HHS Public Access

Author manuscript

Nat Neurosci. Author manuscript; available in PMC 2015 April 01.

Published in final edited form as:

Nat Neurosci. 2014 October ; 17(10): 1395–1403. doi:10.1038/nn.3800.

Encoding and decoding in parietal cortex during sensorimotor decision-making

Il Memming Park[1,2], **Miriam L. R. Meister**[1,2,3], **Alexander C. Huk**[1,2,4], and **Jonathan W. Pillow**[1,2,5]

[1]Center for Perceptual Systems, The University of Texas at Austin, Austin, Texas, USA

[2]Department of Psychology, The University of Texas at Austin, Austin, Texas, USA

[3]Department of Physiology and Biophysics, University of Washington, Seattle, Washington, USA

[4]Department of Neuroscience, The University of Texas at Austin, Austin, Texas, USA

[5]Department of Statistics and Data Science, The University of Texas at Austin, Austin, Texas, USA

Abstract

The lateral intraparietal area (LIP) of macaques has been asserted to play a fundamental role in sensorimotor decision-making. Here we dissect the neural code in LIP at the level of individual trial spike trains using a statistical approach based on generalized linear models. We show that LIP responses reflect a combination of temporally-overlapping task and decision-related signals. Our model accounts for the detailed statistics of LIP spike trains, and accurately predicts spike trains from task events on single trials. Moreover, we derive an optimal decoder for heterogeneous, multiplexed LIP responses that could be implemented in biologically plausible circuits. In contrast to interpretations of LIP as providing an instantaneous code for decision variables, we show that optimal decoding requires integrating LIP spikes over two timescales. These analyses provide a detailed understanding of the neural code in LIP, and a framework for studying the coding of multiplexed signals in higher brain areas.

Introduction

Perceptual decision-making is an important and experimentally-tractable cognitive ability that involves the timely integration of sensory, cognitive, and motor information. Recent work has hypothesized that the lateral intraparietal area (LIP) plays a key role in simple forms of perceptual decision-making. Much of this literature has focused on either normative models, which aim to derive the optimal decision-making strategy for a given task from first principles[1–8], or mechanistic models, which aim to qualitatively reproduce

Author Contributions

J.W.P. and A.C.H. conceived the study, M.L.R.M. and A.C.H. designed the experiments, M.L.R.M. performed the experiments, I.M.P., J.W.P. and A.C.H. designed the analysis, I.M.P. performed the analysis, and A.C.H., I.M.P., M.L.R.M. and J.W.P. wrote the manuscript.

the dynamics governing neural activity in decision-making circuits[9–15]. Although these lines of work have generated many intriguing hypotheses and experiments, both of these approaches start with strong assumptions about the function of LIP and a limited view of its functional heterogeneity.

Here we describe a data-driven, statistical approach for investigating the encoding and decoding of information in LIP spike trains during a sensorimotor decision-making task. We begin by formulating a generalized linear encoding model characterizing LIP spike responses as a function of the external variables of interest on the scale of individual trials. Encoding models have been used to describe and quantify information transfer in early sensory areas[16–19] as well as motor cortices[20–22] and rodent hippocampus[23,24], but have so far had limited application to decision-making areas[25–27].

A statistical model based approach differs from other methodologies in that it does not seek a particular mechanistic or normative theory of LIP function. Rather, it aims to develop a rich, descriptive model of the statistical features of LIP spike responses and their dependencies on task and behavioral variables. The primary challenge here is that, in contrast to primary sensory or motor areas, an area like LIP may reflect a panoply of signals, some of which are tightly coupled with known sensory and motor events, and some of which may be the product of more elusive cognitive operations[28–30]. Here we show that the sort of generalized linear model (GLM) previously applied to spike representations in the early visual system can be extended to model LIP activity recorded during a decision-making task, in which many of the key elements of the task were variable, either temporally or in their value, and hence dissociable. The model reveals that LIP responses are best described as an interacting set of temporally-overlapping response components, implying that some of the spikes emitted during decision formation are potentially related to a number of factors, some of which are irrelevant to the accumulation of evidence. The model also provides a framework for understanding statistically optimal readout of various kinds of information from single-trial spike trains. Our analyses show that the superposition of sensory, decision, and motor variables encoded in LIP can be demultiplexed to read out decisions using a simple linear mechanism that spans multiple timescales. This encoding-decoding approach thus both identifies and then resolves a puzzle: LIP spike rates do not purely encode the accumulation of evidence during the formation of decisions, but LIP (or another area with similar responses) could feasibly be decoded to extract more choice-related information than conventional spike-counting analyses have suggested.

Generally, the analyses of LIP based on the GLM provide a detailed account of the time-varying information carried by LIP spikes, and operates at the level of single trials, providing a testing ground for interpreting the functional meaning of LIP spikes on the timescale of individual stimulus events and decisions. This approach provides a platform for quantitatively characterizing the information carried by LIP, for comparing LIP responses across experiments, and for assessing the adequacy of various theories of LIP function. After explaining the statistical model in the context of the moving dots task, we discuss the implications of this framework for a broader range of ongoing debates and future work surrounding LIP's functional role in sensorimotor behaviors.

Results

We analyzed the spiking activity of 80 LIP neurons recorded from two monkeys while they performed a moving-dot direction-discrimination task[31]. In this well-known task[32], monkeys view a random dot kinetogram and make decisions about the net direction of dot motion. They communicate their choice by making a saccadic eye-movement to one of two choice targets on the screen. From trial to trial, the fraction of dots moving coherently in the correct direction is varied, spanning a range of difficulties from easy (high coherence) to hard (low coherence). Fig. 1a illustrates the decision-making task and the variable timings of the four principal task elements: fixation point, choice targets, moving dots, and saccade. We varied the timings (and/or values) of the first three elements independently, and the saccade exhibited intrinsic timing variability as a part of the animal's behavior (see Methods).

Classical analyses of coding in LIP rely on the peri-stimulus time histogram (PSTH) aligned to events such as the onset of the moving dots or the occurrence of the saccade (Fig. 1b). Other approaches might regress binned spike rates against levels of various experimental variables, and are sometimes applied at the level of single trials or single neurons[7,33–35]. Our model works at the resolution of individual spikes, neurons, and trials—quantifying the dependencies of the neural response on multiple task variables by regressing single trial spike trains against the timing and value of each variable represented on each trial—while also capturing the (autoregressive) influence of a spike history on subsequent spikes, and a nonlinearity associated with spike generation from these inputs. It is therefore possible to dissociate firing rate components associated with each element given the decorrelated design of the task components (e.g., trial-to-trial variability in the relative timings of events, and independent variation of intensities such as motion coherence relative to other factors), while also capturing the neuron's own temporal response properties (e.g., refractoriness, burstiness, longer time scale autocorrelations).

Encoding: a description of the neural code in area LIP

An encoding model aims to describe $p(\boldsymbol{r}|\boldsymbol{x})$, the probability of a spike train response $\boldsymbol{r}$ given a set of external variables $\boldsymbol{x}$ on a single trial. Our model defines this probability in terms of a time-varying spike rate λ_t, given by

$$\lambda_t = \exp\left(\sum_i (\boldsymbol{k}_i * x_i)(t) + (\boldsymbol{h} * \boldsymbol{r}^{hist})(t)\right), \; conditional\ spike\ rate \quad (1)$$

where $(\boldsymbol{k}_i * x_i)(t)$ denotes linear convolution of $x_i(t)$, the timecourse for the i'th external event (e.g., it's zero everywhere except the time at which the saccade targets appear), with the linear filter or *kernel* $\boldsymbol{k}_i$, which captures the time-varying relationship between this event and the neuron's probability of spiking. The second term $(\boldsymbol{h} * \boldsymbol{r}^{hist})(t)$ denotes the linear convolution of the neuron's spike history $\boldsymbol{r}^{hist}$ with the *post-spike filter* $\boldsymbol{h}$. This filter allows the model to capture non-Poisson spike-history effects such as refractoriness, bursting, and spike rate adaptation. We illustrate these components in a model diagram shown in Fig. 2a.

Under this model, the probability of a spike train for a single trial is given by a Poisson likelihood:

$$p(\boldsymbol{r}|\boldsymbol{x},\theta)=\prod_{t=0}^{T}p(r_t|\boldsymbol{x},\theta) \propto \prod_{t=0}^{T}(\Delta\lambda_t)^{r_t}e^{-\Delta\lambda_t}, \textit{encoding distribution} \quad (2)$$

where Δ is the time bin size, T is the number of time bins in the trial, r_t is the spike count at time t, and $\theta = \{\{\boldsymbol{k}_i\}, \boldsymbol{h}\}$ are the model parameters.

Fig. 2 illustrates the model fit to data from a typical LIP neuron. To provide intuition for the model's basic capabilities, we plot the three kernels related to the three primary task elements that occur on each trial: (1) the appearance of the choice targets; (2) the moving dots stimulus; and (3) the saccade made by the monkey to indicate a decision (either into or out of the response field of the neuron under study, which we refer to as IN and OUT). At right, we plot the predicted time-varying change in spike rate due to each task element, for each of five possible motion coherences and two possible saccade directions (Fig. 2b). For each component, we show the linear convolution of the kernel with the timecourse of the corresponding task element, and passed through an exponential nonlinearity. The resulting traces show the multiplicative effect of each task element on the neuron's probability of spiking. The product of three such components, forms the predicted spike rate for a single trial (Eq. 1). These predictions closely match the neuron's actual PSTH (Fig. 2c).

Randomness and variability in the timings of experimental events are essential factors for dissociating the different components of the response. For example, if the interval between the onset of the choice targets and the saccade were constant, we could not differentiate spikes time-locked to the targets from those time-locked to the saccade. Similarly, stimulus kernels for each motion coherence can be dissociated thanks to randomized coherences and duration, despite having onset locked to the appearance of targets. The large number of task elements makes for a large number of model parameters; we therefore fit kernels in smooth temporal basis and applied Bayesian regularization methods to prevent over-fitting (see Methods). We verified the fits via PSTH prediction (Fig. 3b), single trial prediction on the test set (Fig. 5a), time-varying spike count variance (Supplementary Fig. 4), and inter-spike interval statistics (Supplementary Fig. 5).

