Encoding and decoding in parietal cortex during sensorimotor decision-making

Il Memming Park[1–3], Miriam L R Meister[3,4], Alexander C Huk[1–3,5] & Jonathan W Pillow[1–3,6]

It has been suggested that the lateral intraparietal area (LIP) of macaques plays a fundamental role in sensorimotor decision-making. We examined the neural code in LIP at the level of individual spike trains using a statistical approach based on generalized linear models. We found that LIP responses reflected 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 derived an optimal decoder for heterogeneous, multiplexed LIP responses that could be implemented in biologically plausible circuits. In contrast with interpretations of LIP as providing an instantaneous code for decision variables, we found that optimal decoding requires integrating LIP spikes over two distinct timescales. These analyses provide a detailed understanding of neural representations in LIP and a framework for studying the coding of multiplexed signals in higher brain areas.

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 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 the dynamics governing neural activity in decision-making circuits[9–15]. Although these studies 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.

We developed 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 that characterizes 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 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 is that, in contrast with primary sensory or motor areas, an area such as 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. The model reveals that LIP responses are best described as reflecting an interacting set of temporally overlapping response components. This implies that some of the spikes emitted during decision formation are potentially related to multiple factors, some of which are irrelevant to the accumulation of evidence. The model also provides a framework for understanding the statistically optimal readout of various kinds of information from single-trial spike trains. Our analyses reveal 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 therefore both identifies and resolves a puzzle: LIP spikes do not purely encode the accumulation of evidence during the formation of decisions, but they can be feasibly decoded to extract more choice-related information than conventional spike-counting analyses have suggested.

Generally, these analyses provide a detailed account of the time-varying information carried by LIP spikes and operate at the level of single trials. Thus, they provide a testing ground for interpreting 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.

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.

[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]Institute for Neuroscience, The University of Texas at Austin, Austin, Texas, USA. [4]Department of Physiology and Biophysics, University of Washington, Seattle, Washington, USA. [5]Department of Neuroscience, The University of Texas at Austin, Austin, Texas, USA. [6]Department of Statistics and Data Science, The University of Texas at Austin, Austin, Texas, USA. Correspondence should be addressed to J.P. (pillow@mail.utexas.edu).

Received 30 April; accepted 29 July; published online 31 August 2014; doi:10.1038/nn.3800

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). **Figure 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 (Online Methods).

Classical analyses of coding in LIP begin with 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 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. It quantifies 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 influence of spike history on subsequent spikes, and a nonlinearity associated with spike generation. It is therefore possible to dissociate firing rate components associated with each task variable given their decorrelated design (for example, trial-to-trial variability in the relative timings of events and independent variation of motion coherence) while also capturing the neuron's own temporal response properties (for example, refractoriness, burstiness, longer timescale 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

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

where $(\boldsymbol{k}_i * x_i)(t)$ denotes linear convolution of $x_i(t)$, the time course for the ith external event (for example, the target kernel is 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 $\boldsymbol{r}^{\text{hist}}$ with the post-spike filter $\boldsymbol{h}$. We illustrate these components in a model diagram shown in **Figure 2a**.

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

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

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.

Figure 2 illustrates the model fit to data from a typical LIP neuron. To provide intuition for the model's basic capabilities, we highlight the kernels related to the three primary task elements that occur on each trial: the appearance of the choice targets, the moving dots stimulus and 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). We plot the predicted time-varying change in spike rate resulting from each task element for each of five possible motion coherences and two possible saccade directions (**Fig. 2b**). For each component, this change is given by the linear convolution of the kernel with the corresponding task element and then passed through an exponential nonlinearity. The product of three such components forms the predicted spike rate for a single trial (equation (1)). These predictions closely match the neuron's actual PSTH (**Fig. 2c**).

Randomness and variability in the timings of experimental events are essential 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 onsets 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 bases and applied Bayesian regularization methods to prevent over-fitting (Online Methods). We verified the fits via PSTH prediction (**Fig. 3b**), single trial prediction on a test set

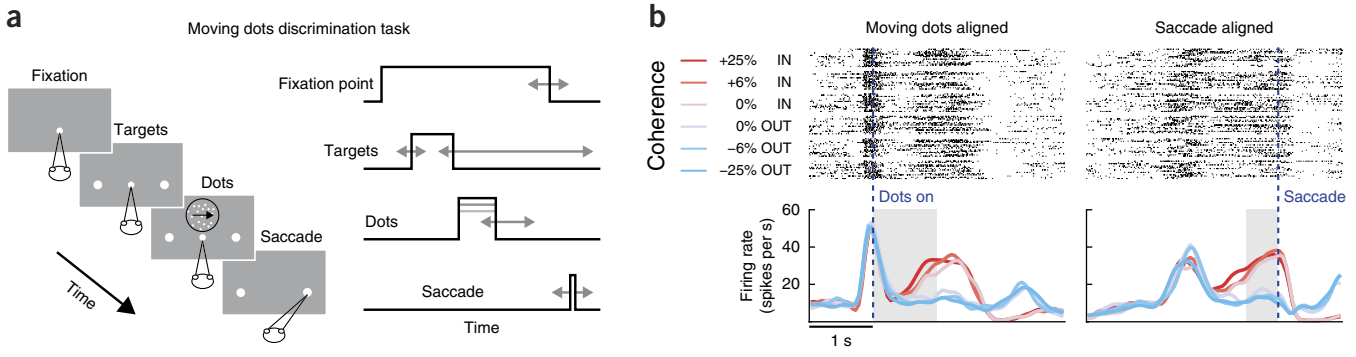

Figure 1 Decision-making task and classical analysis of LIP responses. (**a**) The task requires an observer to fixate, judge the direction of a moving (dots) stimulus and report a decision by moving the eyes to one of two targets. Temporal variability in the task design variables (indicated by gray 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 grouped by the stimulus coherence (the fraction of dots moving toward (+) or away from (–) the neuron's response field), and the saccade into (IN) or opposite (OUT) the response field. These are aligned to the stimulus onset (left) or saccade time (right). Although the spike trains exhibit substantial variability (top), their average time courses (below) exhibits coherence-dependent ramping. The gray area indicates the short portions of trials often considered in prior work[33].

