

Probabilistic Population Codes for Bayesian Decision Making

Jeffrey M. Beck,^{1,7} Wei Ji Ma,^{1,2,7} Roozbeh Kiani,³ Tim Hanks,³ Anne K. Churchland,³ Jamie Roitman,⁴ Michael N. Shadlen,³ Peter E. Latham,⁵ and Alexandre Pouget^{1,6,*}

¹Department of Brain and Cognitive Sciences, University of Rochester, Rochester, NY 14627, USA

²Department of Neuroscience, Baylor College of Medicine, Houston, TX 77030, USA

³Howard Hughes Medical Institute and Department of Physiology and Biophysics, University of Washington, Seattle, WA 98195-7330, USA

⁴Department of Psychology, University of Illinois at Chicago, Chicago, IL 60607-7137, USA

⁵Gatsby Computational Neuroscience Unit, London WC1N 3AR, UK

⁶Theoretical Neuroscience Group, Collège de France, Paris 75005, France

⁷These authors contributed equally to this work

*Correspondence: alex@bcs.rochester.edu

DOI 10.1016/j.neuron.2008.09.021

SUMMARY

When making a decision, one must first accumulate evidence, often over time, and then select the appropriate action. Here, we present a neural model of decision making that can perform both evidence accumulation and action selection optimally. More specifically, we show that, given a Poisson-like distribution of spike counts, biological neural networks can accumulate evidence without loss of information through linear integration of neural activity and can select the most likely action through attractor dynamics. This holds for arbitrary correlations, any tuning curves, continuous and discrete variables, and sensory evidence whose reliability varies over time. Our model predicts that the neurons in the lateral intraparietal cortex involved in evidence accumulation encode, on every trial, a probability distribution which predicts the animal's performance. We present experimental evidence consistent with this prediction and discuss other predictions applicable to more general settings.

INTRODUCTION

Decision making affects all aspects of human behavior, on time scales varying from seconds to hours to days. For instance, imagine you are driving your car toward a busy intersection and your brakes fail. Within a few hundred milliseconds, you have to decide where to steer your car. Although this is a task we handle relatively easily, in fact it involves three separate, and nontrivial, stages. First, sensory evidence must be accumulated over time. Here, the sensory evidence consists of the image of cars and people in the intersection. Second, the accumulation must be stopped at some point (waiting too long can have disastrous consequences in this situation). Third, an action must be selected. This task is difficult because the sensory evidence and the response are continuous variables, the reliability of the sensory evidence is a priori unknown, and it can vary greatly

over time. For instance, as you get closer to the intersection, your ability to distinguish different objects improves. The reliability of the visual information can also vary from day to day: it is much easier to analyze the scene on a sunny day than on a foggy one.

There is currently no neural model that can deal with this type of decision optimally, where by optimal, we mean that the accumulation of evidence is done without loss of information and that the chosen option is the most likely one given the sensory evidence (we do not address the issue of when to make the decision; see *Discussion*). Yet, it is essential to understand optimal decision making in the face of multiple choices and unknown and time-varying reliability, since most decisions we make fall into this category. Most models are concerned only with binary decision making, and even with this limitation, cannot deal optimally with sensory evidence of unknown and continuously changing reliability. This problem is conceptual: these models have no clear probabilistic interpretation or, when they do, are limited to situations in which the evidence has a constant and known reliability over time and over trials. As a result, it is unclear how, or even if, they are related to the general case we consider in this paper.

Here, we present the first neural model of decision making that performs sensory evidence accumulation and response selection optimally when there are multiple or a continuum of possible decisions and the reliability of the sensory input varies over time or across trials. This model is built around the observation that spike counts in the brain are close to what we call "Poisson-like" (Ma et al., 2006; Shadlen and Newsome, 1998; Tolhurst et al., 1983). Given this observation, our main contributions are twofold. First, we show that for Poisson-like distributions, optimal evidence accumulation can be performed through simple integration of neural activities, while optimal response selection can be implemented through attractor dynamics. Second, we show (again for Poisson-like distributions of neural activity) that neurons encode the posterior probability distribution over the variables of interest at all times. This latter contribution has far-reaching implications, since it suggests that neurons implicated in simple perceptual decisions represent quantities that are directly relevant to inference, confidence, and belief.

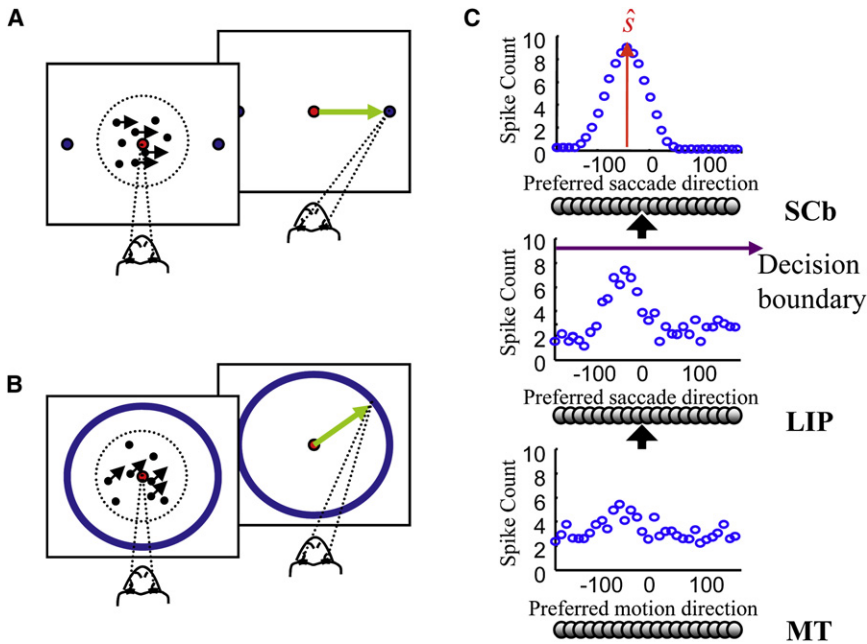


Figure 1. Task and Network Architecture

(A) Binary decision making. The subject must decide whether the dots are moving to the right or to the left. Only a fraction of the dots are moving to the right or the left coherently (black arrows). The other dots move in random directions. The animal indicates its response by moving its eyes in the perceived direction (green arrow).

(B) Continuous decision making, for which the dots can move in any direction. The animal responds by making a saccade to the outside circle in the perceived direction.

(C) Network architecture. The network consists of three interconnected layers of neurons with Gaussian tuning curves. In MT, the tuning curves are for direction of motion, while in LIP and SCb, the tuning curves are for saccade direction. The layers differ by their connectivity and dynamics. The LIP neurons have a long time constant (1 s), allowing them to integrate their input, and lateral connections, allowing them to implement short-range excitation and long-range inhibition. The SCb layer forms an attractor network, for which smooth hills of activity are stable regardless of their position. The blue dots indicate representative patterns of activity 200 ms into a trial for the MT and LIP layer and at the end of the trial for the SCb layer.

When accompanied by a termination rule, our model, like a number of others (Ditterich et al., 2003; Gold and Shadlen, 2007; Laming, 1968; Link, 1992; Link and Heath, 1975; Mazurek et al., 2003; Ratcliff and Rouder, 1998; Reddi and Carpenter, 2000; Smith and Ratcliff, 2004; Stone, 1960; Usher and McClelland, 2001; Wald, 1947; Wald and Wolfowitz, 1948) accounts for the speed-accuracy tradeoff reported in humans and monkeys for binary choices. However, it goes beyond previous neural models in three ways. First, it captures the speed-accuracy trade-off and the physiology of LIP cells in experiments involving four choices. Second, as previously indicated, it predicts that neural activity in LIP encodes a probability distribution over actions. This is a new prediction about the response of LIP neurons, which we have tested and verified using data from area LIP recorded while monkeys engaged in a decision among two or four alternatives. Third, it makes predictions for the responses of cells in LIP and SC when there are multiple choices, a continuum of choices, and when the reliability of the cue varies over time.

RESULTS

Task and Model Architecture

For concreteness, we consider a motion direction task that has been extensively used to study decision making in humans and animals. In this task, an observer sees a random-dot kinematogram in which a fraction of the dots move coherently in a particular direction while all the other dots move randomly (Figure 1A). The task of the observer is to report the direction of motion with a saccadic eye movement to a choice target that is associated with that direction of motion. The reliability of the sensory evidence can be controlled by changing the percentage of dots moving coherently. In most experiments, this task is

restricted to binary decision making (right versus left) and constant coherence over the course of a trial. We also consider a more general setting in which the mean direction of moving dots and the direction of the saccade can take any value (Figure 1B) and the reliability of the motion information (the coherence) can vary not only across trials, but also during a trial.