Previous work has shown that LIP neurons are heterogeneous[36], with diverse response characteristics during the moving-dots task[31]. We fit the model to each LIP neuron in the dataset, and found that it captures the responses of both conventional and radically unconventional LIP neurons with high accuracy (Fig. 3a). The fitted model parameters reveal that LIP neurons carry information about a variety of task elements, and that the output of each LIP neuron reflects a roughly multiplicative combination of signals (Supplementary Fig. 6 and Supplementary Fig. 7). Furthermore, individual cells may encode these elements in distinct ways, both in terms of overall magnitude and in more nuanced aspects of the time course. This cell-by-cell analysis suggests that the striking differences in PSTHs may arise from the combination of heterogeneous task-related components that can now be considered in isolation at the level of single neurons.

In addition to capturing the average timecourse of neural activity for different stimulus and choice conditions, the model can predict spiking activity on single trials from the timings

and values of task elements (Fig. 4). Despite the diversity of responses across trials and across neurons, the model captures the details of single-trial spike rate modulations very well, on par or better than the model's account of the full PSTH. Note that these predictions are unique to each neuron, and differ for each trial because the task elements have randomized times.

Single-trial prediction accuracy improves even further when the model includes a post-spike filter, which captures the effects of spike history on a neuron's probability of firing (Fig. 5a). The autocorrelation functions of spike trains in LIP vary substantially across neurons (Fig. 5b), and frequently exhibit fine timescale structure that is inconsistent with a Poisson response model (or a Poisson model with refractory period). A model without post-spike filters cannot account for the detailed shape of these autocorrelation functions, yet the full model captures them accurately. The fitted post-spike filters (Fig. 5c) reveal detailed and diverse shapes that are not obvious from the autocorrelation function. One striking observation is that for these neurons, the probability of firing is enhanced for more than 200 ms after a spike (but not much longer; see Supplementary Fig. 8). Note that such a long timescale of self-excitation is a statistical description (as opposed to a proposed recurrent mechanism), and reflects a form of temporal integration above and beyond that explained by the components linked to external task elements. Despite their contribution to spike train prediction, it is reassuring that the inclusion of the post-spike filters in the model exerts only a modest scaling effect on the other temporal kernels (i.e., the median correlation coefficient of the kernels fit with or without the post-spike filter is 96%, and the median scaling effect is 75%).

We can validate the use of an exponential nonlinearity by non-parametrically reconstructing the nonlinear relationship between filter output and firing rate, and we find a close match between the non-parametric estimate and an exponential (Supplementary Fig. 6a). This is consistent with multiplicative interaction among components (but see Supplementary Fig. 6b,c for comparison to rectified linear function), which is critical for the statistically optimal linear decoding mechanism we will describe below.

Together, these parts of the encoding model—linear kernels for task variables, a post-spike filter, and an approximately exponential nonlinearity—work together to capture the statistical relationship between LIP responses and various external and internal variables relevant to a sensorimotor decision-making task, enabling the model to predict spike trains on single trials and capture the fine structure of each cell's autocorrelation function. These results suggest that, despite LIP's cognitive function and more distant relation to simple sensory and motor processing, simpler neural signals may make up the bulk of LIP responses and can be captured by a type of statistical model already shown to be successful in sensory and motor circuits. However, the value of the encoding model goes beyond its ability to account for what makes an LIP neuron respond during the dots task, as we demonstrate below with an application for decoding (or de-multiplexing) of decision-related information from single-trial responses.

Decoding: readout of decision-related information from LIP spike trains

Encoding models specify an explicit probability distribution over neural activity given a set of external task variables. They also provide a powerful tool for analyzing the readout of information from spike trains. In this particular application, we can decode decisions from spiking activity in LIP by using the fitted model to evaluate the probability of the spikes on a single trial under both possible choices (i.e., saccades to one target or the other). For each trial, the model provides a prediction in the form of a Poisson spike rate hypothesized to have generated the activity on that trial. Intuitively, decoding amounts to evaluating whether the spikes are more probable under the rate function consistent with a saccade IN or an OUT of an LIP neuron's response field.2

Fig. 6a shows spike trains and corresponding rate predictions for pairs of randomly-selected IN and OUT trials from three different neurons. By considering the spikes up to each point in time, we can obtain a time-varying estimate of the animal's eventual decision (shown below each trial's spike train and rate predictions). Note that these probabilistic estimates diverge from 0.5 (the prior probability of an IN decision) as soon as the rate predictions diverge. For the six trials shown, the decoder achieves near certainty (probability of an IN choice close to one or zero) by the time of saccade, although this does not occur for all trials or in all neurons.

Formally, model-based decoding of the animal's decision relies on the posterior distribution over choice given the spikes, which can be derived from the encoding distribution using Bayes' rule. If we assume that IN and OUT choices are equally probable *a priori*, the posterior probability is given by

$$p(x_c{=}\mathrm{IN}|\boldsymbol{r},\tilde{\boldsymbol{x}}){=}\frac{p(\boldsymbol{r}|x_c{=}\mathrm{IN},\tilde{\boldsymbol{x}})}{p(\boldsymbol{r}|x_c{=}\mathrm{IN},\tilde{\boldsymbol{x}}){+}p(\boldsymbol{r}|x_c{=}\mathrm{OUT},\tilde{\boldsymbol{x}})},\ \textit{decoding distribution} \quad (3)$$

where $\tilde{\boldsymbol{x}}$ denotes all other covariates besides the choice direction component x_c. The posterior probability of an OUT choice is simply $1 - p(x_c = \mathrm{IN}|\boldsymbol{r},\tilde{\boldsymbol{x}})$, as the posterior must sum to 1.

Fortuitously, these posterior probabilities can be computed very simply under our model. For a GLM with Poisson noise and exponential nonlinearity, Bayesian decoding can be achieved with a linear weighting of the spike response. Specifically, the log-likelihood ratio (LLR), which is also the log of the ratio of posterior probabilities, is given by

$$\mathrm{LLR}{=}\log\frac{p(x_c{=}\mathrm{IN}|\boldsymbol{r},\tilde{\boldsymbol{x}})}{p(x_c{=}\mathrm{OUT}|\boldsymbol{r},\tilde{\boldsymbol{x}})}{=}\boldsymbol{w}^{\mathbf{T}}\boldsymbol{r}{+}\textit{constant},\ \textit{linear decoding} \quad (4)$$

where $\boldsymbol{w}$ denotes a set of linear decoding weights over time, and the additive constant does not depend on the response $\boldsymbol{r}$ (see Methods). The LLR leads to a simple decoding rule for predicting the animal's choice: whenever LLR > 0, an IN saccade is more probable; whenever LLR < 0, an OUT saccade is more probable. The optimal decoding weights are in fact given by $\boldsymbol{w} = \boldsymbol{k}_{\mathrm{IN}} - \boldsymbol{k}_{\mathrm{OUT}}$, the difference of the fitted saccade kernels for IN and OUT choices (see Fig. 6a).

Bayesian decoding of decisions from LIP spike trains can therefore be implemented with a set of time-varying linear weights. The shape of each neuron's weight profile tells us how much information that neuron's spikes carry about the animal's decision, as a function of time before the movement. The weights determined empirically from data challenge the view that instantaneous spike rate encodes the animal's accumulated evidence for a decision, as they frequently extend out to 1.5 s before the movement, indicating that spikes from LIP should themselves be integrated over a relatively long window to optimally predict the decision. Note that this linear readout mechanism would not be statistically optimal if the nonlinearity were not exponential, since our model technically sits within the family of log-linear models with "Poisson-like" noise known as *probabilistic population codes* (PPC)[6,37,38].

This decoding analysis supports a model-based extension of the basic concept of *Choice Probability* (CP), a metric for quantifying the information that neural responses carry about an animal's decision[38]. Conventional CP applies to scalar quantities such as spike count, and assumes that the optimal rule for reading out a pair of neuron and "anti-neuron" responses is to choose the option with larger response (a "max" decoding rule). *Model-based CP*, on the other hand, can be defined using any model-based decoding rule, and applies to non-scalar representations of neural activity such as a binned spike train. If we assume an anti-neuron for the neuron we are decoding to have exactly the same model weights, except with swapped choice related decoding weights $\boldsymbol{k}_{\mathrm{IN}}$ for $\boldsymbol{k}_{\mathrm{OUT}}$, and vice versa), then a simple comparison rule for the projection to both neurons suffices for decoding, because the threshold is exactly the same for both neurons. We can therefore decode choices using the model-based projection of a spike train, instead of simple spike count.

The key difference between conventional CP and model-based CP is that the model weights provide the relative importance of different time bins for reading out the choice. Fig. 6b–d shows a comparison of model-based and conventional choice probability for quantifying decoding performance. Model-based CP outperforms conventional CP as the time window grows towards the saccade (Fig. 6c). Moreover, model-based CP increases almost monotonically, while conventional CP fluctuates more around chance levels. Conventional CP is clearly inadequate for decoding LIP neurons like example cell 3, where the polarity of rate associated with IN trials and OUT trials are reversed through time. But even for the more canonical example cell 1, model-based decoding extracts significantly more decision-related information than the conventional approach. Model-based CP is on average 6.2% point higher than conventional CP on average (mean model-based CP, 75.2%; mean traditional CP, 69.9%). Fig. 6d illustrates the advantage of model-based decoding for a population of 80 neurons.

Implementation: low-dimensional readout of LIP population activity

So far we have considered the problem of decoding choices from spikes in LIP using an ideal observer with access to the saccade time and other task variables for the trial in question. This perspective is useful for determining theoretical limits on the accuracy with which downstream neurons could read out choice-related information from LIP, but it does not necessarily appl to neurally plausible readout mechanisms. At first glance, statistically

optimal readout seems difficult due to the fact that each neuron has a unique decoding-weight profile, and these weights exhibit significant cell-to-cell heterogeneity.

