1999; Treue and Maunsell, 1996, 1999), and the strength of this effect would be expected to vary from trial to trial independently of the visual stimulus. The probability of a correct choice by the monkey should also be higher on trials with more focused spatial attention. Thus, we might expect a correlation between the monkey's performance (proportion correct responses) and response modulation due to spatial attention. This, in turn, could induce a correlation between MT responses and behavioral choices. Importantly, however, the sign of any choice effect due to spatial attention would depend on stimulus disparity (preferred or null). For a preferred stimulus, increased attention would tend to produce a larger neuronal response and a preferred choice (hence, choice probability >0.5). By contrast, increased spatial attention to a null stimulus would tend to produce a larger neuronal response and a null choice (hence, choice probability <0.5). There would be no prediction for the 0% correlation case, since increased spatial attention would not tend to favor preferred or null choices. The data of Figure 4B show clearly that choice probabilities do not depend on the sign of the stimulus disparity, thus ruling firmly against spatial attention as a source of the choice effect.

Another possibility is that choice probabilities reflect a choice bias. For example, if the monkey expects a near disparity to appear, this might selectively enhance the activity of Near-tuned neurons and increase the probability of an eventual near choice. If this were the case, then we might observe significant choice probabilities in the prestimulus period. In contrast, we found that the average grand choice probability (0.50) was not larger than 0.5 in the 200 ms period between fixation

and stimulus onset. Moreover, choice probabilities were also absent in the prestimulus period during 0% correlation trials (average 0.50), for which any prestimulus choice bias would tend to be most strongly coupled to the eventual choice, due to the ambiguity of the visual stimulus. These findings are consistent with the results of previous analyses of prestimulus activity in a direction discrimination task (Britten et al., 1996; Seidemann et al., 1998). Thus, the data do not seem consistent with an explanation based on choice bias. We note, however, that choice bias signals elsewhere in the brain (cf. Basso and Wurtz, 1998; Coe et al., 2002; Shadlen and Newsome, 2001) could possibly give rise to choice probabilities (in the form of featural attention, Treue and Martinez Trujillo, 1999) by changing the response gain of MT neurons after stimulus onset.

With the above considerations, there are two basic explanations for our data. The bottom-up explanation is that MT neurons provide sensory input to a decision mechanism, such that trial-to-trial variability in MT responses drives trial-to-trial fluctuations in behavioral choices (Britten et al., 1996; Shadlen et al., 1996). The top-down explanation is that the activity of decision-related neurons elsewhere in the brain (cf. Gold and Shadlen, 2001; Horwitz and Newsome, 1999; Kim and Shadlen, 1999; Shadlen and Newsome, 2001) is fed back to modulate MT responses (Dodd et al., 2001; Parker et al., 2002). At present, it is difficult to distinguish between these two possibilities, and they are not mutually exclusive. Nevertheless, it is useful to frame our results using these two basic models.

By the bottom-up explanation, the dependence of choice probabilities on the shape of the disparity-tuning curve (Figure 6A) would arise because signals from Tuned-Zero/Tuned-Inhibitory neurons do not feed into the decision mechanism used to judge depth. Presumably, this pattern of selective connectivity was instantiated during training based on the specific requirements of our task, namely, that monkeys were trained to discriminate between disparities symmetrically placed around the plane of fixation. The bottom-up model also readily explains the link between choice probabilities and slow fluctuations in neuronal and psychophysical thresholds (Figure 7). Previous work has shown that MT neurons exhibit slow fluctuations in responsiveness having a roughly similar time course to our observed fluctuations in neuronal sensitivity (Bair et al., 2001). Moreover, these slow fluctuations are generally uncorrelated between nearby MT neurons (Bair et al., 2001). Thus, for neurons that the monkey monitors to perform the task (i.e., those with large choice probabilities), slow fluctuations in response would contribute to slow fluctuations in psychophysical sensitivity (Figure 7B). In contrast, slow variations in a higher-level cognitive process, such as motivation, would be likely to affect all MT neurons and would not explain the significant correlation seen in Figure 7B.