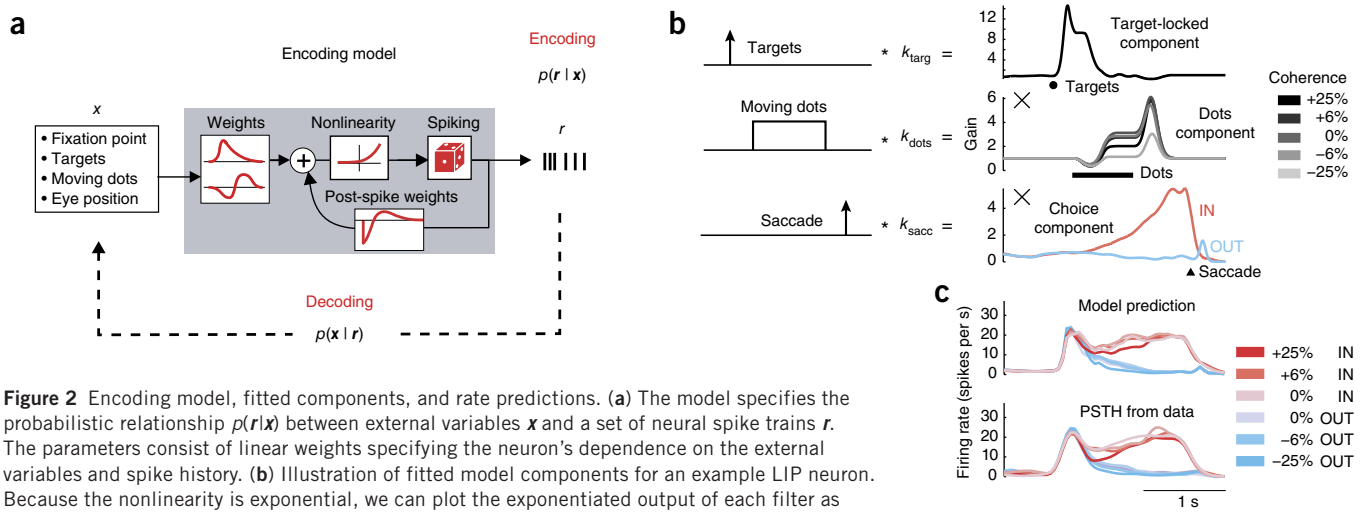

Figure 2 Encoding model, fitted components, and rate predictions. (**a**) The model specifies the probabilistic relationship $p(\boldsymbol{r}|\boldsymbol{x})$ between external variables $\boldsymbol{x}$ and a set of neural spike trains $\boldsymbol{r}$. The parameters consist of linear weights specifying 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 of 12 total kernels shown; Online Methods and **Supplementary Figs. 6** and **7**). (**c**) PSTHs predicted by the model (above) and computed from real data (below). Each trace reflects a different coherence level and saccade direction.

(**Figs. 4** and **5a**), time-varying spike count variance (**Supplementary Fig. 1**) and interspike interval statistics (**Supplementary Fig. 2**).

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 data set and found that it captured 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 Figs. 3** and **4**). Furthermore, individual cells encode these elements in distinct ways, both in terms of overall magnitude and in nuanced aspects of the time course. This cell-by-cell analysis suggests that the marked differences in PSTHs may arise from the combination of heterogeneous task-related components that can now be examined in isolation at the level of single neurons.

In addition to capturing the average time course 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 accurately, 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 model. A model without a post-spike filter cannot account for the detailed shape of these autocorrelation functions, whereas 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. For these neurons, the probability of firing was enhanced for more than 200 ms after a

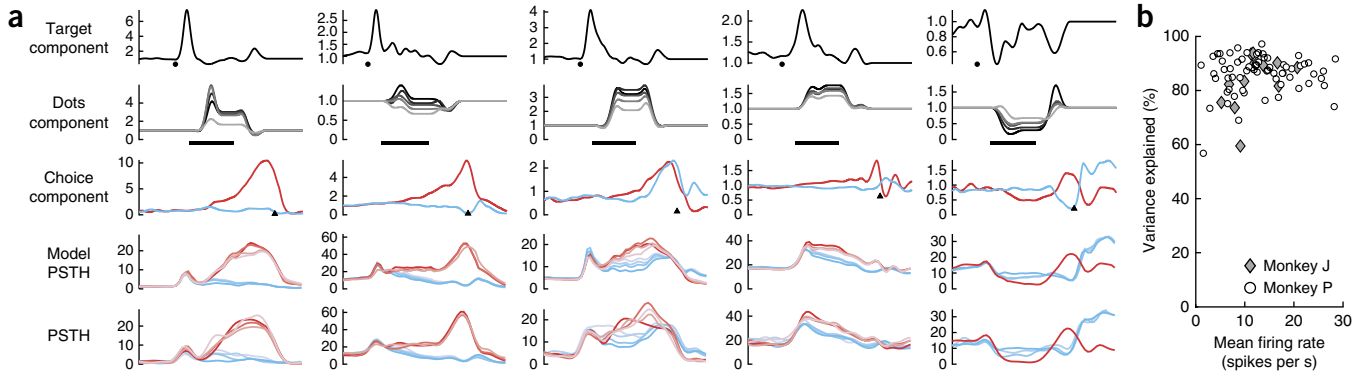

Figure 3 Encoding model captures responses across a heterogeneous population. (**a**) Fitted model components for five additional neurons (columns 1–5) showing the relative contributions of targets, moving dots and saccades to the predicted spike rate of each neuron (rows 1–3) and predicted and observed PSTH (rows 4 and 5) as in **Figure 2**. The population exhibited substantial heterogeneity across neurons; for example, the fourth neuron's PSTH peaked early and then declined, regardless of choice (column 4, bottom row), yet the model still extracted a classical 'ramping' choice-related component (third row). More unusually, the fifth neuron's choice-dependent component (fifth column, row 3) exhibited a time-dependent reversal, meaning that, early in the trial, the neuron fired 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.

