

One-Dimensional Dynamics of Attention and Decision Making in LIP

Surya Ganguli,^{1,7,*} James W. Bisley,² Jamie D. Roitman,³ Michael N. Shadlen,⁴ Michael E. Goldberg,^{5,6} and Kenneth D. Miller^{5,7}

¹Sloan-Swartz Center for Theoretical Neurobiology, University of California, San Francisco, San Francisco, CA 94143, USA

²Department of Neurobiology, University of California, Los Angeles, Los Angeles, CA 90025, USA

³Department of Psychology, University of Illinois, Chicago, IL 60607, USA

⁴Howard Hughes Medical Institute, Department of Physiology and Biophysics, University of Washington, Seattle, WA 98195, USA

⁵Department of Neuroscience, Columbia University, New York, NY 10032, USA

⁶Departments of Neurology and Psychiatry, Columbia University, and New York State Psychiatric Institute, New York, NY 10032, USA

⁷Center for Theoretical Neuroscience, Columbia University, New York, NY 10032, USA

*Correspondence: surya@phy.ucsf.edu

DOI 10.1016/j.neuron.2008.01.038

SUMMARY

Where we allocate our visual spatial attention depends upon a continual competition between internally generated goals and external distractions. Recently it was shown that single neurons in the macaque lateral intraparietal area (LIP) can predict the amount of time a distractor can shift the locus of spatial attention away from a goal. We propose that this remarkable dynamical correspondence between single neurons and attention can be explained by a network model in which generically high-dimensional firing-rate vectors rapidly decay to a single mode. We find direct experimental evidence for this model, not only in the original attentional task, but also in a very different task involving perceptual decision making. These results confirm a theoretical prediction that slowly varying activity patterns are proportional to spontaneous activity, pose constraints on models of persistent activity, and suggest a network mechanism for the emergence of robust behavioral timing from heterogeneous neuronal populations.

INTRODUCTION

We live in a complex visual world with multiple stimuli continually vying for our attention. In the face of this complexity, we must reliably shift our loci of attention over time to achieve our goals. Furthermore, we continually receive noisy sensory cues about objects in the world. We must reliably integrate such noisy evidence over time in order to make perceptual decisions. Both attentional shifting and decision making require robust temporal control and hence place stringent constraints on the dynamics of any neural circuitry mediating these behaviors.

Both behaviors are controlled by cortical circuits in which the underlying dynamics of single neurons are often highly heterogeneous and noisy. The synaptic connections between these neurons are unreliable and not precisely specified in development.

Moreover, these neurons can display significant levels of spontaneous activity, so that their responses reflect the interaction of stimulus-driven inputs with this ongoing activity. How can robust behavioral dynamics emerge naturally out of the biophysics of spontaneously active, imprecisely specified networks of unreliable and heterogeneous neuronal elements?

We address this fundamental question by examining the dynamics of single neurons in the lateral intraparietal area (LIP) of the posterior parietal cortex. A number of studies suggest that this brain region plays a role both in the allocation of visual attention (Colby et al., 1996; Gottlieb et al., 1998; Robinson et al., 1995; Gottlieb and Goldberg, 1999; Powell and Goldberg, 2000) and in perceptual decision making when the decision is reported via a saccade (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Hanks et al., 2006). In particular, this paper is motivated by a remarkable and puzzling correlation between single-neuron dynamics in LIP and shifts in spatial attention found recently by two of the authors (Bisley and Goldberg, 2003, 2006). We arrive at a circuit explanation of this puzzle that not only suggests a more general answer to the question of how heterogeneous cortical dynamics can give rise to robust behavioral dynamics but also yields insights into the nature of persistent and slowly integrating activity in LIP. We show that the puzzle can be explained by circuitry that makes LIP dynamics one-dimensional on slow timescales. This means that slowly varying patterns of activity, including spontaneous activity during fixation, persistent activity during the delay period before a planned saccade, and slowly integrating activity during a perceptual decision-making task, are all scaled versions of a single pattern of relative firing rates across neurons. This simple scaling of a single pattern does not describe more rapidly varying firing patterns, such as the transient response to a visual stimulus, so that this scaling does not arise simply from variations in neural excitability. We uncover direct evidence for this hidden one-dimensional signal underlying LIP responses both in tasks involving attentional allocation (Bisley and Goldberg, 2003, 2006) and in tasks involving perceptual decision making (Roitman and Shadlen, 2002). Functionally, we propose that the reduction of neural signals in LIP to a single dimension on the timescale of hundreds of milliseconds can yield the robust and

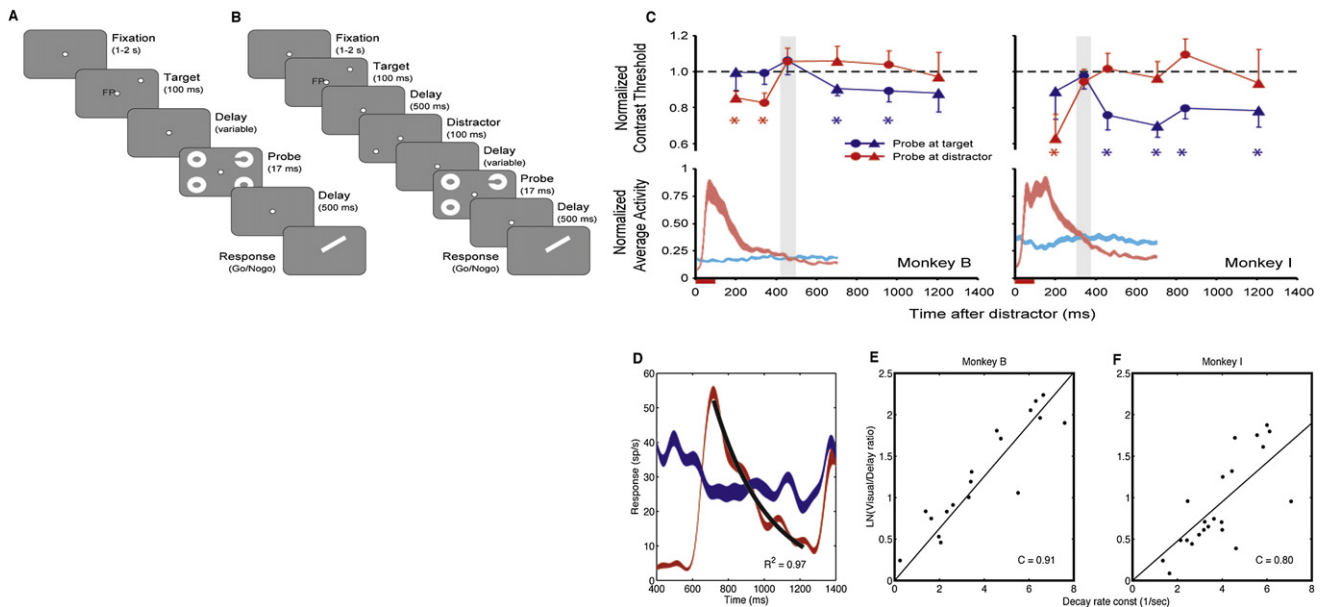


Figure 1. The Experiments of Bisley and Goldberg (2003, 2006)

(A) While a single neuron was recorded, the monkey initiated a trial by fixating a small spot. After a short delay, the saccade target appeared for 100 ms either in the RF of the recorded neuron or opposite. At varying times after the saccade target disappeared, and at varying contrasts, a Landholt ring (the probe) and three complete rings flashed briefly. The probe appeared either at the saccade target or opposite. Five hundred milliseconds after the probe appeared, the monkey had to indicate the orientation of the probe by either maintaining fixation or making a saccade to the target.