By the top-down explanation, feedback connections to MT from decision-related neurons would have to selectively target Near and Far cells in order to explain the results of Figure 6A. Such a pattern of connectivity could also be shaped during training. However, it is more difficult to explain the results of Figure 7 in terms of the top-down model. If slow covariations in neuronal and

psychophysical sensitivity were driven by slow fluctuations in top-down decision signals, then slow fluctuations in choice probabilities should be correlated with changes in neuronal sensitivity. Moreover, this effect should be strongest for neurons with the largest choice probabilities, but our data show that this is not the case.

The data of Figure 9 show that choice effects appear at response onset and reach a plateau within the first several hundred milliseconds of the trial, consistent with a bottom-up explanation of choice probabilities. This result is similar to that found by Newsome and colleagues for direction discrimination in areas MT and MST (Britten et al., 1996; Celebrini and Newsome, 1994), although our time course reaches a plateau a bit later than theirs. The data of Figure 9 would argue against a top-down explanation, provided that the monkeys made their decisions relatively late in the trial. However, we do not know precisely when monkeys made their decisions (as ours was not a reaction time task), and we cannot exclude the possibility that decisions were made early (see Uka and DeAngelis, 2003). If monkeys do make their decisions early, one might expect choice probabilities to decline later in the trial if the bottom-up model holds. However, we have previously reported that MT neurons exhibit fairly strong serial correlations in response variability as a function of time (Uka and DeAngelis, 2003), and this could prolong the time course of choice probabilities. Thus, the fairly flat time course of choice probabilities that we observe does not clearly distinguish between bottom-up and top-down explanations for the choice effect.

Our time course data (Figure 9) do appear to be somewhat different from those reported by Dodd et al. (2001) for MT neurons studied during a bi-stable structure-from-motion task. Whereas our choice effect clearly reached a plateau by 700–800 ms after stimulus onset, they found a steadily increasing time course of choice probabilities throughout their 2 s trial period. This difference might be explained by a stronger contribution of top-down signals in the structure-from-motion task (see next section).

Considering all of the data and arguments laid out above, the bottom-up model seems to provide a more parsimonious explanation for our results than the top-down model, but we cannot exclude a contribution from top-down processes based on our experiments.

Comparison with Other Studies

Correlations between MT activity and behavioral choices have been explored in several previous studies (Bradley et al., 1998; Britten et al., 1996; Dodd et al., 2001; Grunewald et al., 2002; Thiele et al., 1999; Williams et al., 2003). Three of these studies (Britten et al., 1996; Dodd et al., 2001; Williams et al., 2003) have calculated choice probabilities in a two-alternative forced choice task, thus allowing a direct comparison with our results. Britten et al. (1996) reported an average choice probability of 0.55 for MT neurons tested in a direction discrimination task, Dodd et al. (2001) reported an average choice probability of 0.67 in a bi-stable structure-from-motion task, and Williams et al. (2003) reported an average choice probability of 0.50 in an ambiguous (bi-stable) apparent motion task. Given that our average choice

probability (0.59) lies in between these three previous values, it is worth considering differences among these studies.

Our experimental methods and neuronal selection criteria (see Uka and DeAngelis, 2003, for details) are very similar to those used by Britten et al. (1996), thus our results should be comparable. The only substantive difference is that our stimuli were tailored to match both the disparity and velocity preferences of MT neurons, whereas Britten et al. (1996) always presented their stimuli at zero disparity. Since the majority of MT neurons prefer nonzero disparities (DeAngelis and Uka, 2003), most neurons were not maximally activated in the Britten et al. (1996) study, and this may have resulted in a slight overestimation of neuronal thresholds. A mismatch between the stimulus disparity and the disparity preference of the neuron could make other neurons (those responding maximally) most informative to the monkey and may explain why the average choice probability in our study (0.59) is somewhat higher than what Britten et al. (1996) found for direction discrimination (0.55). Indeed, the 99% confidence interval for our mean choice probability (0.558 to 0.613) does not include the mean value of the Britten et al. study, indicating that this difference is significant.