However, there exists a practical mechanism for population readout that requires no knowledge of the saccade (or decision) timing, is insensitive to other task variables, and flexibly accommodates the heterogeneity of the temporal decoding weight profiles—but which still performs very close to theoretically optimal decoding. First, we can compute the linear projection of the spike trains onto the decoding weights in continuous time via linear filtering: at each time the output of the linear filter is exactly the log-likelihood ratio (Eq. 4) for a "preferred" saccade at the current time. Second, instead of projecting to the time-varying decoding weights associated with each neuron—which are highly diverse Fig. 7a (top row)—we can approximately decompose each decoding weight as a linear combination of a few principal components of the collection of decoding weights (Fig. 7a bottom row), and then sum the results. Indeed, we find that the constellation of decoding weights are well spanned by 2 principal components (explaining 93.3% of the variance). To represent the total log-likelihood of LIP neurons corresponding to the same decision (i.e., with overlapping RFs), the filter outputs corresponding to each RF can simply be summed. Third, we find that the "decoding subspace" spanned by these two principal components can be well approximated by a pair of exponential filters. This simplifies the problem of implementation because, while single neurons may not be able to perform arbitrary time-dependent weightings of spikes on the timescale of a second, a neuron or a population of neurons can naturally implement a leaky integrator. We find that two exponential filters $e^{-t/\tau}$ with τ = 237,410 ms explain 92.1% of the variance of the seemingly diverse decoding weights (dashed lines in Fig. 7a). Finally, we posit the existence of two competing populations, each integrating evidence for one of the choice alternatives (i.e., two opposite RF locations), with balanced contributions from the other response components so that the influence of the other task variables cancels.

Fig. 7b shows a schematic of the proposed decoding circuit: (1) Each subpopulation of LIP neurons associated with the same choice/RF project to two readout populations with appropriate weights, (2) each readout population acts as a simple first-order temporal filter with different time constants, (3) the overall activity of the readout population can be instantaneously summed to represent choice activity, and (4) through a winner-take-all mechanism the choice information is selected, then transmitted to the motor system that controls the saccade.

The proposed scheme yields compelling improvements over traditional decoders (Fig. 7c): Decoding performance using the model-derived weights is substantially better than simply adding the spike counts in a fixed time window (i.e., a boxcar). Moreover, decoding with two leaky integrators performs virtually as well as with the seemingly complex individual decoding weights.

In summary, this decoding scheme lays out one plausible way to read out the choice given the population activity of LIP neurons. It characterizes how decision-related information is represented in the temporal evolution of population spiking activity. The dimensionality reduction implementation shows that this information could be exploited by a biologically-

reasonable neural circuit. This exercise suggests that statistically-optimal readout of decisions could be implemented by neural circuits that read out the activity of LIP, even in the face of multiplexed signals and considerable heterogeneity across neurons.

Invariance to target-induced changes in firing

In addition to its neural plausibility, the decoding mechanism we have proposed achieves performance that is invariant to firing rate modulations induced by variables not directly relevant to the task. In particular, we manipulated the visual saccade targets so that on half of the trials they persisted until the end of the trial ("targets-ON" trials), while on the other half they flashed briefly and then disappeared ("targets-FLASH" trials; the animal made a saccade to the remembered, stereotyped target location). Although this manipulation did not affect choice behavior, it produced large changes in spike rate in many neurons[31]. This presents a major challenge to models in the prior literature that regard the firing rate of an LIP neuron as a direct neural correlate of log-likelihood ratio or some other normative quantity[5]. Our descriptive model, however, capture the effects of target removal through a multiplicative interaction with a target-related kernel. (Although saccade endpoints were affected by the targets-FLASH manipulation, the single target-based kernel appears sufficient to account for the changes in LIP response; see also[40]).

Fig. 8a shows an example neuron that reduces its response drastically when visual targets are extinguished early in the trial (i.e., before the onset of the moving dots). The model captures the difference with a single filter aligned to the time when targets disappear. Furthermore, model-based decoding is unaffected by the inclusion of this component, and there is negligible difference in decoding performance for the two kinds of trials (Fig. 8b). This robustness indicates that the superposition of the target-related kernels and the decision-related (saccade-locked) kernels is appropriate. More generally, it demonstrates that additional decision-irrelevant factors that affect LIP responses in other tasks might similarly be isolated so as to preserve the readout of decisions.

Discussion

We have developed a generalized linear model to describe LIP responses during a complex perceptual decision-making task. This framework yields several new insights into the coding of information in LIP spike trains. First, it reveals that representations in LIP are multiplexed: they can be decomposed into separate components related to the targets, the visual stimulus, and the eventual saccadic decision; the superposition of these components, in turn, provides accurate prediction of spike responses on single trials. Second, LIP neurons exhibit notably long-time-scale self-excitation that is statistically separable from the effects of sensory and motor drive, an effect that can be captured with a spike-history dependent model component. Third, despite substantial population-level heterogeneity in LIP response characteristics, decision-related information is well captured by a two-dimensional linear projection of the spike trains, which leads to an implementation of statistically optimal decoding using a pair of leaky integrators with distinct timescales.

Our model-based approach assumes that the spike rate at any time depends on a combination of multiple task components. By exploiting the randomized structure of the experimental design on a single-trial basis, our analysis shows that LIP activity can be well modeled as reflecting a superposition of contributions from different sources. This result mirrors insights recently reported in rodent[41] and other primate[28] decision-making tasks. The dominant normative model in the context of the moving dots task—drift diffusion to bound—takes the coherence-dependent ramping signals seen in averages as a neural correlate of an evolving decision variable. This framework assumes that the instantaneous firing rate in LIP should reflect the eventual choice. Our analysis adds a dimension to this perspective, suggesting a readout rule involving much longer temporal integration of LIP response, on the time scale of seconds. In other words, if LIP is the critical encoder of the accumulation of evidence, an optimal mechanism would still need to decode it using a time-varying weighting function, as opposed to simply thresholding its instantaneous spike rate. Time-varying decoding weights may be more generally applicable in other contexts for which a fixed threshold on LIP responses has also been shown to be untenable[42].

Multiplexing applies even to the decision formation period, where spikes are influenced by the moving dots stimulus, but also by the placement of the choice targets and the upcoming saccade. Although LIP clearly carries decision-related signals, the mixture of sensory, cognitive, and motor signals that simultaneously affect it argues against the assumption that it is a pure integrator underlying the accumulation of sensory evidence[28–30]. Instead, our encoding analysis suggests that LIP should be regarded as the recipient of multiple task-related signals, only some of which are directly related to the formation of decisions. Our decoding analyses show that it is still possible to read out decisions in the face of this multiplexing, with higher accuracy than indicated by conventional choice probability. This perspective raises the possibility that the difficulties in distinguishing between various functional roles of LIP (e.g., attention/salience, decision-making, motor intention) may have arisen because LIP does not have single functional role, but because it simultaneously encodes multiple signals that are at least partially separable downstream. The application of descriptive statistical models in other tasks and task variants may thus provide the means for further integration of visual, attentive, decisional, and motor function in LIP. Further causal manipulations of the circuit are likely necessary for a definitive resolution, but our statistical analyses have shown that the information is multiplexed at encoding and can be demultiplexed by a plausible readout stage. One implication is that if LIP responses are critical for decision-making, future empirical work should shift focus from demonstrating the encoding of various cognitive factors in LIP, to how these complex signals are de-multiplexed by later stages of neural processing.

Although our model provides an illuminating perspective on LIP function, it is not without limitations. It makes very few assertions about the function of LIP, but does require making some assumptions in the structure of the model. However, these assumptions amount to decisions about how to represent external, observable variables within a regression framework, and thus involve far less conceptual baggage, and a tighter reliance on the data, than normative approaches. In this initial report, we attempted to use the simplest approach, modeling brief events as impulses and prolonged events as boxcars. Future work will benefit from more nuanced manipulation and modeling of finer-grained temporal structure on each

trial. Additionally, prior applications of generalized linear models have provided insights into the form of dependencies between neurons and the relevance of correlated firing to information carried by a neural population; such analyses motivate follow-up work that involves recording multiple cells simultaneously[43]. Our findings about neurally plausible implementations of decoding also motivate future investigations into how or where the brain might accomplish readout involving simultaneous recordings in multiple brain areas. Indeed, work in related structures and paradigms has already been to explore how neurally-constrained models of the oculomotor system can be applied[44,45].

Dynamical models provide another kind of approach to understanding the function of area LIP[13], and such models have been elaborated to include multiple factors that drive LIP, such as the response to visual targets[11]. A key difference is that such models start with a semi-biological circuit structure that constrains the neural dynamics, and the framework for adding other components is rather flexible. Statistical models of the kind we propose could complement these models by providing a principled technique for assessing which components to incorporate; these models may thus provide a key tool not just for interpreting LIP activity, but for linking normative and mechanistic models of the sorts already proposed in this context[6,13,46,47].

More generally, we note that a variety of both classical and recent controversies regarding LIP's functional role have relied on attempts to dissociate or isolate various response components. For example, a recent paper reported that LIP responses were modulated by the magnitude of either appealing or aversive outcomes (consistent with a "salience" account) as opposed to being sensitive to value (in the neuroeconomic sense) per se[48]. A follow-up debate then focused on whether the data were generalizable, owing to the lack of strong persistent activity seen during the task[49,50]. Our analysis could provide a path towards resolution of this sort of debate, by testing whether persistent activity can be modeled as an independent driver of the LIP response (in which case, its presence or absence is irrelevant to the salience/value issue), or whether it indeed interacts with salience/valuation. The latter outcome would not just demonstrate that persistent activity is necessary to observe value signals in LIP, but would suggest intriguing computations that link these two components.

More consideration will be required to create families of experimental designs amenable to analysis via the generalized linear model. However, we feel that the framework promises both to enrich interpretations within well-studied paradigms and to pave the way for more ambitious and direct tests of hypotheses about higher brain function[27].

Online Methods

The data analyzed in this paper were initially described in a previous empirical report[31]. Here, we briefly summarize the experimental methods as most important to the analyses described here.

Preparation, neurophysiology, and tasks

All procedures were performed in accordance with National Institutes of Health guidelines and the Institutional Animal Care and Use Committee at the University of Texas at Austin.

Two male rhesus macaque monkeys underwent surgery for implantation of a head-holder and a recording chamber over the posterior parietal cortex. A single electrode was advanced into LIP based on anatomical references, and was further located by signature neural activity observed during an instructed saccade task, which was also used to locate the response field (RF) of a neuron. Only neurons that were spatially selective during an instructed saccade task are included in this paper[31]. After the RF was determined using an instructed saccade task, the neuron's spiking responses were recorded during the decision task.