A minimal model of this task (and, in fact, any decision-making task that involves integrating evidence over time) requires three distinct populations of neurons: an input layer, an evidence accumulation layer, and a readout layer where motor output is generated (Figure 1C). Here, we label these MT (middle temporal), LIP (lateral intraparietal), and SCb (superior colliculus, in particular those cells that exhibit a motor burst; hence the index “b”), based on what is known about the functions of these regions. These labels are used for convenience only: it is quite likely that the sensory integration involves many other cells beside the ones in LIP, and that the motor burst is not generated solely in the SC.

Bayesian Formulation

We denote the population activity of M neurons in area MT at time t_n by a vector $\mathbf{r}^{\text{MT}}(t_n)$ (see Figure 1C for an example), where $\mathbf{r}^{\text{MT}} \equiv \{r_1^{\text{MT}}, \dots, r_M^{\text{MT}}\}$ and $r_i^{\text{MT}}(t_n)$ is the spike count of neuron i in the time interval $[(n-1)\delta t, n\delta t]$. In our simulations we set δt to 50ms, although our results are insensitive to that choice.

The stimulus is characterized by a direction of motion, s , and task-irrelevant variables such as contrast and motion coherence, which we refer to as nuisance parameters and collectively denote \mathbf{c} (where $\mathbf{c} = \{c(t_1), c(t_2), \dots, c(t_N)\}$). When a stimulus (s, \mathbf{c}) is presented, MT generates a series of patterns of activity over time, denoted $\mathbf{r}^{\text{MT}}(t_1:t_N) \equiv \{\mathbf{r}^{\text{MT}}(t_1), \dots, \mathbf{r}^{\text{MT}}(t_N)\}$. Because of neural variability, $\mathbf{r}^{\text{MT}}(t_1:t_N)$ is not the same on every presentation of

(s, \mathbf{c}), but follows a probability distribution $p(\mathbf{r}^{\text{MT}}(t_1:t_N)|s, \mathbf{c})$. If we assume that the activity is uncorrelated on timescales of 50 ms, this distribution can be written as a product over time,

$$p(\mathbf{r}^{\text{MT}}(t_1:t_N)|s, \mathbf{c}) = \prod_{n=1}^N p(\mathbf{r}^{\text{MT}}(t_n)|s, c(t_n)).$$

Given a series of activity patterns $\mathbf{r}^{\text{MT}}(t_1:t_N)$ and assuming that one knows \mathbf{c} , the optimal strategy for inferring the direction of motion is to apply Bayes' rule to compute a probability distribution over s , given $\mathbf{r}^{\text{MT}}(t_1:t_N)$. If the prior on s is flat, this so-called posterior distribution is given by

$$p(s|\mathbf{r}^{\text{MT}}(t_1:t_N), \mathbf{c}) \propto \prod_{n=1}^N \frac{p(\mathbf{r}^{\text{MT}}(t_n)|s, \mathbf{c})}{p(\mathbf{r}^{\text{MT}}(t_n)|\mathbf{c})}. \quad (1)$$

This distribution captures everything there is to know about s given all the data from MT since the beginning of the trial, and as such, it retains all the information in the MT activity. Therefore, if the brain uses a Bayesian approach to decision making, the goal of the accumulation layer (LIP) should be to generate a pattern of activity at time t_n that encodes this distribution (Equation 1). An even better solution would be to encode a posterior distribution $p(s|\mathbf{r}^{\text{MT}}(t_1:t_N))$ that does not depend on \mathbf{c} , the nuisance parameters (or in the jargon of probabilistic inference, a posterior in which \mathbf{c} has been marginalized out; $p(s|\mathbf{r}^{\text{MT}}) = \int d\mathbf{c} p(s|\mathbf{r}^{\text{MT}}, \mathbf{c}) p(\mathbf{c}|\mathbf{r}^{\text{MT}})$). This would allow downstream areas to perform optimal computations over LIP activity without having to estimate the nuisance parameters. In other words, we should seek a set of feedforward connections between MT and LIP, and lateral connections within LIP, such that

$$p(s|\mathbf{r}^{\text{LIP}}(t_N)) = p(s|\mathbf{r}^{\text{MT}}(t_1:t_N)). \quad (2)$$

It is critical to note that the approach we have just outlined requires that neural responses in MT and LIP represent probability distributions. In MT, $\mathbf{r}^{\text{MT}}(t_n)$ represents $p(s|\mathbf{r}^{\text{MT}}(t_n))$, which is obtained from the response distribution, $p(\mathbf{r}^{\text{MT}}(t_n)|s)$ (sometimes called the noise distribution) through Bayes' rule: $p(s|\mathbf{r}^{\text{MT}}(t_n)) \propto p(\mathbf{r}^{\text{MT}}(t_n)|s)$, (we are assuming a flat prior over s for this encoding step; nonflat priors can be incorporated into our approach but are not central to the current argument). The same idea also applies to LIP. We refer to populations that represent probability distributions in this way as probabilistic population codes (Ma et al., 2006). The existence of such codes is central to our approach: neurons represent probability distributions via Bayes' rule and, as a result, neural computations, such as accumulation of evidence, can be optimized by tailoring neural operations to the encoded distributions.

Once the accumulation is stopped, an action must be selected. The optimal strategy under many reasonable cost functions is to choose the action corresponding to the most likely stimulus. This value, denoted \hat{s} , is given by

$$\hat{s} = \underset{s}{\text{argmax}} p(s|\mathbf{r}^{\text{MT}}(t_1:t_{\text{stop}})), \quad (3)$$

where t_{stop} is the stopping time. Note that, for simplicity, t_{stop} use the same variable s to refer to both the direction of motion and the direction of a saccade, since they are indistinguishable in this experiment.

In a minimal optimal network, the third layer should encode the estimate, \hat{s} , as a stereotyped motor command. This is a task for which attractor networks are ideally suited, because they can take a noisy hill of activity as input and produce a smooth hill of stereotyped shape and height as output (Zhang, 1996) (see top layer in Figure 1C). Stereotyped hills like these are observed, for instance, in the motor layer of the superior colliculus (SC), where the position of the peak of a hill determines the direction and amplitude of the upcoming saccade (Lee et al., 1988). The fact that the hill is smooth in Figure 1C might appear unrealistic, but see the Supplemental Data available online for why this is in fact not a significant concern.

The question we address in the rest of the paper is how to implement optimal accumulation (Equation 2) and optimal response selection (Equation 3) in neural hardware.

Optimal Network

Not surprisingly, the network connectivity needed to achieve optimality depends strongly on how the information about the stimulus is represented in MT, which in turn depends on the structure of the neuronal variability. Here, we assume that the variability in MT conditioned on the value of a stimulus belongs to the exponential family with linear sufficient statistics (Ma et al., 2006). This choice is a natural one, since it is consistent with experimental measurements in a wide range of cortical areas (Ma et al., 2006). Specifically, we assume that $p(\mathbf{r}^{\text{MT}}(t_n)|s, c(t_n))$ has the form

$$p(\mathbf{r}^{\text{MT}}(t_n)|s, c(t_n)) = \Phi(\mathbf{r}^{\text{MT}}(t_n), c(t_n)) \exp(\mathbf{h}(s) \cdot \mathbf{r}^{\text{MT}}(t_n)) \quad (4)$$

where $\Phi(\mathbf{r}^{\text{MT}}(t_n), c(t_n))$ is an arbitrary function of $\mathbf{r}^{\text{MT}}(t_n)$ and $c(t_n)$, and “ \cdot ” is the standard dot product: $\mathbf{h}(s) \cdot \mathbf{r}^{\text{MT}}(t_n) = \sum_i h_i(s) r_i^{\text{MT}}(t_n)$. Note that the nuisance parameter, $c(t_n)$, does not appear in the kernel \mathbf{h} . In the rest of the paper, we refer to distributions with the property that \mathbf{h} depends only on s as Poisson-like.