(B) In half the trials, a task-irrelevant distractor, identical to the saccade target, was flashed 500 ms after the target offset in a location opposite to the saccade target.

(C) (Top) The monkey's visual sensitivity was measured by the probe contrast threshold (Bashinski and Bacharach, 1980; Ciaramitaro et al., 2001) at which the monkey performed this task with 75% accuracy. This threshold depended on the spatial location and time at which the probe appeared. Lower thresholds indicate higher sensitivity. (Bottom) The population average response of LIP neurons to a distractor appearing in the RF (red) and the population average delay activity in response to a target appearing earlier in the RF with the distractor flashed elsewhere (blue). The standard deviation of these rates (as well as those in panel [D]) is given by the width of the traces. The center of the vertical gray column marks both the center of a temporal window of ambiguity in which there is no significant difference between the activity in both populations (bottom) and the time at which psychophysical measurements found no significant attentional advantage at either the target or distractor sites (top).

(D) The response of a single neuron in trials in which a distractor appears in its RF from 600 to 700 ms (red), a typical exponential fit to this transient (black), and the delay period activity of the same neuron in trials in which the target appears in its RF from 0 to 100 ms (blue). The mean of this last response yields the delay period activity D , while the peak of the black curve yields the visual response V , and its decay rate yields k .

(E and F) For each monkey, there is a relationship across all neurons recorded between the visual/delay (V/D) ratio and the rate of decay of activity k following the distractor. The solid black line shows a linear fit to the data. The slope t_c of each fit is the common neuronal crossing time for each monkey.

reliable temporal signal needed for LIP function in attentional allocation, perceptual decision making, and related oculomotor behaviors, despite substantial heterogeneity in many aspects of single-neuron responses.

RESULTS

Single Neurons and Attention: The Puzzle

The Bisley and Goldberg results are reviewed in more detail in Figure 1. The authors manipulated the locus of attention in a top-down manner by instructing the monkey to plan a delayed saccade to the site of a flashed saccade target. During the delay period, the target site captures the monkey's attention, in keeping with psychophysical studies in humans showing that attention is allocated to the endpoint of a planned saccade (Shepherd et al., 1986; Deubel and Schneider, 1996). Attention was also manipulated in a bottom-up manner by flashing a behaviorally irrelevant distractor at a different site during the delay period. This was designed to transiently capture the monkey's attention,

in keeping with prior studies showing that a flashed object can also attract attention (Yantis and Jonides, 1984, 1996; Egeth and Yantis, 1997). Measurements of the locus of attention, defined as the visual field location with enhanced contrast sensitivity, confirmed that attention was initially captured by the distractor but then moved back to the target (Figure 1C, top).

Single neurons in LIP were recorded during this task. On different trials, either (1) the target flashed in the neuron's receptive field (RF) and then the distractor appeared elsewhere or (2) the distractor flashed in the RF after the target had appeared elsewhere. In the first trial type, an LIP neuron displayed a transient visual response to the target, followed by elevated levels of persistent activity during the delay period, signifying that the monkey was planning a saccade to the RF of the recorded neuron. In the second trial type, the same neuron would display a transient visual response to the distractor that decayed back toward baseline firing levels. Figure 1C (bottom) shows the average population response, averaged over these two trial types. Reinterpreting these trial types as the simultaneous activity of

two populations of neurons, one at the target and one at the distractor, one sees that at all points during the task, attention is allocated to the RF of the population with highest firing rates. For each monkey there is, however, a brief window of ambiguity in which the two population averages are statistically indistinguishable. The authors probed the center of this window psychophysically and found that at this time (455 ms in monkey B and 340 ms in monkey I), there was no attentional advantage at either location. This time was defined as the attentional switching time. This correlation between the population response and the psychophysics led the authors to the interpretation that LIP plays the role of a salience map in which attention is allocated to the site of greatest salience, as represented by greatest LIP activity (Bisley and Goldberg, 2003).

However, the single-neuron data analyzed by Bisley and Goldberg (2006) reveal our motivating puzzle. The response of a given neuron i , averaged over trials when the distractor was in the neuron's RF, was fit as a simple exponential decay of the firing rate (Figure 1D). This response began from the peak of the visual response to the distractor, V_i , and fell with decay rate k_i . For the same neuron i , the average delay period activity on trials in which the target was in the RF was denoted as D_i . Individually, each of these response properties, V_i , k_i , and D_i , displayed almost a 10-fold variation across neurons. Furthermore, there was no stereotyped relation between any pair of these three quantities (Bisley and Goldberg, 2006). But remarkably, when considered together, the three quantities were not independent of each other. They obeyed a relation

$$\ln \frac{V_i}{D_i} \approx t_c k_i \quad (1)$$

as shown in Figures 1E and 1F, where t_c is a neuron-independent constant for a given monkey. Rearranging (1), we find

$$V_i e^{-k_i t_c} \approx D_i. \quad (2)$$

This is the mathematical statement that for every neuron i , the exponential decay, $V_i e^{-k_i t}$, of the response to the distractor crosses the delay level activity, D_i , of the response to the target at approximately the same time, t_c . Furthermore, for each monkey, this common neuronal crossing time t_c , which is 421 ms and 375 ms for monkeys B and I, respectively, is within 34 ms and 35 ms, respectively, of the attentional switching time, the time when psychophysical measurement found that neither the distractor nor target site had an attentional advantage (Figure 1C, top).

The observation in Equations 1 and 2 of a common single-neuron crossing time within a given monkey is both important and nontrivial for several reasons. First, while V_i and k_i are presumably determined by bottom-up mechanisms via the response to a new visual input, D_i is reflective of top-down salience mechanisms induced by the planned saccade to the target. Thus, the surprising relation obeyed by these three quantities in Equation 1 suggests the existence of a mechanism within LIP that mediates interactions between the top-down and bottom-up salience systems. Furthermore, this mechanism must be very robust in order to enforce Equation 1 despite the large degree of heterogeneity in the individual quantities V_i , D_i , and k_i . Also, the single-neuron

data in Figures 1E and 1F were collected over a number of months, indicating that this mechanism is not only robust but stable. Finally, the common neuronal crossing time t_c is different for each monkey, but nevertheless in both monkeys t_c is close to the attentional switching time, despite significant differences in the attentional switching time of the two monkeys. Following the interpretation in Bisley and Goldberg (2003), this suggests that the putative neural mechanism within LIP that enforces Equation 1 actually plays a functional role in determining the attentional switching time. However, despite its correlation with behavior, the crossing time of any neuron cannot be available on a single trial; average responses from two *different* trial types are required for its calculation. This only adds to the puzzle: why would LIP be organized to keep invariant a quantity that is not available on single trials while allowing substantial heterogeneity in all other aspects of the response?