The decision task was a standard variable-duration moving-dot direction-discrimination that has been used previously. In order to achieve fixation at a trial's start, the monkey's eye position had to register within a window around the central fixation point within 3 seconds of fixation point appearance. The monkey then needed to maintain eye position within that window until the fixation point vanished, which was the "go-signal" for the monkey to make a saccade to a target.

500 ms after the monkey achieved fixation, two saccadic target choices appeared. One choice target was located in the RF of the neuron (the IN choice), whereas the other choice target was diametrically opposite (the OUT choice). Randomly, on half of the trials, the targets disappeared after 100 ms (in figures showing PSTHs or single trials, we show only the "targets-ON" trials). 200 ms after the appearance of the choice targets, a motion stimulus appeared in a circular aperture and the monkey had to decide the net direction of motion and communicate their decision with a saccade to one of the two choice targets.

The algorithm for generating the dot motion display was identical to that used in prior LIP studies by Shadlen and colleagues[31,33]. Motion coherence strengths of 0, 3.2, 6.4, 12.8, 25.6 and 51.2% were used. At 0% coherence (no motion on average), the monkey was randomly rewarded with 50% probability. After the motion was displayed for a specified time (500–1000 ms, uniform distribution), the motion stimulus vanished. The monkey continued to maintain fixation for another 500 ms until the fixation point also disappeared, thus signaling that he could now saccade to the correct target location to obtain a reward. The monkey received a liquid reward 200 ms after eye position entered the spatial window around the correct target. The monkey was considered to have made a saccade to a target location if eye position registered within the spatial window around the target location within 100 ms after leaving the fixation point window. Entry into the spatial window around the target also had to occur within 450 ms after the go-signal to be considered a complete trial. About 800–1200 trials were collected per recording session (per neuron).

Plotting neural response

PSTHs were smoothed with a gaussian filter (standard deviation, 75 ms). In Figs. 2 and 3, each conditional PSTH was stitched together by averaging temporally overlapping local PSTHs of length 500–1000 ms that were obtained by aligning at the median event time. The spike trains corresponding to fixed windows around the aligned time are collected and averaged to obtain the local PSTHs. For example, for the dots motion event, we aligned to the onset and offset separately with −25–600 ms and −400–500 ms time windows, respectively. Conditional PSTHs with less than 20 trials are not shown and were excluded from the variance explained analysis.

Trials in Fig. 4 were chosen by sorting the mean square error between the predicted spike rate and the boxcar smoothed spike trains within the 2 second window around saccade time. Only the correct decision trials with high coherence in the cross-validation set were used.

The autocorrelation function $R(\tau)$ was normalized by the mean firing rate m to quantify excess spike rate:

$$R(\tau)=\frac{1}{m}\left(\frac{1}{N(\tau)}\sum_t r(t)r(t-\tau)\right)-m,$$

where t is over all bins, $r(t)$ is the binned spike train, and $N(\tau)$ is the number of bins such that both $r(t)$ and $r(t-\tau)$ are valid. We have removed the 0-th lag component from the autocorrelation plots in Fig. 5b.

Model parameter representation

The spike trains were discretized into 1 ms bins. Each event was represented as a delta-function over time, and convolved with a filter (or a kernel, see Supplementary Fig. 1). We used smooth temporal basis functions defined by raised cosine bumps separated by $\pi/2$ radians (50 ms) to parameterize the filters (see Supplementary Fig. 10). Each event kernel was represented as a linear combination of basis functions that cover a specified range of time: 2 s window after fixation point onset, 1 s window after target onset, 4 s window after target disappearance, and 500 ms window after the dots disappearance. For the saccade response, we had two kernels, one for each direction which was anti-causal for 2504 ms for monkey J, and 2662 ms for monkey P and with total duration of 5.3 s. We grouped the directional coherence levels into 5 groups (to reduce the number of parameters): one for zero-coherences (0%), and two each for high (51.2%, 25.6%, 12.8%), and low (6.4%, 3.2%) in each direction. The dot-motion stimulus was represented as a boxcar of corresponding duration, and the filter duration for each coherence level was 800 ms. The post-spike history filter was parameterized by 20 linear weights; ten 1 ms uniform basis (to represent the fast refractory effects) followed by ten raised cosine basis stretched in a logarithmic scale that spanned 265 ms[17] (see Supplementary Fig. 10). The total number of parameters for each neuron was 402 (or 422 with post-spike history filter) from the 11 (or 12) filters above. To facilitate visual interpretation, a rank-2 parametrization was used for the 5 coherence kernels for Figs. 1 and 2.

Fitting

The encoding model was fit by maximizing the log posterior, where we used a ridge prior to regularize the weights:

$$\mathscr{L}(\theta)=\sum_{t=0}^{T}(r(t)\log(\lambda_t\Delta)-\lambda_t\Delta)-\xi\|\theta\|^2,$$

where θ is a vector representing the weights on the basis functions, and ξ is chosen from a grid to maximize the marginal likelihood. The marginal likelihood $p(\boldsymbol{r}|\xi)=\int p(\boldsymbol{r}|\theta,\xi)p(\theta|\xi)d\theta$

was computed using Laplace approximation. For each neuron, we divided the trials randomly into five equally-sized sets to perform 5-fold cross-validation. Maximizing marginal likelihood for selecting the hyperparameter ξ is known as "evidence optimization", and it does not require cross-validation, hence we only use cross-validation sets for evaluating the resulting fits.

Goodness-of-fit in Fig. 3 was computed on the stitched PSTH smoothed with a 25 ms Gaussian for each stimulus coherence and behavioral choice pair. Conditions with less than 30 trials were discarded from the comparison. Spike prediction accuracy in Fig. 5a were computed by taking the difference between the full model log-likelihood and the log-likelihood of a (single parameter) homogeneous Poisson model normalized by the number of spikes on the cross-validation set[17].

Decoding

We derive the log-likelihood ratio (Eq. 4) decoding rule. Let t_c be the time when the decision is read out.

$$
\begin{aligned}
\mathrm{LLR} &= \log \frac{p(\boldsymbol{r}|\tilde{\boldsymbol{x}}, x_c{=}\mathrm{IN})}{p(\boldsymbol{r}|\tilde{\boldsymbol{x}}, x_c{=}\mathrm{OUT})} \\
&= \log\left[\prod_{t=0}^{T} \Delta\lambda_t(x_c{=}\mathrm{IN})^{r_t} e^{-\Delta\lambda_t(x_c=\mathrm{IN})}\right] - \log\left[\prod_{t=0}^{T} \Delta\lambda_t(x_c{=}\mathrm{OUT})^{r_t} e^{-\Delta\lambda_t(x_c=\mathrm{OUT})}\right] \\
&= \sum_{t=0}^{T} r_t \log\left(\Delta\lambda_t(x_c{=}\mathrm{IN})\right) - \Delta\lambda_t(x_c{=}\mathrm{IN}) - r_t \log\left(\Delta\lambda_t(x_c{=}\mathrm{OUT})\right) + \Delta\lambda_t(x_c{=}\mathrm{OUT}) \\
&= \sum_{t=0}^{T} r_t \left(\sum_{\mathrm{i}\neq\mathrm{c}} (\boldsymbol{k}_i * x_i)(t) + \left(\boldsymbol{k}_c^{\mathrm{IN}} * x_c\right)(t) + \left(\boldsymbol{h} * \boldsymbol{r}^{hist}\right)(t) - \sum_{\mathrm{i}\neq\mathrm{c}} (\boldsymbol{k}_i * x_i)(t) + \left(\boldsymbol{k}_c^{\mathrm{IN}} * x_c\right)(t) + \left(\boldsymbol{h} * \boldsymbol{r}^{hist}\right)(t)\right) + constant \\
&= \sum_{t=0}^{T} r_t \left(\left(\boldsymbol{k}_c^{\mathrm{IN}} * x_c\right)(t) - \left(\boldsymbol{k}_c^{\mathrm{OUT}} * x_c\right)(t)\right) + constant \\
&= \sum_{t=0}^{T} r_t \left(\boldsymbol{k}_c^{\mathrm{IN}}(t - t_c) - \boldsymbol{k}_c^{\mathrm{OUT}}(t - t_c)\right) + constant = \boldsymbol{w}^{\mathbf{T}} \boldsymbol{r} + constant
\end{aligned}
$$

where the constant terms do not depend on $\boldsymbol{r}$.

Note that the dot product can be computed continuously at every time point by a simple linear filter since $\boldsymbol{w}^{\boldsymbol{T}}\boldsymbol{r}=\sum_{i=1}^{L} w(i)r(t-L+i)=\tilde{\boldsymbol{w}}*\boldsymbol{r}(t)$ where L is the length of the weight vector, $\boldsymbol{w}$ ĩs the time reversed weight, and $\boldsymbol{r}(t)$ is the length L vector of response approaching up to t.

Choice probability (CP) quantifies the dependence between the spike counts of a neuron and a binary decision variable while keeping all sensory information constant. Given a randomly selected IN trial, and an OUT trial, the probability that the higher spike count belongs to the IN trial is CP. A high CP implies that the spike count contains much “information” about the decision. CP in Fig. 6c,d was computed by randomly drawing 1000 pseudorandom (IN, OUT) pairs of trials. We estimated repeated this procedure 40 times to obtain the error bars. For the traditional CP, we used the spikes count in the −1500–50 ms window before the saccade detection (for Fig. 6c, the beginning of the window was fixed at −1500 ms, and x-axis shows the end of the window). For the model-based CP, spike trains on the cross-validation set were projected with the weights obtained from the corresponding training set.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank J. Yates, K. Latimer, and L. Cormack for discussions, K. Eastman for assistance with data collection, and C. Rorex and S. Winston for technical support.

This project was supported by NEI EY017366 (A.C.H.), NIMH MH099611 (J.W.P. & A.C.H.), Sloan Foundation (J.W.P.), McKnight Foundation (J.W.P.), and NSF Career award IIS-1150186 (J.W.P.).