Independent Poisson variability is a special case of the Poisson-like family, with $h_i(s)$ being the log of the tuning curve of neuron i . Importantly, correlated neuronal responses (as observed in the brain) are also in the Poisson-like family, although there are restrictions on the nuisance parameters (Ma et al., 2006). These restrictions arise because $\mathbf{h}(s)$ is not independent of the tuning curves and the covariance matrix, but is related via

$$\mathbf{h}'(s) = \Sigma^{-1}(s, c(t_n)) \mathbf{f}'(s, c(t_n)) \quad (5)$$

where $\mathbf{f}(s, c(t_n))$ is the tuning curve (the mean of \mathbf{r} as a function of s), a prime denotes a derivative with respect to s , and $\Sigma(s, c(t_n))$ is the covariance matrix of \mathbf{r} . Since the right-hand side of Equation 5 depends on $c(t_n)$ and the left-hand side does not, satisfying this equation is not trivial. There is, however, a rather natural condition under which it is satisfied: $c(t_n)$ is contrast (Sclar and Freeman, 1982). This is because contrast has a multiplicative effect on both tuning curves (Anderson et al., 2000; Sclar and Freeman, 1982) and covariance (Gershon et al., 1998; Kohn and Smith, 2005; Tolhurst et al., 1983), so $\mathbf{f}'(s, c(t_n))$ is proportional to some monotonic function $g(c(t_n))$, $\Sigma^{-1}(s, c(t_n))$ is proportional to $1/g(c(t_n))$, and thus the $c(t_n)$ dependence disappears from the right-hand side. (This is also the case

when $c(t_n)$ is stimulus intensity for tactile stimuli.) Whether or not Equation 5 is satisfied for other nuisance parameters must be checked on a case-by-case basis.

If the activity in MT satisfies Equation 4, then we can insert Equation 4 into Equation 1, and we see that the right-hand side is independent of \mathbf{c} . Thus, the probability of the stimulus given the entire history of MT activity is given by

$$p(s|\mathbf{r}^{\text{MT}}(t_1 : t_N)) \propto \exp\left(\mathbf{h}(s) \cdot \sum_{n=1}^N \mathbf{r}^{\text{MT}}(t_n)\right) \quad (6)$$

where the constant of proportionality depends on MT activity but not on s . Consequently, when the prior is flat, we see that Equation 2 is satisfied if the LIP activity is constructed by simply adding the MT activity,

$$\mathbf{r}^{\text{LIP}}(t_N) = \sum_{n=1}^N \mathbf{r}^{\text{MT}}(t_n). \quad (7)$$

The problem with this simple summing operation, however, is that LIP activity would saturate very quickly. Fortunately, it is possible to show that global inhibition can be used to alleviate this problem while preserving optimality (see [Supplemental Note](#)).

Finally, to guarantee that the SCb activity peaks at the optimal location (at \hat{s} in Equation 3 and in [Figure 1C](#)), we must introduce recurrent connectivity so that the SCb layer can support a hill of activity without input. In addition, input to the SCb must be gated, so that it receives no input until decision time. Once a decision is made, the instantaneous activity in the LIP layer is used to initialize the SCb activity. After initialization, the LIP activity is removed, and the SCb layer evolves under its own dynamics. As shown in the [Supplemental Note](#), if the neuronal variability is Poisson-like in LIP, the SCb layer peaks at \hat{s} when

$$\mathbf{v}^{\text{SC}}(s) \propto \mathbf{h}'(s) \quad (8)$$

where $\mathbf{v}^{\text{SC}}(s)$ is the left null eigenvector of the Jacobian evaluated at the attractor implemented by the SCb and $\mathbf{h}'(s)$ is the same function that appears in Equation 5. Importantly, $\mathbf{v}^{\text{SC}}(s)$ can be tuned to satisfy Equation 8 by adjusting network parameters such as the weights of lateral connections in SCb. Consequently, when the neuronal variability is Poisson-like in LIP, there exists a set of parameters for which the superior colliculus generates a maximum-likelihood estimate.

Note that if the variability in LIP is not Poisson-like, attractor dynamics is no longer guaranteed to be optimal. In fact, there is no known optimal network for most distributions. It is therefore quite remarkable that, of all distributions, the cortex appears to exhibit those for which attractor dynamics can be tuned to be optimal.

Implications of Optimality

There are several important features of our network that are somewhat hidden by the above analysis. First, if the neurons are Poisson-like, Equation 7 leads to optimal accumulation of evidence (i.e., Equation 2 is satisfied) even when the reliability of the sensory information varies from trial to trial or over the course of a single trial. This might sound counterintuitive at first. Consider, for example, an image whose contrast increases over

time. Since the data become progressively more reliable, the decision should be based more strongly on the information acquired at later times. A way to implement this would be to boost the weights from MT onto LIP as the contrast increases. However, this reweighting would have to be done on very short timescales and would require a constantly updated and reliable estimate of contrast. With Poisson-like variability, there is no need to reweight the input over time, because both MT and LIP represent probability distributions at all times and in a manner which is invariant to the value of contrast. This is easy to see in the case of contrast: as contrast increases, the reliability of the sensory evidence increases, but so does the amplitude of the population activity in MT. Since the MT activity is added on top of LIP activity, its impact scales with its amplitude, and therefore, in proportion to its reliability.

A second feature of our network is that the reliability of the data (encoded in the nuisance parameters, \mathbf{c}) plays no role in estimating the stimulus, in the sense that even if we knew \mathbf{c} our estimate of the posterior over the stimulus would not improve. This is a strong result, and one that is highly unusual in Bayesian inference. Much more typical is that the nuisance parameters are either estimated, or integrated out of the posterior, both of which introduce additional uncertainty in the inference process. The ramification of this is that, assuming no loss of information in processing after the SCb layer, the posterior in LIP exactly reflects the behavioral performance. For instance, if the posterior in LIP on a given trial is Gaussian with a standard deviation of 10° at decision time, and the decision involves computing the maximum-likelihood estimate, the discrimination threshold of the animal should be around 10° as well, across multiple trials of the same type. As we will discuss later, this prediction can be tested with existing data. Importantly, this prediction does not apply to the SCb layer in our model: instead, variability in this region, estimated on a trial-by-trial basis, would encode the motor error for the saccadic eye movement.

Evidence Accumulation: Simulation Results

We have shown so far that if the responses of MT neurons have Poisson-like statistics, optimal evidence accumulation can be performed by adding spikes over time, and optimal action selection can be performed with a single attractor network. Importantly, the attractor network can extract the maximum-likelihood estimate of the stimulus, s , without any need to know either the nuisance parameters, \mathbf{c} , or how much time has elapsed since the start of the trial (M.N. Shadlen et al., 2006, *Soc. Neurosci.*, abstract).

These results are important but they are based on assumptions that are not necessarily exactly true in vivo. For instance, real neurons do not simply add spikes over time. Moreover, the response of MT neurons to random-dot kinematograms may not be exactly Poisson-like (in fact, it is not exactly Poisson-like according to the current models of MT; see [Experimental Procedures](#)). It is therefore essential that we test our theory in biologically realistic networks. In particular, we want to address two critical questions in the simulated network: (1) Does the LIP layer accumulate evidence optimally? (2) Can a single attractor network extract the maximum-likelihood estimate from LIP activity, for all coherences and at all times?

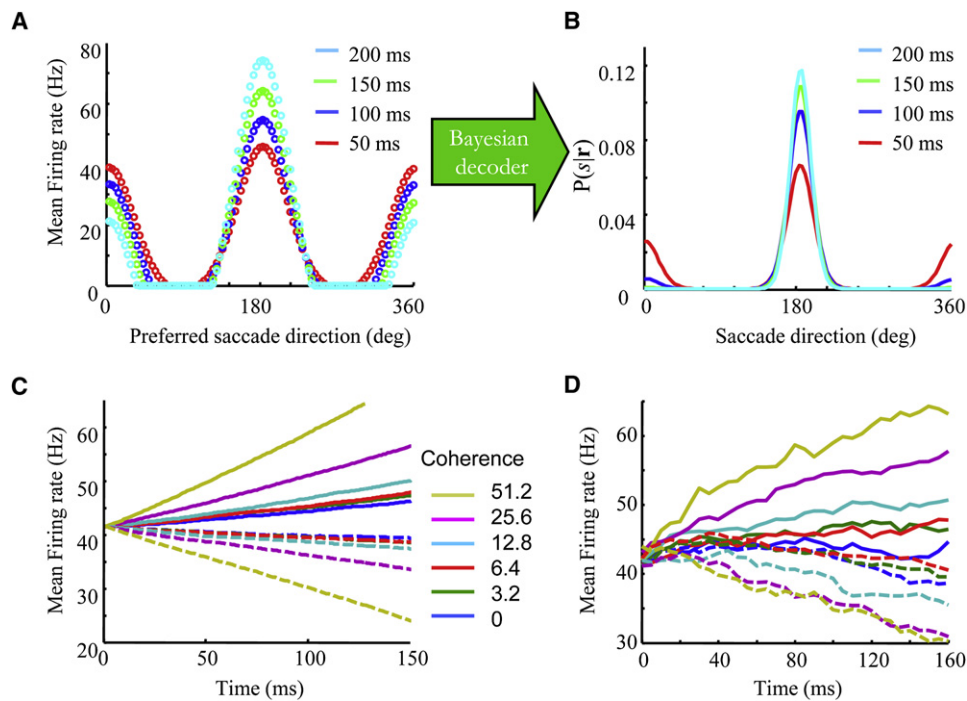


Figure 2. Binary Decision Making (as Illustrated in Figure 1A)

(A–C) model; (D) data.

