

# A Role for Neural Integrators in Perceptual Decision Making

Mark E. Mazurek, Jamie D. Roitman, Jochen Ditterich and Michael N. Shadlen

Howard Hughes Medical Institute, Department of Physiology and Biophysics, and National Primate Research Center, University of Washington, Seattle, WA 98195-7290, USA

**Decisions based on uncertain information may benefit from an accumulation of information over time. We asked whether such an accumulation process may underlie decisions about the direction of motion in a random dot kinetogram. To address this question we developed a computational model of the decision process using ensembles of neurons whose spiking activity mimics neurons recorded in the extrastriate visual cortex (area MT or V5) and a sensorimotor association area of the parietal lobe (area LIP). The model instantiates the hypothesis that neurons in sensorimotor association areas compute the time integral of sensory signals from the visual cortex, construed as evidence for or against a proposition, and that the decision is made when the integrated evidence reaches a threshold. The model explains a variety of behavioral and physiological measurements obtained from monkeys.**

## Introduction

A hallmark of higher brain function is a capacity to act and sense on different time scales. This capacity allows organisms to acquire sensory information from one or more sources, combine information across time, and alter their behavior some time later, perhaps contingent on other intervening input. This contrasts with simpler reflexive behaviors in which sensory information precipitates an immediate motor response. From this perspective, the capacity for higher brain function would seem to rest on the brain's ability to accumulate, combine and preserve information over time.

The persistent spike activity that is commonly observed from neurons in the association cortex is thought to provide a neural substrate by which information can span such time gaps. It has been implicated in working memory (Fuster and Alexander, 1971; Fuster, 1973; Miyashita and Chang, 1988; Funahashi *et al.*, 1989), motor planning (Evarts and Tanji, 1976; Bruce and Goldberg, 1985; Gnadt and Andersen, 1988) and decision making (Kim and Shadlen, 1999; Shadlen and Newsome, 2001; Hernandez *et al.*, 2002). A reasonable hypothesis is that such persistent activity represents the outcome of a process of integration with respect to time—that is, the accumulation and storage of information. By analogy to neural integrators in the brainstem that convert eye velocity to position (Robinson, 1989; Fukushima *et al.*, 1992), our hypothesis is that cortical neural integrators combine evanescent sensory data to generate an evolving conception about the state of the world, which can then be used to plan appropriate behaviors.

In this paper we explore the idea that the accumulation of sensory information underlies a simple perceptual decision. We review evidence suggesting that rhesus monkeys make decisions about the direction of random dot motion by integrating, with respect to time, the information they receive through direction-selective neurons in the visual cortex. We test this idea by developing a model of a neural circuit that we believe may perform

such a computation. The model encompasses three stages of neural processing: (i) the representation of visual motion by ensembles of direction-selective neurons; (ii) the representation of a decision variable by ensembles of neurons that accumulate input from the first stage; and (iii) a comparison to threshold, which terminates the process and determines the model's choice. The neurons that comprise the first stage are based on the known properties of neurons in area MT. The second and third stages of the model predict the response properties of neurons in LIP and behavioral measurements obtained from monkeys. Our results show that temporal integration can explain a fairly wide range of behavioral and physiological observations, suggesting that this simple computation may play an essential role in sensorimotor decisions.

## Background

The processes underlying sensorimotor decisions in humans and monkeys have been investigated using a motion discrimination task (Fig. 1) (Britten *et al.*, 1992; Britten, 2003). The sensory stimulus in the task consists of a dynamic display of dots which appear and disappear at random locations within a 5–10° circular aperture. A fraction of these dots are displaced at some fixed offset to impart an overall sense of motion in one direction or its opposite (e.g. left versus right). The subject's goal is to discriminate between these two alternatives. Rhesus monkeys are trained to indicate their choice of direction by making a saccadic eye movement to one of two targets, corresponding to the two motion directions. If the fraction of moving dots, termed the percent coherence, is sufficiently high, it is easy to identify the correct direction. By controlling the percent coherence, the task can be made arbitrarily difficult or easy.

In this study we refer to two versions of the random dots task. In the reaction time (RT) task (Fig. 1A), subjects are allowed to report their decision as soon as they are ready, once the random dots appear. In the fixed duration (FD) task (Fig. 1B), subjects view the random dot motion stimulus for a specified duration (usually 1 or 2 s) that is under the control of the experimenter. In the FD task, motion viewing is typically followed by a memory/delay period before the subject is allowed to respond.


The motion information in the random dots task is represented by direction selective neurons in the extrastriate visual cortex, especially area MT (also known as area V5). When random dot motion appears in the receptive field of an MT neuron, there is a large initial burst followed by a sustained response whose magnitude depends on the strength and direction of motion (Fig. 2A). Evidence from microstimulation, lesion, and neural recording experiments demonstrate that such responses underlie the monkey's judgment in the random dots task (Newsome and Paré, 1988; Salzman *et al.*, 1990; Britten *et al.*, 1992, 1993, 1996; Bisley *et al.*, 2001; Britten, 2003). In

particular, the animal's ability to discriminate weak motion appears to be limited by variability in the response of MT neurons (Parker and Newsome, 1998). The responses of MT neurons provide the evidence upon which the monkey bases its decision about direction, but they do not gather this evidence together to form a decision. Their activity fluctuates with each passing random dot and the evidence they provide is as evanes-

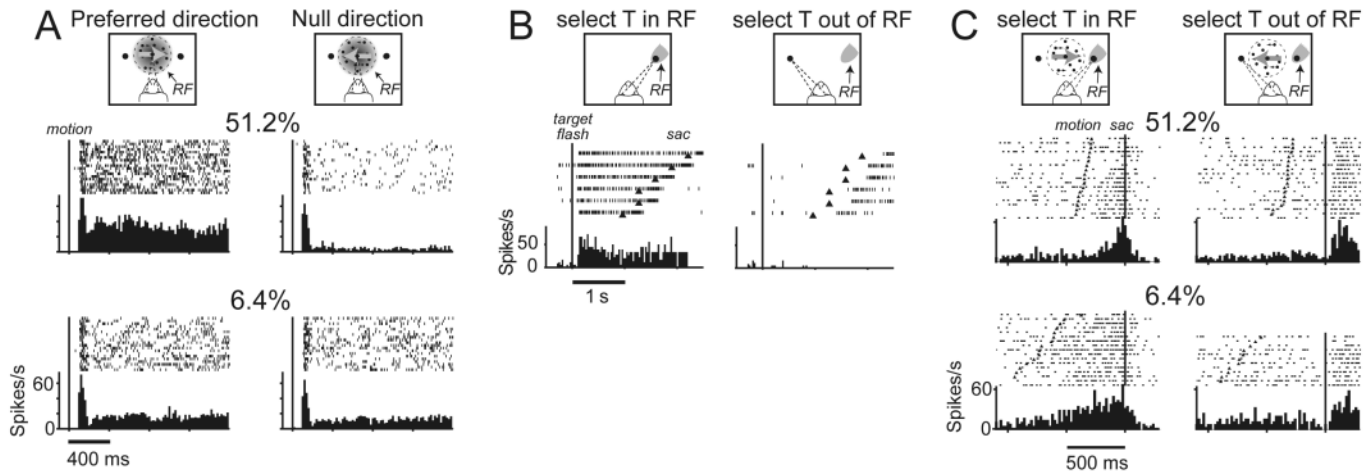
cent. To reach a decision, other neurons must 'read out' the evidence from area MT.

Less is known about how sensory responses are read out and interpreted in order for the brain to reach a categorical decision. In the motion discrimination task described above, a decision about motion is linked to an eye movement response. Therefore, neurons that represent a plan to move the eyes to a particular place in space are potential sites for the representation of a decision about motion direction. Such neurons have been found in the parietal cortex, prefrontal cortex and superior colliculus (Wurtz and Albano, 1980; Gnadt and Andersen, 1988; Funahashi *et al.*, 1989; Glimcher and Sparks, 1992). These neurons exhibit persistent activity when the monkey is instructed to plan an eye movement to a restricted part of the visual field, termed the response field (RF) (Fig. 2B). In the random dots task, these neurons have been shown to represent information about the monkey's decision (Shadlen *et al.*, 1996; Kim and Shadlen, 1999; Horwitz and Newsome, 2001; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002). Here, the direction of motion effectively instructs an eye movement to one of a pair of choice targets (Fig. 1). When the monkey decides that the direction is the one indicated by an eye movement to the RF, the spike rate gradually increases (Fig. 2C). For the opposite decision, the spike rate declines during the period of motion viewing. Such activity represents the evolving plan of how to answer the question, 'Which direction is the motion?'

In this study we use simulated neural responses to test the notion that temporal integration of sensory signals from area MT forms the basis for direction discrimination in the random dots task, and that neurons with persistent activity preceding eye



**Figure 1.** Direction discrimination tasks. Monkeys are trained to discriminate the direction of motion in a dynamic random dot display. An eye movement signals the choice of direction. Task difficulty is governed by the fraction of coherently moving random dots. (A) Reaction time (RT) task. The subject is allowed to respond at any time after the onset of random dot motion to indicate the direction of perceived motion. The reaction time (RT) is measured as the time between motion onset and the initiation of the saccade. (B) Fixed duration (FD) task. The experimenter controls the duration of the random dot motion, which is typically 1 or 2 s. In some experiments, the subject maintains fixation through a variable delay period after the motion stimulus is extinguished. The instruction for the response is the disappearance of the fixation point, after which the monkey makes an eye movement to one of the targets to indicate the direction of perceived motion. In Roitman and Shadlen (2002), RT and FD tasks used different fixation point color and were performed in alternating blocks.