The observation in Equation 1 is all the more striking because there is no obvious biophysical explanation for the emergence of a common crossing timescale t_c out of such a heterogeneous neuronal population. Any explanation involving the dynamics of isolated single neurons would suffer from severe fine-tuning problems. The crossing time of a single neuron would depend sensitively on the strength of top-down inputs and bottom-up visual inputs as well as intrinsic membrane and synaptic time constants that in turn determine the neuron's decay rate. The observed heterogeneity in the single-neuron response properties V_i , k_i , and D_i implies that these various parameters themselves are heterogeneous across neurons. The tight correlation between these three response properties seen in Figures 1E and 1F would then require unrealistic, finely-tuned relations between these heterogeneous parameters (see Supplemental Data for a mathematical analysis of these claims).

Moreover, the relation of Equation 1 also cannot be simply explained by a functional coupling between the different neurons that were recorded in LIP. There is no evidence for any such coupling. The RF's of the recorded neurons were at a variety of locations within the visual field. In trials when the target appeared in the RF of a recorded neuron, large transient firing rates induced in other regions of LIP by distractors flashed far from the RF seemed to have no effect on the low levels of delay period activity of the recorded neuron. Since many of the recorded neurons had RFs far enough from one another that, by this assay, they showed no functional coupling, any explanation for the observed common crossing time cannot rely on all of the recorded neurons participating in a common network.

A Common Crossing Time from One-Dimensional Dynamics

In the following, we find a simple and robust explanation for the fundamental observation in Equation 1 that circumvents all the above problems. We begin by explaining a key conceptual insight: how a rapid reduction in multineuronal dynamics to a single dimension would be sufficient to realize a common crossing time.

We first imagine that each recorded neuron i is part of a local, connected network of N neurons all sharing the same RF. We will later address the actual case in which the recorded neurons come from different subnetworks that are not connected to




Figure 2. Concept of One-Dimensional Dynamics

An example of firing-rate space for $N = 2$ neurons. Visual (V), delay (D), and spontaneous activity (S) are points in this space, while the visual transient response to a distractor is a curve starting at V and traveling to S . This dynamics can be viewed as motion down an energy valley, shown above, which is determined by the connectivity between neurons. The floor could be curved, as shown here, due to nonlinear effects. But regardless of its shape, its critical property of one dimensionality forces the decaying transient response to intersect the delay period activity. Furthermore, to the extent to which S and D do lie along a single line through the origin, the delay activity of each neuron will be proportional to its spontaneous activity.

one another. The collective firing-rate dynamics of the local population of neurons can be described by an N -dimensional firing-rate vector $\vec{r}(t)$, whose i^{th} component $r_i(t)$ represents the firing rate of neuron i at time t . At any given fixed time t , we can think of the collection of firing rates $r_i(t)$ as a point in N -dimensional firing-rate space, by plotting each value of $r_i(t)$ as i ranges from 1 to N on a different axis; see Figure 2 for a two-dimensional example. Then as time progresses, the firing-rate vector $\vec{r}(t)$ traces out a curve in firing-rate space. The set of delay activities D_i for each neuron i can be thought of as a special point, or vector \vec{D} in this firing-rate space. In the interpretation of Bisley and Goldberg (2003), this point represents the network's steady-state activity when salience is assigned to the common RF of all the neurons.

Now consider the response to a transient visual stimulus, corresponding to the distractor, flashed in this common RF. The firing rate of each neuron i will rise to a peak value V_i and then subsequently decay to its spontaneous firing rate during fixation S_i . Again we can think of each collection V_i and S_i as special points or vectors \vec{V} and \vec{S} in the same firing-rate space. Since all of the neurons peak at approximately the same time (Bisley et al., 2004), the transient response $\vec{r}(t)$ traces out a curve that starts at the point \vec{V} just after the peak response to the distractor and eventually decays back down to \vec{S} . Now in this picture, the constraint in Equation 2, that each neuron's transient response crosses its own delay activity at approximately the same time, is equivalent to requiring that the transient decay curve $\vec{r}(t)$

intersect the delay activity vector \vec{D} as it travels from \vec{V} to \vec{S} . The crossing time t_c is the time at which this intersection occurs. At this time, the transient response $r_i(t_c)$ equals D_i for all i , reproducing the observation in Equation 2.

As illustrated in Figure 2, a dynamical scenario in which this crossing will occur is if (1) all directions but one in firing-rate space decay away quickly, represented by rapid flow down the walls of the valley in Figure 2, and (2) the one direction that decays away more slowly, represented by a slower flow down the valley floor, contains the delay activity vector \vec{D} . That is, after the visual transient induced by the distractor, the network rapidly settles into a single dimension in firing-rate space, which represents a particular pattern of activity across neurons that is proportional to \vec{D} . The amplitude of this pattern then decays more slowly back down to spontaneous activity levels. We will subsequently explain how this dynamical scenario can be simply and robustly realized by the LIP circuit and how this realization then explains why even neurons that are not part of the same local network should have a common crossing time. But first, we examine the data for evidence of this one-dimensional dynamics.

Direct Experimental Evidence for One-Dimensional Dynamics

These dynamics lead to strong experimental predictions. It is possible that strongly nonlinear dynamics could build a valley floor in Figure 2 that is curved, so that although the delay activity \vec{D} and spontaneous activity \vec{S} both lie on the valley floor, they lie in very different directions in the firing-rate space. But let us make the stronger assumption on the dynamics that the valley floor is relatively straight. Then, independent of any other details of the circuit that realizes these dynamics, the dynamical scenario leads to the following predictions. First, it predicts that the multineuronal data in Bisley and Goldberg (2006) is one dimensional over long timescales, that is, it lies along a single direction in firing-rate space except during transient activity. This is not at all obvious from plots of the single-neuron data shown in Figures 3A and 3B, which display a high degree of heterogeneity. Second, it predicts that this one preferred dimension in firing-rate space corresponds to the spontaneous pattern of activity \vec{S} across neurons, as shown in Figure 2. Third, since the delay activity also settles into this preferred activity pattern, it predicts that the delay activity pattern \vec{D} across neurons should simply be a scaled-up version of the spontaneous pattern \vec{S} across neurons. This predicts a new fundamental constant in the LIP circuit: the ratio of delay activity of any given cell to its own spontaneous activity should be approximately identical across all cells. Fourth, on short timescales, just after large input perturbations, due for example to the distractor, the dynamics should not be confined to one dimension; it should depart from the spontaneous pattern of activity. Thus, in contrast to the delay activity \vec{D} , we predict that the peak visual response \vec{V} should show little relationship to the spontaneous activity pattern \vec{S} .