In most other respects, our findings are quite similar to those of Britten et al. (1996). Specifically, the independence of choice probability on signal strength, the correlation between neuronal thresholds and choice probabilities, the independence of choice probability on stimulus variations, and the time course of choice effects are all comparable between the two studies. Unlike Britten et al. (1996), however, we did not perform a control in which the stimulus was moved off of the receptive field, or a nonpreferred stimulus was used.

Dodd et al. (2001) have reported remarkably strong choice probabilities for MT neurons recorded during a structure-from-motion (rotating cylinder) task. Given that the ambiguous stimuli in their task were perceptually bi-stable, one speculation is that these large choice probabilities may have resulted from a stronger top-down decision signal in the structure-from-motion task versus our depth discrimination task. Even in the ambiguous case, the structure-from-motion percept is compelling and stable during a single trial (Dodd et al., 2001; Parker et al., 2002), whereas the depth percept in our 0% correlation stimulus is more tenuous and fleeting. Thus, an observer is probably much more confident in his/her decision when viewing the ambiguous structure-from-motion stimulus than when viewing our 0% correlation stimulus, potentially leading to a much stronger top-down decision signal. Neuronal sampling bias may also have contributed to the difference between our results and those of Dodd et al. (2001). Dodd et al. (2001) excluded 132/322 (41%) of neurons from their sample due to poor disparity tuning, whereas we excluded only 10/170 (6%) of neurons for this reason (see Uka and DeAngelis, 2003). Neurons with poor tuning will have high neuronal thresholds and thus will tend to have lower choice probabilities (see our Figure 6B and Figure 6 of Parker et al., 2002). If we exclude an additional 35% of neurons with the highest thresholds, our mean choice probability rises to 0.61. Thus, sampling bias probably accounts for a modest portion of the difference in choice

probabilities between our study and that of Dodd et al. (2001).

Williams et al. (2003) have reported a complete lack of significant choice probabilities for MT neurons recorded during an ambiguous (bi-stable) apparent motion task. This is surprising considering that a few other studies (including ours) have found large choice probabilities in MT during various tasks, including a structure-from-motion task using directionally bi-stable stimuli (Dodd et al., 2001). Williams et al. (2003) suggest that their choice probabilities may have been small because the visual stimulus was not optimized for each MT neuron. Indeed, the average differential response between preferred and null direction motion with unambiguous stimuli was less than 10 spikes/s (see Figure 2a of Williams et al., 2003), which is quite low for MT neurons. Weak responses might explain why Williams et al. (2003) don't see large choice probabilities, since neurons with poor tuning will tend to have lower choice probabilities. However, it seems somewhat unlikely that this factor could account for the complete absence of choice effects in the Williams et al. (2003) study.

Finally, Thiele et al. (1999) recorded from various visual areas (including MT) while monkeys were engaged in a hybrid detection/discrimination task using moving grating stimuli presented near contrast threshold. MT neurons gave stronger responses when monkeys correctly detected motion in the neuron's preferred direction than when they incorrectly reported motion in a different direction (for the same visual stimulus). Although they quantified these effects as "choice probabilities," Thiele et al. (1999) noted that the effects were strong for stimulus motion in the preferred direction but not for motion in nonpreferred directions (quite unlike what we show in Figure 4 for our task). This led them to propose that their effects were due mainly to attention and/or stimulus expectation (i.e., choice bias), whereas our data cannot be explained by these factors. Presumably, trial-to-trial variations in spatial attention and/or expectation had much stronger effects on the responses recorded by Thiele et al. (1999) because the visual stimuli were presented near contrast threshold. This comparison highlights the importance of examining choice effects using both preferred and nonpreferred stimuli, since genuine choice effects should be observable regardless of the polarity of the stimulus.

In closing, choice probabilities have previously been used to demonstrate that single MT neurons are functionally linked to perceptual judgments of motion direction or structure-from-motion (Britten et al., 1996; Dodd et al., 2001). Our findings extend this line of work by showing that MT neurons are functionally linked to judgments of depth. More generally, our finding that choice probabilities depend on conventional tuning parameters (disparity-tuning shape) implies that these types of measurements can be used to infer the strategies by which animals extract task-relevant information from populations of neurons. This finding opens up many possibilities for experiments that examine how specific types of neurons are dynamically recruited into population codes that are appropriate for solving different types of tasks.