References

1. Platt ML, Glimcher PW. Neural correlates of decision variables in parietal cortex. Nature. 1999 Jul; 400(6741):233–238. [PubMed: 10421364]
2. Shadlen MN, Newsome WT. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J. Neurophysiol. 2001 Oct; 86(4):1916–1936. [PubMed: 11600651]
3. Huk AC, Shadlen MN. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J. Neurosci. 2005 Nov; 25(45):10420–10436. [PubMed: 16280581]
4. Kable JW, Glimcher PW. The neural correlates of subjective value during intertemporal choice. Nat. Neurosci. 2007 Dec; 10(12):1625–1633. [PubMed: 17982449]
5. Yang T, Shadlen MN. Probabilistic reasoning by neurons. Nature. 2007; 447(7148):1075–1080. [PubMed: 17546027]
6. Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, Shadlen MN, Latham PE, Pouget A. Probabilistic population codes for Bayesian decision making. Neuron. 2008; 60(6):1142–1152. [PubMed: 19109917]
7. Seo H, Barraclough DJ, Lee D. Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game. J. Neurosci. 2009; 29(22):7278. [PubMed: 19494150]
8. Cisek P, Kalaska JF. Neural mechanisms for interacting with a world full of action choices. Annu. Rev. Neurosci. 2010; 33:269–298. [PubMed: 20345247]

9. Mazurek ME, Roitman JD, Ditterich J, Shadlen MN. A role for neural integrators in perceptual decision making. Cereb. Cortex. 2003; 13(11):1257. [PubMed: 14576217]
10. Lo CC, Wang XJ. Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. Nat. Neurosci. 2006 Jul; 9(7):956–963. [PubMed: 16767089]
11. Wong KF, Huk AC, Shadlen MN, Wang XJ. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. Front Comput Neurosci. 2007; 1:6. [PubMed: 18946528]
12. Fusi S, Asaad WF, Miller EK, Wang X-J. A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple timescales. Neuron. 2007 Apr; 54(2):319–333. [PubMed: 17442251]
13. Wang X-J. Decision making in recurrent neuronal circuits. Neuron. 2008 Oct; 60(2):215–234. [PubMed: 18957215]
14. Ganguli S, Bisley JW, Roitman JD, Shadlen MN, Goldberg ME, Miller KD. One-dimensional dynamics of attention and decision making in LIP. Neuron. 2008 Apr; 58(1):15–25. [PubMed: 18400159]
15. Soltani A, Wang XJ. Synaptic computation underlying probabilistic inference. Nat. Neurosci. 2010 Jan; 13(1):112–119. [PubMed: 20010823]
16. Pillow JW, Paninski L, Uzzell VJ, Simoncelli EP, Chichilnisky EJ. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. J. Neurosci. 2005; 25:11003–11013. [PubMed: 16306413]
17. Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EPEJ. Simoncelli. Spatio-temporal correlations and visual signaling in a complete neuronal population. Nature. 2008; 454:995–999. [PubMed: 18650810]
18. Jacobs AL, Fridman G, Douglas RM, Alam NM, Latham PE, Prusky GT, Nirenberg S. Ruling out and ruling in neural codes. Proc. Natl. Acad. Sci. U.S.A. 2009 Apr; 106(14):5936–5941. [PubMed: 19297621]
19. Fernandes HL, Stevenson IH, Phillips AN, Segraves MA, Kording KP. Saliency and saccade encoding in the frontal eye field during natural scene search. Cerebral Cortex. 2013 Jul.
20. Paninski L, Fellows M, Shoham S, Hatsopoulos N, Donoghue J. Superlinear population encoding of dynamic hand trajectory in primary motor cortex. J. Neurosci. 2004; 24:8551–8561. [PubMed: 15456829]
21. Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects. J. Neurophysiol. 2005; 93(2):1074–1089. [PubMed: 15356183]
22. Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. J. Neurophysiol. 2009; 102(1):614. [PubMed: 19357332]
23. Brown E, Frank L, Tang D, Quirk M, Wilson M. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. J. Neurosci. 1998; 18:7411–7425. [PubMed: 9736661]
24. Barbieri R, Wilson MA, Frank LM, Brown EN. An analysis of hippocampal spatio-temporal representations using a Bayesian algorithm for neural spike train decoding. IEEE T. Neur. Sys. Reh. 2005 Jun; 13(2):131–136.
25. Rorie AE, Gao J, McClelland JL, Newsome WT. Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. PLoS One. 2010; 5(2):e9308. [PubMed: 20174574]
26. Park J, Zhang J. Sensorimotor locus of the buildup activity in monkey lateral intraparietal area neurons. J. Neurophysiol. 2010; 103(5):2664–2674. [PubMed: 20164399]
27. Jenison RL, Rangel A, Oya H, Kawasaki H, Howard MA. Value encoding in single neurons in the human amygdala during decision making. J. Neurosci. 2011; 31(1):331–338. [PubMed: 21209219]
28. Rishel CA, Huang G, Freedman DJ. Independent category and spatial encoding in parietal cortex. Neuron. 2014 Jul; 77(5):969–979. [PubMed: 23473325]
29. Huk AC, Meister ML. Neural correlates and neural computations in posterior parietal cortex during perceptual decision-making. Frontiers in integrative neuroscience. 2012; 6

30. Gottlieb J, Goldberg ME. Activity of neurons in the lateral intraparietal area of the monkey during an antisaccade task. Nat Neurosci. 1999 Oct; 2(10):906–912. [PubMed: 10491612]
31. Meister MLR, Hennig JA, Huk AC. Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. J. Neurosci. 2013 Feb; 33(6):2254–2267. [PubMed: 23392657]
32. Newsome WT, Pare EB. A selective impairment of motion perception following lesions of the middle temporal visual area (MT). J. Neurosci. 1988 Jun; 8(6):2201–2211. [PubMed: 3385495]
33. Roitman JD, Shadlen MN. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J. Neurosci. 2002; 22(21):9475. [PubMed: 12417672]
34. Sugrue LP, Corrado GS, Newsome WT. Matching behavior and the representation of value in the parietal cortex. Science. 2004 Jun; 304(5678):1782–1787. [PubMed: 15205529]
35. Foley NC, Jangraw DC, Peck C, Gottlieb J. Novelty enhances visual salience independently of reward in the parietal lobe. The Journal of Neuroscience. 2014 Jun; 34(23):7947–7957. [PubMed: 24899716]
36. Premereur E, Vanduffel W, Janssen P. Functional heterogeneity of macaque lateral intraparietal neurons. J. Neurosci. 2011 Aug; 31(34):12307–12317. [PubMed: 21865473]
37. Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. Nat. Neurosci. 2006; 9:1432–1438. [PubMed: 17057707]
38. Jazayeri M, Movshon JA. Optimal representation of sensory information by neural populations. Nat Neurosci. 2006 May; 9(5):690–696. [PubMed: 16617339]
39. Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA. A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis. Neurosci. 1996; 13(01):87–100. [PubMed: 8730992]
40. Steenrod SC, Phillips MH, Goldberg ME. The lateral intraparietal area codes the location of saccade targets and not the dimension of the saccades that will be made to acquire them. Journal of neurophysiology. 2013 May; 109(10):2596–2605. [PubMed: 23468388]
41. Brunton BW, Botvinick MM, Brody CD. Rats and humans can optimally accumulate evidence for Decision-Making. Science. 2013 Apr; 340(6128):95–98. [PubMed: 23559254]
42. Gottlieb J, Balan P. Attention as a decision in information space. Trends in Cognitive Sciences. 2010 Jun; 14(6):240–248. [PubMed: 20399701]
43. Bollimunta A, Totten D, Ditterich J. Neural dynamics of choice: Single-trial analysis of decision-related activity in parietal cortex. J. Neurosci. 2012 Sep; 32(37):12684–12701. [PubMed: 22972993]
44. Schall JD, Purcell BA, Heitz RP, Logan GD, Palmeri TJ. Neural mechanisms of saccade target selection: gated accumulator model of the visual-motor cascade. The European journal of neuroscience. 2011 Jun; 33(11):1991–2002. [PubMed: 21645095]
45. Purcell BA, Schall JD, Logan GD, Palmeri TJ. From salience to saccades: multiple-alternative gated stochastic accumulator model of visual search. The Journal of neuroscience : the official journal of the Society for Neuroscience. 2012 Mar; 32(10):3433–3446. [PubMed: 22399766]
46. Gold JI, Shadlen MN. The neural basis of decision making. Annu. Rev. Neurosci. 2007; 30:535–574. [PubMed: 17600525]
47. Churchland AK, Kiani R, Chaudhuri R, Wang X, Pouget A, Shadlen MN. Variance as a signature of neural computations during decision making. Neuron. 2011; 69(4):818–831. [PubMed: 21338889]
48. Leathers ML, Olson CR. In monkeys making value-based decisions, LIP neurons encode cue salience and not action value. Science. 2012 Oct; 338(6103):132–135. [PubMed: 23042897]
49. Newsome WT, Glimcher PW, Gottlieb J, Lee D, Platt ML. Comment on “In monkeys making value-based decisions, LIP neurons encode cue salience and not action value”. Science. 2013 Apr. 340(6131):430. [PubMed: 23620037]
50. Leathers ML, Olson CR. Response to comment on “In monkeys making value-based decisions, LIP neurons encode cue salience and not action value”. Science. 2013 Apr.340(6131):430. [PubMed: 23620038]

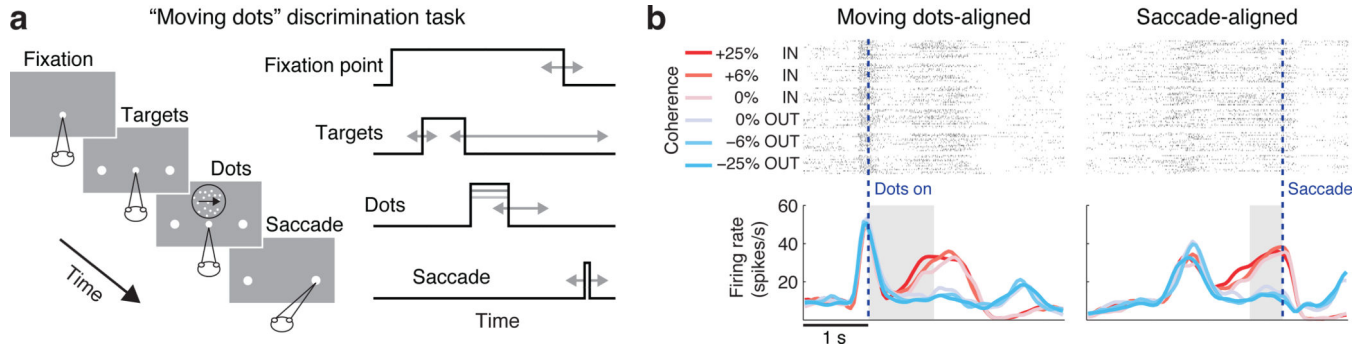

Figure 1. Decision-making task and classical analysis of LIP responses

(a) The task requires an observer to fixate, judge the direction of a moving (dots) stimulus, and report a decision by moving the eyes to one of two targets. Temporal variability in the task design variables (indicated by gray arrows) allows for statistical dissociation of the effects of different extrinsic variables on neural responses. **(b)** Standard analysis of spike responses from an LIP neuron: Spike trains are conditioned on the stimulus coherence, the fraction of dots moving towards ('+') or away from ('−') the neuron's response field, and aligned to the stimulus onset (left) or saccade time (right). Although the spike trains exhibit substantial variability (top), their average time-courses (below) exhibit coherence-dependent ramping. The gray area indicates the short portions of trials often considered in prior work[33].