We verify all four of these predictions at once in Figures 3C and 3D by replotting the single-neuron data in a new coordinate system in firing-rate space. We plot the component of neural activity in the direction of the average spontaneous activity and the magnitude of the neural activity orthogonal to this direction in firing-rate space (see Supplemental Data). As predicted, we see that,

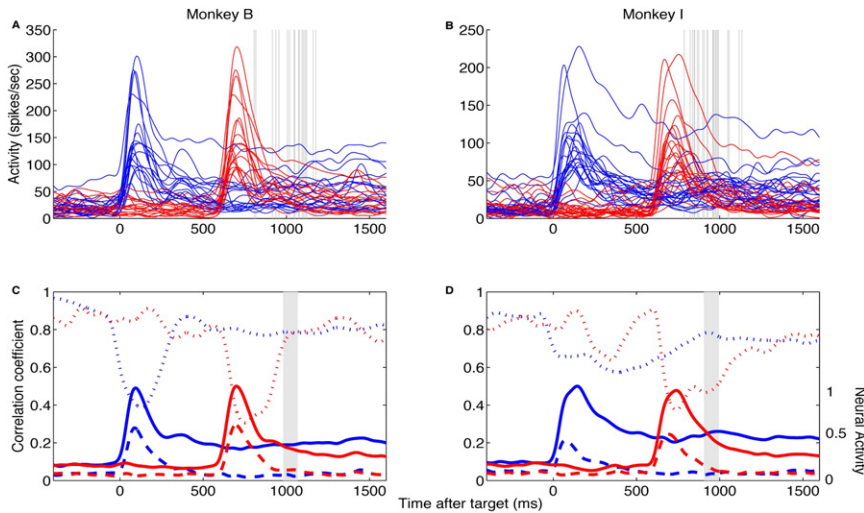


Figure 3. Reduction of Neural Activity to a Single Dimension

(A and B) Single-neuron responses in trials when the target (blue) or distractor (red) appeared in the RF. The vertical gray lines mark the times at which each neuron's decaying response to the distractor crosses its own delay activity in response to the target. Although there are outliers, the standard deviation of the crossing times are only 109 ms for monkey B and 90 ms for monkey I. Both these deviations are about the width of the window of neuronal ambiguity inherent in the experiment (Bisley and Goldberg, 2003).

(C and D) Solid curves represent the component of the target response (blue) and distractor response (red) along the direction of the spontaneous vector, in normalized units. The dashed curves represent the norm of the residual activity orthogonal to this direction. The dotted curves show the correlation coefficient, over 100 ms bins, between the instantaneous activity and the fixed spontaneous vector. The vertical gray bar represents a window centered at the time at which behaviorally there is no attentional advantage either at the distractor or the target.

throughout the duration of both trial types, the activity pattern across neurons is almost always identical to that of the spontaneous activity (solid lines). The activity orthogonal to this dimension (dashed lines) is suppressed at all times, except during the visual transient in response to either a target or a distractor. Ultimately, by the time attention switches back to the target, as measured behaviorally, activity orthogonal to the spontaneous activity has died away, and only a single dimension is left. One can also see the one-dimensional nature of the time course by considering a plot of the correlation coefficient between the instantaneous activity and the fixed average spontaneous activity (dotted lines). As predicted, the multineuronal activity is highly correlated with the spontaneous activity, except during visual transients, when feed-forward inputs presumably reduce this correlation.

In particular, the spontaneous activity is highly correlated with the delay activity, but not the visual transient activity. To demonstrate explicitly our third and fourth predictions, we have shown the scatterplots of spontaneous activity against delay and visual activity in Figures 4A and 4B. Of course, taken alone, a cell-by-cell correlation between average spontaneous and average delay activity could be easily explained by a trivial

excitability argument: more excitable cells fire more in both situations. However, this argument would also predict that these highly excitable cells would fire more during the visual transient. Thus, it would predict a similarly strong correlation between spontaneous and visual activity. This correlation is in fact much weaker (Figure 4B), arguing against such a simple excitability argument. Instead, the correlation between spontaneous and delay, but not visual activity, is explained naturally as a consequence of the reduction of neural activity to a single dimension. Thus, overall, the data in Figures 3 and 4 provide direct experimental evidence for our dynamical scenario and, moreover, reveal a fundamental new constant in the LIP network, namely a roughly fixed ratio between the delay and spontaneous activity across all cells in LIP.

The Robustness of One-Dimensional Dynamics

Having found evidence for our dynamical scenario, we now address its implementation at a network level, focusing on the crucial issue of why two neurons, in two different local subnetworks, or patches, of LIP each processing widely separated regions of visual space, would exhibit the same crossing time. To address this issue, we consider the worst-case situation, in which each local patch




Figure 4. Nontrivial Relationship between Spontaneous and Delay Activity

(A) The spontaneous activity of each cell is plotted on the horizontal axis and is the average of activity during fixation over 400 ms before target or distractor onset, averaged both over distractor-in-RF and target-in-RF trials. The delay activity of the same cell on the vertical axis is the average of activity between 800 and 1200 ms after target appearance in target-in-RF trials.

(B) The same spontaneous activity is plotted against peak visual activity in distractor-in-RF trials.




Figure 5. Model Reproduction

(A) From each of 100 different randomly chosen sparsely coupled excitatory networks of 100 neurons each, a single neuron is chosen at random and its visual/delay ratio is plotted against its decay rate. The solid black line shows a linear fit to the data.

(B) The same is done for 100 randomly chosen isolated neurons.

(C and D) For these same neurons, the relationship between the spontaneous versus delay activity and spontaneous versus peak visual activity are also plotted.

processes a region of visual space, without talking to other patches processing distant regions of visual space. This situation is consistent with the observed lack of strong interactions between neurons with substantially different RFs (Bisley and Goldberg, 2006).

To model the dynamics of each local patch, we use a phenomenological linear rate model because it is the simplest model that can reproduce the data and because it can be analyzed exactly (see Supplemental Data for a precise mathematical analysis of all the claims made in this section), yielding considerable intuition about the robustness of our proposed mechanism. However, we emphasize that the linear assumption is not essential to our approach; it serves simply as a particularly tractable realization of the essential dynamical idea proposed in Figure 2.