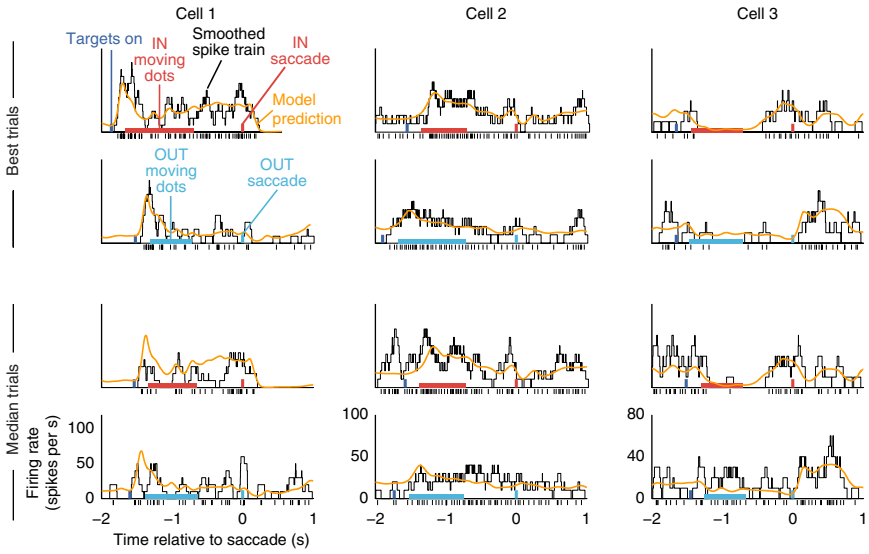

Figure 4 Model-based single trial predictions. Each column shows four example trials from a single neuron. Black ticks at the bottom indicate the observed spike train and the black trace shows an estimate of spike rate obtained by smoothing the spike train with a 100-ms boxcar. Colored rectangles show 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 and the bottom two rows show example trials with median prediction accuracy. More trials are shown in **Supplementary Figure 8**.

spike (but not much longer; **Supplementary Fig. 5**). Note that such a long timescale of self-excitation is a purely statistical description (as opposed to a proposed recurrent mechanism) and reflects a form of temporal integration above and beyond that explained by the external task elements. Despite its contribution to spike train prediction, inclusion of the post-spike filter in the model exerts only a modest scaling effect on the other temporal kernels (the median correlation coefficient between the kernels fit with and without the post-spike filter was 96%, and the median scaling effect was 75%).

We validated the use of an exponential nonlinearity by comparing it to a nonparametric estimate of the nonlinear relationship between filter output and firing rate (**Supplementary Fig. 3a**). This implies a multiplicative interaction among components (but see **Supplementary Figs. 3b**,**c** and **4** for comparison to rectified linear function) and is critical for the statistically optimal linear decoding mechanism we describe below.

Together, the distinct pieces of the encoding model—linear kernels for task variables, a post-spike filter and an approximately exponential nonlinearity—jointly 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, simple task- and behavior-related signals may explain the bulk of LIP responses. 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.

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, but they also provide a powerful tool for analyzing the readout of information from spike trains. In this 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 (that is, saccades to one target or the other). For each trial, the model provides a prediction in the form of a Poisson spike rate underlying 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 OUT of an LIP neuron's response field.

Figure 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. Note that these probabilistic estimates diverge from 0.5 (the prior probability of an IN decision) as soon as the rate predictions diverge.

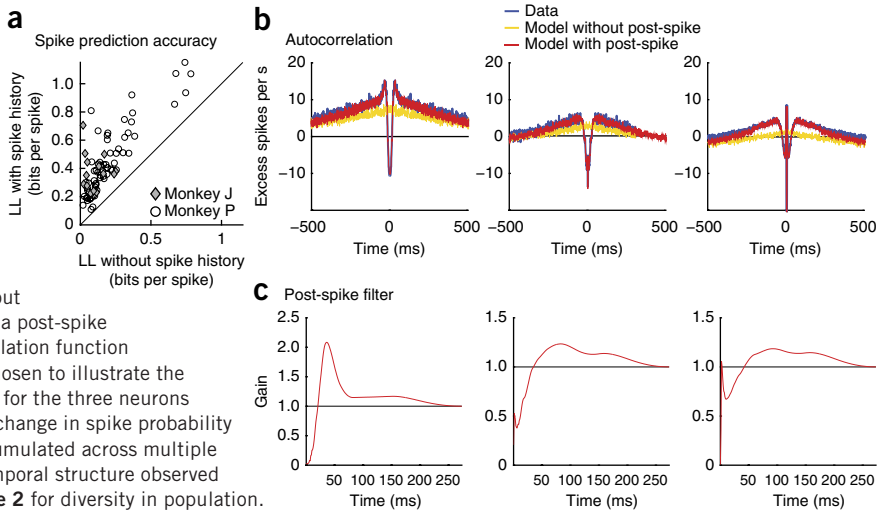

Figure 5 Spike history effects captured by the post-spike filter. (**a**) Spike train prediction accuracy, quantified by the log-likelihood per spike under the model fit with and without spike history (cross-validation data). Positive values indicate how much better each model is than a baseline model (homogeneous Poisson). The model with post-spike filters captured 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 (gold) and with post-spike filter (red). The model with a post-spike filter almost perfectly captured the empirical autocorrelation function (red curve overlaps blue curve). Example cells were chosen to illustrate the diversity of history effects. (**c**) Fitted post-spike filters for the three neurons shown in **b**. These filters expressed the multiplicative change in spike probability as a function of time since a spike. These effects accumulated across multiple spikes, allowing the model to capture the complex temporal structure observed in the autocorrelations in **b**. See **Supplementary Figure 2** for diversity in population.