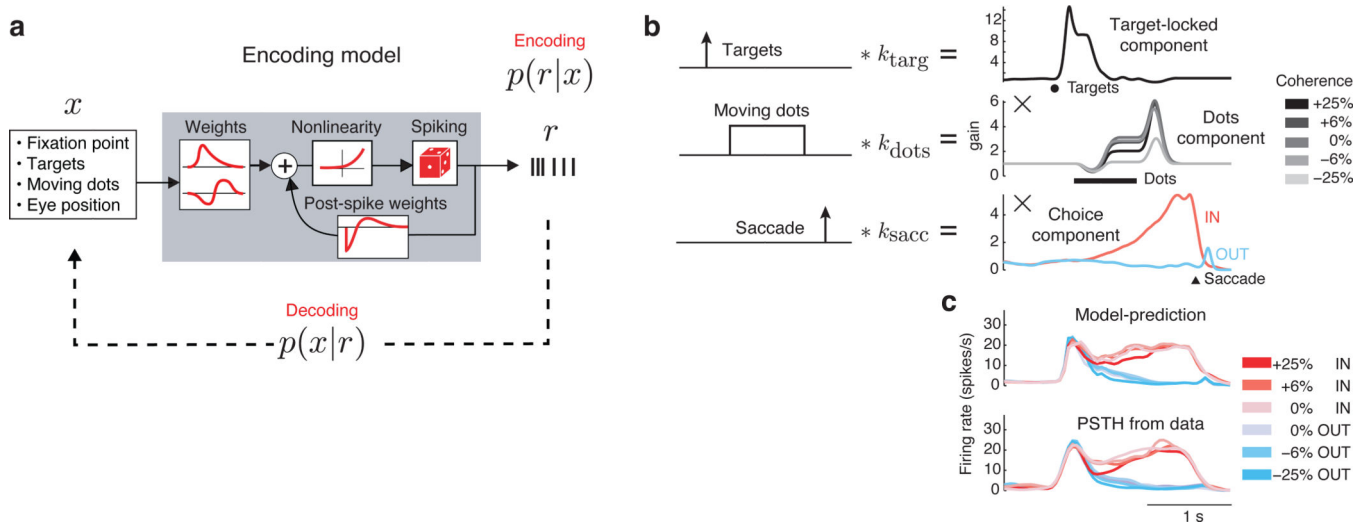

Figure 2. Encoding model, fitted components, and rate predictions

(a) The model specifies the probabilistic relationship $p(\boldsymbol{r}|\boldsymbol{x})$ between external variables $\boldsymbol{x}$ and a set of neural spike trains $\boldsymbol{r}$. The parameters consist of linear weights specifying the neuron's dependence on the external variables and spike history. **(b)** Illustration of fitted model components for an example LIP neuron. Because the nonlinearity is exponential, we can plot the exponentiated output of each filter as a gain signal reflecting the influence of each task element on the time-varying spike rate. These signals are combined multiplicatively to obtain the instantaneous spike rate, which drives spiking via a conditionally Poisson process with feedback (8 out of 12 total kernels shown; see Methods, Supplementary Figs. 1 and 2 for more detail). **(c)** Peri-stimulus time histograms (PSTHs) predicted by the model (above) and computed from real data (below). Each trace reflects a different coherence level and saccade direction.

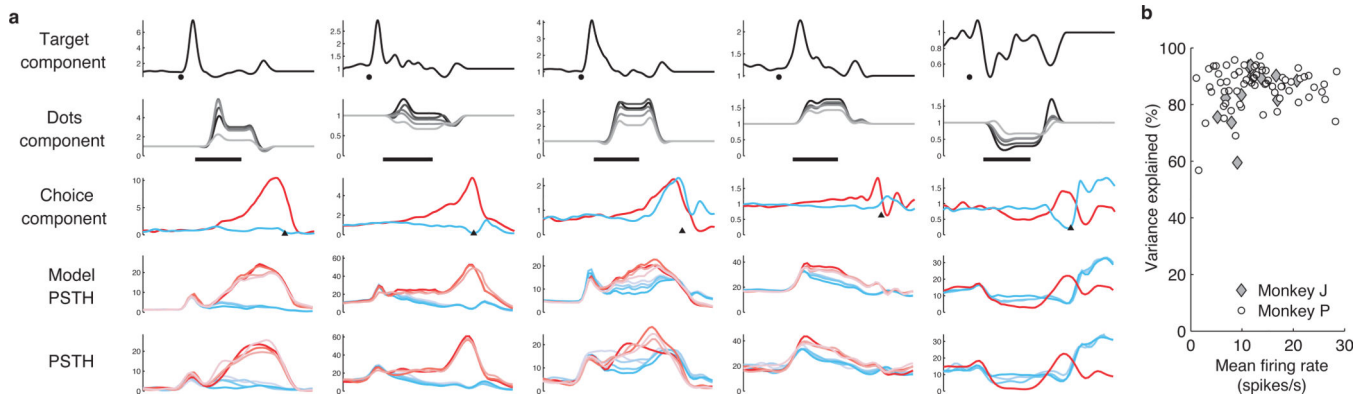

Figure 3. Encoding model captures responses across a heterogeneous population

(a) Fitted model components for five additional neurons (columns 1–5), showing the relative contributions of targets, moving dots, and saccades to the predicted spike rate of each neuron (rows 1–3) and predicted and observed PSTH (rows 4–5) as in Fig. 2. The population exhibits substantial heterogeneity across neurons: e.g., the fourth neuron's PSTH peaks early and then declines, regardless of choice (column 4, bottom row), yet the model still extracts a classical "ramping" choice-related component (3rd row). More unusually, the fifth neuron's choice-dependent component (5th column, row 3) exhibits a time-dependent reversal, meaning that early in the trial, the neuron fires more spikes before saccades to the *anti-preferred* (OUT) target. **(b)** Goodness-of-fit across the population. Each point corresponds to the percent of variance accounted for in the PSTH of each neuron.

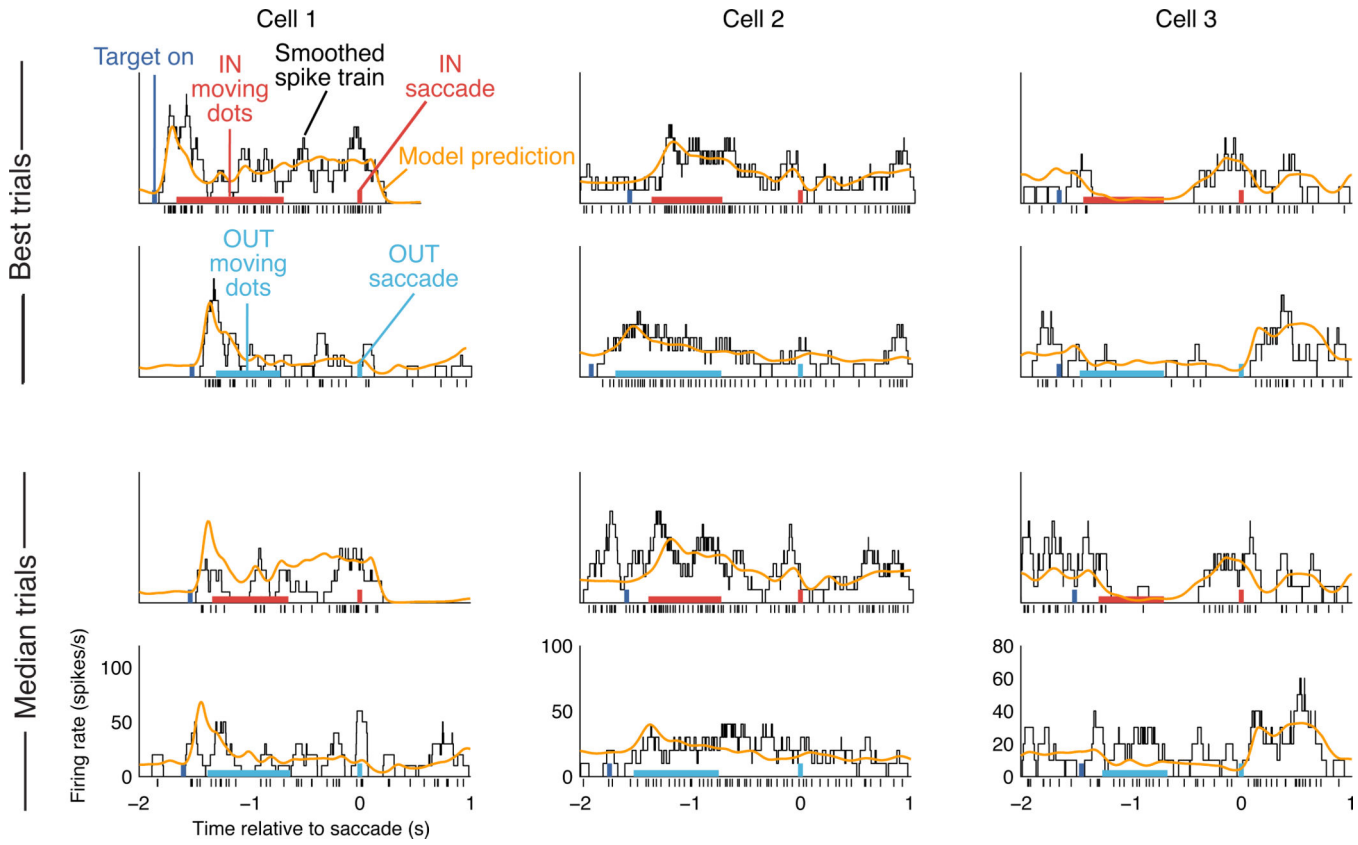

Figure 4. Model-based single trial predictions

Each column shows 4 example trials from a single neuron. Black ticks at the bottom indicate the observed spike train, and the black trace shows an estimate of spike rate obtained by smoothing the spike train with a 100 ms boxcar. Colored rectangles show the timing and identity of task elements for the trial in question (targets, moving dots, and saccade), and the orange trace shows the model-based spike rate prediction for that trial. The top two rows show example cross-validation trials for which prediction accuracy was highest, while the bottom two rows show example trials with median prediction accuracy. More trials are shown in Supplementary Fig. 3.