**Figure 2.** Responses of neurons in areas MT and LIP. (A) Response of a direction selective MT neuron in the direction discrimination task. The random dot stimulus was centered on the neuron's receptive field (RF). The direction of motion was either in the best direction for activating the neuron, called the preferred direction, or in the opposite direction, the null direction. Spike rasters from individual trials and peristimulus time spike rate averages are shown for a strong motion stimulus (51.2% coherence) and a weak motion stimulus (6.4% coherence). Mean spike rate is calculated in 20 ms bins. The solid line in each panel shows the time of onset of the random dot motion. These data were recorded in a 2 s FD version of the task (Britten *et al.*, 1992). Data obtained from the Newsome database (<http://www.cns.nyu.edu/~wyeth/data/newsome/newsome.html>). (B) Response of an LIP neuron in a memory guided eye movement task. In this task a target is flashed on and off in the periphery. The monkey must make an eye movement (sac) to the remembered location of the target after a variable delay. In some locations, the onset of the target evokes a response from the neuron that is maintained until the time of the saccade, even when the target is absent. This location is designated as the neuron's response field (RF). The panels show the responses of an LIP neuron when the target appears in the neuron's RF and when it appears in the opposite hemifield. The data format is the same as in (A). The solid line indicates the onset of the target, and the filled triangles show the times of the saccade. Mean spike rate is calculated in 30 ms bins. Adapted from Shadlen and Newsome (2001). (C) Response of an LIP neuron in the direction discrimination task. The visual objects were arranged so that one of the two choice targets appears in the neuron's response field (RF) and the other target and the random dot motion appear outside the RF. The panels show responses when the motion instructed eye movements into the RF or away from the RF, for strong motion (51.2% coherence) and weak motion (6.4% coherence). The data are shown in the same format as in (A) except that the spikes and the mean spike rate are aligned relative to the time of the saccade (indicated by the solid vertical line). The filled triangle shows the time of motion onset in the trial. These data were collected while the monkey performed the RT motion discrimination task. The trials are arranged by ascending reaction time, which is the time between motion onset and saccade initiation. Mean spike rate is calculated in 20 ms bins. Adapted from Roitman and Shadlen (2002).

movements represent this integration. For the latter, we focus on neurons in the lateral intraparietal area (LIP), which have been studied using both FD and RT versions of the task. We begin by simulating the well-described responses of MT neurons to random dot motion, and we model LIP responses as a simple time-integral of the difference from opposing pools of MT neurons. We simulate two ensembles of LIP neurons, one for each of the possible eye movement responses. The decision is made when the activity of one of the LIP ensembles exceeds a critical value, that is, when the evidence reaches a threshold. The model furnishes several novel interpretations of behavioral and physiological measurements obtained using the random dots task.

The electrophysiological and behavioral data in this paper have been described in earlier publications (Britten *et al.*, 1992; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002).

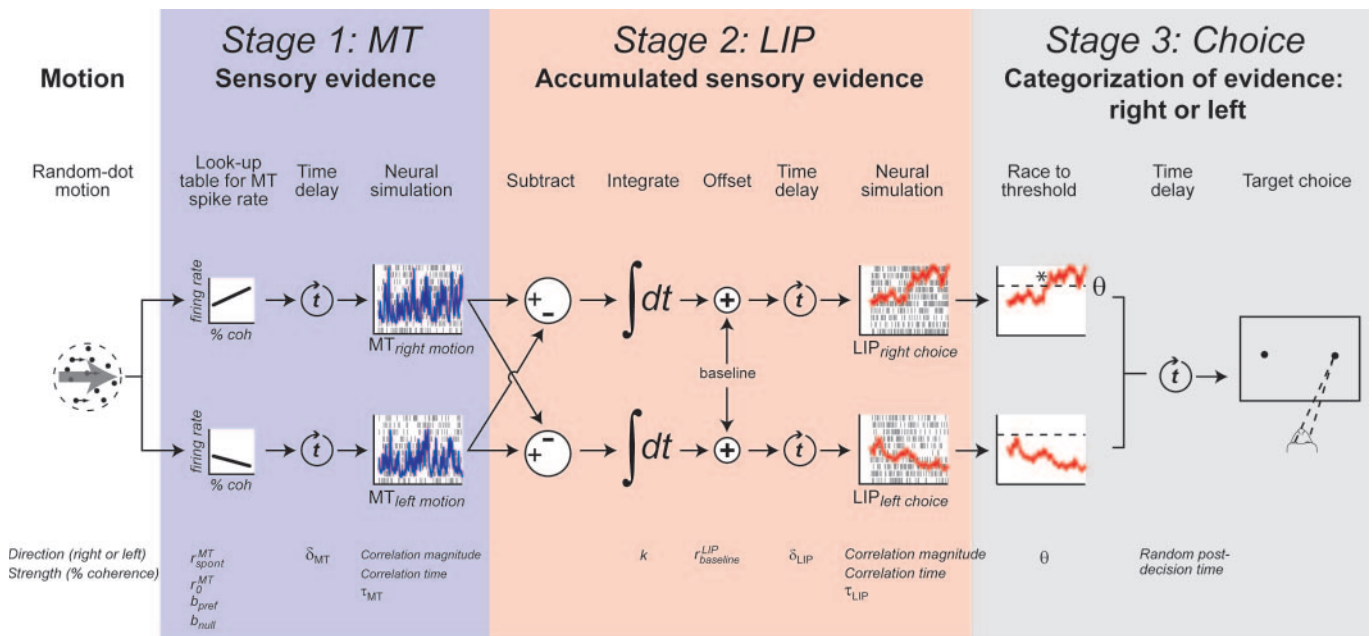
### Model Design

The goal of this modeling exercise is to investigate how well integration of sensory signals can account for behavioral and neural data obtained in the random-dot motion task. We address this question by simulating the responses from two brain regions: the middle temporal area (MT or V5) and the lateral intraparietal area (LIP) (Fig. 3). Each brain region consists of an ensemble of simulated cortical neurons whose firing rates are determined by the model but whose actual spike times are random.

The model is divided into three stages: representation of the evidence, accumulation of the evidence into decision variable,

and comparison of the decision variable to threshold (choice). In the first stage, we simulate the responses of two opposed pools of direction selective MT neurons to random dot motion of varying motion strength. For simplicity, all simulations assume that the motion is either to the right or to the left. Each neuron in the two pools produces a sequence of spikes with an expected rate proportional to the strength of motion, based on values from Britten *et al.* (1993). The averaged spike rate from the two pools furnishes the output of stage 1.

In the second stage we simulate the responses of two pools of LIP neurons. One of the LIP pools represents a plan to choose the right choice target (Fig. 3,  $LIP_{\text{right choice}}$ ), whereas the other represents the opposite plan (Fig. 3,  $LIP_{\text{left choice}}$ ). Unlike the MT neurons, the expected firing rates of these neurons are time-dependent, determined by the time integral of the difference in the output of the right and left MT pools. The expected LIP firing rate is calculated by integrating the difference in spike rate signals from MT starting from when the coherence-dependent MT response begins ( $\delta_{\text{MT}}$ ), then delaying the result by  $\delta_{\text{LIP}}$ ; we do not simulate the LIP response before the integration begins at time  $\delta_{\text{MT}} + \delta_{\text{LIP}}$ . Because of the temporal fluctuations in the responses of the MT pools [ $r_{\text{right}}^{\text{MT}}$  and  $r_{\text{left}}^{\text{MT}}$ ; see equations A1 and A2, Supplementary Material], the expected LIP responses [ $\langle r_{\text{right}}^{\text{LIP}}(t) \rangle$  and  $\langle r_{\text{left}}^{\text{LIP}}(t) \rangle$ , equations A3 and A4, Supplementary Material] exhibit random deviations around the motion-dependent drift, similar to a biased diffusion process. As with the MT stage, the averaged spike rate from the two pools of LIP neurons furnishes the output of this stage of the model.



**Figure 3.** Model of the decision process. The input to the model is a two-parameter description of the random dot motion stimulus shown on one experimental trial: direction and strength (% coherence). The model predicts the behavioral choice (left or right) and the response time. It also represents neural activity which is intended to simulate the responses of direction selective neurons in area MT and decision related neurons in area LIP, as shown in Figure 2. The elements shown in the schematic correspond to mathematical operations that are believed to underlie the neural and behavioral responses. The order of several (linear) operations can be combined or exchanged, but the grouping into stages (colored background) corresponds to sensory representation, accumulation of evidence for two competing hypotheses, and comparison of the accumulated evidence to a threshold. Labels below the elements in the schematic show the variable names which appear in the text equations; corresponding values are given in Table A1 (Supplementary Material). The schematic demonstrates a trial with strong rightward motion. Each Neural simulation block shows simulated spikes from 10 neurons (tic marks) and the smoothed average spike rates from the full ensemble of 100 simulated neurons (solid curves). The Race to threshold elements show the smoothed average rates from the two LIP ensembles racing against each other toward the decision threshold; \* indicates the time of decision in the winning LIP.

For simulations of the reaction time (RT) task, the LIP responses are compared to a decision threshold. The two ensemble average spike rates are smoothed using a first order filter with time constant  $\tau_{LIP} = 0.1$  s in order to tame the moment-to-moment fluctuations. The smoothed LIP signals race against each other to provide the weight of evidence for their preferred choice direction. The first to reach the decision threshold,  $\theta$ , determines the target choice and the decision time. A random time is then added before saccade initiation. Note that the reaction time in a trial is the sum of decision time plus the non-decision intervals [saccadic latency,  $\delta_{MT}$  and  $\delta_{LIP}$ , amounting to  $\sim 300$  ms non-decision time on average (Luze, 1986)].

For simulations of the fixed duration (FD) task, the decision is generated in one of two ways. In the first method, LIP signals are allowed to integrate for the full duration of the simulation. The decision in this case is determined by simply comparing the value of the smoothed LIP signals corresponding to the two possible choices: whichever one is higher at the end of the trial determines the choice. In the second method, a decision threshold ( $\theta$ ) is applied, and if one of the LIP signals crosses  $\theta$ , the decision is determined as in the RT simulation. If neither LIP signal crosses  $\theta$  by the end of the simulation, the decision is determined by whichever of the LIP signals is higher.

