

1-1-2012

Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas

Vinod Rao

Washington University School of Medicine in St. Louis

Gregory C. DeAngelis

Washington University School of Medicine in St. Louis

Lawrence H. Snyder

Washington University School of Medicine in St. Louis

Follow this and additional works at: http://digitalcommons.wustl.edu/open_access_pubs

 Part of the [Medicine and Health Sciences Commons](#)

Recommended Citation

Rao, Vinod; DeAngelis, Gregory C.; and Snyder, Lawrence H., "Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas." *The Journal of Neuroscience*.32,29. 10063-10074. (2012).
http://digitalcommons.wustl.edu/open_access_pubs/1141

This Open Access Publication is brought to you for free and open access by Digital Commons@Becker. It has been accepted for inclusion in Open Access Publications by an authorized administrator of Digital Commons@Becker. For more information, please contact engeszer@wustl.edu.

Neural Correlates of Prior Expectations of Motion in the Lateral Intraparietal and Middle Temporal Areas

Vinod Rao,^{1,2} Gregory C. DeAngelis,^{1,3} and Lawrence H. Snyder¹

¹Department of Anatomy and Neurobiology, Washington University School of Medicine, Saint Louis, Missouri 63110, ²Department of Psychiatry, Massachusetts General Hospital, Boston, Massachusetts 02114, ³Department of Brain and Cognitive Sciences, University of Rochester, Rochester, New York 14627

Successful decision making involves combining observations of the external world with prior knowledge. Recent studies suggest that neural activity in macaque lateral intraparietal area (LIP) provides a useful window into this process. This study examines how rapidly changing prior knowledge about an upcoming sensory stimulus influences the computations that convert sensory signals into plans for action. Two monkeys performed a cued direction discrimination task, in which an arrow cue presented at the start of each trial communicated the prior probability of the direction of stimulus motion. We hypothesized that the cue would either shift the initial level of LIP activity before sensory evidence arrived, or it would scale sensory responses according to the prior probability of each stimulus, manifesting as a change in slope of LIP firing rates. Neural recordings demonstrated a clear shift in the activity level of LIP neurons following the arrow cue, which persisted into the presentation of the motion stimulus. No significant change in slope of responses was observed, suggesting that sensory gain was not strongly modulated. To confirm the latter observation, middle temporal area (MT) neurons were recorded during a version of the cued direction discrimination task, and we found no change in MT responses resulting from the presentation of the directional cue. These results suggest that information about an immediately upcoming stimulus does not scale the sensory response, but rather changes the amount of evidence that must be accumulated to reach a decision in areas that are involved in planning action.

Introduction

Both prior knowledge and sensory input determine how we perceive the external world. Although much research has examined how sensory representations are used to form perceptual decisions (Gold and Shadlen, 2007), considerably less work has examined how prior knowledge is incorporated into the decision-making process. One proposed mechanism for perceptual decision making is the integration-to-bound model (Link and Heath, 1975; Mazurek et al., 2003; Smith and Ratcliff, 2004; Gold and Shadlen, 2007). In this account, each possible choice is associated with an integrator that accumulates the value of sensory evidence in favor of that option. When the total evidence favoring one of the options reaches a threshold, the evidence-gathering process is halted, and that option is selected. Various studies have suggested that activity in the lateral intraparietal area (LIP) reflects such a decision-making process (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Huk and Shadlen, 2005; Hanks et al., 2006; Kiani et al., 2008). Specifically, the activity of an LIP neuron is thought to represent the cumulative evidence in favor of a choice indicated by an eye movement into that neuron's

response field (RF) (Gold and Shadlen, 2007; Yang and Shadlen, 2007).

If LIP activity reflects a decision variable that determines the choice, then LIP activity should reflect relevant prior knowledge in one of at least two ways. First, less evidence might be required to choose the more likely option. This could manifest as an initial offset of LIP activity toward a bound, or as a change in bound height for that choice option (Gold and Shadlen, 2007; Yang and Shadlen, 2007). Alternatively, more likely stimuli might have a larger gain on their sensory representation. An integrator of those sensory responses would consequently accumulate greater evidence in favor of the cued option. This could manifest as a steeper rate of ramping (slope) in LIP responses to visual motion. This possibility is plausible given the abundant evidence that behavioral context (e.g., attention) can modify sensory representations (Treue, 2001; Reynolds and Chelazzi, 2004; Cisek, 2007).