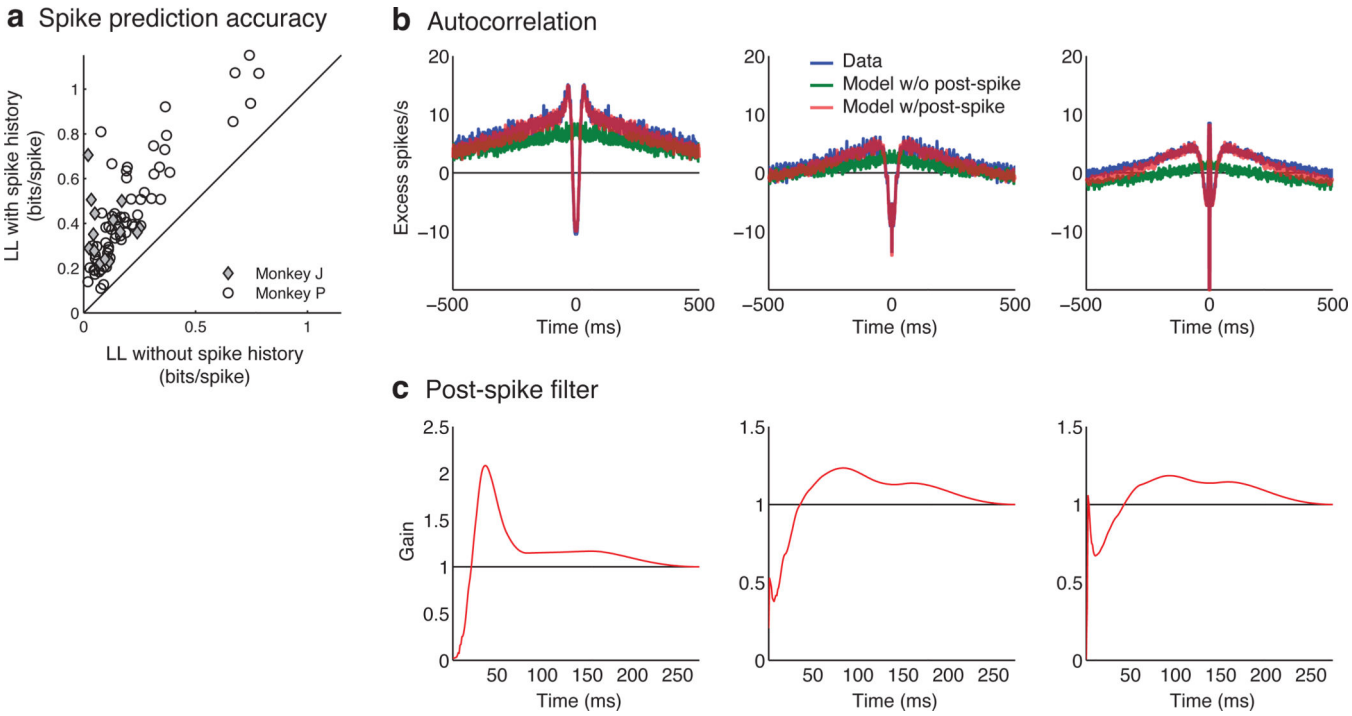

Figure 5. Spike history effects captured by the post-spike filter

(a) Spike train prediction accuracy, quantified by the log-likelihood per spike under the model fit with and without spike history (cross-validation data). Positive values indicate how much better each model is compared to a baseline model (homogeneous Poisson). The model with post-spike filters captures an average of 137% more information about single-trial spike trains. **(b)** Spike train autocorrelations computed on neural data (blue), and on spike trains simulated from the model fit without (green) and with post-spike filter (red). The model with a post-spike filter almost perfectly captures the empirical autocorrelation function (red curve overlaps blue curve). Example cells chosen to illustrate the diversity of history effects. **(c)** Fitted post-spike filters for the three neurons shown in **b**. These filters express the multiplicative change in spike probability as a function of time since a spike. These effects accumulate across multiple spikes, allowing the model to capture the complex temporal structure observed in the autocorrelations in **b**. See Supplementary Fig. 2 for diversity in population.