A complete description of the model parameters along with a justification of the values we used can be found in the Supplementary Material. Importantly, although there are many parameters, most are constrained by measurement. The threshold,  $\theta$ , is the only parameter in the model that was fit to match behavioral accuracy (see Supplementary Material).

## Results

We developed a model consisting of ensembles of spiking neurons in order to simulate a process of decision formation in a two-alternative forced choice direction discrimination task. The model represents motion ‘evidence’ in the response of ensembles of simulated MT neurons, and it represents the accumulation of this evidence in the response of ensembles of simulated LIP neurons. The commitment to one or the other alternative is based on the read out of the ensemble spike rate signals in area LIP. Figure 3 shows examples of the smoothed ensemble-average firing rates for MT (blue) and LIP (red). The filters used to smooth the firing rates in this figure, as well as the relative height of  $\theta$ , are the same as those used to produce the predictions in this study. Note that the smoothed LIP signals, which function to determine the time and direction of the decision, possess a sizeable degree of temporal noise resulting from both the variability of neural spiking and a weak tendency for pairs of neurons to emit spikes within  $\pm 10$  ms of one another (Mazurek and Shadlen, 2002). This temporal noise is not a consequence of variability in the visual stimulus or in the input from MT, but rather it is inherent in the decision signal.

The results are presented in three parts. First, we compare the predictions of the model to behavioral and physiological data obtained in the reaction time (RT) direction discrimination task, where the subject responds as soon as they arrive at a decision. We then test whether the assumptions of the model hold in the fixed duration (FD) task, where subjects must view the motion stimulus for a specified duration. Finally, we examine the predictions of the model for the relationship between single unit responses and behavior.

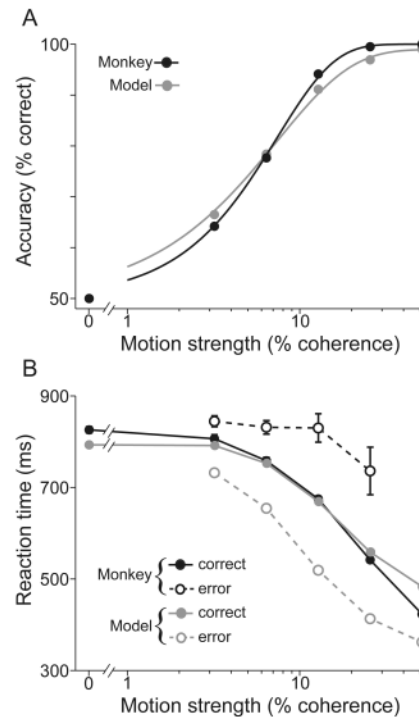
## Reaction-time Direction Discrimination

The combined measurement of RT and perceptual choice permits an examination of the neural responses that underlie decision formation in the time frame used by the animal to solve the task. The task furnishes two behavioral measures, choice and response time, which must be accounted for by a model of the decision process.

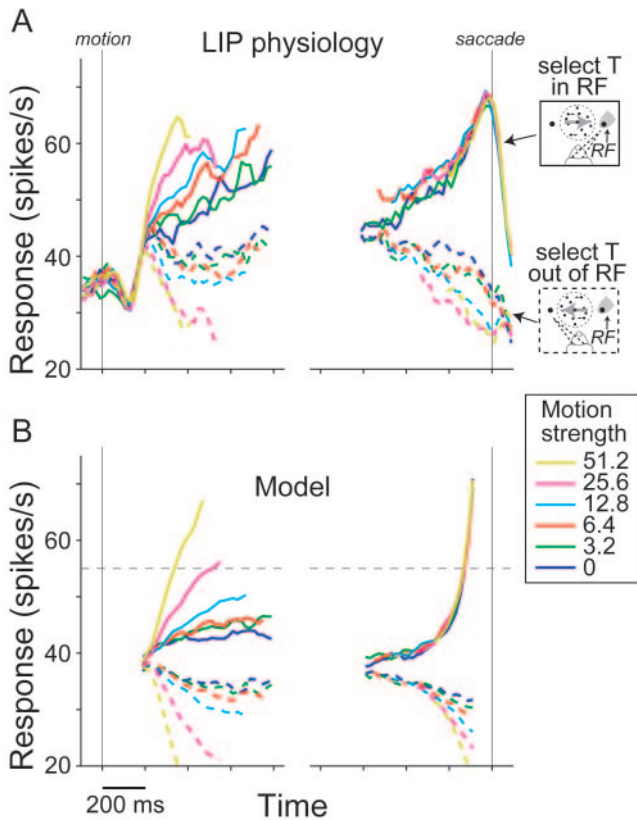
### Behavior: Accuracy and Decision Time

The model in Figure 3 works by integrating the sensory evidence—the difference in ensemble spike rates from area MT representing the alternative directions—toward a threshold level ( $\theta$ ). The level of this threshold controls the model’s accuracy: a higher threshold requires more evidence to be accumulated before a response, leading to fewer errors than would be obtained with a lower threshold. Importantly,  $\theta$  is the only parameter in the model that was fit to behavioral data. We adjusted  $\theta$  to match the overall accuracy seen in the monkey. At this level, the model produces a psychometric function with similar shape to that measured experimentally (Fig. 4A).

To achieve this level of accuracy in the model,  $\theta = 55$  spikes/s, which is 15 spikes/s above the simulated baseline (see Fig. 5B).



**Figure 4.** Behavior of monkey and model in the RT task. The data from the monkey show averages from all experiments for which there are LIP recordings in the RT task. (A) Psychometric functions. The percentage of correct choices is plotted as a function of motion strength (% coherence). The black points show data from the monkey. The gray points show predicted accuracy from the model. Sigmoid curves are best fitting Weibull functions,  $p = 1 - 0.5 \exp(-[C/\alpha]^\beta)$ , where  $C$  is the motion strength and  $\alpha$  and  $\beta$  are the threshold and slope of the Weibull function. The average accuracy across all conditions is 80% correct for the monkey. This is the target for setting the decision threshold ( $\theta$ ) in the model (see Model Design). (B) Chronometric functions. Reaction time is plotted as a function of motion strength. Reaction time is defined as the time from motion onset to saccade initiation. Black points show data from the monkey; correct and error trials are as indicated. Error bars ( $\pm 1$  SEM) are smaller than the points for the correct data. Gray points show predicted mean reaction times from the model.



**Figure 5.** Time course of LIP responses in the RT direction discrimination task. (A) Average spike rate from 54 LIP neurons. Responses are grouped by motion strength and choice as indicated by color and line type, respectively. The left portion of the graph shows responses aligned in time with the onset of random dot motion (vertical line). The responses in this portion of the graph are plotted up to the time of the mean reaction time at that coherence, excluding any activity within 100 ms of saccade initiation. The right portion of the graph shows responses aligned in time with the initiation of the saccadic eye movement response (vertical line). The neural responses in this portion of the graph are plotted back in time from the saccade up to the mean reaction time for that coherence, excluding any activity within 200 ms of motion onset. The spike rates are smoothed with a running mean ( $\pm 30$  ms). Error trials are not included in this analysis. Adapted from Roitman and Shadlen (2002). (B) Responses of model LIP neurons. The simulated responses are graphed using the same conventions as in (A). The model output is undefined in the first 200 ms after onset of motion and in the epoch after the decision is made. The horizontal dashed line shows the height of the decision threshold ( $\theta$ ), 55 spikes/s.

The particular values for spike rate are not critically important for explaining accuracy since they reflect our choice for the baseline spike rate ( $r_{\text{baseline}}^{\text{LIP}}$ ) and the scaling of the integral ( $k$ ), which were chosen to mimic the response of neurons in LIP (see next section). More fundamentally, the excursion from baseline to threshold amounts to a net excess of  $\sim 3$  spikes per MT neuron from the rightward- or leftward-preferring pool. This characterization of the decision threshold holds regardless of the particular values for  $r_{\text{baseline}}^{\text{LIP}}$  and  $k$ . Importantly, the same value for  $\theta$  is applied for all motion strengths.

With the decision threshold fixed to approximate the monkey's overall accuracy, the model predicts the amount of time it takes to reach a decision. Reaction times for monkey and model are depicted in Figure 4B as a function of motion strength (the chronometric function). For correct choices (Fig. 4B, solid lines), the predicted mean RTs closely match the mean RTs observed in the monkey. Of course, the total RT includes delays

besides the decision time (i.e. visual latencies  $\delta_{\text{MT}}$  and  $\delta_{\text{LIP}}$  and saccade preparation time) whose properties are not well understood, so the precise values of RT predicted by the model should be taken as approximations. More significant is the observation that RTs in the model decrease with motion strength in a similar manner to that seen in the data, despite the fact that the non-decision times in the model are independent of coherence. This suggests that the dependence of RT on motion strength arises primarily from different rates of evidence accumulation at different coherences.

While integration to threshold accounts for accuracy and RT on correct choices, it fails to account for the pattern of RT on error trials. As shown in Figure 4B (gray curves), the model predicts that error RTs will be faster than correct RTs at each coherence, which is inconsistent with data (black curves). These fast errors appear to be caused by early threshold crossings, which are caused by the instantaneous fluctuations in the model LIP activity. We will consider possible solutions to this shortcoming in Discussion. Such errors constitute only  $\sim 14\%$  of trials.