For the six trials shown, the decoder achieved near certainty (probability of an IN choice close to one or zero) by the time of saccade, although this did 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

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

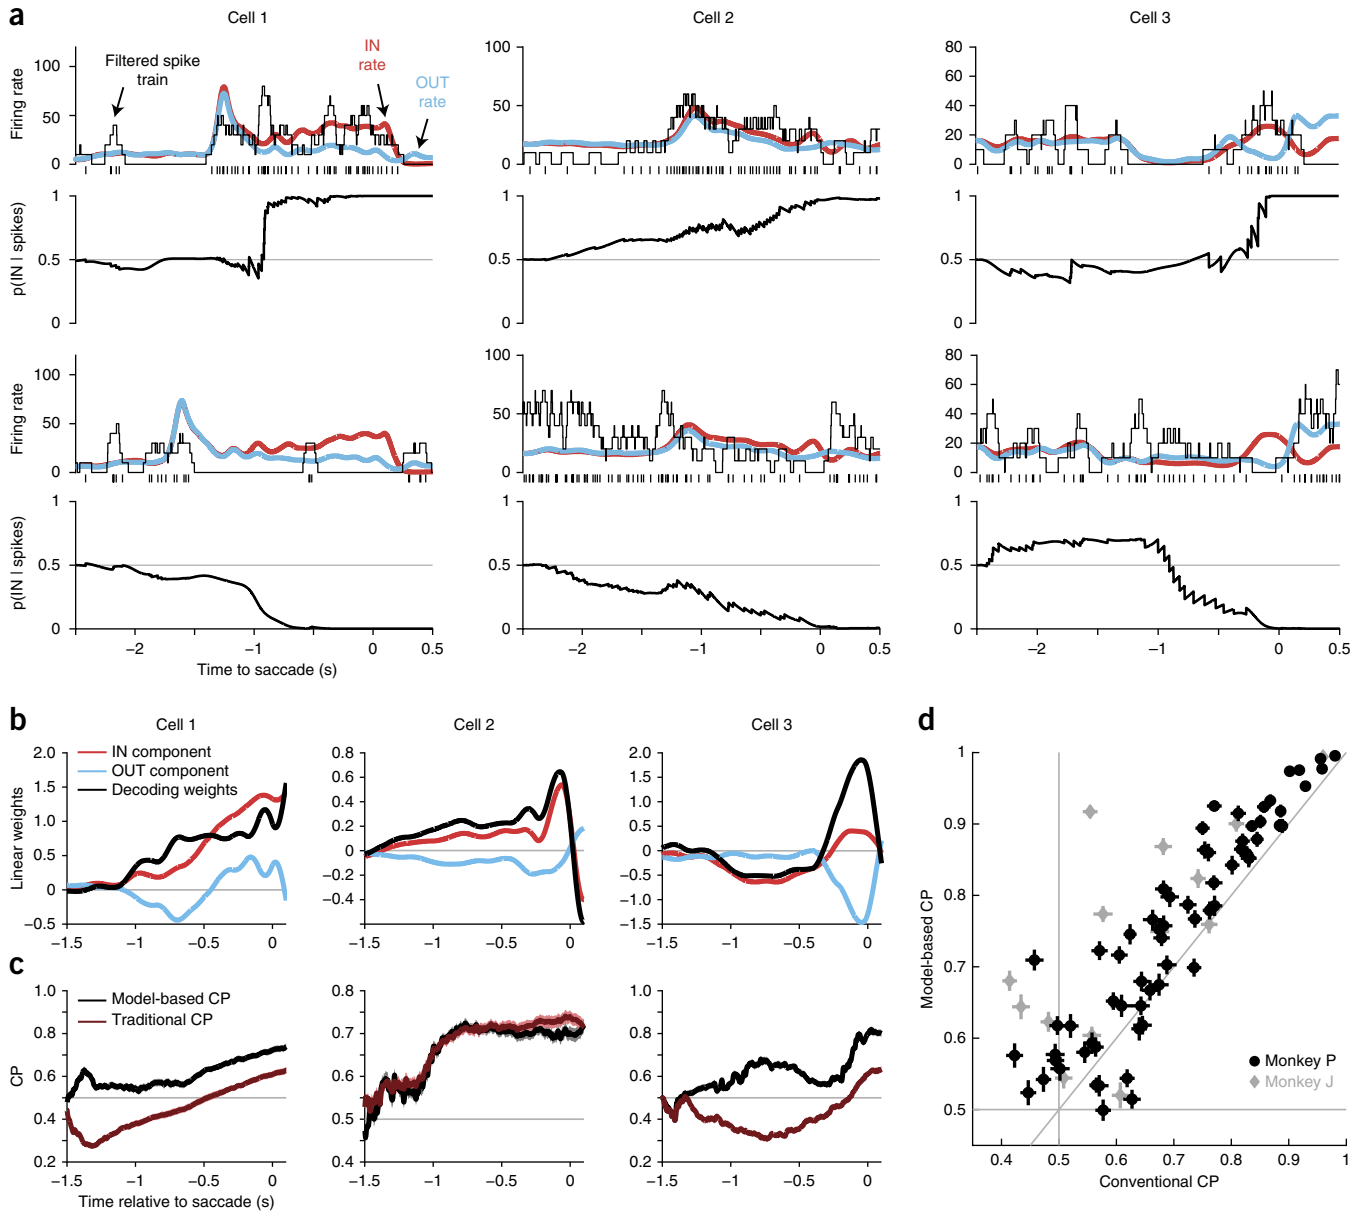

Figure 6 Single-trial prediction and model-based decoding of spike trains. (**a**) Single-trial responses from three neurons for one randomly selected IN (top row) and OUT trial (bottom row) spanning a range of coherences. For each trial, the model predicts time-varying spike rate both under IN (red) and OUT (blue) choices. Model-based decoding amounts to assessing whether the observed spikes on a single trial (black ticks) were more likely generated from the 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 thus 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 to optimally predict the saccade direction. (**c**) Traditional and model-based 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, whereas model-based CP (black) uses the time-dependent weights shown in **b**. Model-based CP exceeded or matched traditional CP for nearly all time windows and accurately deciphered the peculiar time-varying decision signal from Cell 3 (shown in **Fig. 4**, column 5) using weights that change sign midway through the trial. (**d**) Comparison of traditional and model-based CP for all neurons in the population, using spikes at least 50 ms before saccade on zero-coherence trials. Vertical and horizontal lines show ±s.d. 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.

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

$$\mathrm{LLR} = \log \frac{p(x_c = \mathrm{IN}\,|\,\boldsymbol{r},\tilde{\boldsymbol{x}})}{p(x_c = \mathrm{OUT}\,|\,\boldsymbol{r},\tilde{\boldsymbol{x}})} = \boldsymbol{w}^{\mathrm{T}}\boldsymbol{r} + \text{constant} \qquad (4)$$

where $\boldsymbol{w}$ denotes a set of linear decoding weights over time, and the additive constant does not depend on the response $\boldsymbol{r}$ (Online 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 (**Fig. 6b**).

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 eye movement. The weights determined empirically from data suggest that instantaneous spike rate in LIP (or a race between competing pools in LIP[9]) does not fully characterize the amount of decision-related information in the neural response, as the weights 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, as our model formally belongs to the family of log-linear models known as probabilistic population codes[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[39]. Conventional CP, which applies to scalar quantities such as spike count, assumes that the optimal rule for reading out a pair of neuron and 'anti-neuron' responses is to choose the preferred stimulus of the neuron with the larger response (a 'max' decoding