(A) Firing rate in LIP at four different times for a coherence of 51.2%. The direction of the moving dots is 180°.

(B) Probability distributions encoded by the firing rates shown in a averaged over 1000 trials. As expected, the probability of the 180° direction goes up while the probability of the 0° direction goes down.

(C) Firing rate over time for two units tuned to 180° (solid line) and 0° (dotted lines) for six different levels of coherence. These averages were obtained over trials for which the model's choice was 180°.

(D) Same as in (B) but for actual neurons in LIP (n = 45). Data from [Roitman and Shadlen \(2002\)](#). The model and the data show similar trends.

For these simulations, we use a network similar to the one depicted in [Figure 1C](#). For the LIP layer, we use linear-nonlinear-Poisson (LNP) neurons ([Plesser and Gerstner, 2000](#)) with a long time constant (1 s) ([Renart et al., 2003](#)) (see [Experimental Procedures](#) for details). We use LNP neurons because they provide a good approximation to real neurons, while producing spikes with realistic count statistics close to the exponential family ([Paninski, 2004](#); [Plesser and Gerstner, 2000](#)). The LIP layer receives spatially correlated spike trains from area MT ([Britten et al., 1993](#); [Zohary et al., 1994](#)). The feedforward connections from MT to LIP, which are purely excitatory, connect neurons with similar direction preferences using a Gaussian weighting profile. The LIP layer also has lateral connections with short-range excitation and long-range inhibition; i.e., the weights are excitatory between neurons with similar preferred directions and inhibitory otherwise. The inhibition is used to prevent saturation. Finally, rather than constructing the SCb network, we make use of the fact that the line attractor implements a local linear estimator which can be tuned to be optimal ([Deneve et al., 1999](#); [Latham et al., 2003](#); see [Supplemental Note](#)).

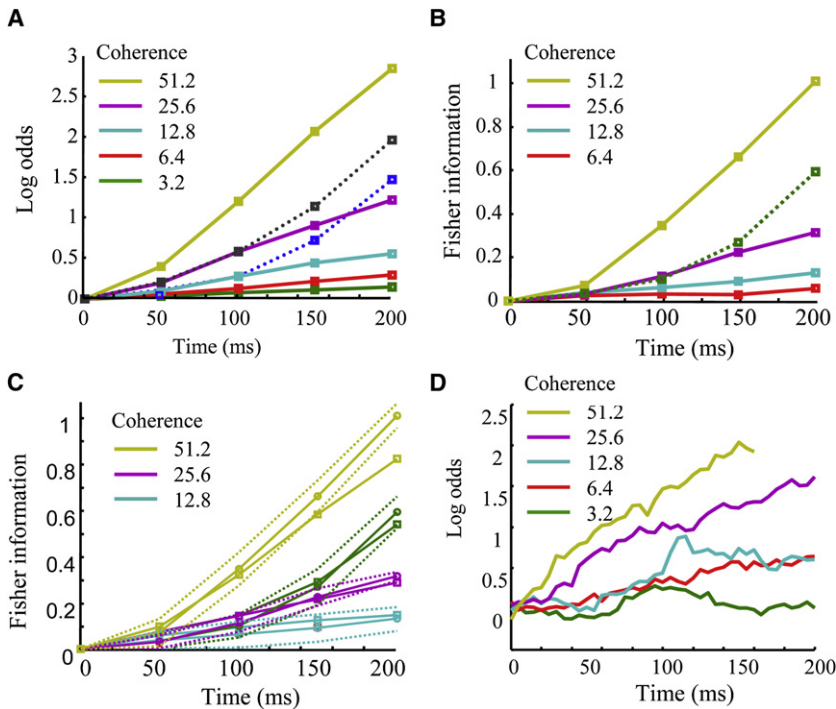
We focus first on the binary case, for which dots move at either 0° or 180°. [Figure 2A](#) shows the average activity in the LIP layer over time for a stimulus moving at 180°. The average posterior distribution encoded by these activity patterns is illustrated in [Figure 2B](#). As expected, the probability corresponding to 180° grows

over time while the probability corresponding to 0° decreases. In [Figure 2C](#), we show firing rate versus time for all coherences and for the neurons optimally tuned to 180° and 0°. The neurons in the model behave quantitatively like actual LIP neurons, as can be seen in [Figure 2D](#) (data from [Roitman and Shadlen \(2002\)](#)).

To determine whether the LIP layer accumulates evidence optimally, we first consider experiments in which coherence is fixed within a trial. Here, we take “optimal” to mean that when LIP updates its estimate of the direction of motion of the moving dots, it takes into account both its own uncertainty about direction and the uncertainty in MT. From a quantitative point of view, this implies that the expected log odds of making a correct choice ($\log[p_{\text{correct}}/(1-p_{\text{correct}})]$ where p_{correct} is the probability of making a correct choice) grows linearly with time, because the evidence is provided at a constant rate (see [Experimental Procedures](#)). Moreover, the slope should increase with coherence, and if the coherence changes during the trial, so should the slope.

[Figure 3A](#) shows that the log odds do indeed grow linearly with time, and the larger the coherence, the faster it increases. Furthermore, if we double or quadruple the coherence at time $t = 100$ ms, the slope of the log odds changes to the correct slope within 100 ms ([Figure 3A](#), dotted lines).

We repeated these simulations for the continuous case, where the stimulus can move in any direction. [Figure 4A](#) shows the time evolution of the firing rate in LIP and [Figure 4B](#) shows the average



information shows that decoding LIP can be done nearly optimally without any knowledge of time or coherence. 25.6% and then switches to 51.2% at 100 ms.

(D) Same as in (A) but for actual LIP neurons ($n = 45$; data from Roitman and Shadlen, 2002). The results are quantitatively similar to the model. The y axis is arbitrary up to a multiplicative factor and a DC offset.

posterior distributions encoded by this activity. As evidence accumulates in favor of 180° , the activity at 180° increases and the probability distribution becomes narrower. To determine whether this accumulation process is optimal, we can run the same test as in the binary case, except this time we use the average of the inverse of the variance of the posterior distributions (the Fisher information [Papoulis, 1991]) rather than the log odds. Figure 3B shows that, indeed, Fisher information increases linearly with time and the slope is an increasing function of coherence. Furthermore, when the coherence increases during the trial, so does the slope.

We now turn to the second question: can the maximum-likelihood estimate be computed from LIP activity, for all coherences and at all times, with a single attractor network? Because attractor networks are mathematically equivalent to local linear estimators (Deneve et al., 1999), this question can be rephrased as: is the performance of a single local linear estimator similar to the performances of a family of estimators, each specialized for one time and one coherence? Figure 3C shows that the Fisher information recovered by the specialized linear estimators is indeed very similar to the information recovered by a single one, hence demonstrating that a single attractor network can optimally decode LIP for all coherences and at all times.

Finally, we performed another test, now at decision time. With our framework, the network encodes a probability distribution at all times and in particular at decision time. This distribution reflects the quality of the data that have been accumulated and, consequently, the performance of the animal. Hence, for

Figure 3. Log Odds and Fisher Information as a Function of Time

The origin ($t = 0$) on all plots corresponds to the start of the integration of evidence, which about 220 ms after stimulus onset in the experimental data. (A–C) model; (D) data.