#### LIP Physiology

According to our hypothesis, neurons that represent the read-out of area MT compute the integral of the difference between opposing direction signals. Here we examine how well this idea predicts neural responses seen in area LIP. Figure 5 shows averaged responses from 54 neurons in area LIP recorded during the combined direction-discrimination RT task (Roitman and Shadlen, 2002). Recall that the task is arranged so that one of the choice targets is in the neuron's RF. The responses increase gradually before saccades to the target in the RF, and they decrease gradually before saccades to the target outside the RF (Fig. 5A). The rate of increase or decrease initially depends on the motion strength: when responses are aligned to the onset of the motion (Fig. 5A, left panel), the spike rates diverge more rapidly for strong motion than for weak motion. This part of the graph emphasizes the epoch of decision formation. When we focus instead on the epoch immediately preceding the behavioral choice (Fig. 5A, right panel), the dependence of the responses on motion strength appears greatly reduced. Indeed, on trials when the monkey chose the target in the RF (solid curves), the responses appear to converge to a common point, independent of the motion coherence. In summary, the data show that (i) different strengths of motion result in different rates of accumulation; and (ii) the responses converge to a common point before decisions that motion is toward the RF but not before decisions in favor of the opposite direction. These properties are explained by the model.

The model produces a family of predicted LIP responses that displays many of the same properties as seen in neural recordings from LIP. Figure 5B shows the simulated responses of LIP neurons that would signal a rightward choice. In other words, the RF of these model neurons is aligned to the rightward choice target. When aligned to the onset of motion (Fig. 5B, left panel), the average spike rates of these model neurons trace out ramp-like linear trajectories. The slopes of the ramps are positive (increasing) for trials leading to rightward choices and negative (decreasing) for leftward choices, and the ramps are steeper for stronger motion than for weaker motion. When the responses are aligned to the time of the motor response (Fig. 5B, right panel), the increasing responses for rightward choices look nearly identical, whereas responses associated with leftward



choices do not achieve a stereotyped level of activity but instead retain their dependency on motion strength.

These properties are explained by the underlying computations in the model. Responses aligned to the beginning of the trial reflect the integral of the difference between rightward- and leftward-preferring MT ensemble responses. Because the magnitude and sign of the MT difference signal varies with the strength and direction of motion, the average rate of rise or decay in the LIP response also depends on motion strength and direction. In contrast, during the time period immediately preceding the motor response, the responses of rightward-choice LIP neurons during rightward-choice trials appear similar for all motion strengths. This is because the responses are constrained to cross a common decision threshold  $\sim 100$  ms before saccade initiation regardless of whether the threshold was approached slowly or quickly. Indeed, the precise moment of this crossing is often determined by a brief upward fluctuation in the ensemble LIP signal. On the other hand, rightward-choice neurons are not constrained to cross a common threshold during leftward choices because decision time in that case is governed by the opposing population of LIP neurons, whose RFs align with the left-choice target. Thus, the average responses from rightward-choice neurons (dashed curves) retain a dependency on motion strength even at the end of the decision process culminating in a leftward choice. The model thus explains this puzzling asymmetry that is evident in the data (Fig. 5A).

One somewhat surprising observation is that the simulated response rates achieve peak values well above the decision threshold  $\theta$  (Fig. 5B, right panel). This occurs because the LIP ensemble activity is smoothed in the formation of the decision signal, which is then compared to decision threshold. Such smoothing delays the growth of the decision signal relative to the spike rates of the neurons within LIP. As a result, the spike rates of single LIP neurons continue their trajectory beyond the threshold by the time the decision is made. The degree of overshoot depends on the value of  $\tau_{LIP}$ , but the effect is seen with even modest smoothing. For  $\tau_{LIP} = 100$  ms, the model achieves spike rates of 65–70 spikes/s at the end of the simulation,  $\sim 100$  ms before the saccade, similar to the spike rates seen in the data at this time.

### Fixed-duration Direction Discrimination

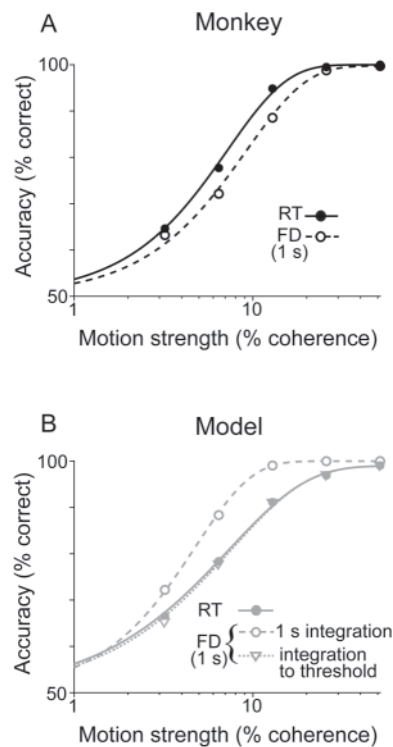
#### Accuracy

Up to this point we have mainly discussed the reaction time version of the random dot motion task, where subjects report their decision as soon as possible. Random dot motion discrimination has been studied more extensively under conditions in which the subject views the motion stimulus for a duration that is controlled by the experimenter. In these ‘fixed duration’ (FD) experiments, it is commonly assumed that the subject makes a decision using all the available motion information—that is, that the decision occurs after the motion is turned off. This idea receives experimental support, especially for short stimulus durations in the range of 100–800 ms (Gold and Shadlen, 2000; see also Britten *et al.*, 1992; Burr and Santoro, 2001). This assumption leads to one obvious prediction: accuracy in 1 s FD experiments should be better than the accuracy in RT experiments because the animals typically view motion for less than a second in the RT task.

However, the experimental evidence contradicts this prediction. Figure 6A shows psychometric functions obtained from experiments in which monkeys alternated blocks of trials in the RT task with blocks of trials in a 1 s FD task. The fits are very similar; in fact, performance on the FD task is slightly worse ( $P < 0.05$ ). This is remarkable considering that the monkey took an average of only 680 ms to initiate an eye movement response in the RT task. Subtracting out visual latencies and motor preparation time, this finding implies that the monkey performed about equally well in the RT task as it did in the FD task even though it had about half the viewing time. Clearly, this is at odds with the most straightforward version of the integration model. If the monkey has twice the viewing time, we would expect an improvement in signal to noise of  $\sim \sqrt{2}$ , which would translate to a left shift of the PMF by about this same amount (Fig. 6B).

The remarkable similarity between PMFs on the RT and 1 s FD tasks raises the possibility of a common mechanism. Suppose that instead of using all available evidence during the 1 s motion viewing period, the brain accumulates evidence to a threshold, just as in the RT task, and thereafter stops incorporating motion evidence into the decision variable. This idea, first proposed by Link (Link and Heath, 1975; Link, 1992), leads to several predictions about behavior and neural activity.

First, the hypothesis that both RT and FD tasks use the race-to-threshold mechanism predicts that the performance in the FD



**Figure 6.** Comparison of accuracy in RT and FD tasks. (A) Psychometric functions from alternating blocks of RT and 1 s FD experiments. Data are averages from all experiments in two monkeys (see Roitman and Shadlen, 2002). Accuracy was slightly better on the RT task ( $P < 0.05$ ). Sigmoid functions are best fitting Weibull functions. (B) Model predictions. The dashed curve shows predicted performance based on the assumption that sensory information is integrated throughout the 1 s period of motion viewing. The two curves that nearly superimpose are predictions based on the integration-to-threshold model. For the FD task (dotted curve), a decision is made as if the monkey were surreptitiously performing the RT task (see Model Design). The solid curve is the same as in Figure 4A.

task should be similar to that obtained in the RT task so long as the decision threshold ( $\theta$ ) is the same (Fig. 6*B*; compare RT prediction against the FD prediction for integration to threshold). In fact, accuracy should be slightly worse on the FD task because a decision must be rendered from 1 s of stimulus viewing, even if the evidence is  $<\theta$ . In principle, there is no reason to set the threshold to the same level for both FD and RT tasks, but we suspect that this occurs because the monkeys performed RT and FD experiments in alternating blocks of trials on the same day. The critical point is that even when given a fixed amount of time to solve the task, the monkey might operate in the same fashion as in the RT task. We next consider the implications of this hypothesis for neural activity in LIP.

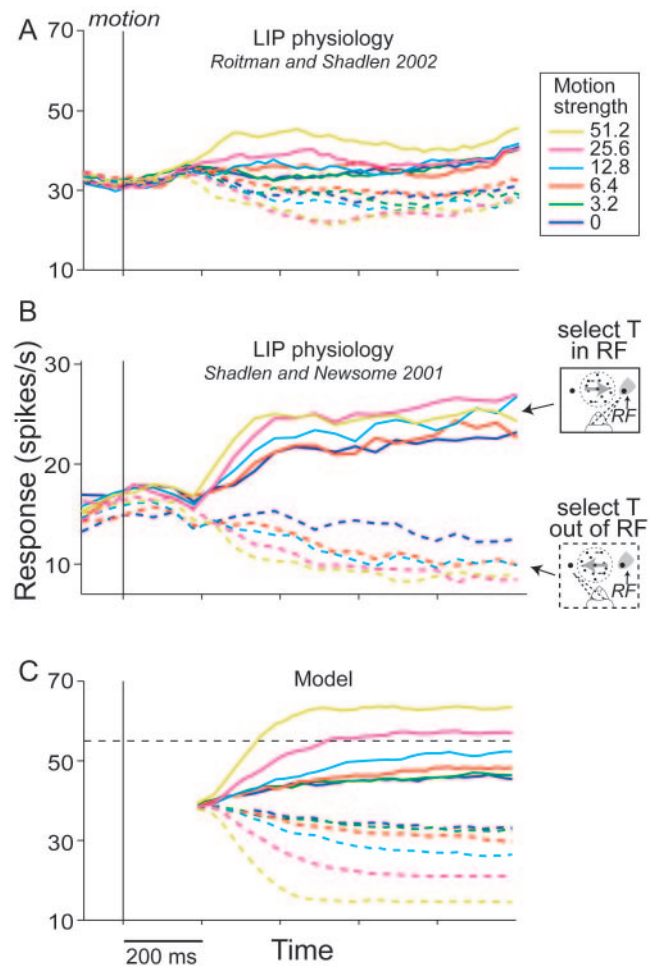
### LIP Physiology

Figure 7*A* (adapted from Roitman and Shadlen, 2002) shows the responses of LIP neurons recorded in 1 s FD experiments; these are the same neurons whose responses in the RT task are plotted in Figure 5*A*. The responses during FD trials begin with trajectories that are similar to those recorded in RT trials, but as time ensues, the response begins to flatten, eventually reaching a plateau rate whose magnitude depends on the strength of the motion. The pattern is best seen in the larger data set of Shadlen and Newsome (2001), shown in Figure 7*B*, using viewing durations of 0.5, 1 and 2 s.