Sparse, random, net excitatory connectivity among the neurons in a local patch is sufficient to ensure the basic requirement of the dynamics we propose—that one pattern of activity decay much more slowly than all other patterns. Furthermore, assuming net inputs are excitatory, this connectivity also ensures that the delay activity \bar{D} lies along the slowly decaying direction. We assume that different patches are different random instantiations of such connectivity. Intrinsic decay rates are assigned randomly to each neuron in each patch. We assume that spontaneous activity is the response of the network to weak tonic bottom-up input, which we model as positive input that is randomly chosen for each neuron in each patch. We assume that the visual transient is driven by a strong pulse of bottom-up input, lasting 100 ms and with amplitude chosen randomly for each neuron. We assume that delay-period activity is the response of the network to both the ongoing spontaneous bottom-up input and to a positive top-down input representing the assignment of salience to the RF of a patch. This top-down input is randomly chosen for

each neuron in the patch. We assume tonic top-down input underlies delay activity, because it is unlikely that LIP alone can maintain elevated levels of persistent activity. For example, if a distractor flashes away from a saccade target, the population of neurons whose RF is at the distractor has much greater activity than any other location in LIP, yet this activity cannot maintain itself and decays back down to spontaneous levels. Instead, elevated levels of persistent activity in

response to a saccade target are more likely to arise from cooperative effects distributed across various nodes in the oculomotor circuit, including the dorsolateral prefrontal cortex (DLPFC), frontal eye fields (FEF), and the superior colliculus (SC), each of which also show persistent activity in the delayed saccade task (Bruce and Goldberg, 1985; Funahashi et al., 1989; Glimcher and Sparks, 1992; Kustov and Robinson, 1996). Alternative scenarios for LIP persistent activity are addressed in the Discussion.

By the above arguments, neurons within a single patch will have a common crossing time, but there is no a priori reason to expect neurons chosen from different patches to have a common crossing time. Both the local circuitry within each patch and the inputs to each patch are chosen randomly and independently of the circuitry and inputs to any other patch. The only thing different patches have in common is the overall distributions from which their circuit and input parameters are sampled. Nevertheless, as seen in the model results shown in Figure 5A, the responses of the neurons in different patches are correlated and show a common crossing time. We have simulated $k = 100$ local patches of $N = 100$ neurons each. For each patch, we chose a single neuron at random and recorded its spontaneous, delay, and peak visual activities and its decay rate in response to the various randomly chosen inputs for that patch. Despite the fact that none of the neurons whose response characteristics are plotted in Figure 5A were ever coupled to one another, the model reproduces the experimentally observed correlation across neurons observed in Figures 1E and 1F, representing a common crossing time. Furthermore, Figures 5C and 5D reproduce the essential experimental observation in Figures 4A and 4B, namely a high correlation between spontaneous and delay activity, but a low correlation between spontaneous and visual activity.

A key reason why such correlations exist in the model is uncovered when we reduce the number of neurons N in each local patch: in the extreme case, when $N = 1$, the correlation largely disappears (Figure 5B). Thus, the correlations arise specifically due to network dynamics. Due to the variability in inputs, cellular time constants, and connectivity, the different neurons chosen from different patches show realistic levels of variability in their effective time constants, in their spontaneous, delay, and visual activities, and in their ratios of visual to delay activity, and this is true whether $N = 1$ or $N = 100$. But despite this heterogeneity, the crossing times of neurons in different patches can remain relatively invariant as follows. Because the crossing time of any neuron within a patch occurs after all the fast modes have decayed away, this crossing time depends predominantly on the dynamics of the slow mode of that patch. In particular, the common crossing time of all neurons within a patch is governed by the time at which the patch's slow mode amplitude in response to the visual transient crosses the slow mode amplitude in the delay activity. In turn, this time depends on only three attributes of the slow mode for that patch: the amplitude of the slow mode excitation in response to visual inputs, the rate of decay of the slow mode, and the amplitude of the slow mode during delay activity. For sufficiently large N , each of these three attributes of the slow mode depends only on the statistical properties of each patch's connectivity and input strengths, not on their detailed realizations. Thus, as long as each patch has similar statistical properties, the three critical properties of the slow mode dynamics will remain invariant across different patches. As a result, the single-neuron crossing time is not only common to all the neurons within a patch, but it varies little from patch to patch, again provided that N is sufficiently large. As seen in Figure 5A, as few as 100 sparsely connected neurons per patch is sufficient to reproduce the original, puzzling experimental observations.

In summary, so long as LIP connectivity is statistically homogeneous and shows a connectivity pattern such as sparse local excitatory connectivity that leads to a single slow mode in each local patch, neurons in distant regions of LIP will have the same crossing time even if they do not talk to each other, because they will be embedded in two different local networks that have the same robust slow dynamical properties, despite having otherwise highly heterogeneous response characteristics.

One-Dimensional LIP Dynamics in a Decision-Making Task

Is one-dimensional slow dynamics a robust feature of LIP circuitry, or is it something that only occurs in the particular context of persistent delay-period activity? To address this question, we examine data from monkeys performing a completely different task involving perceptual decision making (Roitman and Shadlen, 2002). In these experiments, LIP neurons showed slowly changing activity (rather than persistent activity as in the delay-period task) that was correlated with the integration of sensory evidence in the formation of a perceptual decision. Specifically, monkeys were trained to perform a random-dot-motion discrimination task, where the direction of motion of the dots was reported by making a saccade either into or out of the RF of a recorded neuron in LIP. When the neuron's activity was averaged over trials in which the monkey decided to make a saccade into

the RF, the firing rate showed a ramping dynamics during the stimulus presentation. The initial slope of the ramp was positively correlated with the motion strength, and in the reaction time version of the task, the average firing rate was stereotyped at the end of the decision-making process, independent of the motion strength and response time. These results led to the interpretation by Roitman and Shadlen (2002) that LIP firing-rate dynamics reflect the stochastic integration of sensory evidence, originating in MT (Salzman, et al., 1992; Ditterich, et al., 2003; Hanks, et al., 2006), during the decision-making process and that a decision is made after the firing rate reaches a fixed threshold.

Despite the differences in tasks and neuronal behaviors, the two experiments shared similar criteria for selecting the neurons studied. In Bisley and Goldberg (2003), neurons that had visual responses to saccade targets in the delayed saccade task were chosen for further study. Of these neurons, 83% also had delay period activity in the same task. In Roitman and Shadlen (2002), neurons that showed delay activity in the delayed saccade task were chosen for further study. Given that similar classes of neurons were involved, we decided to test whether the framework of slow one-dimensional dynamics in LIP applies to the Roitman and Shadlen (2002) experiments.

We repeat, on the decision-making data, the same analysis that was used to extract the one-dimensional structure in the attentional data shown in Figures 3C and 3D. For each monkey, we estimate the spontaneous firing rate S_i for each neuron i using data from the fixation period, before the presentation of the random-dots stimulus. Then, thinking of this collection of rates as a vector \vec{S} , we plot in Figure 6 the ramping activity in a new coordinate system, with activity in the direction of the average spontaneous activity vector \vec{S} in solid lines and the magnitude of all activity orthogonal to this mode in dashed lines. We do this for both monkeys in each of 12 trial types, specified by six possible coherence levels and two possible choices. On each of these plots we also superimpose the time course of the correlation coefficient between the ramping activity pattern and the fixed spontaneous activity pattern \vec{S} (dotted lines).