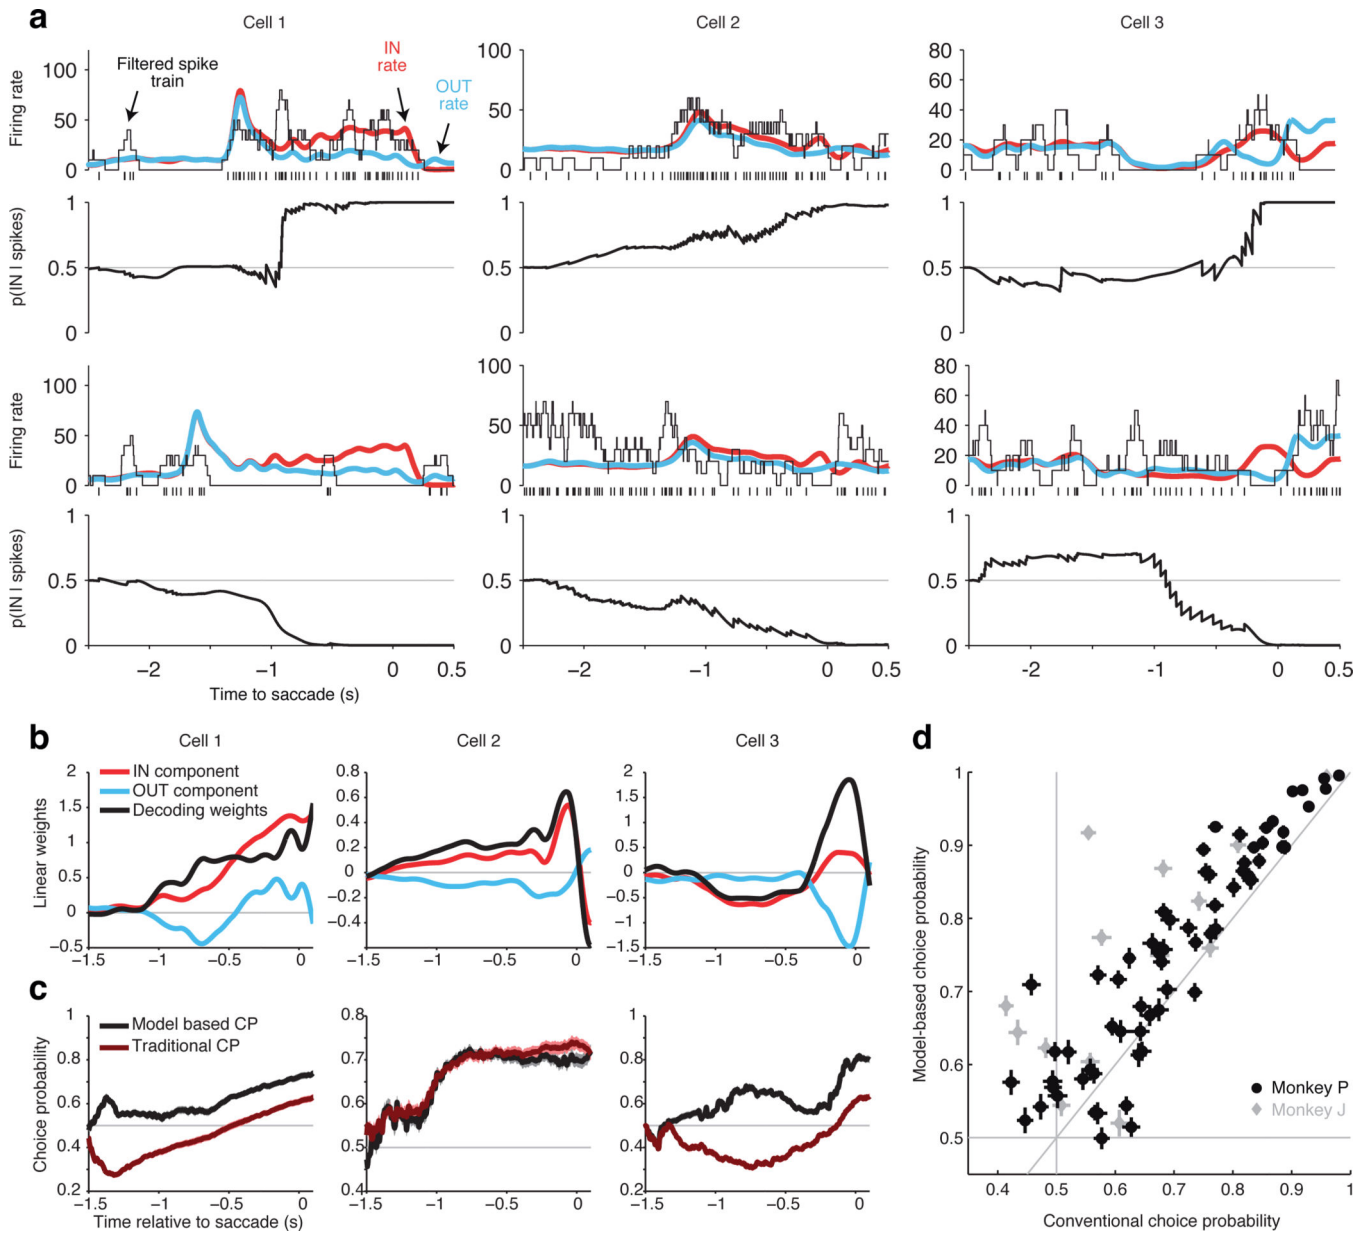

Figure 6. Single-trial prediction and model-based decoding of spike trains

(a) Single trial responses from three neurons, for one randomly selected IN (top row) and OUT trial (bottom row) spanning a range of coherences. For each trial, the model predicts time-varying spike rate both under IN (red) and OUT (blue) choices. Model-based decoding amounts to assessing whether the observed spikes on a single trial (black ticks) are more likely generated from IN or the OUT spike rate. (Thin black traces show estimated rate by filtering spikes with a 100 ms boxcar, for visualization purposes only). Thick black traces (below) indicate the probability that the monkey will saccade IN given the spikes observed so far in the trial, as computed under the model. When the model predicts a higher rate for the IN choice, every spike evokes a jump in the probability of an IN choice, whereas silence evokes a gradual decrease. **(b)** Red and blue traces show saccade-related model components

for movements IN or OUT of the neuron's RF. The difference of these components gives linear decoding weights (black), which determine how spikes from each neuron should be integrated in order to optimally predict the saccade direction. **(c)** Traditional and model-based choice probability (CP), as a function of time interval used to integrate spikes, for zero-coherence trials (cross-validation data). Traditional CP (red) weights all spikes in the decoding window equally, while model-based CP (black) uses the time-dependent weights shown in **b**. Model-based CP exceeds or matches traditional CP for nearly all time windows, and accurately deciphers the peculiar time-varying decision signal from Cell 3 (shown previously in Fig. 4, column 5) using weights that change sign midway through the trial. **(d)** Comparison of traditional and model-based choice probability for all neurons in the population, using spikes at least 50 ms prior to saccade on zero-coherence trials. Vertical and horizontal lines show SD bootstrap confidence intervals. For the majority of neurons in the population, model-based decoding supports more accurate readout of the animal's decisions from LIP spike trains.

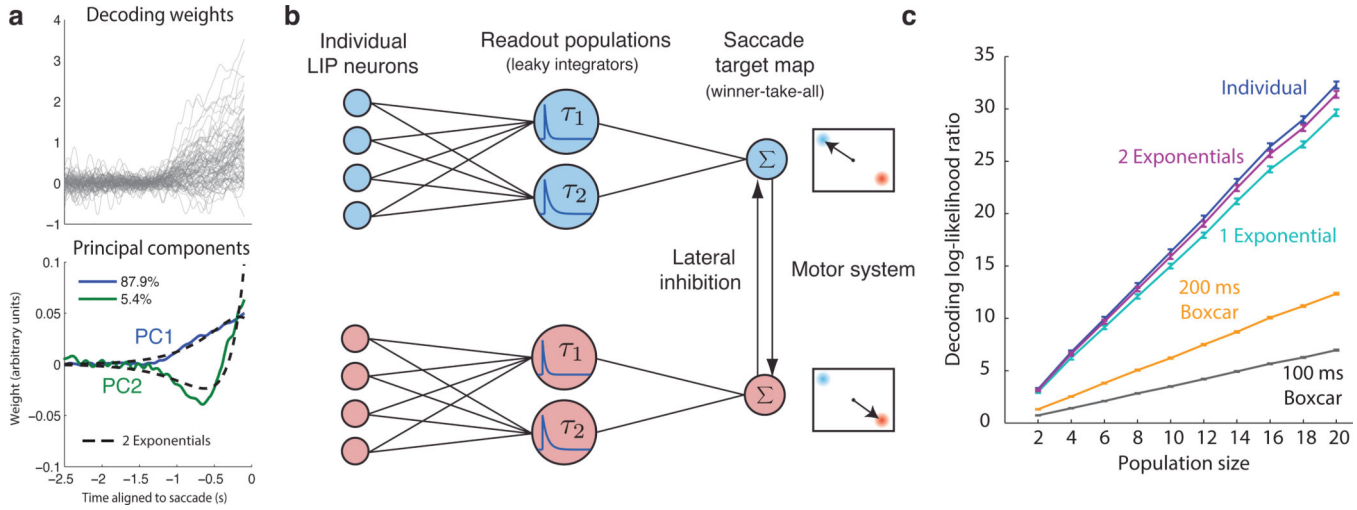

Figure 7. Decoding from low dimensional decoding weight space

(a) Top, decoding weights from all cells. Bottom, corresponding principal components (PC) that span the weights (numbers indicates variance explained). Dashed lines show the 2 PCs approximated with 2 optimized exponential functions (τ = 237,410 ms). For more details, see Supplementary Fig. 9 and Supplementary Table 1. **(b)** Schematic for decoding LIP neurons. Two populations with opposing RFs have corresponding readout populations. Each neuron projects to a subpopulation that filters their spiking activity with first order dynamics. The population response is averaged within the two pools of neurons that compete with each other to control the saccade. **(c)** Decoding performance comparisons on zero-coherence trials assuming neuron-antineuron pools. Higher log-likelihood ratio implies easier distinction between the choices given a pair of trials.

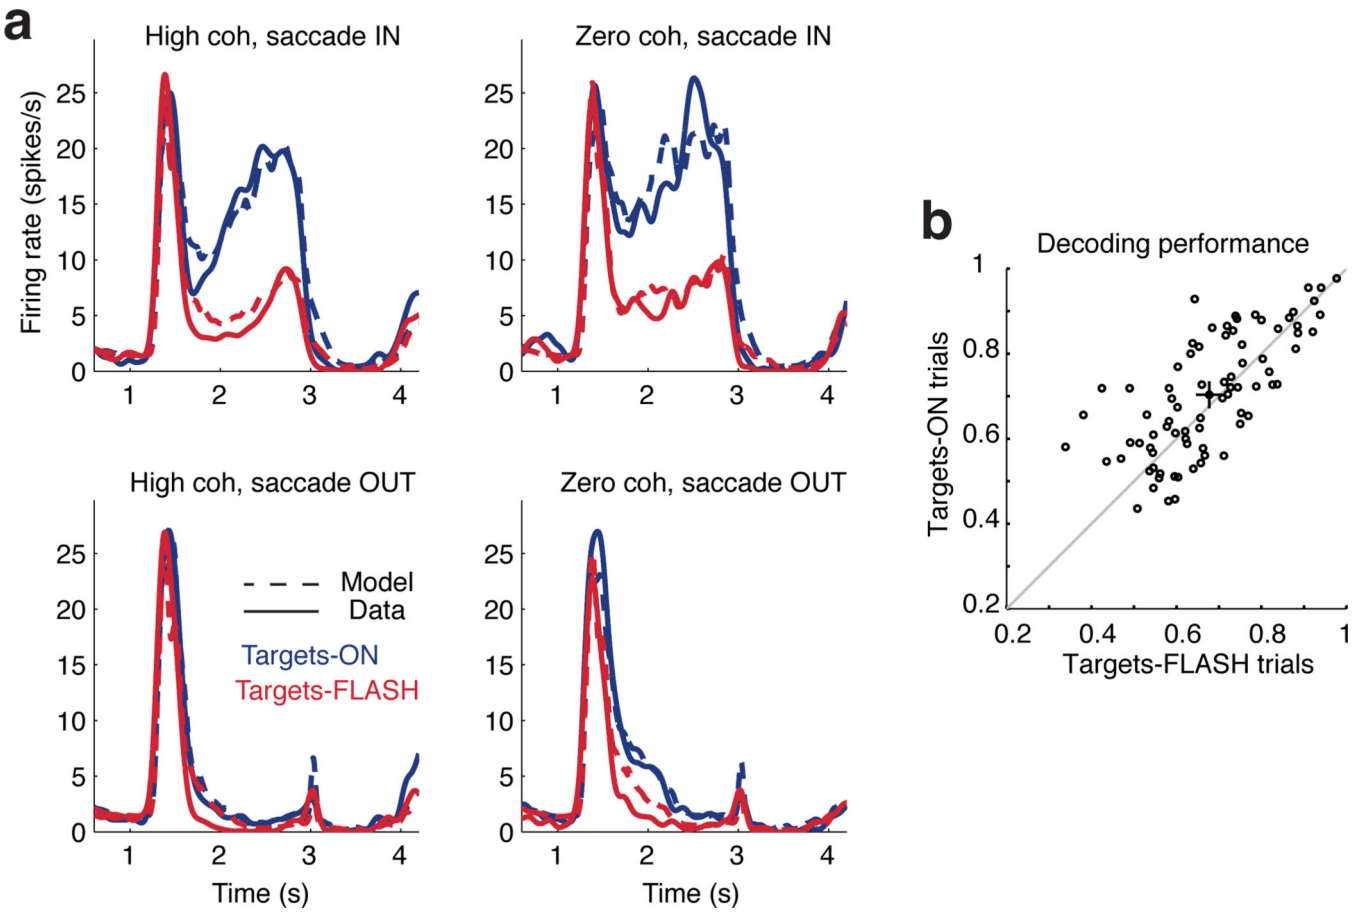

Figure 8. Targets absent manipulation

(a) On half of the trials, the targets disappeared early (Targets-FLASH condition), which often resulted in large changes in firing rate (blue vs. red curves). The encoding model captures this effect well with a single kernel aligned to target offset. The four sub-panels show high (left), zero (right) coherence conditions for IN (top) and OUT (bottom) eventual saccades. **(b)** Decoding performance of zero-coherence trials with and without targets (80 neurons). Using the same decoding kernel, the trials can be decoded almost as well. Filled circle indicates the population average with standard errors. Decoding targets-ON trial is 2.6% easier (p-value 0.003; paired t-test), a significant but small effect.