Although estimates of prior probability are often built up over long periods of time and change gradually, contextual cues sometimes alert us to rapid shifts in statistical structure. Recent physiological studies that explored how prior probabilities influence decision processes have generally used tasks with fixed or slowly-changing priors (Fischer and Pena, 2011; Hanks et al., 2011). In contrast, we examined how priors that change on a trial-by-trial basis influence neural correlates of choices. In our cued direction discrimination task, monkeys were given a partially predictive cue about the direction of motion of an upcoming stimulus and learned to incorporate cue information into direction judgments. Recordings in LIP showed that when monkeys expected motion toward the neuron's response field, activity was increased before and during the motion display. Changes in the slope and end-

Received Nov. 28, 2011; revised May 21, 2012; accepted June 4, 2012.

Author contributions: V.R., G.C.D., and L.H.S. designed research; V.R. performed research; V.R., G.C.D., and L.H.S. analyzed data; V.R., G.C.D., and L.H.S. wrote the paper.

This work was supported by National Eye Institute Grants EY013644, EY012135 and EY002687. V.R. was supported by National Institute for Mental Health F30 Grant MH077368-01A1. Special thanks to Donna Lalor, Amy Wickholm, and Jonathon Tucker for technical assistance. We thank Dora Angelaki, Maurizio Corbetta, and David Van Essen for helpful discussions.

Correspondence should be addressed to Dr. Lawrence H. Snyder, Campus Box 8108, 660 South Euclid Avenue, St. Louis, MO 63110. E-mail: larry@eye-hand.wustl.edu.

DOI:10.1523/JNEUROSCI.5948-11.2012

Copyright © 2012 the authors 0270-6474/12/3210063-12\$15.00/0

point of activity buildup could not account for the cueing effects on behavior. Furthermore, recordings in the middle temporal area (MT), which provides motion input to LIP, showed that sensory responses were unaffected by the cue.

Materials and Methods

Subjects and surgery. Experiments were performed on two male rhesus macaques (*Macaca mulatta*). General procedures have been described previously (Uka and DeAngelis, 2003). Each animal was fitted with either a CILUX or stainless-steel head holder (Crist Instruments), and a scleral search coil was implanted under the conjunctiva for monocular eye position monitoring. All experimental procedures were approved by the Institutional Animal Care and Use Committee at Washington University and conformed to National Institutes of Health guidelines.

After behavioral training, a beveled CILUX recording chamber (Crist Instruments) was surgically attached to the monkey's skull. Chambers for LIP experiments were centered ~12 mm lateral to the midline and 21 mm anterior to the occipital ridge. For one monkey, a follow-up structural MRI was used to guide placement of electrodes within the chamber to target LIP. The chamber for the MT experiment was centered 17 mm lateral to the midline and 14 mm anterior to the occipital ridge. All chambers were oriented at an angle of ~25° above stereotaxic horizontal in a parasagittal plane.

We selected neurons using the same criteria that have been used in previous studies of evidence accumulation, as LIP neurons are heterogeneous, and the selection criteria could affect our results (Premereur et al., 2011). LIP and MT were identified functionally as areas of cortex, located beneath contiguous grid holes, having the expected neural responses: for LIP, expected properties include spatially tuned responses during execution of delayed saccades and frequent presence of memory activity; for MT, expected properties include robust direction-selective responses and appropriate receptive field size/eccentricity ratios. Although we targeted our recordings to LIP-ventral, we recorded in all contiguous grid holes for which neurons exhibited memory-saccade responses, and we cannot rule out that some of our recordings were in LIP-dorsal.

Stimulus control and presentation. Task events and data acquisition were controlled by TEMPO software (Reflective Computing). Online behavioral and neural analyses were performed using Matlab software (Mathworks). Visual stimuli were generated using an OpenGL accelerator board (Oxygen GVX1; 3Dlabs) using custom software written in Visual C++. In the LIP experiments, the monkey viewed the stimuli on a 22 inch flat-screen 100 Hz CRT monitor. At a viewing distance of 22 cm, the monitor subtended $84 \times 68^\circ$ of visual angle. In the MT experiments, the images were displayed using a stereoscopic projector (Mirage 2000; Christie Digital) onto a $56 \times 46^\circ$ projection screen positioned 57 cm in front of the monkey. To generate stereoscopic depth in the MT experiment, stimuli for each eye were presented on alternate frames and viewed through liquid-crystal ferroelectric shutter goggles, which were synchronized to the refresh of the projector. This setup had no measurable stereo cross talk.