rule). Model-based CP, on the other hand, can be defined using any model-based decoding rule and non-scalar representations of neural activity such as a binned spike train. If we assume an anti-neuron for the neuron that we are decoding to have exactly the same model weights, except with opposite 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, as the threshold is exactly the same for both neurons. We can therefore compute model-based CP using projections of spike trains instead of simple spike counts.

The key difference between conventional and model-based CP is that model weights specify the relative importance of different time bins for reading out the choice. **Figure 6b–d** shows a comparison of model-based and conventional CP for quantifying decoding performance. Model-based CP outperforms conventional CP as the time window grows toward the saccade (**Fig. 6c**). Moreover, model-based CP increases almost monotonically, whereas conventional CP remains closer to chance levels. Conventional CP is clearly inadequate for decoding LIP neurons such as example cell 3, where the polarity of rate associated with IN and OUT choices reverse during the trial. Even for the more canonical example cell 1, model-based decoding extracts substantially more decision-related information than the conventional CP. Model-based CP was, on average, 6.2% higher than conventional CP (mean model-based CP, 75.2%; mean traditional CP, 69.9%). **Figure 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 apply to neurally plausible mechanisms. At first glance, statistically optimal readout seems difficult, as each neuron has a unique decoding-weight profile and these weights exhibit substantial 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 that 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 moment in time, the output of the linear filter is the LLR (equation (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—we can approximately decompose each decoding weight as a linear combination of a few principal components of the collection of decoding weights, and then

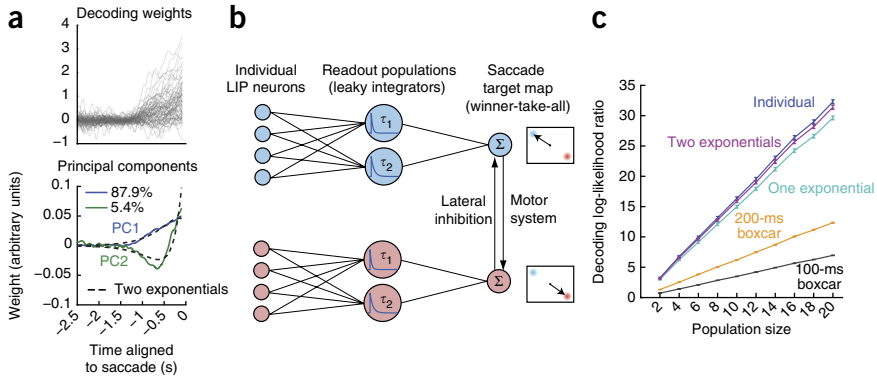

Figure 7 Decoding from low-dimensional decoding weight space. (**a**) Top, decoding weights from all cells. Bottom, corresponding principal components (PCs) that span the weights (percentages indicates variance explained). Dashed lines show the two PCs approximated with two optimized exponential functions (τ = 237,410 ms). For more details, see **Supplementary Figure 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 its spiking activity with first-order dynamics. The population response is averaged in the two pools of neurons that compete with each other to control the saccade. (**c**) Decoding performance comparisons on zero-coherence trials assuming neuron–anti-neuron pools. Higher LLR implies easier distinction between the choices given a pair of trials.