Experimental Procedures

Details regarding our experimental preparation and task are given in two recent publications (DeAngelis and Uka, 2003; Uka and

DeAngelis, 2003). Here, we briefly review these methods with an emphasis on the portions most relevant to the present study.

Subjects and Surgery

Physiological experiments were performed using two male rhesus monkeys (*Macaca mulatta*) weighing 5–6 kg. All animal care and experimental procedures were approved by the Institutional Animal Care and Use Committee at Washington University and were in accordance with NIH guidelines. Monkeys were prepared for behavioral training sessions by attaching a CILUX head post receptacle (Crist Instruments, Hagerstown, MD) to the monkey's skull and implanting a coil of wire under the conjunctiva of one eye for monitoring eye position (Judge et al., 1980). To reduce slippage, the eye coil was sutured to the sclera using either a permanent or long-lasting dissolvable suture (8-0 Nylon or 7-0 Dexon). Animals were allowed to recover fully for 4–8 weeks before the first behavioral training session took place. Following 3–6 months of training on the depth discrimination task (described below), a beveled CILUX recording chamber (Crist Instruments) was attached to the monkeys' skull at an angle of 25° above the horizontal and was located over occipital cortex roughly 17 mm lateral and 14 mm dorsal to the occipital ridge. A second eye coil was also implanted to allow measurements of vergence posture. After 1–2 weeks of recovery time, the animal underwent an additional training period in which vergence angle was monitored and enforced to be accurate to within $\pm 0.25^\circ$; subsequently, we started electrophysiological recordings in MT.

Visual Stimuli

The monkeys viewed visual stimuli presented on a flat-screen 22 inch color monitor (Sony GDM-F500) placed at a distance of 57 cm. The display subtended a visual angle of $40^\circ \times 30^\circ$, had a resolution of 1152×864 pixels, and was refreshed at 100 Hz. Visual stimuli were generated by an OpenGL accelerator board with hardware stereo support (Oxygen GVX1, 3D Labs, Milpitas, CA). Stereoscopic images were displayed by presenting the left and right half-images alternately at a refresh rate of 100 Hz. The monkey viewed the display through a pair of ferroelectric liquid crystal shutters (DisplayTech, Longmont, CO) that were synchronized to the video refresh such that one shutter was closed while the other was open. This produced stereo images with minimal crosstalk ($<3\%$) due solely to phosphor persistence. Precise disparities and smooth motion were achieved by plotting dots with subpixel resolution using the hardware anti-aliasing capabilities of the OpenGL accelerator board (see DeAngelis and Uka, 2003, for additional details).

All stimuli were random-dot stereograms (RDS) that were presented within a circular aperture that was matched to the size of the receptive field of the recorded neuron. The location and size of the circular RDS aperture did not vary with the disparity or binocular correlation of the dots, thus eliminating monocular cues to depth. All dots within the RDS moved coherently (100% motion coherence) at a velocity tailored to each MT neuron. Thus, dots did not disappear until they reached the boundary of the circular aperture, after which point motion resumed from the opposite side of the aperture. Dot density was 64 dots per square degree per second, with each dot subtending $\sim 0.1^\circ$. The starting position of each dot within the aperture was newly randomized for each trial ("VAR" condition), except for some trials, specifically noted in the text, in which the dot patterns were identical across trials ("NOVAR" condition, see Uka and DeAngelis, 2003, for details). Stationary background dots (in fixation trials) or flickering background dots (in discrimination trials) were presented at zero disparity to help anchor the monkey's vergence posture (gray dots in Figure 1A).

For the discrimination task, the strength of the disparity signal was varied by manipulating the percentage of binocularly correlated dots in the RDS. Correlated (i.e., "signal") dots were assigned one of two fixed disparities (crossed versus uncrossed) during each trial, and the remaining ("noise") dots were assigned random disparities within the range from -2° to 2° (Figure 1C). Dots retained their identities (signal or noise) throughout a trial; hence, the distribution of noise disparities was fixed within a given trial.