Tasks and training procedures. All experiments required the monkey to perform some version of a cued direction discrimination task. The task was designed both to require the monkey to make a perceptual judgment about a noisy visual stimulus and to encourage the monkey to bias these choices in accordance with a prior cue. In our task, the monkey had to judge the direction of a moving random-dot pattern and was biased by a partially predictive arrow-shaped cue.

The timing of the task in the main experiment is as follows (see Fig. 1). The task began when the monkey fixated upon a 0.2° diameter, yellow fixation point. Fixation windows were 2.2° in diameter for Monkey 1, and 3° for Monkey 2. Immediately after the monkey acquired the fixation point, two 0.5° diameter targets appeared on opposite sides of the fixation point. One target was placed in an LIP neuron's RF and was denoted T_{in} ; the target opposite the RF was called T_{out} . After the monkey held fixation for 1050 ms, a green, 3° arrow-shaped cue appeared at the fovea for 200 ms. On most trials the arrow pointed toward one of the two choice targets; however, on some trials, a double-headed arrow pointed to both targets simultaneously and is referred to as a neutral cue. The

arrow cue was displayed for 200 ms, followed by a variable delay of 150–800 ms (flat hazard function; mean, 290 ms).

Next a random-dot motion stimulus was displayed in a 5° diameter patch centered on the fixation point. The motion stimulus consisted of a field of dots moving at 8° per second toward one of the two targets, chosen randomly on each trial. Noise was added to the stimulus by varying the motion coherence. Coherence refers to the percentage of dots moving smoothly toward one of the two targets; the remaining dots were replotted at random positions to generate a "snowlike" masking noise. The motion coherence (c) on a given trial was chosen from the set 0, c , $2c$, $4c$, and $8c$, where $c = 3\%$ for Monkey 1 and $c = 2\%$ for Monkey 2. These different ranges of coherence assured that each monkey's choices would be appropriately biased by the arrow cue. Note that a 0% coherent motion stimulus consists entirely of randomly plotted dots and so is ambiguous with respect to the direction of motion. The motion stimulus was presented for 1000 ms on 75% of trials, chosen randomly, or for 250 ms on the remaining 25% of trials. The interleaved short duration stimuli encouraged the monkey to make use of motion information from the beginning of the presentation; however, these trials were too few to analyze. Instead, we include the results of a separate version of the task, run on one monkey, in which only 250 ms motion stimuli were presented.

After the motion stimulus disappeared, the monkey maintained fixation during a 250 ms delay period. The disappearance of the fixation point instructed the monkey to initiate a saccade to one of the two peripheral targets and to maintain eye position within 3.5° of that target for 150 ms. The trial was scored as correct if the chosen target matched the direction of the moving dots, regardless of the direction of the arrow cue. On correct trials, the monkey was rewarded with a drop of juice, whereas on incorrect trials, the monkey endured a 1000 ms time-out before proceeding to the next trial. Trials in which the monkey broke fixation or failed to make a saccade into either target window were aborted, and the monkey was required to wait 3500 ms before proceeding to the next trial.

Critically, the correspondence between the direction of the arrow cue and the direction of random-dot motion was chosen to make the arrow cue partially predictive of motion direction. Specifically, trials in which the arrow-cue direction matched the motion direction ("valid cue" trials) were twice as frequent as trials in which the arrow-cue direction was opposite the motion direction ("invalid cue" trials). Neutral-cue trials, in which the arrow cue had two heads and thus contained no directional information, were equally as frequent as valid cue trials. Thus, an ideal observer attending only to the cue would be rewarded at 67% on all non-neutral-cue trials. An ideal observer attending only to the motion stimulus would be rewarded at 50% on 0%-coherence trials and up to 100% on other trials. To promote attention to the cue, Monkey 1 was presented with occasional cue-only trials. These trials were equivalent to valid cue trials with the motion stimulus omitted, so that the monkey was rewarded for reporting the direction of the arrow cue. Given two directions of motion, five coherence levels, and five cue validities (two valid, two neutral, one invalid), the task contained 50 distinct stimulus conditions (not including Monkey 1's cue-only trials, which were not included in the analysis).

Training the monkeys to perform this task consisted of two main stages: (1) training to discriminate between two opposite directions of motion and (2) learning to incorporate the cue information. The first stage was thoroughly described previously (Britten et al., 1992). Briefly, one-target trials were intermixed with two-target trials early in training to encourage the animal to form an association between full-coherence motion and a saccade to the corresponding target. When this association was stabilized, staircase procedures were used to gradually decrease the motion coherence as the monkey became more sensitive. During this procedure, the monkeys were exposed to all possible axes of discrimination and eccentricities of targets ranging from 5 to 20° from the fovea.