The model predictions are shown in Figure 7*C*. To obtain these predictions we used the same architecture as in the RT simulations, with one exception: when the evidence reaches the decision threshold ( $\theta$ ), the accumulation stops, and the spike rate is fixed at the level it has achieved. Under these assumptions, the responses of model neurons demonstrate the features described above (Fig. 7*A,B*). The fact that LIP responses plateau at different rates for different coherences may seem counterintuitive: why are the trajectories not fixed at the value of the threshold ( $\theta$ )? First, the decision is derived from the smoothed output of LIP. The inherent delay between instantaneous measured activity in LIP and the detection of a threshold crossing allows the LIP response to continue to rise beyond the threshold before the decision is made. This effect is more pronounced the higher the motion strength. Secondly, the last phase of the decision is often dominated by transient noise impulses. These transients occur with different latencies from motion onset and are thus concealed within the response averages aligned to the beginning of the trial. That is why, at low motion strengths, the averaged responses remain below threshold despite the fact that on any single trial, this threshold was crossed. Together, the results in Figures 5 and 7 illustrate that the same computation—integration to threshold—can produce different patterns of single-unit responses depending on the demands of the task.

### Relationships between Single Neurons and Behavior on Single Trials

One of main dividends of the style of model that we have developed is that it portrays the activity of neurons as sequences of spikes that resemble spike trains from neurons in areas MT and LIP, incorporating their associated variability. We can therefore use the model to formulate quantitative predictions about the relationship between neural activity and decisions on a trial-by-trial basis. A basic assumption is that neural signals are represented by the average spike rate from a large ensemble of



**Figure 7.** Time course of LIP responses in the FD direction discrimination task. (*A*) Average spike rate from 38 LIP neurons recorded in 1 s FD experiments which alternated with blocks of RT experiments. Responses are grouped by motion strength and choice as indicated by color and line type, respectively. The responses are aligned in time with the onset of random dot motion (vertical line). The scale on the ordinate is the same as in the model, panel (*C*), for comparison. The spike rates are averaged using a running mean ( $\pm 30$  ms). Adapted from Roitman and Shadlen (2002). (*B*) Average spike rate from 104 LIP neurons recorded in 0.5, 1 and 2 s FD experiments. Same conventions as in (*A*). Spike rates are averaged using 50 ms bins. Note, ordinate scale is different than in (*A*) and (*C*). Adapted from Shadlen and Newsome (2001). (*C*) Responses of model LIP neurons. The simulated response averages are plotted using the same conventions as in (*A*). The model output is undefined in the first 200 ms after onset of motion (see Model Design). The horizontal dashed line indicates the height of the decision threshold ( $\theta$ ), 55 spikes/s. Error trials are not included in these graphs.

neurons with similar response properties. Consequently, the impact of any one neuron on the animal's behavior is relatively small. In this section, we examine this assumption quantitatively. First, we consider how trial-to-trial variability in the response of single MT neurons relates to the decisions on those trials and to the monkey's overall performance. Then, we examine how the variability of single neurons in LIP relates to the monkey's RT on individual experimental trials and to the monkey's overall performance.

### Single Neurons in Area MT

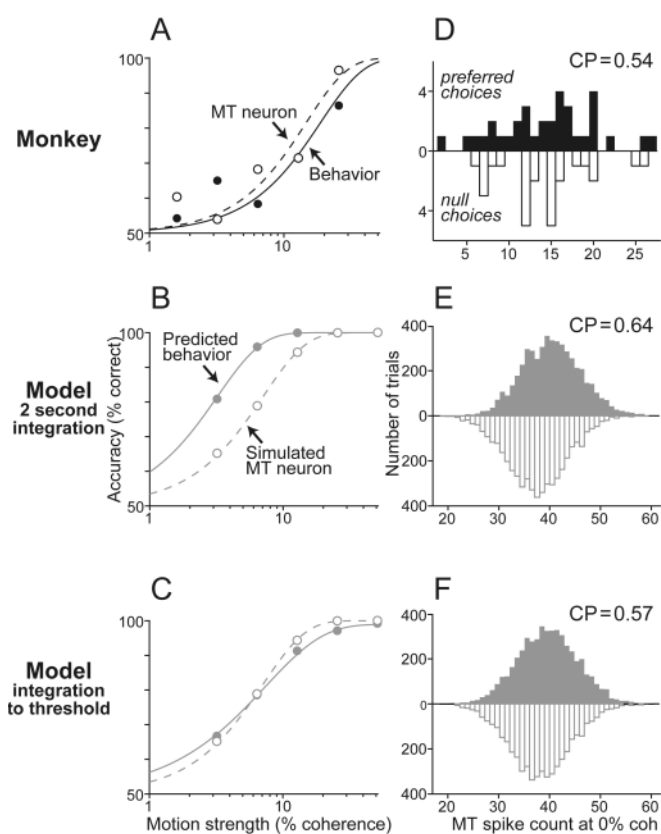
In their original studies, Newsome and colleagues made two quantitative comparisons between MT neurons and behavior. First, they compared the sensitivity of single neurons to the

monkey's behavioral sensitivity in a 2 s FD task (Newsome *et al.*, 1989; Britten *et al.*, 1992). To quantify the neural sensitivity, they compared the distributions of spike counts measured in response to preferred- and null-direction motion. They calculated the probability that a spike count drawn from the distribution of preferred direction responses would exceed a spike count drawn from the distribution of null direction responses. They thus predicted the level of performance that would be achieved if the monkey could compare responses from a pair of neurons tuned to opposite directions of motion, that is, the recorded neuron and one just like it but for the opposite direction preference. This analysis showed that single neurons are surprisingly sensitive to random dot motion: the predicted accuracy versus coherence function from one neuron, which they termed the *neurometric function*, was strikingly similar to the monkey's psychometric function on average (Fig. 8A). This result implies that the monkey could achieve observed levels of accuracy using only a single pair of MT neurons with opposed direction preferences.

In the second comparison, the authors measured the correlation between the response of an MT neuron and the monkey's choice on a single trial (Britten *et al.*, 1996). The trial-by-trial correlation between MT responses and choices was quantified using a measure called the *choice probability* [CP; also called *sender operating characteristic*, or SOC, in earlier publications (Newsome *et al.*, 1989; Celebrini and Newsome, 1994)]. To obtain the CP, the spike counts are compiled from each repetition of the same stimulus (i.e. same motion coherence and direction). The spike counts are then partitioned into two distributions, based on whether the animal chose the direction corresponding to the preferred direction of the motion (Fig. 8C). The separation of these distributions is indexed by the CP, which is the probability that a random draw of values from the two distributions would show the right ordinal relationship (i.e. larger for preferred direction choices). A CP value of 0.5 would indicate that the distributions overlap perfectly, hence the neuron's response is uncorrelated with the monkey's decision, while a CP value of 1 would mean that every spike count associated with a preferred choice was larger than every spike count associated with a null choice. The studies showed a weak but significant relationship between the variable spike counts from the neurons and the monkey's choices. Across the population of MT neurons recorded in the task, the average CP was 0.54–0.59, depending on monkey.

Together, these two results presented a paradox. The similarity between the psychometric function and the neurometric function based on a pair of neurons seemed to suggest that the brain does not combine signals from many MT neurons: more neurons would improve the monkey's accuracy to an unrealistic level. On the other hand, if this were truly the case, the few neurons underlying the decision should be highly correlated with the monkey's choice. A decision based on a comparison between two neurons would yield a CP of ~0.85 (Newsome *et al.*, 1989; Shadlen *et al.*, 1996), whereas the average CP was closer to 0.5, suggesting a very weak relationship. Thus the data appear incompatible with either a population-coding model or a model where a few MT neurons determine the decision.

Two solutions to this paradox were proposed by Newsome and colleagues (Shadlen *et al.*, 1996). First, one could assume that the brain uses many MT neurons whose preferred directions are not as well matched to the random dot motion as the



**Figure 8.** Relationships between single MT neurons and behavior in FD experiments. Left column: comparison of neurometric and psychometric functions. The neurometric function estimates discrimination accuracy from the neuron's responses to motion in its preferred and null direction (see text). Sigmoid curves are best fitting Weibull functions. (A) Data from a representative MT neuron in the Newsome database using a 2 s FD task. Behavioral accuracy is similar to the accuracy of the neuron. (B) Predicted neurometric and psychometric functions based on the assumption that the brain integrates motion information for the 2 s of motion viewing. The predicted behavior is superior to the accuracy estimated from a single neuron. (C) Predicted psychometric function based on the assumption that the brain integrates to a decision threshold. The neurometric function is the same as in (B). The accuracy of the predicted behavior is similar to the accuracy of the neuron. Right column: responses of MT neurons on single trials, sorted in accordance with the behavioral choice. Solid and open histograms show responses to the 0% coherent motion on trials in which the monkey (or model) chooses the preferred and null direction, respectively. The choice probability (CP) indexes the separation of these distributions. It indicates how well a single MT neuron predicts the monkey's choices. (D) Choice probability (CP) for the neuron whose data are shown in (A). (E) Predicted CP based on the assumption that the brain integrates motion information for the 2 s of motion viewing. (F) Predicted CP based on the assumption that the brain integrates to a decision threshold.

neurons that were studied by Newsome *et al.* (1989), where the motion was precisely aligned to the preferred direction of the neurons. Pooling from such a mixture of optimal and sub-optimal neurons would lead to less improvement in accuracy. However, this resolution was not sufficient to explain the CP. Even with arbitrarily large pools, and regardless of the mix of optimal and sub-optimal neurons, single neurons would be expected to retain a higher degree of correlation with the monkey's decision than observed (the predicted CP was ~0.65). This is because the neurons in the pools are thought to covary weakly in their responses. To reduce the predicted CP, they assumed that some additional noise is added during read-out of the MT signal; they referred to this as pooling noise. By adjusting



the mixture of neurons and pooling noise, the model could account for both the accuracy of the monkey and the trial-by-trial relationship between single neurons and decisions.