Tasks and Training

Behavioral tasks and data acquisition were controlled by a commercially available software package (Reflective Computing, Saint

Louis, MO). Monkeys were first trained on a fixation task, in which they were required to fixate on a yellow spot ($0.15^\circ \times 0.15^\circ$) within a $1.6^\circ \times 1.6^\circ$ electronic window. Monkeys received a water or juice reward for maintaining fixation throughout a 1.5 s trial. Conjugate eye position (the average of the left and right eye positions) was used to enforce fixation, and trials were aborted without reward when the monkey left the fixation window prematurely.

After fixation training, monkeys were subsequently trained on the depth discrimination task (Figure 1). An RDS containing signal dots at one of two fixed disparities was presented, and monkeys were required to report whether the signal dots were near (crossed) or far (uncrossed) by making a saccade to one of two targets (located 5° below and above the fixation point, respectively) that appeared 200 ms after offset of the RDS. Correct responses were rewarded with a drop of water or juice.

Discrimination training began with 100% binocular correlation trials, and lower correlations were gradually introduced after monkeys reached at least 75% correct. The range of correlation levels was then lowered gradually over many weeks of training until the monkey's performance reached a plateau. To discourage choice biases in the early stages of training, we used a staircase procedure in which the stimulus probabilities could be altered based upon the recent history of the monkey's choices (see Uka and DeAngelis, 2003). After choice biases were minimized, we transitioned each animal to the *method of constant stimuli*, in which a fixed set of disparities and correlation levels was presented in blocks of randomly interleaved trials. Although it was occasionally necessary to return to the staircase procedure in the days following this transition, all recording experiments were performed using the method of constant stimuli. Monkeys were extensively trained using stimuli with various directions, speeds, disparities, and locations in the visual field. This allowed us to tailor the stimulus to the preferences of each neuron under study.

Electrophysiological Recordings

Extracellular activity of single neurons was recorded using tungsten microelectrodes (Frederick Haer Inc., Bowdoinham, ME, tip diameter 7–15 μm , impedance 0.2–1 $\text{M}\Omega$ at 1 kHz) that were advanced into cortex through a transdural guide tube. Single units were isolated using a conventional amplifier, band-pass filter (500–5000 Hz), and window discriminator (Bak Electronics, Mount Airy, MD). Spike times and behavioral event markers were stored to disk with 1 ms resolution.

Area MT was recognized based on extensive experience with interpreting the pattern of gray matter and white matter regions encountered during electrode penetrations and based on the physiological response properties (direction, speed, and disparity tuning; receptive field location and size) of both single neurons and multiunit clusters (see DeAngelis and Uka, 2003). All data included in this study were derived from recordings that were confidently assigned to area MT.

Experimental Protocol

After isolating an MT neuron and carefully mapping its receptive field by hand, we conducted a set of quantitative preliminary tests to measure the direction tuning, speed tuning, and size tuning of each MT neuron (see DeAngelis and Uka, 2003, for details). Tuning curves constructed online during these tests were used to optimize the size of the RDS stimulus as well as the direction and speed of the moving dots. Next, we quantitatively measured the horizontal disparity tuning of each neuron using the optimal stimulus parameters. In most cases, disparities were tested from -1.6° to 1.6° in steps of 0.4° ; however, these parameters were adjusted as necessary based on our initial qualitative assessment of the breadth of disparity tuning. A disparity-tuning curve was plotted online based on the responses to five repetitions of each different disparity. The peak of the disparity tuning curve was located to define the "preferred" disparity, and the location of the trough of the curve was estimated to define the "null" disparity (see Figure 2A, arrowheads).