One monkey was also trained in a modified version of the task that included only short-duration (250 ms) trials (see Fig. 2D, 5B). All other timing parameters were consistent with the main task. The animal's behavior was sufficiently consistent that all sessions used a constant set of five motion strengths: 0, 4, 8, 16, and 32% coherence.

The direction discrimination task was slightly modified for the MT recordings, performed in one animal (see Fig. 7A). Most importantly, the

parameters of the motion stimulus were chosen to maximally drive MT neuron responses (see below, Experimental protocol). The position and size of the motion patch were selected to fill the MT neuron's receptive field, and the axis of discrimination was aligned to the neuron's preferred–null axis. Furthermore, the speed and horizontal disparity of the dots were chosen to match the neuron's preferences. Only 300 ms elapsed between the acquisition of the fixation point and the display of the arrow cue. After the motion presentation, which lasted 1000 ms on all trials, the choice targets immediately appeared, and the fixation point disappeared, signaling the monkey to report its choice.

Incorporation of the arrow cue into the direction discrimination task began by including only valid cue trials for which the cue overlapped in time with the motion stimulus or for which there was no interstimulus interval (ISI) between the offset of the cue and the onset of the motion. After the monkey demonstrated consistent use of the cue on trials having either low coherence or zero coherence (>80% correct), we gradually separated the cue and the motion stimulus in time, until eventually reaching the final ISI distribution. Next, we introduced neutral-cue trials and interleaved them with the valid cue trials. Finally, we gradually incorporated invalid cues starting at ratios of five valid/one invalid and slowly increasing in frequency until we reached the final ratio of two valid/one invalid.

In addition to the cued direction discrimination task, monkeys were trained to perform other tasks needed to identify LIP neurons and characterize their response fields. In the delayed saccade task, the monkey first fixed his gaze on the fixation point, and 300 ms later a red 0.5° peripheral target was presented. After 200 ms, the target luminance was dimmed (30% original luminance for Monkey 1, 20% for Monkey 2) for 800 ms. Finally, the fixation point and target were extinguished, and the monkey had to make a saccade to the location of the target and was rewarded for landing within 3.5° of the target. Memory saccade trials were almost identical, except that the target was completely extinguished after the initial 200 ms presentation. The monkey had to remember the target location for 800 ms and then make a saccade to that location when the fixation point disappeared. Finally, the monkey was trained to perform a receptive field mapping task. In this task, the monkey fixated while a peripheral target was flashed in a series of one to five positions. The targets were presented at full luminance for 300 ms each. The monkey was rewarded for making a saccade to the target that was on the screen when the fixation point was turned off.

Experimental protocol. We recorded responses from single neurons in two monkeys. Fine tungsten microelectrodes (1–2 M Ω ; Alpha Omega) were inserted into the brain via a transdural guide tube held in a plastic grid (Crist Instruments) and were advanced through the cortex with a micromanipulator (Narishige) mounted on the recording chamber. Action potentials were isolated using an amplifier (A-M Systems or FHC), bandpass filter (Krohn-Hite), and window discriminator (Bak Electronics). Action potential and event timings were recorded at 1 ms temporal resolution.

While searching for neurons during the LIP recordings, the monkey performed delayed saccades to targets at an eccentricity of 15° in each of eight directions. After a single unit was isolated, the location of its response field (polar angle and eccentricity) was estimated using the delayed saccade and receptive field mapping tasks, run iteratively as needed. Next, the monkey performed memory saccades to a target located in the response field and a target positioned diametrically opposite the fixation point. LIP neurons that failed to show tuned memory responses ($p < 0.05$, Mann–Whitney U test) were not studied further.

For the LIP versions of the cued direction discrimination task, the center of the response field and the position diametrically opposite the fixation point defined the two target locations, referred to as T_{in} and T_{out} , respectively (see Fig. 1, top left). Coherent dot motion, and thus the axis of discrimination, was along the line defined by these two positions. For a neuron to be included in the analysis, single-unit isolation was maintained for a minimum of five repetitions of each of the 50 trial types, with a median of seven repetitions.