The model in Figure 3 offers a simpler account of these results. When the model is allowed to integrate for a full 2 s, it yields the same problematic predictions as the Shadlen *et al.* (1996) pooling model: a level of sensitivity that is about two times better than single neurons and the monkey (Fig. 8B), and CP values that are higher than those observed in the monkey (Fig. 8E). Suppose, however, that instead of accumulating evidence for the full duration of the motion stimulus, the monkey reaches a decision when the evidence reaches a threshold,  $\theta$ , as in the RT task (see Fixed-duration Direction Discrimination). According to this idea, MT spikes would only contribute to the decision during the portion of the trial before the threshold crossing. Importantly, the experimenter does not know when this time occurs and thus includes all spikes from the full 2 s recording into calculations of neural sensitivity and CP.

This alternative decision model resolves the paradox posed by the sensitivity of single MT neurons and the weak trial-by-trial correlation between their variable responses and the monkeys' decisions. First, the model no longer attains the 2-fold improvement over the measured sensitivity of single neurons because the decision process uses only a fraction of the 2 s viewing time. In fact, the predicted psychometric function is now approximately the same as the neurometric function calculated from just one typical MT neuron using 2 s of spike discharge (Fig. 8C). This correspondence reflects the fact that the decision takes ~500 ms on average, which is only one-quarter of the total viewing time, so the model's accuracy is factor of 2 lower than the theoretical maximum. Secondly, the predicted CP in this model is reduced to 0.57 (Fig. 8F), which is within the range of values calculated by Britten *et al.* (1996). Again, this occurs because many of the MT spikes used to compute the CP do not actually contribute to the decision, and thus they show no correlation with the choice on a given trial. Hence, this model advanced here can account for all sources of noise between the sensory representation in MT and the decision. Moreover, as shown in Figure 7, it explains features of the LIP responses recorded in FD experiments.

The model makes one prediction that is contradicted by the data. According to our idea, spikes recorded late in the trial are less likely to affect the monkey's choices because they are often recorded after the monkey has made its decision. Assuming, as we have, that these spikes are independent of those occurring earlier in the trial, we would expect that they would exhibit no correlation with the monkey's choices. The data from Britten *et al.* (1996) appear to contradict this prediction (see their figure 11). We will return to this issue in Discussion.

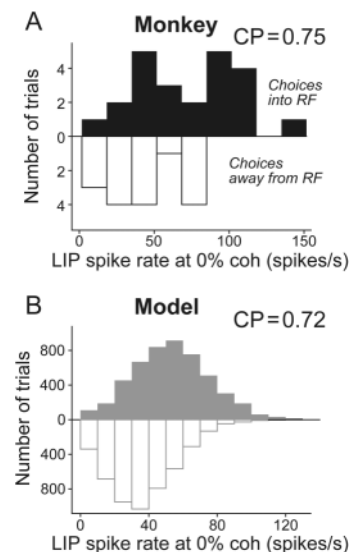
#### Single Neurons in Area LIP

Our model instantiates the idea that LIP represents the accumulation of evidence for choosing the direction of motion associated with the choice target in the neuron's response field. As with MT, we assume that ensembles of noisy, weakly correlated, spiking neurons supply the signals used for making decisions about the direction of motion. Therefore, the same intuitions should apply: responses from single LIP neurons should covary weakly with the ensemble signals to which they contribute and therefore with the monkey's behavioral response.

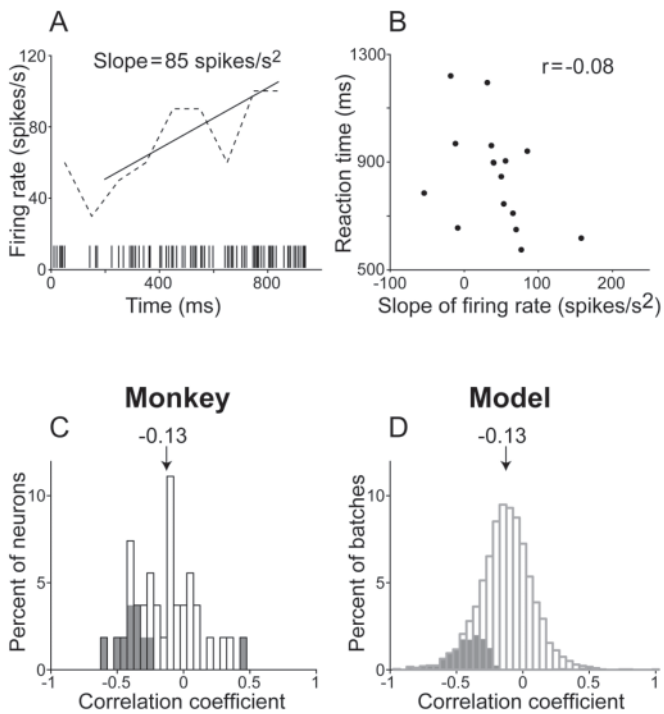
The LIP neurons studied by Shadlen and Newsome (2001) and Roitman and Shadlen (2002) exhibit responses that can be interpreted as motor planning signals. Not surprisingly, in the 100 ms epoch just before the monkey makes an eye movement, the responses are clearly associated with the eye movement response. Indeed, for many neurons, it is possible to predict with 100% reliability the monkey's choice based on the spikes in that trial. A more interesting quantity is the spike rate in the epoch ending 100 ms before eye movement initiation in the RT task (Fig. 9A), about the time when the evidence reaches the threshold for commitment. Although the ensemble response in LIP determines the monkey's choice, individual neurons predict the monkey's choices with only modest reliability (mean CP = 0.73). Interestingly, the model yields a very similar CP value for single neurons in the same epoch (Fig. 9B).

The model also predicts an association between the responses of LIP neurons and the monkey's reaction time on a given trial. We would expect the rate of rise of the LIP spike rate to be inversely correlated with RT: steeper changes in LIP should lead to shorter RT. The challenge is to estimate the rate of rise of the spike rate from the sequence of spikes on a single trial. Roitman and Shadlen (2002) did this by finding the ramp that best explained the data (spike train) in the maximum likelihood sense. They used only spikes occurring from 200 ms after motion onset to 100 ms before the saccade. An example of one of their fits is shown in Figure 10A. This analysis yields a spike rate slope value on each trial, which can then be compared to the reaction time on that trial.

Figure 10B shows a scatter plot of RTs and the best estimate of spike rate slopes from one neuron using 15 trials in which the monkey viewed an intermediate motion strength (6.4% coher-



**Figure 9.** Choice probability in area LIP near the end of the decision process. Spike rates were measured in the epoch 200 to 100 ms before saccade initiation in the RT task. Solid and open histograms show responses to the 0% coherent motion on trials in which the monkey (or model) chooses the target in or out of the neurons response field, respectively. The choice probability (CP) indexes the separation of these distributions. It indicates how well a single LIP neuron predicts the monkey's choices. (A) Data from one representative neuron. Across the population of LIP neurons studied by Roitman and Shadlen (2002), the average CP = 0.73 ( $P < 0.05$ ;  $n = 54$  neurons). (B) Model prediction. The expected CP based on the integrate-to-threshold model is 0.72 in this epoch.



**Figure 10.** Trial-by-trial correlation between LIP response and reaction time. (A) Example of a spike train and the estimated rate function from one trial. The response was obtained during exposure to 6.4% coherence motion toward the neuron's RF, leading to a correct response. The raster shows the times of action potentials with respect to the time of motion onset ( $t = 0$ ). The dashed line indicates the average firing rate from the trial calculated in 100 ms bins. The solid line indicates the maximum likelihood fit of a linear rate function to the spike train. The fit assumes that the rate changes linearly from 200 ms after motion onset until 100 ms before the saccade (for details, see Roitman and Shadlen, 2002). It provides an estimate of the slope of the 'firing-rate-versus-time' function. (B) Relationship between reaction time and slope of the firing-rate-versus-time function. Scatter plot shows RT and slopes for correct choices from one experiment using the 6.4% coherent motion toward the neuron's RF ( $n = 15$  trials). The weak inverse correlation is not significant owing to the small number of samples, but the value is representative. (A) and (B) are adapted from Roitman and Shadlen (2002). (C) Distribution of correlation coefficients from 36 LIP neurons. Each neuron contributes one correlation value to the distribution based on all correct choices that end in an eye movement to the RF. The correlation value is obtained after standardizing the firing rate slopes and reaction times within each motion strength. 27% of the neurons (10/36) showed individually significant correlation ( $P < 0.05$ , Fisher's  $z$ ; gray bars). Across the population of neurons, the average correlation coefficient was  $r = -0.13$  ( $P < 0.05$ ,  $H_0: r = 0$ ). (D) Model prediction. We grouped simulated trials into batches where each batch contained the same number of trials as one of the neurons in the data set. These simulated data batches show the same average correlation between LIP firing rate slope and reaction time as seen in the neural data ( $r = -0.13$ ), with comparable spread in the observed  $r$  values. 16% of the batches showed individually significant correlation (gray bars).

ence) in the direction that led to a correct choice (an eye movement to the choice target in the neuron's RF). A weak but measurable relationship is apparent between the slope of the spike rate and the reaction time ( $r = -0.08$ ), with larger slopes leading to shorter reaction times. We repeated this analysis for each neuron, using only those trials in which the monkey made a correct choice to the RF. The histogram of  $r$  values obtained in this fashion is shown in Figure 10C. Across the population of 36 LIP neurons, the average correlation was  $r = -0.13$  ( $P < 0.05$ ). Using data sets of comparable size to those obtained in LIP, the model yields a similar inverse correlation between RT and esti-

mated ramp slope (mean  $r = -0.13$ ; Fig. 10D). The relatively close match between the predicted and observed correlation levels suggests that the ramping responses seen in single LIP neurons reflect an accumulating decision signal being carried in the average rate of a local ensemble of similar neurons.