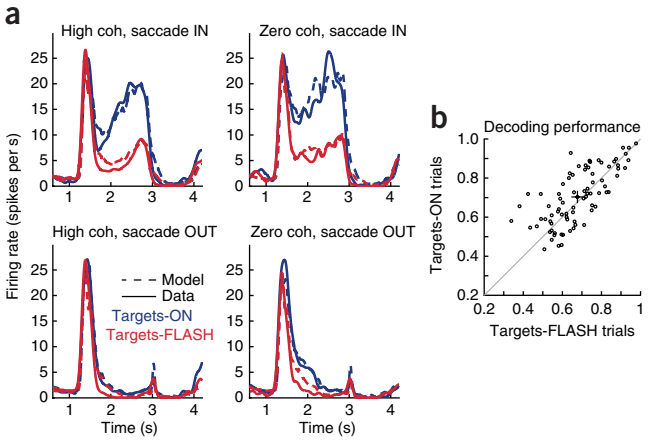

Figure 8 Targets absent manipulation. (**a**) On half of the trials, the targets disappeared early (targets-FLASH condition), which often resulted in large changes in firing rate (blue versus red curves). The encoding model captures this effect well with a single kernel aligned to target offset. The four panels show high (left) and 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 ($P = 0.003$, paired t test), a significant, but small, effect.

sum the results (**Fig. 7a**). Indeed, we found that the constellation of decoding weights are well described by two principal components (explaining 93.3% of the variance). To represent the total log-probability of a decision given the spikes from a population of LIP neurons, the filter outputs corresponding to each RF can simply be summed. Third, we found 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, although 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[14]. We found that two exponential filters, $e^{(-t/\tau)}$, with $\tau = 237$ and 410 ms, explain 92.1% of the variance of the seemingly diverse decoding weights (**Fig. 7a**). Finally, we posit two competing populations, each integrating evidence for one of the choice alternatives (that is, two opposite choice target locations), with balanced contributions from the other response components, so as to nullify the effects of other task variables.

Figure 7b shows a schematic of the proposed plausible decoding circuit: LIP neurons associated with the same choice project with appropriate weights to a pair of subpopulations functioning as leaky integrators, each with a different time constant. The instantaneous log-probability of a choice can be obtained by summing activity of the two subpopulations. Finally, the preferred target is read out through a winner-take-all mechanism and transmitted to the motor neurons that generate 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 counting spikes in a fixed time window. Moreover, decoding with two leaky integrators performs virtually as well as with the complex individual decoding weights.

In summary, this decoding scheme offers one plausible way to read out decisions from the population activity of LIP neurons. It characterizes how decision-related information is represented in population spiking activity in LIP and suggests that statistically optimal decoding could be implemented by low-dimensional dynamics in biologically plausible circuits, even in the face of multiplexed signals and considerable heterogeneity across neurons.

Invariance to target-induced changes in firing

The decoding mechanism we have proposed achieves performance that is invariant to firing rate modulations induced by variables not directly relevant to the task. To illustrate this invariance, we manipulated the visual saccade targets so that on half of the trials they persisted until the end of the trial (targets-ON trials), whereas 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 motivates a substantial refinement to previously described models that regard the firing rate of an LIP neuron as a direct neural correlate of LLR or some other normative quantity[5]. Our descriptive model, however, captures the effects of visual target removal through a multiplicative interaction with an additional target-related kernel (although saccade endpoints were affected by the targets-FLASH manipulation, the single target–based kernel is sufficient to account for the changes in LIP response; also see ref. 40).

Figure 8a shows an example neuron that reduced its response markedly when visual targets are extinguished early in the trial (that is, 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 suggests 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 developed a GLM to describe LIP responses during a complex perceptual decision-making task. This framework yielded several new insights into the coding of information in LIP spike trains. First, we found that the multiplexed representations in LIP 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 exhibited notably long-timescale self-excitation that was 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 was well captured by a two-dimensional linear projection of the spike trains, which led to an implementation of statistically optimal decoding using a pair of leaky integrators with distinct timescales.

Our model allows the spike rate at any time to depend on a combination of multiple task components. We exploited the randomized structure of the experimental design to identify these dependencies and showed that LIP activity could 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 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 and often appeals to a race between competing integrators to generate decisions[9]. 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 timescale 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 in 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 analyses suggest 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 reveal that it is still possible to read out decisions in the face of this multiplexing, with higher accuracy than indicated by conventional CP. This perspective raises the possibility that the difficulties in distinguishing between various functional roles of LIP (for example, attention/salience, decision-making, motor intention) may not have arisen because LIP has a 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 task variants may therefore 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 determining how these complex signals are demultiplexed 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 in a regression framework, and therefore involve far less conceptual baggage and a tighter reliance on the data than normative approaches. 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. In addition, prior applications of GLMs 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 could 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 explored 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 that we propose could complement these models by providing a principled technique for assessing which components to incorporate; these models may therefore provide a tool not just for interpreting LIP activity, but for linking normative and mechanistic models of the sorts that have already been 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 study 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 toward resolving 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 GLM. However, we feel that the framework promises both to enrich interpretations in well-studied procedures and to pave the way for more ambitious and direct tests of hypotheses about higher brain function.

METHODS

Methods and any associated references are available in the online version of the paper.

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

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 grants from the National Eye Institute (EY017366 to A.C.H.) and the National Institute of Mental Health (MH099611 to J.W.P. and A.C.H.), by the Sloan Foundation (J.W.P.), McKnight Foundation (J.W.P.), and a National Science Foundation CAREER award (IIS-1150186 to J.W.P.).

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. A.C.H., I.M.P., M.L.R.M. and J.W.P. wrote the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

1. Platt, M.L. & Glimcher, P.W. Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238 (1999).
2. Shadlen, M.N. & Newsome, W.T. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* **86**, 1916–1936 (2001).