The combined data in Figure 6 largely verify the predictions of our theory. In almost all cases, the ramping activity occurs along a single dimension that is well correlated with the spontaneous activity. This effect is more pronounced in monkey N (Figures 6B and 6D) than in monkey B (Figures 6A and 6C) and is particularly weak for monkey B for trials on which saccades were made out of the RF of the recorded neuron (Figure 6C). This may simply reflect poorer sampling in monkey B. After dividing all the trials into each of the 12 trial types, for any given trial type only a subset of the neurons had enough data to reliably estimate the firing rate. As a result, of 54 cells recorded in monkey N, an average of 34 neurons contributed to each curve in Figures 6B and 6D. In contrast, of 33 cells recorded in monkey B, an average of only 11 neurons contributed to each curve in Figures 6A and 6C. Comparatively, in the data from the attentional task, 18 neurons from monkey B contributed to the curves in Figure 3C, and 23 neurons from monkey I contributed to those in Figure 3D. With the smaller number of neurons, the analysis for any given trial type is more sensitive to fluctuations in single-neuron dynamics. In addition, the lower firing rates on trials involving saccades out of the RF might lead these fluctuations to have relatively more




Figure 6. Decision-Making Data of Roitman and Shadlen (2002)

(A and B) For each monkey, the traces describe the ramping activity at various coherence levels (percentage of dots moving in a single direction), indicated by color, on trials in which the monkey made a saccade into the RF of the recorded neuron. The solid curves show the component of the multineuron rate vector along the direction of the spontaneous vector. The dashed curves show the norm of the residual activity orthogonal to this. Both these activity levels in the lower half of each plot are in normalized units varying from 0 to 1, with the scale shown on the right-hand axes. Superimposed on these plots are dotted curves that show the correlation coefficient between the instantaneous ramping activity and the spontaneous vector. The scale of the correlation coefficient is shown on left-hand axes.

(C and D) The same data for each monkey and each coherence level, but on trials in which the monkey made a saccade away from the RF of the recorded neuron.

effect. Nevertheless, the combined data from both monkeys exhibit evidence for our theoretical framework, here in the context of decision making, despite the significantly different dynamical behavior of LIP neurons in this context compared to their simpler, steady-state behavior in the attentional task.

DISCUSSION

Both the functional role and mechanistic origins of persistent and integrating activity in LIP have been a topic of debate in the literature. This activity has not only been implicated in attention and perceptual decision making but also in motor planning (Snyder et al., 1997; Platt and Glimcher, 1997), economic decision making (Platt and Glimcher, 1999; Dorris and Glimcher, 2004), and visual categorization (Freedman and Assad, 2006) as well as the representation of value (Sugrue et al., 2004), time (Leon and Shadlen, 2003; Janssen and Shadlen, 2005), and inferred motion (Assad and Maunsell, 1995; Eskandar and Assad, 1999). We have shown, in two different experimental paradigms, that local LIP dynamics rapidly reduces to a single dimension and that the persistent and integrating activity underlying LIP functional responses are simply a scaled version of the spontaneous activity during fixation. These observations may in some sense transcend the debates on LIP function; they suggest a functional mechanism by which LIP could exert robust temporal control over a wide variety of oculomotor behaviors on the timescale of 100s of milliseconds and also suggest a distributed origin of persistent activity in the oculomotor system.

Robust Behavioral Dynamics from Heterogeneous Biophysical Dynamics

The variety of behaviors implicated in LIP function require fine temporal control. For example, in the context of attentional shifting, top-down information about behavioral relevance allows us to focus on a fixed location to accomplish goal-directed tasks, while the bottom-up salience of new or changing stimuli can shift our attention elsewhere in the visual field, allowing us to deal with

new threats or opportunities. There must be a delicate balance between the two mechanisms: if the former is too strong, we will be either unable or too slow to attend to important new stimuli, whereas if the latter dominates, then we will be rapidly distracted and unable to attend to any one location.

Similarly, in the context of decision making, neural circuitry must make a dynamical transition from an uncommitted state to a final decision by temporally integrating noisy evidence. If the circuitry performs this transition too quickly, we may arrive at an incorrect decision because it does not take enough time to integrate all available evidence. If the circuit performs this transition too slowly, a decision, although more likely to be correct, may come too late to be of behavioral relevance. This tradeoff is well known in the theory of decision making as the speed-accuracy tradeoff (Gold and Shadlen, 2002) and is analogous to the delicate balance faced in allocating attention between top-down and bottom-up salience signals, in the sense that both behaviors require fine temporal control.

The required, and exhibited, robustness of attentional and decision-making dynamics stands in sharp contrast to the unreliability and heterogeneity of single-neuron data in LIP, a brain region that is thought to mediate both behaviors. We have proposed and found evidence for a simple dynamical principle that can help span the bridge between heterogeneous single neurons and robust behavior. Although single neurons in LIP can have highly heterogeneous response properties, the temporal properties of each slow mode in each local patch of LIP, such as its level of excitation and decay time constant, are highly reliable from trial to trial and uniform from patch to patch because they depend only on statistical properties of LIP connectivity and inputs. For example, as shown in the [Supplemental Data](#), the decay time constant of this slow mode is a particularly robust quantity that depends only on the mean level of recurrent excitation and is remarkably insensitive to statistical fluctuations in synaptic strengths. Overall, these slow mode dynamics can reliably drive behavior on a trial-by-trial basis. This picture explains the functional significance of a common crossing time across

neurons, a quantity that is not available on a single trial; this common crossing time is really a signature of the robust decays of slow modes in local regions of LIP, which *are* available on a single trial.

Consider the above requirement of a delicate balance between top-down and bottom-up attention. If, as suggested in [Bisley and Goldberg \(2003\)](#), the locus of attention is determined by the RF of neurons with maximal firing rates, then a robust attentional switching time would arise on a trial-by-trial basis, simply because the slow mode activity at the distractor would cross the slow mode activity at the target at approximately the same time trial after trial, due to the emergent, robust temporal dynamics of each slow mode. No other special mechanism would be required to fine tune the balance between top-down and bottom-up salience systems. Consider also the above requirement of tightly controlling the rate of integration in decision making in order to implement a particular speed-accuracy tradeoff. If, as suggested in [Roitman and Shadlen \(2002\)](#), the decision making is implemented by LIP neurons via an integration-to-threshold mechanism, it is crucial that the integration occur with the same dynamics across neurons processing distant regions of the visual space. If this does not happen, then there would be a nonuniform implementation of the speed-accuracy tradeoff across the visual field. However, if in each local region the dynamics is one dimensional, then this dynamics is determined not by heterogeneous single-neuron properties but by the robust dynamical properties of the slow mode, which would be the same across all regions of LIP, thereby allowing a uniform implementation of the speed-accuracy tradeoff without any fine tuning.