(A) Log odds for a binary decision as a function of time for four different levels of coherence (solid lines). Blue and black dotted lines: the coherence increases to 51.2% at $t = 100$ ms. After 100 ms, the slope matches the 51.2% coherence trials, as expected if the model is Bayes optimal.

(B) Fisher information as a function of time for continuous decision making (as in Figure 1B). Fisher information rises linearly with time, with higher slopes for higher coherences, as expected for Bayesian optimality. Dotted line: trial in which the coherence increases from 25.6% to 51.2%. In both (A) and (B), the kink at $t = 50$ ms is due to the discretization of time. (C) Squares: Fisher information estimated by a single local linear estimator across all times and all coherences. Circles: Fisher information estimated by a local optimal estimator trained separately for each time and each coherence. Dotted lines: for each coherence, the upper line corresponds to the information estimated from the training set, while the lower trace is the information obtained from the testing set. The solid line is the average of the upper and lower dotted lines. The fact that both estimators return similar values of Fisher

information shows that decoding LIP can be done nearly optimally without any knowledge of time or coherence. Green line: trials in which the coherence starts at

both two- and four-choice experiments, the log odds estimated in the LIP layer should be higher at high coherence than at low coherence, since the performance of the animal is better in the former case. Figures 5A and 5B show that our model behaves as predicted. Note the important distinction with single-race bounded accumulation models (Bogacz et al., 2006; Huk and Shadlen, 2005; Link, 1992; Link and Heath, 1975; Ratcliff and Rouder, 1998). In such models, the state of the system is characterized by the value of the accumulation process. When the bound is hit, this value is always the same (Gold and Shadlen, 2001; Link, 1992; Shadlen et al., 2006a). Thus, there is no principled way to recover the probability that the decision is correct. An ad hoc solution has been proposed for two race models (Vickers, 1979), but it was not derived from probabilistic principles, and does not readily generalize to more than two choices.

Speed-Accuracy Tradeoff

When monkeys are tested on our decision making task in which they are free to choose when to respond, their psychometric and chronometric functions follow the profiles shown in Figures 6A and 6B. To obtain these curves with our model, we used a stopping rule similar to the one used in most models: a fixed bound on the maximum activity in the network (see Experimental Procedures). As can be seen, our model readily captures the performance and reaction time reported in monkeys whether the task involves two or four choices (data from Churchland et al., 2008). Moreover, the rate at which activity builds up on average

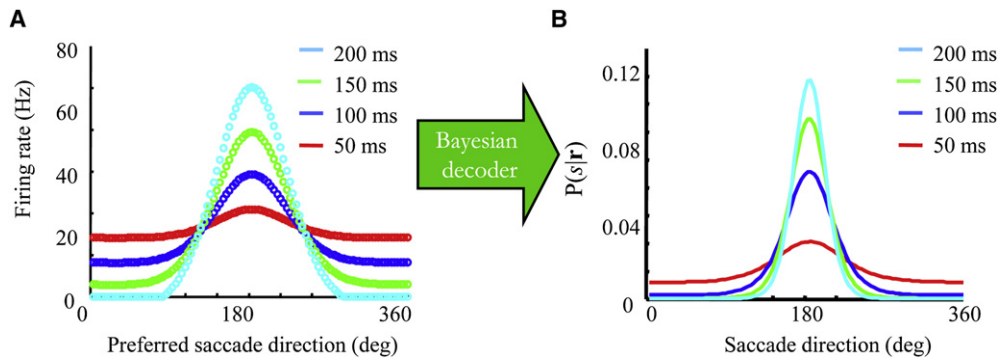


Figure 4. Continuous Decision Making (as Illustrated in Figure 1B)

(A) Firing rates of model neurons in LIP at four different times for a coherence of 51.2%. The direction of the moving dots is 180°.

(B) Probability distributions encoded by the firing rates shown in (A) averaged over 1000 trials. As expected, the peak of the distribution is close to 180° and the variance of the distribution decreases over time.

in the LIP layer of the model as a function of coherence and number of choices is similar to what has been reported in vivo (see Figure 7).

Experimental Predictions

Our model makes two experimentally testable predictions. The first is that the population response in LIP encodes a probability distribution over the stimulus and, more importantly, that this distribution reflects both the reliability of the evidence and the performance of the animal. Therefore, we predict that if the population activity in LIP is decoded with the same method used in our simulations, the results will match those shown in Figures 3A and 3B and Figure 5A. Rigorously testing this prediction requires multiunit recordings in LIP, which are not currently available. However, we can test it qualitatively with the spike trains

obtained from single cell recordings (Roitman and Shadlen, 2002). If the spikes trains in LIP reflect the quality of the sensory data, the expected log odds computed from these spike trains should grow linearly with time, and the rate of growth should be proportional to coherence. We have performed this analysis, and this is indeed what we found, as illustrated in Figure 3D. Furthermore, if these odds reflect the performance of the animal, we should find that the log odds in LIP at decision time grows with coherence for both the two- and four-choice experiment (since performance improves with coherence). Again, this is what we observed (Figure 5B).

Recent experiments suggest that a similar property may hold for buildup cells in the superior colliculus (Kim and Basso, 2008; Ratcliff et al., 2006). For instance, Kim and Basso (2008) have recorded simultaneously from neurons responding to the selected

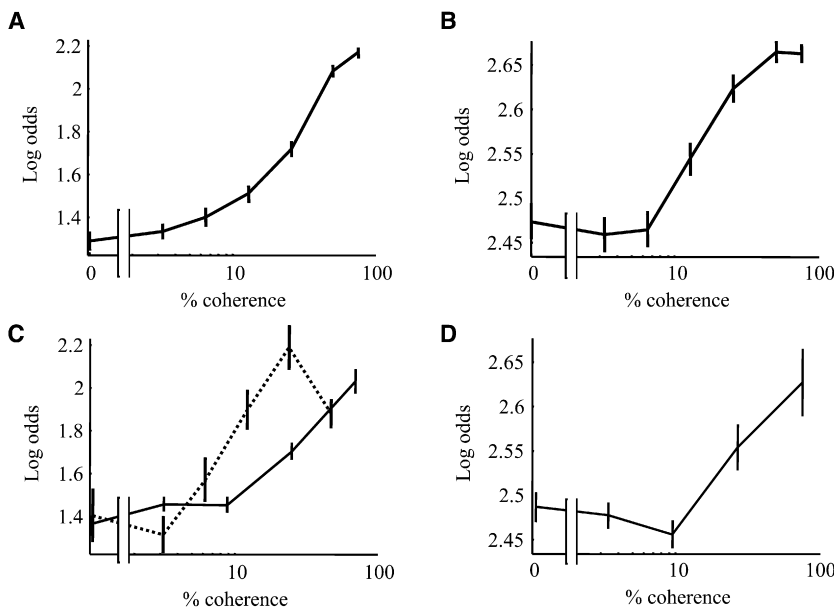


Figure 5. Average Log Odds at Decision Time Computed from the Model and Data for the Two- and Four-Choice Experiments

(A) Average log odds at decision time for a two-choice experiment estimated from two neurons in the LIP layer of the model tuned to 0° and 180° on trials for which the model selected 180°. The average log odds is defined as the log of the ratio of the probability that the direction is equal to 180° to the probability that it is equal to 0° averaged over trials.

(B) Same as in (A) but for the four-choice experiment (for consistency with the two-choice experiment, we use log odds in the four-choice experiment).

(C) Same as in (A) but for actual LIP neurons (n = 45) in the two-choice experiment (dotted line, data from Roitman and Shadlen, 2002; solid line, data from Churchland et al., 2008).

(D) Same as in (B) but for actual LIP neurons (n = 51–70) in the four-choice experiment (data from Churchland et al., 2008). In both (C) and (D), the log odds increases with coherence. Since higher coherence also implies higher performance, log odds also increases with performance. This is indeed what is expected if the posterior encoded in LIP reflects the quality of the data and, at decision time, the performance of the animal. On these plots, the y axis is arbitrary up to a multiplicative factor.

The method used to obtain the error bars is described in the Supplemental Data.