## Discussion

Integration with respect to time is thought to play a role in neural computations involved in navigation and motor control. In the brainstem oculomotor system, integration converts a velocity 'pulse' to a persistent 'step' in spike rate, which is used to hold the gaze in an eccentric position (Robinson, 1989). Similarly, neurons that represent head position are thought to incorporate the time integral of angular velocity of the head (Fukushima *et al.*, 1992; Seung *et al.*, 2000; Aksay *et al.*, 2001; Goldman *et al.*, 2002). Indeed, wherever one observes persistent neural activity in the brain, it seems reasonable to hypothesize that it represents the integration of a transient input. However, unlike the previous examples, persistent activity in higher brain centers can be elicited under a wide range of conditions, and is subject to modulation by a number of factors. How does integration contribute to higher brain function? The experiments described and simulated in this paper represent an attempt to begin to answer that question. For the problem of discriminating between two categorical interpretations of a stochastic sensory stimulus, it seems likely that integration allows the brain to accumulate evidence over time, thereby improving the fidelity of perception and accuracy of decision making.

Our model instantiates the idea that the decision in the random dots task is based on the time integral of motion evidence represented by neurons in area MT. The simulations described in this paper do not test the idea, but instead explore its implications. In that sense, the model exposes important features of the data, furnishes novel interpretations, and makes predictions about future experiments (see Supplementary Material). In what follows, we provide some intuition for the structure of the model, explain the rationale behind several assumptions, and explore its successes, limitations and alternatives.

## Structure of the Model

A substantial body of experimental results indicates that neurons in area MT represent the critical sensory signals that monkeys use to base their judgment of random dot motion (for reviews, see Parker and Newsome, 1998; Britten, 2003). This evidence fluctuates in time, reflecting both the variability in the dynamic stimulus and the noise in the neural response (Mazurek and Shadlen, 2002). It must be accumulated in order for the brain to reach an accurate decision. The model suggests that this integrated motion evidence is carried in the responses of sensorimotor association areas such as LIP. When the accumulated evidence reaches a threshold level, the decision process terminates. The higher the threshold, the longer it will take to reach a decision but the more accurate the decision will be, on average. The model thus takes us from the representation of motion to a commitment to a choice.

The model's structure incorporates several assumptions about the nature of the neural signals in these computations. These are detailed in the Supplementary Material along with a description of the mathematical operations used to perform the simulations.

### **What Do We Learn from the Model?**

Integration of MT signals to a decision threshold explains the effect of motion strength on the monkey's behavior in the reaction time (RT) task (Fig. 4). It reconciles the accuracy of the monkey's decisions with the amount of time required to make them. Stronger motion produces a rapid accumulation of evidence toward the correct decision threshold. Weaker motion produces a more random accumulation that meanders between either decision threshold, leading to a mixture of correct and incorrect responses. For the 0% coherent motion, the accumulation is equally likely to 'diffuse' toward either threshold.

Accumulation to threshold also accounts for the psychometric function observed in the fixed duration (FD) task, where the experimenter controls the stimulus duration (Fig. 6). This idea resolves a puzzling observation: monkeys achieve similar level of accuracy on RT and 1 s FD tasks (Fig. 6A), despite the longer viewing times provided in the FD experiment. We propose that this surprising behavior occurs because the monkey employs a threshold of evidence for committing to a decision in the FD task, just as in the RT task (Link and Heath, 1975; Meyer *et al.*, 1988; Ratcliff, 1988; Link, 1992). The similarity in performance suggests that the monkeys used the same level of the decision threshold ( $\theta$ ) in the two tasks.

This idea does not imply that monkeys cannot improve their performance with increasing viewing times. If the viewing duration is relatively short, the integrated evidence may fail to reach threshold. In that case the optimal solution is to respond based on the available evidence, which is expected to increase in reliability as a function of  $\sqrt{t}$ . Behavioral evidence for such integration can be found in FD experiments (Britten *et al.*, 1992; Gold and Shadlen, 2000; Gold and Shadlen, 2003).

Integration to threshold explains the time course of the responses from LIP neurons in both RT and FD tasks (Figures 5 and 7). This observation reconciles a remarkable discrepancy in the experimental data: the responses during motion viewing exhibit different time courses depending on whether the monkey is performing the RT or the FD task (Roitman and Shadlen, 2002). Our simulations demonstrate that integration to threshold can account for the time course of the responses in both of these situations, predicting both ramp-like increases and decreases in the RT task and saturating functions of time in the FD task. The model also reconciles some puzzling observations about the LIP activity: (i) the convergence of rising responses preceding eye movements to the RF in the RT task (Fig. 5A,B, solid curves, right panels); (ii) the absence of convergence of declining responses before eye movements to away from the RF in the RT task (Fig. 5A,B, dashed curves, right panels); (iii) the absence of convergence of rising responses (i.e. the coherence-dependent saturation) in FD trials (Fig. 7); and (iv) the apparent failure of averaged responses in FD experiments to reach a common threshold.

Because the model explicitly represents the activity of single MT and LIP neurons, it lends insight into the relationship between behavioral measurements and single neuron recordings obtained in the laboratory. By mimicking the spike discharge of single MT neurons, the model explains the observed relationship between MT responses and the monkey's decisions on a trial-to-trial basis (choice probability; Fig. 8D-F). By mimicking the spike discharge of single LIP neurons, the model explains the observed relationship between LIP responses and behavior in the RT task (Figures 9 and 10). The explanation of these single-trial comparisons, which are only

available from combined psychophysics-neurophysiology experiments, are the main dividend of our modeling exercise. Finally, the idea that evidence is integrated to a threshold, even when given a fixed viewing duration, serves to reconcile the exquisite sensitivity of single MT neurons with the overall sensitivity of the monkey to random dot motion (Fig. 8A-C).

### **Model Limitations and Failures**

A major limitation of our model is that it fails to specify how key computations are achieved by neurons. It offers no insight into the neural mechanisms that would underlie integration, or direction selectivity, or even addition and subtraction. Nor does the model attempt to explain the mechanism by which the LIP activity is compared to the decision threshold. The biophysical and circuit properties underlying operations such as integration are an active topic of investigation at the theoretical and experimental level (Camperi and Wang, 1997, 1998; Seung *et al.*, 2000; Aksay *et al.*, 2001; Koulakov *et al.*, 2002). Our model circumvents these issues by performing simple mathematical calculations on spike rate signals, but obviously this leaves open the question of how the brain actually accomplishes the calculations.

Similarly, our model provides little insight into the control of the integration process. For example, we do not attempt to explain the early dip and recovery of neural activity that occurs just after onset of random dot motion. These 'dips' are a common feature of neurons with persistent activity (Sato *et al.*, 2001), and we suspect that this phenomenon might play a role in initiating the integration, perhaps serving as a reset. Similarly, we assume that LIP stops receiving sensory input after the decision threshold is crossed, even when motion information remains present as in the FD task. This may be plausible in the context of the random dot motion task, but it is undoubtedly an oversimplification. Clearly, neural integration for high-level cognitive behavior requires complex controls and is susceptible to numerous influences, and our model makes no attempt to capture these various factors.

Finally, neither our data nor our model address the question of whether integration actually occurs in LIP or elsewhere, only to be relayed to LIP. Several brain areas show decision related activity similar to what is observed in LIP, including the prefrontal cortex and the superior colliculus (Kim and Shadlen, 1999; Horwitz and Newsome, 2001). It is unknown whether any or all of these areas are responsible for integration of the evidence from MT, or to what degree their responses simply reflect this process.

In addition to its limitations, the model makes several predictions that are known to be incorrect. The most serious discrepancy between model and data is in the RT associated with error trials. The model predicts relatively fast error responses, whereas monkeys tend to take slightly more time to respond on error trials than on correct trials at each coherence level (Fig. 4). Diffusion or random walk models, which are related to ours, predict that errors and correct trials share the same chronometric function (RT versus motion strength) (Luce, 1986). Our model produces fast errors because of variability in the LIP signal, resulting in early threshold crossings. How might these fast errors be remedied? Ratcliff and Rouder (1998) demonstrated that the inclusion of trial-to-trial variability in the effective stimulus strength produces slower errors in diffusion models sharing features with ours. This source of variability is not incorporated in our model: the expected MT response is

identical across trials sharing the same motion strength, which is clearly an oversimplification. Remedying this might bring the predicted error RTs closer to those observed in the data. In addition, we should note that monkeys performing the discrimination task probably make errors for a wide variety of reasons, many of which are not captured by the integration model (e.g. lapses in attention, blinks). Hence, error RTs may be harder to explain than other aspects of the behavior.

There is at least one other prediction of the model that is contradicted by data. In simulations using a fixed 2 s duration, the model posits that MT inputs are ignored after LIP crosses the decision threshold. This predicts that MT spikes occurring early in the viewing period should correlate with the monkey's decision more strongly than spikes occurring later in the trial, which is contrary to the findings of Britten *et al.* (1996). One possible explanation may be that the monkeys in the Britten *et al.* study began integrating MT spikes at a variety of latencies during the 2 s motion viewing period. This is consistent with our observations that the animals require much less than the 2 s to reach a decision. Some support for this idea can be found in Shadlen and Newsome's study of LIP (Shadlen and Newsome, 2001). They found that exposure to viewing durations shorter than 2 s caused an acceleration in the build up of LIP activity for trials of all durations. Direct measurements of MT neurons in monkeys trained to perform both FD and RT tasks may help to clarify these issues.