It remains unclear precisely how the slow mode dynamics are used in integration. The slow mode time constant is a few hundred milliseconds in the experiments of [Bisley and Goldberg \(2003, 2006\)](#). The time constant of integration in the experiments of [Roitman and Shadlen \(2002\)](#) seems likely to be closer to 1 s (see also [Huk and Shadlen, 2005](#)). One possibility is that the time constant can be modulated, for example by factors that control the overall strength of recurrent excitation across LIP. Another is that integration is performed elsewhere and transmitted via top-down input that drives the LIP slow modes. Alternatively, it has been argued that the neural time constant of integration in the [Roitman and Shadlen \(2002\)](#) experiments actually could be a few hundred milliseconds or even faster ([Ditterich, 2006](#)).

More generally, the considerable temporal robustness achieved in the reduction of local neural signals to a single dimension is a phenomenon that would allow LIP, as well as other brain regions, to tightly control behavioral dynamics, regardless of the particular behavior involved. Stronger recurrent excitation in a given cortical area would yield slower timescales, so one could conjecture that the observed gradient across cortical areas in the strength of recurrent excitation ([Elston et al., 2005](#)) might correspond to a gradient in the timescales over which cortical areas control behavior. Our direct observation of a collective mode of neural activity in [Figures 3C and 3D](#) decaying over hundreds of milliseconds then suggests that LIP circuitry may specialize in the robust control of a variety of oculomotor behaviors on this timescale.

Implications for the Origins of Persistent Activity in LIP

The dominant theoretical paradigm for the explanation of persistent neural activity in the absence of a stimulus is “attractor dynamics,” in which persistent activity arises as an attractor or self-maintaining state of the neural dynamics that maintains itself through reverberating positive feedback. There exists a large class of attractor models employing such reverberating feedback either at the level of the single cell ([Booth and Rinzel, 1995](#); [Lowenstein and Sompolinsky, 2003](#)) or the local network ([Ben-Yishai, et al., 1995](#); [Seung, 1996](#); [Amit and Brunel, 1997](#); [Compte et al., 2000](#); [Wang, 2001](#)). A common property of simple attractor models is that a transient response to a stimulus selects one attractor out of many possible ones by pushing the network’s dynamical state into “the basin of attraction” for that attractor (the set of states that will evolve over time into that attractor state). Thus, if the transient response to a stimulus is not dependent on stimulus history, then the current attractor state of the network is essentially determined by the last presented stimulus.

This theoretical scenario is insufficient to explain the observed LIP data ([Bisley and Goldberg, 2003, 2006](#)). The saccade target and distractor are identical stimuli that appear to evoke identical transient visual responses. Yet the target leads to a localized state of persistently elevated activity among neurons that respond to it, while the distractor does not. Although attractor models exist that are capable of ignoring distractors to some extent ([Brunel and Wang, 2001](#)), they rely on a suppression of the transient response to the distractor by the active attractor. We are unaware of any attractor model that can ignore a distractor whose transient neural response is identical to the one that leads to the attractor state or involves much larger firing rates than that of the attractor state, both of which are seen in the experimental data.

A crucial theoretical ingredient that is missing in simple attractor models is some form of gating mechanism that, for example, would allow LIP activity to either shift to persistent activity or not, depending on the behavioral relevance of the visual stimulus. Such a gating mechanism is likely to arise from the interactions among different brain regions active during a task. Indeed, in the delayed saccade task, not only LIP, but many other brain regions, including the DLPFC ([Funahashi, et al., 1989](#)), FEF ([Bruce and Goldberg, 1985](#)), and SC ([Glimcher and Sparks, 1992](#)), display persistent activity during the delay period. From the above considerations, it seems clear that LIP alone cannot function as an attractor network. Top-down inputs from FEF ([Chafee and Goldman-Rakic, 2000](#)) or other areas must contribute significantly to LIP persistent activity, whether as transient inputs that gate LIP’s state or as ongoing inputs during the persistent activity. At the same time, the observation that LIP delay activity is tightly correlated with LIP spontaneous activity, when LIP presumably is not receiving large inputs from other brain regions, strongly suggests that during delay activity, LIP is not simply passively receiving inputs from other brain regions; LIP also plays a role in sculpting its own persistent activity. In essence, persistent states in LIP are a relatively weak perturbation of spontaneous states, which in turn are determined by intrinsic LIP properties. Thus, persistent activity in LIP is likely to originate from cooperative feedback between multiple brain regions,

allowing for a rich variety of gating mechanisms that would act by modulating this cooperative feedback.

The feedback could be as simple as some region other than LIP acting alone as an attractor, and then, when it is in the attractor state, providing persistent input to LIP that maintains persistent activity in LIP. In this case, the localized attractor model would still be basically correct, but the localized attractor would not be located in LIP (and the mechanism of behavioral gating in this other area would remain to be worked out). But the slow modes we have seen in LIP activity suggest an alternative hypothesis that may be worth exploring. Suppose that none of the areas involved can act alone as attractors, but that all have slow modes. The attractor state representing persistent activity could then be built out of coupling between the slow modes of the different areas. A slow mode is basically a leaky attractor, which does not have strong enough recurrent excitation to maintain activity in the absence of an input. If the slow modes of different areas provide input to one another, they could “bootstrap” themselves into an attractor: each area would provide persistent input to the others, and this persistent input would in turn enable each area’s slow mode to maintain persistent activity. By coupling together, they create a network with more recurrent excitation than any one alone, and this multiarea network could maintain activity without an input. The gating decisions would then be made in the decisions of how much feedback the areas send to one another in a given circumstance. Since there is rich cortical circuitry between the inputs that a cortical area receives and the output it sends to other areas, this provides rich possibilities for gating, though of course the precise mechanisms constitute a key research question. Such an added level of complexity underlying persistent activity may be essential to bringing existing attractor models closer to experimental data, by allowing them to incorporate the necessary gating mechanisms so that the current attractor state of the distributed oculomotor system can be driven by behavioral relevance rather than simply by the last seen stimulus.

SUPPLEMENTAL DATA

The Supplemental Data for this article can be found online at <http://www.neuron.org/cgi/content/full/58/1/15/DC1/>.

ACKNOWLEDGMENTS

We would like to thank Larry Abbott for helpful comments on the manuscript. This work was supported by the Human Frontiers Science Program, HHMI, the J.S. McDonnell, Swartz, Whitehall, and McKnight Foundations, and NIH grants EY-11001, EY-15634, EY-014978, EY-05603, RR-00166, and EY-11378.