3. Huk, A.C. & Shadlen, M.N. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* **25**, 10420–10436 (2005).
4. Kable, J.W. & Glimcher, P.W. The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* **10**, 1625–1633 (2007).
5. Yang, T. & Shadlen, M.N. Probabilistic reasoning by neurons. *Nature* **447**, 1075–1080 (2007).
6. Beck, J.M. *et al.* Probabilistic population codes for Bayesian decision making. *Neuron* **60**, 1142–1152 (2008).
7. Seo, H., Barraclough, D.J. & Lee, D. Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game. *J. Neurosci.* **29**, 7278 (2009).
8. Cisek, P. & Kalaska, J.F. Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* **33**, 269–298 (2010).
9. Mazurek, M.E., Roitman, J.D., Ditterich, J. & Shadlen, M.N. A role for neural integrators in perceptual decision making. *Cereb. Cortex* **13**, 1257 (2003).
10. Lo, C.C. & Wang, X.J. Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. Neurosci.* **9**, 956–963 (2006).
11. Wong, K.F., Huk, A.C., Shadlen, M.N. & Wang, X.J. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front. Comput. Neurosci.* **1**, 6 (2007).
12. Fusi, S., Asaad, W.F., Miller, E.K. & Wang, X.-J. A neural circuit model of flexible sensorimotor mapping: learning and forgetting on multiple timescales. *Neuron* **54**, 319–333 (2007).
13. Wang, X.-J. Decision making in recurrent neuronal circuits. *Neuron* **60**, 215–234 (2008).
14. Ganguli, S. *et al.* One-dimensional dynamics of attention and decision making in LIP. *Neuron* **58**, 15–25 (2008).
15. Soltani, A. & Wang, X.J. Synaptic computation underlying probabilistic inference. *Nat. Neurosci.* **13**, 112–119 (2010).
16. Pillow, J.W., Paninski, L., Uzzell, V.J., Simoncelli, E.P. & Chichilnisky, E.J. Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *J. Neurosci.* **25**, 11003–11013 (2005).
17. Pillow, J.W. *et al.* Spatio-temporal correlations and visual signaling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
18. Jacobs, A.L. *et al.* Ruling out and ruling in neural codes. *Proc. Natl. Acad. Sci. USA* **106**, 5936–5941 (2009).
19. Fernandes, H.L., Stevenson, I.H., Phillips, A.N., Segraves, M.A. & Kording, K.P. Saliency and saccade encoding in the frontal eye field during natural scene search. *Cereb. Cortex* published online, doi:10.1093/cercor/bht179 (17 July 2013).
20. Paninski, L., Fellows, M., Shoham, S., Hatsopoulos, N. & Donoghue, J. Superlinear population encoding of dynamic hand trajectory in primary motor cortex. *J. Neurosci.* **24**, 8551–8561 (2004).
21. Truccolo, W., Eden, U.T., Fellows, M.R., Donoghue, J.P. & Brown, E.N. A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects. *J. Neurophysiol.* **93**, 1074–1089 (2005).
22. Yu, B.M. *et al.* Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J. Neurophysiol.* **102**, 614 (2009).
23. Brown, E.N., 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.* **18**, 7411–7425 (1998).
24. Barbieri, R., Wilson, M.A., Frank, L.M. & Brown, E.N. An analysis of hippocampal spatio-temporal representations using a Bayesian algorithm for neural spike train decoding. *IEEE Trans. Neural Syst. Rehabil. Eng.* **13**, 131–136 (2005).
25. Rorie, A.E., Gao, J., McClelland, J.L. & Newsome, W.T. Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS ONE* **5**, e9308 (2010).
26. Park, J. & Zhang, J. Sensorimotor locus of the buildup activity in monkey lateral intraparietal area neurons. *J. Neurophysiol.* **103**, 2664–2674 (2010).
27. Jenison, R.L., Rangel, A., Oya, H., Kawasaki, H. & Howard, M.A. Value encoding in single neurons in the human amygdala during decision making. *J. Neurosci.* **31**, 331–338 (2011).
28. Rishel, C.A., Huang, G. & Freedman, D.J. Independent category and spatial encoding in parietal cortex. *Neuron* **77**, 969–979 (2013).
29. Huk, A.C. & Meister, M.L. Neural correlates and neural computations in posterior parietal cortex during perceptual decision-making. *Front. Integr. Neurosci.* **6**, 86 (2012).
30. Gottlieb, J. & Goldberg, M.E. Activity of neurons in the lateral intraparietal area of the monkey during an antisaccade task. *Nat. Neurosci.* **2**, 906–912 (1999).
31. Meister, M.L.R., Hennig, J.A. & Huk, A.C. Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. *J. Neurosci.* **33**, 2254–2267 (2013).
32. Newsome, W.T. & Pare, E.B. A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J. Neurosci.* **8**, 2201–2211 (1988).
33. Roitman, J.D. & Shadlen, M.N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475 (2002).
34. Sugrue, L.P., Corrado, G.S. & Newsome, W.T. Matching behavior and the representation of value in the parietal cortex. *Science* **304**, 1782–1787 (2004).
35. Foley, N.C., Jangraw, D.C., Peck, C. & Gottlieb, J. Novelty enhances visual salience independently of reward in the parietal lobe. *J. Neurosci.* **34**, 7947–7957 (2014).
36. Premereur, E., Vanduffel, W. & Janssen, P. Functional heterogeneity of macaque lateral intraparietal neurons. *J. Neurosci.* **31**, 12307–12317 (2011).
37. Ma, W.J., Beck, J.M., Latham, P.E. & Pouget, A. Bayesian inference with probabilistic population codes. *Nat. Neurosci.* **9**, 1432–1438 (2006).
38. Jazayeri, M. & Movshon, J.A. Optimal representation of sensory information by neural populations. *Nat. Neurosci.* **9**, 690–696 (2006).
39. Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S. & Movshon, J.A. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* **13**, 87–100 (1996).
40. Steenrod, S.C., Phillips, M.H. & Goldberg, M.E. The lateral intraparietal area codes the location of saccade targets and not the dimension of the saccades that will be made to acquire them. *J. Neurophysiol.* **109**, 2596–2605 (2013).
41. Brunton, B.W., Botvinick, M.M. & Brody, C.D. Rats and humans can optimally accumulate evidence for decision-making. *Science* **340**, 95–98 (2013).
42. Gottlieb, J. & Balan, P. Attention as a decision in information space. *Trends Cogn. Sci.* **14**, 240–248 (2010).
43. Bollimunta, A., Totten, D. & Ditterich, J. Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex. *J. Neurosci.* **32**, 12684–12701 (2012).
44. Schall, J.D., Purcell, B.A., Heitz, R.P., Logan, G.D. & Palmeri, T.J. Neural mechanisms of saccade target selection: gated accumulator model of the visual-motor cascade. *Eur. J. Neurosci.* **33**, 1991–2002 (2011).
45. Purcell, B.A., Schall, J.D., Logan, G.D. & Palmeri, T.J. From salience to saccades: multiple-alternative gated stochastic accumulator model of visual search. *J. Neurosci.* **32**, 3433–3446 (2012).
46. Gold, J.I. & Shadlen, M.N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
47. Churchland, A.K. *et al.* Variance as a signature of neural computations during decision making. *Neuron* **69**, 818–831 (2011).
48. Leathers, M.L. & Olson, C.R. In monkeys making value-based decisions, LIP neurons encode cue salience and not action value. *Science* **338**, 132–135 (2012).
49. Newsome, W.T., Glimcher, P.W., Gottlieb, J., Lee, D. & Platt, M.L. Comment on "In monkeys making value-based decisions, LIP neurons encode cue salience and not action value". *Science* **340**, 430 (2013).
50. Leathers, M.L. & Olson, C.R. Response to comment on "In monkeys making value-based decisions, LIP neurons encode cue salience and not action value". *Science* **340**, 430 (2013).