### Model Extensions and Alternatives

The analyses presented in this paper show that time integration of sensory evidence to a threshold accounts for a wide variety of behavioral and physiological observations in monkeys trained to discriminate the direction of random dot motion. Several extensions to the basic idea deserve consideration. Variability in the parameters of integration, such as the starting time, the baseline spike rate, and the drift rate, may be helpful in remedying the failures and improving the generality of the model, especially when incorporating prior biases into the decision (Carpenter and Williams, 1995; Ratcliff and Rouder, 1998). Similarly, the decision threshold may be variable and is likely to change as a function of time. A dynamic decision threshold embodies the idea that commitment to one or another behavioral option may need to occur by some time point or with some degree of urgency (Reddi and Carpenter, 2000). In general, the decision threshold is likely to incorporate knowledge of the accuracy that has been achieved, the passage of time, and the rate at which reward is attained (Gold and Shadlen, 2002). These factors will need to be incorporated into a more complete computational model.

The main alternative to temporal integration is a decision process that is not based on an accumulation of information but instead on extreme values in the sensory data, considered as a passing stream. This type of process is often termed 'probability summation' (Watson, 1979). Like integration, probability summation predicts that accuracy and decision time will depend on motion strength. However, probability summation would not predict the existence of neurons like those in LIP that represent a decision variable that evolves gradually. Rather, it predicts that at any moment, the brain is either committed to a decision or it has yet to detect the decisive sensory data. The pattern of neural responses seen in LIP fails to support this alternative, but individual neurons could in principle undergo the kind of state

change predicted by this model on individual trials. The idea is that the average of many such switches from an intermediate spike rate to a high spike rate (or a low spike rate for the opposite choice) at different times could give rise to the ramp-like trajectories seen in Figure 5. The evidence at present favors integration over probability summation (see Gold and Shadlen, 2000; Roitman and Shadlen, 2002), but further experiments will be necessary to clarify the matter

### Concluding Remarks

We have examined the proposal that neural integration of time-varying sensory signals underlies the formation of decisions about the direction of random dot motion. The model explains a variety of behavioral and physiological observations, and its deficiencies are relatively minor. The model provides no insight into how integration comes about, but it does help to answer two fundamental questions: what is integrated and to what end? The answer to the first question is sensory data construed as evidence for versus against a proposition. The answer to the second question is that the evidence is integrated to a threshold level, and the crossing of the threshold signals a commitment to a proposition or behavioral response. For the random dots task, the integral is represented by neurons in the parietal cortex and other visual-oculomotor association areas. A common property of these neurons is their ability to discharge in a persistent fashion and thus carry their representation of space or motor intention forward through time. Persistent activity is observed in many cortical areas, where it has been associated with memory traces (Fuster and Alexander, 1971; Fuster, 1973; Miyashita and Chang, 1988; Funahashi *et al.*, 1989) and motor planning (Evarts and Tanji, 1976; Tanji and Hoshi, 2001). It is tempting to view such activity as a hallmark of integration, furnishing the brain with the capacity to evaluate propositions about the state of the world in order to plan appropriate behavior.

### Notes

Supported by HHMI, NIH grants RR00166 and EY11378. M.E.M. was supported by an NIH training grant (GM07108), a Poncin grant, and an ARCS fellowship. We thank Carlos Brody, Alex Huk, Joshua Gold, and John Palmer for helpful comments and discussion.

Address correspondence to Michael N. Shadlen, Department of Physiology and Biophysics, University of Washington Medical School, Box 357290, Seattle, WA 98195-7290, USA. Email: shadlen@u.washington.edu.

### Supplementary Material

Supplementary material can be found at: <http://www.cercor.oupjournals.org>

### References

- Aksay E, Gamkrelidze G, Seung HS, Baker R, Tank DW (2001) In vivo intracellular recording and perturbation of persistent activity in a neural integrator. *Nat Neurosci* 4:184-193.
- Bisley JW, Zaksas D, Pasternak T (2001) Microstimulation of cortical area MT affects performance on a visual memory task. *J Neurophysiol* 85:187-196.
- Britten K (2003) The middle temporal area: motion processing and the link to perception. In: *The visual neurosciences* (Chalupa WJ, ed.). Boston, MA: MIT Press.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745-4765.

- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10:1157-1169.
- 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.
- Bruce CJ, Goldberg ME (1985) Primate frontal eye fields I Single neurons discharging before saccades. *J Neurophysiol* 53:603-635.
- Burr DC, Santoro L (2001) Temporal integration of optic flow, measured by contrast and coherence thresholds. *Vision Res* 41.
- Camperi M, Wang X-J (1997) Modeling delay-period activity in the prefrontal cortex during working memory tasks. In: *Computational neuroscience* (Bower, ed.), pp. 273-279. New York: Plenum Press.
- Camperi M, Wang X-J (1998) A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability. *J Comput Neurosci* 5:383-405.
- Carpenter R, Williams M (1995) Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377:59-62.
- 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.
- Evarts EV, Tanji J (1976) Reflex and intended responses in motor cortex pyramidal tract neurons of monkey. *J Neurophysiol* 39:1069.
- Fukushima K, Kaneko C, Fuchs A (1992) The neuronal substrate of integration in the oculomotor system. *Prog Neurobiol* 39:609-639.
- Funahashi S, Bruce C, Goldman-Rakic P (1989) Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *J Neurophysiol* 61:331-349.
- Fuster JM, Alexander GE (1971) Neuron activity related to short-term memory. *Science* 173:652-654.
- Fuster JM (1973) Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. *J Neurophysiol* 36:61-78.
- Glimcher PW, Sparks DL (1992) Movement selection in advance of action in the superior colliculus. *Nature* 355:542-545.
- Gnadt JW, Andersen RA (1988) Memory related motor planning activity in posterior parietal cortex of macaque. *Exp Brain Res* 70:216-220.
- Gold JI, Shadlen MN (2000) Representation of a perceptual decision in developing oculomotor commands. *Nature* 404:390-394.
- Gold JI, Shadlen MN (2002) Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36:299-308.
- Gold JI, Shadlen MN (2003) The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J Neurosci* 23:632-651.
- Goldman MS, Kaneko CR, Major G, Aksay E, Tank DW, Seung HS (2002) Linear regression of eye velocity on eye position and head velocity suggests a common oculomotor neural integrator. *J Neurophysiol* 88:659-665.
- Hernandez A, Zainos A, Romo R (2002) Temporal evolution of a decision making process in medial premotor cortex. *Neuron* 33:959-972.
- Horwitz GD, Newsome WT (2001) Target selection for saccadic eye movements: prelude activity in the superior colliculus during a direction-discrimination task. *J Neurophysiol* 86:2543-2558.
- Kim J-N, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176-185.
- Koulakov AA, Raghavachari S, Kepecs A, Lisman JE (2002) Model for a robust neural integrator. *Nat Neurosci* 5:775-782.
- Link SW (1992) *The wave theory of difference and similarity*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Link SW, Heath RA (1975) A sequential theory of psychological discrimination. *Psychometrika* 40:77-105.
- Luce RD (1986) *Response times: their role in inferring elementary mental organization*. New York: Oxford University Press.
- Mazurek ME, Shadlen MN (2002) Limits to the temporal fidelity of cortical spike rate signals. *Nat Neurosci* 5:463-471.
- Meyer DE, Irwin DE, Osman AM, Kounios J (1988) The dynamics of cognition and action: mental processes inferred from speed-accuracy decomposition. *Psychol Rev* 95:183-237.
- Miyashita Y, Chang H (1988) Neuronal correlate of pictorial short-term memory in the primate temporal cortex. *Nature* 331:68-70.
- Newsome WT, Paré EB (1988) A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J Neurosci* 8:2201-2211.
- Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. *Nature* 341:52-54.
- Newsome WT, Britten KH, Movshon JA, Shadlen M (1989) Single neurons and the perception of visual motion. Neural mechanisms of visual perception. In: *Proceedings of the Retina Research Foundation* (Lam DM-K, Gilbert CD, eds), vol. 2, pp. 171-198. The Woodlands, TX: Portfolio Publishing Company.
- Parker AJ, Newsome WT (1998) Sense and the single neuron: probing the physiology of perception. *Annu Rev Neurosci* 21:227-277.
- Ratcliff R (1988) Continuous versus discrete information processing: modeling accumulation of partial information. *Psychol Rev* 95:238-255.
- Ratcliff R, Rouder JN (1998) Modeling response times for two-choice decisions. *Psychol Sci* 9:347-356.
- Reddi BA, Carpenter RH (2000) The influence of urgency on decision time. *Nat Neurosci* 3:827-830.
- Robinson DA (1989) Integrating with neurons. *Annu Rev Neurosci* 12:33-45.
- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475-9489.
- Salzman CD, Britten KH, Newsome WT (1990) Cortical microstimulation influences perceptual judgements of motion direction. *Nature* 346:174-177.
- Sato T, Murthy A, Thompson KG, Schall JD (2001) Search efficiency but not response interference affects visual selection in frontal eye field. *Neuron* 30:583-591.
- Seung HS, Lee DD, Reis BY, Tank DW (2000) Stability of the memory of eye position in a recurrent network of conductance-based model neurons. *Neuron* 26:259-271.
- Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86:1916-1936.
- 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.
- Tanji J, Hoshi E (2001) Behavioral planning in the prefrontal cortex. *Curr Opin Neurobiol* 11:164-170.
- Watson AB (1979) Probability summation over time. *Vision Res* 19:515-522.
- Wurtz RH, Albano JE (1980) Visual-motor function of the primate superior colliculus. *Annu Rev Neurosci* 3:189-226.