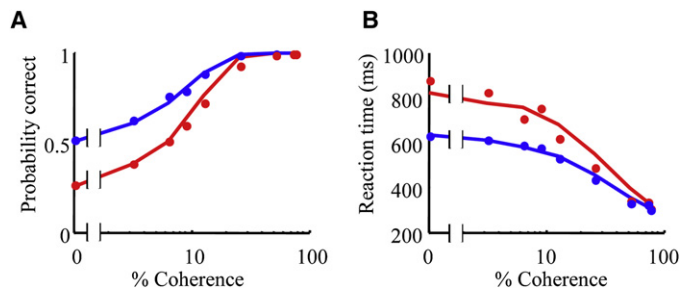


Figure 6. Performance and Reaction Time for the Model versus Monkeys

(A) Probability of correct responses as a function of coherence. Blue: two-choice experiment. Red: four-choice experiment. Solid lines: model. Closed circles: data from Churchland et al. (2008).

(B) Reaction time as a function of coherence. Legend as in (A).

target and neurons responding to the distractors in a four-choice experiment. They reported that the difference in activity between these neurons increases with performance. Under the assumption of Poisson-like neural variability, this difference would lead to an increase in the posterior probability assigned to the target as a function of performance.

As additional multiunit data become available, it will be interesting to test our predictions more quantitatively. In particular, it will be important to determine whether the population of firing rates representing evidence for competing directions affects confidence judgments. It will also be important to determine whether a decoder that has knowledge of time and coherence performs better than a decoder that does not have such knowledge, and whether or not this additional information is accessible to the animal. As shown in Figure 3C, our model predicts that there should be little difference.

The second experimental prediction concerns the time evolution of the population activity in LIP. As can be seen in Figure 4A, the width of the population activity does not change over time (once the curves are normalized for height and the baseline is removed), in contrast to the decoded probability distribution, which gets narrower as time progresses (Figure 4B). This prediction is slightly weaker, since a population code with an invariant width is a sufficient but not a necessary condition for our proposed model. Nonetheless, the finding that the width of the population activity is invariant over time would be consistent with our model, while ruling out codes in which neural activities are proportional to probability (Barber et al., 2003; Eliasmith and Anderson, 2003).

DISCUSSION

We have shown that when the variability in spike count is Poisson-like, integration of evidence and action selection can be performed optimally in networks of spiking neurons, even when the variables involved are continuous and the reliability of the data changes over time. This result might explain why spike counts appear to follow Poisson-like distributions throughout most of the cortex: this particular format greatly simplifies optimal Bayesian inference for decision making.

We have also shown that performance is near-optimal even when the distribution of spike counts is not exactly Poisson-like (see Figure 3C) but instead follows the experimentally observed distribution in MT in response to random-dot kinematograms. It would be interesting to explore how far one has to be from the Poisson-like family before there is significant departure

from optimal Bayesian inference. If a stimulus could trigger such non-Poisson statistics in the brain, we could test whether subjects' performance degrades as predicted by our model.

At first glance, it might appear that our model does not differ much from previous neural models of decision making (Machens et al., 2005; Mazurek et al., 2003; Ratcliff and Rouder, 1998; Reddi and Carpenter, 2000; Smith and Ratcliff, 2004; Usher and McClelland, 2001; Wang, 2002; Wong and Wang, 2006). Previous neural models have indeed shown that a neural integrator can capture the behavior of subjects in a binary decision task, as can a point attractor network. They have even provided a probabilistic interpretation of the neural integration in terms of accumulation of log odds. It is important to emphasize, however, that these models, and their probabilistic interpretations, apply under very restrictive conditions and do not generalize to real-world problems. In particular, in the context of estimating motion direction, they cannot handle decision making over continuous choices, or time- or trial-varying coherence. For example, the notion that LIP neurons are effectively accumulating log odds when they integrate the difference in activity of MT neurons with opposite preferences is true only for binary decisions and fixed coherence (Gold and Shadlen, 2001). This notion does not generalize easily to multiple directions (Bogacz and Gurney, 2007; McMillen and Holmes, 2006) and does not generalize at all to time- and trial-varying coherence.

The general case requires that we deal with the difficult problem of hidden variables: how do we extract information about a variable (e.g., direction) from neural activity which is influenced by other, hidden variables (e.g., coherence) whose value is unknown and varies over time? This is one of the hardest problems faced by the brain, and no general solution has been provided in the context of decision making. Here, however, we have found a solution that can be implemented with biologically plausible mechanisms. Moreover, this solution led to a strong prediction which is that the log odds (or the posterior distribution in the case of multiple or continuous choice) are available on a trial-by-trial basis in LIP at all times and in particular at decision time (without any knowledge of coherence or time). As shown in Figure 5, the responses of LIP neurons in vivo are consistent with this prediction.

Our probabilistic framework also helps to clarify the benefits and limitations of using point attractor (Machens et al., 2005; Wang, 2002; Wong and Wang, 2006) or line attractor dynamics (Furman and Wang, 2008) for accumulation of evidence in decision making. Line attractor dynamics are a good way to perform

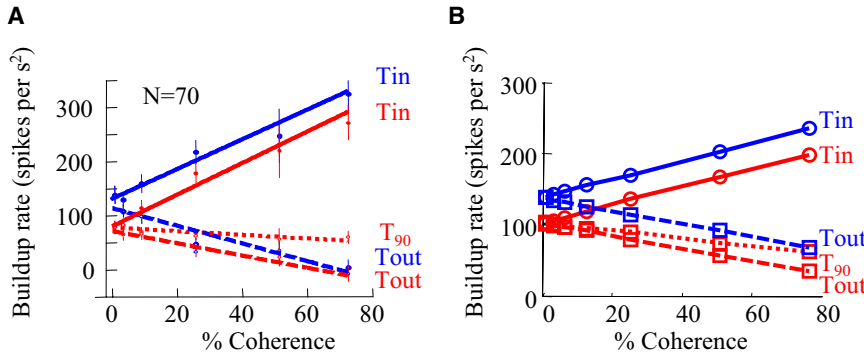


Figure 7. Buildup Rate of LIP Neurons

T_{in} : neurons tuned to the direction of the stimulus. T_{out} : neurons tuned to a direction 180° away from the stimulus direction. T_{90} : neuron tuned to a direction 90° away from the stimulus direction. Blue: two-choice experiment. Red: four-choice experiment. (A) LIP data from Churchland et al., 2008. (B) Model.

optimal action selection (as we do in the SCb layer), but not optimal evidence integration (which is why we do not use it in the LIP layer). Moreover, attractor dynamics can provide an optimal solution for action selection, but, importantly, only for a limited family of distributions, one of which is Poisson-like. This is a critical point, as it emphasizes the strong link between the response distribution and optimal inference.

Our framework is sufficiently powerful that it can be extended in several directions, including incorporating prior information, dealing with time-varying stimuli, and taking into account nontrivial reward functions when selecting actions. This last extension is critical. We have shown how the evidence accumulation and the response selection can be optimized in neural circuits, but we have not shown how to optimize reward rates. Optimizing reward rate is a complex problem that depends crucially on the cost function and the stopping process (Kiani et al., 2008). This lies beyond the scope of the present paper, but it is an important issue, which we intend to explore in future studies. It remains to be seen if it can be incorporated in the PPC framework. We believe that a promising idea is to explore whether LIP encodes the expected reward as a function of saccade direction and amplitude. Recent experimental data suggest that LIP might indeed represent either expected reward for all actions or the probability that an action will maximize reward (Platt and Glimcher, 1999; Sugrue et al., 2004). Either way, our framework should be applicable, since these quantities are similar to probability distributions over stimulus values.

EXPERIMENTAL PROCEDURES

Network Simulations with LNP Neurons

The MT layer contained 100 stochastic spiking neurons with bell-shaped tuning curves to direction of motion. At each time step, the probability of a spike in neuron i was determined according to

$$p(r_i^{MT}(t_n) = 1) = [\delta t(c\Delta d \exp(K_{MT}(\cos(s_0 - s_i) - 1)) + d_{null}c + r_{spont} + n_i)]^+$$

where $r_i^{MT}(t_n)$ is the response of neuron i within the interval $[t_n - \delta t, t_n]$, s_0 is the direction of motion of the random dots, s_i is the preferred direction of neuron i , c is the percentage of dots moving in direction s_0 (the coherence level), Δd is the difference in drive between the preferred and null directions ($d_{pref} - d_{null}$), d_{null} is the drive in the null direction, r_{spont} is the spontaneous firing rate, n_i is a random variable used to induce correlations, and $[\cdot]^+$ is the threshold-linear operator: $[x]^+ = \max(0, x)$. The parameters of the model were: $d_{pref} = 0.4$, $d_{null} = -0.2$, $r_{spont} = 20$, $K_{MT} = 4$ (as reported in MT [Britten et al., 1993]), $\delta t = 1$ ms. Note that in the equation above, the coherence c cannot be factored out of the equation. As a result, the spike statistics in MT are not exactly in the Pois-

son-like family as defined in Equation 4, because the kernel $h(\cdot)$ depends on both s and c .