ONLINE METHODS

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

Preparation, neurophysiology and tasks. All procedures were performed in accordance with US 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 RF of a neuron. Only neurons that were spatially selective during an instructed saccade task were included[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. To achieve fixation at a trial's start, the monkey's eye position had to register within a window around the central fixation point no later than 3 s 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 choice targets 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[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 (500–1,000 ms, uniform distribution), the motion stimulus vanished. The monkey continued to maintain fixation for another 500 ms until the fixation point also disappeared, thereby 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 inside the spatial window around the target location within 100 ms of leaving the fixation point window. Entry into the spatial window around the target also had to occur within 450 ms of the go signal to be considered a complete trial. About 800–1,200 trials were collected per recording session (per neuron).

Plotting neural response. PSTHs were smoothed with a gaussian filter (s.d., 75 ms). In **Figures 2** and **3**, each conditional PSTH was stitched together by averaging temporally overlapping local PSTHs of length 500–1,000 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 **Figure 4** were chosen by sorting the mean square error between the predicted spike rate and the boxcar smoothed spike trains in the 2-s 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 0th lag component from the autocorrelation plots in **Figure 5b**.

Model parameter representation. The spike trains were discretized into 1-ms bins. Each event was represented as a delta or boxcar function over time, and convolved with a filter (or a kernel; **Supplementary Fig. 1**). We used smooth temporal basis functions defined by raised cosine bumps separated by π/2 radians (50 ms) to parameterize the filters (**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 for the targets-FLASH condition and 500-ms window after the dots disappearance. For the saccade response, we had two kernels, one for each direction which was anti-causal (that is, predicting spikes as a function of time prior to the saccade) for 2,504 ms for monkey J, and 2,662 ms for monkey P and with total duration of 5.3 s. We grouped the directional coherence levels into five 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%) coherence 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; 10 1-ms uniform bases (to represent the fast refractory effects) followed by 10 raised cosine bases stretched in a logarithmic scale that spanned 265 ms[17] (**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 five coherence kernels for **Figures 1** and **2**.

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

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

corrected (mistake introduced by copy editor)

where θ is a vector representing the weights on the basis functions, and ξ is chosen from a grid to maximize the marginal likelihood. We computed marginal likelihood $p(\boldsymbol{r}|\xi) = \int p(\boldsymbol{r}|\theta,\xi)\, p(\theta|\xi) \mathrm{d}\theta$ using Laplace approximation. For each neuron, we divided the trials randomly into five equally sized sets to perform fivefold 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[51].

Goodness-of-fit in **Figure 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 **Figure 5a** was 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].

doi:10.1038/nn.3800

Decoding. We derive the LLR (equation (4)) decoding rule. Let t_c be the time when the decision is read out.

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

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 $\boldsymbol{w}^{\mathrm{T}}\boldsymbol{r} = \sum_{i=1}^{L} w(i)r(t - L + i) = \tilde{\boldsymbol{w}} * \boldsymbol{r}(t)$ where $\tilde{\boldsymbol{w}}$ is the time reversed weight vector, L is its length, and $\boldsymbol{r}(t)$ is the length-L vector of response history at time t.

CP quantifies the dependence between the spike counts of a neuron and a binary decision variable under a fixed sensory stimulus. Given a pair of randomly selected IN and OUT choice trials, the probability that the higher spike count belongs to the IN trial is CP. A high CP implies that the spike count covaries strongly with the decision. CP in **Figure 6c**,**d** was computed by randomly drawing 1,000 pseudorandom (IN, OUT) pairs of trials. We repeated this procedure 40 times to obtain the error bars. For traditional CP, we used the spike count in the −1,500 to −50 ms window before the saccade detection (for **Fig. 6c**, the beginning of the window was fixed at −1,500 ms, and x axis shows the end of the window). For the model-based CP, spike trains from the cross-validation set were projected onto the weights obtained from the training set.

A **Supplementary Methods Checklist** is available.

51. Cunningham, J.P., Yu, B.M., Shenoy, K.V. & Sahani, M. Inferring neural firing rates from spike trains using Gaussian processes. *Adv. Neural Inf. Process. Syst.* **20**, 329–336 (2007).