Following these preliminary tests, we recorded while the monkey performed the depth discrimination task. Both the binocular correlation of the stimulus and the binocular disparity (preferred versus null) were varied in blocks of randomly interleaved trials. Binocular correlation values were typically 0%, 1.5%, 3%, 6%, 12%, 24%,

and 48% for monkey B, and 0%, 2%, 4%, 8%, 16%, 32%, and 64% for monkey J. These ranges were determined from extensive psychophysical testing during the advanced stages of training, and the vast majority of the data was collected using this fixed set of parameters for each monkey. In a few cases, it was necessary to increase the range of binocular correlations due to poor disparity selectivity of the recorded neuron. Whenever possible, data were collected for 40 or more repetitions of each unique stimulus condition, and data sets were discarded if isolation was not maintained for at least ten repetitions. Across the range of accepted data sets, the average number of repetitions was 33 ± 10 SD, and the average number of total trials was 461 ± 139 SD.

Data Analysis

Calculation of Choice Probability

We quantified the relationship between MT responses and the animals' choices by computing "choice probabilities" using ROC analysis (Britten et al., 1996). At each disparity and binocular correlation level, the responses of the MT neuron were sorted into two groups based on the choice that the animal made at the end of each trial (preferred choices versus null choices). An ROC curve was calculated from these distributions of responses, and the area under the ROC curve gave the choice probability for that disparity/binocular correlation combination. To arrive at a single "grand" choice probability for each neuron, responses were normalized (using z scores) separately for each disparity/binocular correlation combination, and the normalized responses were then combined across stimulus conditions into a single pair of distributions for preferred and null choices (see Figure 5; Britten et al., 1996). ROC analysis on this pair of distributions yielded the grand choice probability.

The reliability of a choice probability depends on the amount of data that goes into the calculation; thus, it is important to determine the statistical significance of choice probabilities relative to the chance expectation of 0.5. We used a permutation test for this purpose. A "bootstrap" choice probability, CP_b , was calculated after we randomly reassigned each MT response to either a preferred or null choice, thus disrupting any correlation between responses and choices. This was repeated 1000 times, and we measured the absolute value of the difference between each bootstrap choice probability and the chance value of 0.5, $|CP_b - 0.5|$. The distribution of these values was compared to the absolute value of the difference between the measured choice probability, CP_m , and 0.5, $|CP_m - 0.5|$, and the proportion of bootstrap values that exceeded the measured value was taken as the p value for significance.

Measurement of Symmetry of Disparity-Tuning Curves

To measure the symmetry of each disparity-tuning curve around zero disparity, the curve was fit with a Gabor function that was constrained to have its Gaussian envelope centered at zero disparity:

$$R(d) = R_0 + A \times e^{-0.5 \frac{(d)^2}{\sigma^2}} \times \cos(2\pi f(d) + \phi), \quad (1)$$

where d is the stimulus disparity, R_0 is the baseline response level, A is the amplitude, σ is the standard deviation of the Gaussian, f is the frequency of the sinusoid, and ϕ is the phase of the sinusoid. ϕ was then wrapped into the range from 0 to $\pi/2$ (90°), and this was used as an estimate of the symmetry of the disparity-tuning curve around zero. Values close to 0 indicate even-symmetry, whereas values close to $\pi/2$ indicate odd-symmetry. We tried other measures of disparity-tuning symmetry, but we found this metric to be most reliable.

Statistics

All statistical analyses were done using STATISTICA (StatSoft, Inc.) software. To account for differences between the two monkeys in our study, all correlation analyses were done as within-cell regressions in the context of an analysis of covariance (ANCOVA) with monkey identity as an independent factor. Multiple regression analyses also took into account differences between monkeys, using appropriate dummy variables. For all parametric statistics, we log-transformed variables whenever this made the distributions closer to normal. We also verified that none of our conclusions would change when nonparametric statistics were used (Spearman rank correlations).

Acknowledgments

We thank Amy Wickholm and Heidi Loschen for excellent technical support and monkey training. We are grateful to Larry Snyder for valuable comments on the manuscript. This work was supported by the National Eye Institute (EY-013644), by a Searle Scholar Award from the Kinship Foundation, and by a Career Award in the Biomedical Sciences from the Burroughs-Wellcome Fund. T.U. was supported by the JSPS Research Fellowship for Young Scientists and by a Long-Term Fellowship from the Human Frontier Science Program.