The noise n_i consisted of independent Gaussian noise convolved with a circular Gaussian kernel,

$$n_i = \sum_j A_{\eta_j} \exp(K_{\eta_j}(\cos(s_i - s_j) - 1))\eta_j,$$

with all η_j drawn independently from a zero-mean Gaussian distribution and K_{η_j} and A_{η_j} set to 2 and 10^{-5} , respectively. These were chosen so that the average correlation coefficients in spike trains were approximately 0.2 between neurons whose preferred directions differed by less than 90° , and approximately 0 for neurons whose preferred directions differed by more than 90° . These values are close to the ones that have been reported in vivo (Zohary et al., 1994).

In the LIP layers, we used 100 Linear-Nonlinear-Poisson (LNP) neurons. In the linear step, the membrane potential proxy of neuron i , denoted $m_i(t_n)$, is obtained from

$$m_i(t_{n+1}) = \left(1 - \frac{\delta t}{\tau}\right)m_i(t_n) + \frac{1}{\tau} \left(\sum_j W_{ij}^{ff} r_j^{MT}(t_n) + \sum_j W_{ij}^{rec} r_j^{LIP}(t_n) \right) + u(t_n) \quad (9)$$

where W_{ij}^{ff} and W_{ij}^{rec} are the matrices for the feedforward and recurrent weights, respectively, and $u(t)$ is an urgency signal (see below). The time constant, τ , was set to 1 s. The linear step is followed by a nonlinear one in which the membrane potential proxy, $m_i(t_n)$, is used to determine the probability that neuron i emits a spike between times t_n and $t_n + \delta t$,

$$p(r_i^{LIP}(t_n) = 1) = [m_i(t_n)]^+ \quad (10)$$

We used translation-invariant weights for both the feedforward and lateral connections (W_{ij}^{ff} and W_{ij}^{rec}),

$$W_{ij} = W(s_i - s_j) = a \exp(K(\cos(s_i - s_j) - 1)) + b.$$

For the feedforward weights, W_{ij}^{ff} , we used $a = 0.25$, $K = 5$ and $b = 0$, and for the lateral weights, W_{ij}^{rec} , we used $a = 0.35$, $K = 10$, and $b = -0.11$.

In the experiment of Shadlen and Newsome (2001), each trial starts with the appearance of M visual targets, where M is the number of choices. This triggers a response in the subset of LIP neurons whose receptive fields overlap with the visual targets. To model this activity, we initialize the firing rate of the neurons in the LIP layer according to

$$p(r_i^{LIP}(t_1)) = \frac{2}{(M/2 + 0.5)} \sum_{m=1}^M p_0 \exp(K_0(\cos(s_m - s_i) - 1))$$

where M is the number of possible directions for the moving dots and $\{s_1, \dots, s_M\}$ are the positions of the targets corresponding to the M choices. The parameters were set to $p_0 = 0.042$ (corresponding to a firing rate of 42 spikes/s for $\delta t = 1$ msec) and $K_0 = 4$. Given the width of these tuning curves (determined by the parameter K_0), the resulting population activity is almost perfectly flat for $M \geq 8$. Accordingly, we used $M = 8$ when simulating ‘‘continuous’’ decision making. This visually induced activity served as a starting point for the accumulation of evidence.

The urgency signal (denoted $u(t)$) in Equation 9 was estimated directly from neural responses as described in Churchland et al. (2008) and was parameterized by a hyperbolic function

$$u(t) = u_{\infty} \frac{t}{t + \tau_{1/2}}$$

In the simulations of the two-choice experiment, we used $u_{\infty} = 34.7$ spike s^{-1} and $\tau_{1/2} = 133.3$ ms, and in the four-choice experiment, we used $u_{\infty} = 39$ spike s^{-1} and $\tau_{1/2} = 343.2$ ms. These values were directly obtained from in vivo recordings in LIP as described in Churchland et al. (2008).

Stopping Bound and Action Selection

The psychometric and chronometric curves shown in Figure 6 were obtained by stopping the accumulation of evidence when the firing rate of any of the neurons in LIP, determined from the probability of firing (Equation 10), reached 55 spikes/s for the four-choice experiment and 66 spikes/s for the two-choice experiment. In the two-choice experiment, we determine the action selected by the network using the preferred direction of the neuron that reached the bound first. If this preferred direction was within the interval $[-90^{\circ}, 90^{\circ}]$, the network decision was set to 0° ; otherwise, it was set to 180° . A similar strategy was used for the four-choice experiment, except that we used four quadrants.

Note that the implementation of the bound (as well as saccade selection) is not based on the output spike trains of individual neurons. Indeed, this would not be a robust way to proceed: estimating the rate of a single neuron on a single trial is subject to a very large variability. A more robust approach consists of using spike counts filtered across direction and time. This is effectively what we have done here, since we are using the probability of firing. That probability, which is m_i in Equation 10, is a filtered version of the spike trains both from MT and LIP.

Decoding Probability Distributions from Population Activity

To generate Figures 2B and 4B, we need to compute the posterior, $p(s|r^{LIP}(t_n))$ where here t_n is shorthand for spike count between times t_n and $t_n - \Delta t$ (where $\Delta t = 50$ ms). For this we use Bayes' rule (assuming a flat prior), which gives us

$$p(s|r^{LIP}(t_n)) \propto p(r^{LIP}(t_n)|s). \quad (11)$$

To model the likelihood in LIP, $p(r^{LIP}(t_n)|s)$, we use a distribution that lies in the exponential family with linear sufficient statistics,

$$p(r^{LIP}(t_n)|s) = \Phi(r^{LIP}(t_n), \mathbf{c}(t_n)) \exp(\mathbf{h}^{LIP}(s) \cdot r^{LIP}(t_n)). \quad (12)$$

Note that this is an approximation: as discussed in the main text, the true distribution in LIP does not lie in this class. However, the approximation appears to be a good one, since we fail to find any significant Fisher information in LIP spike count beyond what can be recovered with a local optimal linear estimator, even when that estimator is independent of both coherence and time (see section Estimating Fisher Information).

To estimate $\mathbf{h}^{LIP}(s)$, we took advantage of the fact that $\mathbf{h}^{LIP}(s)$ must satisfy Equation 5 (Ma et al., 2006). Importantly, the right hand-side of Equation 5 is the local optimal linear estimator (LOLE) of LIP activity (Series et al., 2004). Therefore, we can approximate $\mathbf{h}^{LIP}(s)$ by estimating the LOLE of LIP activity and integrating it as a function of s . To obtain the LOLE, we ran the network for 10,000 trials at 51.2% coherence, with each trial lasting 200 ms. We divided each trial into four time windows of 50 ms each and extracted the spike count over each time window. We then trained four LOLEs over each of the four time windows (see Series et al. [2004] for details). This gave us four sets of weights, $\mathbf{W}_i^{LOLE}(s)$, with $i = \{1, 2, 3, 4\}$ referring to the 50 ms time interval. We then integrated the $\mathbf{W}_i^{LOLE}(s)$ with respect to s to obtain an estimate of the kernels, $\mathbf{h}^{LIP}(s, t_i)$, at the four time intervals. The resulting kernels were then averaged to obtain the overall kernel, $\mathbf{h}^{LIP}(s)$. This kernel was used in Equation 12 and 11 to obtain posterior distributions at all times and across all coherences, as illustrated in Figures 2B and 4B.

Estimating Fisher Information

To estimate the Fisher information, we used the kernels computed in the previous section to obtain the maximum-likelihood estimate of the stimulus on 5000 trials, and then computed the variance of those estimates. The Fisher

information is the inverse of the variance. The maximum likelihood estimates were given by

$$\hat{s} = \underset{s}{\operatorname{argmax}} (\exp(\mathbf{h}^{LIP}(s) \cdot r^{LIP}(t_n))).$$

The activity, $r^{LIP}(t_n)$, is the spike count in a 50 ms bins between times t_n and $t_n - 50$ ms. We computed the Fisher information both for $\mathbf{h}^{LIP}(s)$, the average kernel, and for $\mathbf{h}_i^{LIP}(s)$, the individual kernels (see previous section). The results are shown in Figure 3B.