Received: April 9, 2007

Revised: July 23, 2007

Accepted: January 9, 2008

Published: April 9, 2008

REFERENCES

- Amit, D.J., and Brunel, N.P. (1997). Model of global spontaneous activity and local structured activity during delay periods. *Cereb. Cortex* 7, 237–252.
- Assad, J.A., and Maunsell, J.H. (1995). Neuronal correlates of inferred motion in primate posterior parietal cortex. *Nature* 373, 518–521.
- Bashinski, H.S., and Bacharach, V.R. (1980). Enhancement of perceptual sensitivity as the result of selectively attending to spatial locations. *Percept. Psychophys.* 28, 241–248.
- Ben-Yishai, R., Baror, R.L., and Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proc. Natl. Acad. Sci. USA* 92, 3844–3848.
- Bisley, J.W., and Goldberg, M.E. (2003). Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299, 81–86.
- Bisley, J.W., and Goldberg, M.E. (2006). Neural correlates of attention and distractibility in the lateral intraparietal area. *J. Neurophysiol.* 95, 1696–1717.
- Bisley, J.W., Krishna, B.S., and Goldberg, M.E. (2004). A rapid and precise on-response in posterior parietal cortex. *J. Neurosci.* 24, 1833–1838.
- Booth, V., and Rinzel, J. (1995). A minimal, compartmental model for a dendritic origin of bistability of motoneuron firing patterns. *J. Comput. Neurosci.* 2, 299–312.
- Bruce, C.J., and Goldberg, M.E. (1985). Primate frontal eye fields. I. Single neurons discharging before saccades. *J. Neurophysiol.* 3, 603–635.
- Brunel, N.P., and Wang, X.J. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *J. Comput. Neurosci.* 11, 63–85.
- Chafee, M.V., and Goldman-Rakic, P.S. (2000). Inactivation of parietal and prefrontal cortex reveals interdependence of neural activity during memory-guided saccades. *J. Neurophysiol.* 3, 1550–1566.
- Ciaramitaro, V.M., Cameron, E.L., and Glimcher, P.W. (2001). Stimulus probability directs spatial attention: an enhancement of sensitivity in humans and monkeys. *Vision Res.* 41, 57–75.
- Colby, C.L., Duhamel, J.R., and Goldberg, M.E. (1996). Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. *J. Neurophysiol.* 76, 2841–2852.
- Compte, A., Brunel, N.P., and Wang, X.J. (2000). Synaptic mechanisms and network dynamics underlying visuospatial working memory in a cortical network model. *Cereb. Cortex* 10, 910–923.
- Deubel, H., and Schneider, W. (1996). Saccade target selection and object recognition: evidence for a common attentional mechanism. *Vision Res.* 36, 1827–1837.
- Ditterich, J. (2006). Stochastic models of decisions about motion direction: behavior and physiology. *Neural Netw.* 19, 981–1012.
- Ditterich, J., Mazurek, M., and Shadlen, M.N. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat. Neurosci.* 6, 891–898.
- Dorris, M.C., and Glimcher, P.W. (2004). Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 44, 365–378.
- Egeth, H.E., and Yantis, S. (1997). Visual attention: control, representation, and time course. *Annu. Rev. Psychol.* 48, 269–297.
- Elston, G.N., Benavides-Piccione, R., and Defelipe, J. (2005). A study of pyramidal cell structure in the cingulate cortex of the macaque monkey with comparative notes on inferotemporal and primary visual cortex. *Cereb. Cortex* 15, 64–73.
- Eskandar, E.N., and Assad, J.A. (1999). Dissociation of visual, motor and predictive signals in parietal cortex during visual guidance. *Nat. Neurosci.* 2, 88–93.
- Freedman, D.J., and Assad, J.A. (2006). Experience-dependent representation of visual categories in parietal cortex. *Nature* 443, 85–88.
- Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1989). Mnemonic coding of visual space in the monkey’s dorsolateral prefrontal cortex. *J. Neurophysiol.* 61, 331–349.
- Glimcher, P., and Sparks, D. (1992). Movement selection in advance of action in the superior colliculus. *Nature* 355, 542–545.
- Gold, J.I., and Shadlen, M.N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions and reward. *Neuron* 36, 309–322.

- Gottlieb, J.P., and Goldberg, M.E. (1999). Activity of neurons in the lateral intraparietal area of the monkey during and antisaccade task. *Nat. Neurosci.* 2, 906–912.
- Gottlieb, J.P., Kusunoki, M., and Goldberg, M.E. (1998). The representation of visual salience in monkey parietal cortex. *Nature* 391, 481–484.
- Hanks, T., Ditterich, J., and Shadlen, M.N. (2006). Microstimulation of macaque area lip affects decision making in a motion discrimination task. *Nat. Neurosci.* 9, 682–689.
- Huk, A.C., and Shadlen, M.N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* 25, 10420–10436.
- Janssen, P., and Shadlen, M.N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. *Nat. Neurosci.* 8, 234–241.
- Kustov, A.A., and Robinson, D.L. (1996). Shared neural control of attentional shifts and eye movements. *Nature* 384, 74–77.
- Leon, M.I., and Shadlen, M.N. (2003). Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron* 38, 317–327.
- Lowenstein, Y., and Sompolinsky, H. (2003). Temporal integration by calcium dynamics in a model neuron. *Nat. Neurosci.* 6, 961–967.
- Platt, M.L., and Glimcher, P.W. (1997). Responses of intra-parietal neurons to saccadic targets and visual distractors. *J. Neurophysiol.* 78, 2164–2175.
- Platt, M.L., and Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238.
- Powell, K.D., and Goldberg, M.E. (2000). Response of neurons in the lateral intraparietal area to a distractor flashed during the delay period of a memory guided saccade. *J. Neurophysiol.* 84, 301–310.
- Robinson, D.L., Bowman, E.M., and Kertzman, C. (1995). Covert orienting of attention in macaques II: Contributions of parietal cortex. *J. Neurophysiol.* 74, 698–712.
- Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489.
- Saizman, C.D., Murasugi, C.M., Britten, K.H., and Newsome, W.T. (1992). Microstimulation in visual area MT: Effects on direction discrimination in performance. *J. Neurosci.* 12, 2331–2355.
- Seung, H.S. (1996). How the brain keeps the eyes still. *Proc. Natl. Acad. Sci. USA* 93, 13339–13344.
- 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.
- Shepherd, M., Findlay, J., and Hockey, R. (1986). The relationship between eye movements and spatial attention. *Q. J. Exp. Psychol.* 38, 475–491.
- Snyder, L.H., Batista, A.P., and Andersen, R.A. (1997). Coding of intention in the posterior parietal cortex. *Nature* 386, 167–170.
- Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science* 18, 1782–1787.
- Wang, X.J. (2001). Synaptic basis of cortical persistent activity: the importance of NMDA receptors to working memory. *J. Neurosci.* 19, 9587–9603.
- Yantis, S., and Jonides, J. (1984). Abrupt visual onsets and selective attention: evidence from visual search. *J. Exp. Psychol. Hum. Percept. Perform.* 10, 601–621.
- Yantis, S., and Jonides, J. (1996). Attentional capture by abrupt onsets: new perceptual objects or visual masking? *J. Exp. Psychol. Hum. Percept. Perform.* 22, 1505–1513.