Received: October 13, 2003

Revised: January 8, 2004

Accepted: March 5, 2004

Published: April 21, 2004

References

- 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.
- Basso, M.A., and Wurtz, R.H. (1998). Modulation of neuronal activity in superior colliculus by changes in target probability. *J. Neurosci.* *18*, 7519–7534.
- Bradley, D.C., Chang, G.C., and Andersen, R.A. (1998). Encoding of three-dimensional structure-from-motion by primate area MT neurons. *Nature* *392*, 714–717.
- 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.
- 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.
- Coe, B., Tomihara, K., Matsuzawa, M., and Hikosaka, O. (2002). Visual and anticipatory bias in three cortical eye fields of the monkey during an adaptive decision-making task. *J. Neurosci.* *22*, 5081–5090.
- Cook, E.P., and Maunsell, J.H. (2002). Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. *J. Neurosci.* *22*, 1994–2004.
- Dean, A.F. (1981). The variability of discharge of simple cells in the cat striate cortex. *Exp. Brain Res.* *44*, 437–440.
- DeAngelis, G.C., and Uka, T. (2003). Coding of horizontal disparity and velocity by MT neurons in the alert Macaque. *J. Neurophysiol.* *89*, 1094–1111.
- DeAngelis, G.C., Cumming, B.G., and Newsome, W.T. (1998). Cortical area MT and the perception of stereoscopic depth. *Nature* *394*, 677–680.
- 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.
- Gold, J.I., and Shadlen, M.N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* *5*, 10–16.
- 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.
- Horwitz, G.D., and Newsome, W.T. (1999). Separate signals for target selection and movement specification in the superior colliculus. *Science* *284*, 1158–1161.
- Judge, S.J., Richmond, B.J., and Chu, F.C. (1980). Implantation of magnetic search coils for measurement of eye position: an improved method. *Vision Res.* *20*, 535–538.
- Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* *2*, 176–185.
- Maunsell, J.H., and Van Essen, D.C. (1983). Functional properties of neurons in middle temporal visual area of the macaque monkey. II. Binocular interactions and sensitivity to binocular disparity. *J. Neurophysiol.* *49*, 1148–1167.
- Parker, A.J., and Newsome, W.T. (1998). Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* *21*, 227–277.
- 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.
- Poggio, G.F., and Fischer, B. (1977). Binocular interaction and depth sensitivity in striate and prestriate cortex of behaving rhesus monkey. *J. Neurophysiol.* *40*, 1392–1405.
- Poggio, G.F., Gonzalez, F., and Krause, F. (1988). Stereoscopic mechanisms in monkey visual cortex: binocular correlation and disparity selectivity. *J. Neurosci.* *8*, 4531–4550.
- Schiller, P.H., Finlay, B.L., and Volman, S.F. (1976). Short-term response variability of monkey striate neurons. *Brain Res.* *105*, 347–349.
- Seidemann, E., and Newsome, W.T. (1999). Effect of spatial attention on the responses of area MT neurons. *J. Neurophysiol.* *81*, 1783–1794.
- Seidemann, E., Zohary, E., and Newsome, W.T. (1998). Temporal gating of neural signals during performance of a visual discrimination task. *Nature* *394*, 72–75.
- Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* *86*, 1916–1936.
- 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.
- Thiele, A., Distler, C., and Hoffmann, K.P. (1999). Decision-related activity in the macaque dorsal visual pathway. *Eur. J. Neurosci.* *11*, 2044–2058.
- Tolhurst, D.J., Movshon, J.A., and Dean, A.F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* *23*, 775–785.
- 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. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* *382*, 539–541.
- 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.
- Uka, T., and DeAngelis, G.C. (2003). Contribution of middle temporal area to coarse depth discrimination: comparison of neuronal and psychophysical sensitivity. *J. Neurosci.* *23*, 3515–3530.
- 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.
- Zohary, E., Celebrini, S., Britten, K.H., and Newsome, W.T. (1994). Neuronal plasticity that underlies improvement in perceptual performance. *Science* *263*, 1289–1292.