We also tried a variety of nonlinear methods to estimate Fisher information (see Series et al. [2004] for details), but we found no significant information beyond what is recovered by the method described above.

Slopes of Integration In Vivo and in Simulations

Figure 7 shows the slope of integration of LIP neurons in vivo (from Churchland et al., 2008) and in the model as a function of coherence for the two-choice and four-choice experiments. The slopes of integration in the model were obtained by fitting a line in the average probability of firing $m_i(t_{n+1})$ of LIP neurons (Equation 10) over the first 50 ms of the integration period (i.e., 50 ms after the start of the response to the moving dots).

In both the model and in vivo, the slope of the integration for T_{in} (the neuron whose response field corresponds to the chosen target) increases linearly with coherence. Conversely, the slope of integration for T_{out} (the neuron whose response field corresponds to a saccade 180° away from the chosen target) decreases linearly with coherence. In addition, in the four-choice experiment, the slope of integration for T_{90} (the neuron whose response field is 90° away from the chosen target) decreases but less so than for the T_{out} neuron. Finally, for a given coherence, the slope of integration for two choices is always larger than the slope for four choices.

In the case of the model, the slope of integration for T_{90} neurons is determined by the shape of the tuning curves to saccade direction. For very narrow tuning curves, the slope of integration for T_{90} and T_{out} neurons are very similar, while for wide tuning curves, the slope of integration for T_{90} can in fact increase with coherence although always less so than for the T_{in} neurons (not shown). This is a noteworthy result because some LIP neurons show an increase in integration slope as a function of coherence (see for instance Figure 4E in Churchland et al., 2008).

SUPPLEMENTAL DATA

The Supplemental Data include two figures and a Supplemental Note and can be found with this article online at [http://www.neuron.org/supplemental/S0896-6273\(08\)00803-9](http://www.neuron.org/supplemental/S0896-6273(08)00803-9).

ACKNOWLEDGMENTS

P.E.L. is supported by the Gatsby Charitable Foundation and National Institute of Mental Health Grant R01 MH62447 and A.P. by NSF grant # BCS0446730 and MURI grant N00014-07-1-0937. M.N.S. and A.P. are jointly supported by NIDA grants #BCS0346785 and a research grant from the James S. McDonnell Foundation. We thank Daphne Bavelier for her suggestions and comments.

Accepted: September 16, 2008

Published: December 24, 2008

REFERENCES

- Anderson, J.S., Lampl, I., Gillespie, D.C., and Ferster, D. (2000). The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. *Science* 290, 1968–1972.
- Barber, M.J., Clark, J.W., and Anderson, C.H. (2003). Neural representation of probabilistic information. *Neural Comput.* 15, 1843–1864.
- Bogacz, R., and Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput.* 19, 442–477.

- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* 113, 700–765.
- Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Vis. Neurosci.* 10, 1157–1169.
- Churchland, A.K., Kiani, R., and Shadlen, M.N. (2008). Decision-making with multiple alternatives. *Nat. Neurosci.* 11, 693–702.
- Deneve, S., Latham, P., and Pouget, A. (1999). Reading population codes: A neural implementation of ideal observers. *Nat. Neurosci.* 2, 740–745.
- Ditterich, J., Mazurek, M.E., and Shadlen, M.N. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat. Neurosci.* 6, 891–898.
- Eliasmith, C., and Anderson, C.H. (2003). *Neural Engineering: Computation, Representation and Dynamics in Neurobiological Systems* (Cambridge, MA: MIT Press).
- Furman, M., and Wang, X.J. (2008). Similarity effect and optimal control of multiple-choice decision making. *Neuron* 60, this issue, 1153–1168.
- Gershon, E.D., Wiener, M.C., Latham, P.E., and Richmond, B.J. (1998). Coding strategies in monkey V1 and inferior temporal cortices. *J. Neurophysiol.* 79, 1135–1144.
- Gold, J.I., and Shadlen, M.N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* 5, 10–16.
- Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574.
- 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.
- Kiani, R., Hanks, T.D., and Shadlen, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.* 28, 3017–3029.
- Kim, B., and Basso, M.A. (2008). Saccade target selection in the superior colliculus: a signal detection theory approach. *J. Neurosci.* 28, 2991–3007.
- Kohn, A., and Smith, M.A. (2005). Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* 25, 3661–3673.
- Laming, D.R.J. (1968). *Information Theory of Choice-Reaction Times* (London: Academic Press).
- Latham, P.E., Deneve, S., and Pouget, A. (2003). Optimal computation with attractor networks. *J. Physiol. (Paris)* 97, 683–694.
- Lee, C., Rohrer, W.H., and Sparks, D.L. (1988). Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature* 332, 357–360.
- Link, S.W. (1992). *The Wave Theory of Difference and Similarity* (Hillsdale, NJ: Lawrence Erlbaum Associates).
- Link, S.W., and Heath, R.A. (1975). A sequential theory of psychological discrimination. *Psychometrika* 40, 77–105.
- Ma, W.J., Beck, J.M., Latham, P.E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438.
- Machens, C.K., Romo, R., and Brody, C.D. (2005). Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* 307, 1121–1124.
- Mazurek, M.E., Roitman, J.D., Ditterich, J., and Shadlen, M.N. (2003). A role for neural integrators in perceptual decision making. *Cereb. Cortex* 13, 1257–1269.
- McMillen, T., and Holmes, P. (2006). The dynamics of choice among multiple alternatives. *J. Math. Psychol.* 50, 30–57.
- Paninski, L. (2004). Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems* 15, 243–262.
- Papoulis, A. (1991). *Probability, Random Variables, and Stochastic Process* (New York: McGraw-Hill, inc).
- Platt, M.L., and Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238.
- Plesser, H.E., and Gerstner, W. (2000). Noise in integrate-and-fire neurons: from stochastic input to escape rates. *Neural Comput.* 12, 367–384.
- Ratcliff, R., and Rouder, J.N. (1998). Modeling response times for two-choice decisions. *Psychol. Sci.* 9, 347–356.
- Ratcliff, R., Hasegawa, Y.T., Hasegawa, R.P., Smith, P.L., and Segraves, M.A. (2006). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. *J. Neurophysiol.* 97, 1756–1774.
- Reddi, B.A., and Carpenter, R.H. (2000). The influence of urgency on decision time. *Nat. Neurosci.* 3, 827–830.
- Renart, A., Song, P., and Wang, X.J. (2003). Robust spatial working memory through homeostatic synaptic scaling in heterogeneous cortical networks. *Neuron* 38, 473–485.
- 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.
- Scar, G., and Freeman, R. (1982). Orientation selectivity in the cat's striate cortex is invariant with stimulus contrast. *Exp. Brain Res.* 46, 457–461.
- Series, P., Latham, P., and Pouget, A. (2004). Tuning curve sharpening for orientation selectivity: coding efficiency and the impact of correlations. *Nat. Neurosci.* 7, 1129–1135.
- Shadlen, M.N., and Newsome, W.T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J. Neurosci.* 18, 3870–3896.
- Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* 86, 1916–1936.
- Shadlen, M.N., Hanks, T.D., Churchland, A.K., Kiani, R., and Yang, T. (2006a). The speed and accuracy of a simple perceptual decision: a mathematical primer. In *Bayesian Brain: Probabilistic Approaches to Neural Coding*, K. Doya, S. Ishii, R. Rao, and A. Pouget, eds. (Cambridge: MIT Press).
- Smith, P.L., and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci.* 27, 161–168.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika* 25, 251–260.
- Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782–1787.
- Tolhurst, D., Movshon, J., and Dean, A. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* 23, 775–785.
- Usher, M., and McClelland, J.L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* 108, 550–592.
- Vickers, D. (1979). *Decision Processes in Visual Perception* (New York: Academic Press).
- Wald, A. (1947). *Sequential Analysis* (New York: Wiley).
- Wald, A., and Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *Ann. Math. Stat.* 19, 326–339.
- Wang, X.J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36, 955–968.
- Wong, K.F., and Wang, X.J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* 26, 1314–1328.
- Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *J. Neurosci.* 16, 2112–2126.
- Zohary, E., Shadlen, M., and Newsome, W. (1994). Correlated neuronal discharge rate and its implication for psychophysical performance. *Nature* 370, 140–143.