During MT recordings, receptive field position and size as well as preferred stimulus parameters (direction, speed, and horizontal disparity) were first estimated by hand and then quantitatively characterized

during fixation trials. The two possible motion directions in the cued direction discrimination task were selected to match the direction that maximized the neuron's response (the “preferred” direction) and the direction 180° opposite (the “null” direction). The remaining parameters of the motion stimulus (position, size, speed, and horizontal disparity) were matched to each neuron's preferences to maximize the relevance of the neuron to the task (Britten et al., 1992; DeAngelis and Newsome, 2004).

Behavioral data analysis. All offline analyses were conducted using Matlab software (Mathworks). We constructed psychometric curves to assess whether the monkeys' reports were influenced by the direction of the arrow cue. First, the trials were split into groups based on the direction of the arrow cue [toward T_{in} , toward T_{out} , or toward both (neutral cue)]. For each cue group, we plotted the proportion of T_{in} choices as a function of signed coherence, where positive coherences denote motion toward T_{in} , negative values denote motion toward T_{out} , and zero-coherence trials are ambiguous. These data were fit with logistic functions of the following form:

$$y = 1 / (1 + e^{(-\alpha(x-\beta))}), \quad (1)$$

where y is the proportion of T_{in} choices, x is the signed coherence, and α and β are free parameters reflecting the slope and bias of the curves, respectively. The cue-related bias for a given cue direction was computed as the difference in β between trials cued toward that direction (either T_{in} or T_{out}) and β in neutral-cue trials. The average cue-related bias for a behavioral session was therefore the average of the bias for each of the two cue directions.

To determine whether the monkey's reports were influenced by the validity of the arrow cue, we performed a logistic regression analysis on the monkey's percent correct in the various conditions, with motion duration, motion coherence, cue validity, and all interaction terms as predictors. For a session to be included in the final data set, we required a significant ($p < 0.05$) effect of cue validity, or a significant effect of at least one of the three interaction terms that include cue validity. Furthermore, the sign of the significant terms was required to be consistent with the predicted effect of the cue. For example, a significant interaction between cue validity and motion coherence must be such that increasing coherence decreases the effect of the cue, since the motion determines the response. Data from 10 of 91 recording sessions in the main LIP task were excluded based on these behavioral criteria. This was done to ensure that we sought neural correlates of prior expectation under conditions in which behavior was clearly affected by the arrow cue.

Neural data analysis: LIP. Analyses were based on unsmoothed spike counts. For figures, action potential trains were smoothed with a Gaussian kernel (20 ms SD, equivalent to a low-pass filter with a -3 dB point at 6.8 Hz) to generate spike density functions (SDFs) before plotting. Population SDFs were generated by first averaging neural responses for the appropriate trial types in the relevant time interval within each cell and then averaging SDFs across neurons. By first averaging within conditions, we avoided artifacts from unequal numbers of T_{in} and T_{out} responses.

To estimate the rate of LIP buildup activity, the raw (unsmoothed) spike counts were grouped in 10 ms bins, and the slope was computed by linear regression of the binned spike counts over a specified time window (Churchland et al., 2008). This slope was computed separately at each coherence for each neuron.

To test for consistent ordering of these firing rate slopes according to coherence, we first computed the Spearman rank correlation between response slope and coherence for each neuron. We then asked whether the average correlation for our sample of neurons was significantly different from zero using a Student's t test. Only correct neutral-cue trials with motion in one direction were used in this analysis (see Fig. 3).

To compute the latency of responses to motion, we focused on correct, neutral-cue trials. To determine the onset time of directional specificity, LIP responses to motion toward T_{out} were subtracted from responses to motion toward T_{in} . This difference was computed separately for each level of coherence, and a bootstrapping procedure was used to compute confidence intervals around the time course of the mean activity differ-

ence. For each level of coherence, we selected the first time bin for which the Bonferroni-corrected 95% confidence interval did not include zero. The directional latency was chosen as the earliest time (290 ms) that a significant directional signal was detected at any coherence (which unsurprisingly occurred at the largest coherence). Other procedures, such as applying a piecewise linear fit to the time course of the response variance, produced similar results (Kiani et al., 2008). Note that this estimate represents an upper bound on the true latency for directional signals to arrive in LIP.

The time course of the cue effect (see Fig. 5) was generated using overlapping 50 ms windows, stepped through the data in 25 ms increments. We averaged the raw spike counts within each time window for each task condition, which we define as a unique combination of cue direction, motion direction, coherence, and choice. Next, for each task condition, we subtracted off activity from the corresponding neutral-cue version of that condition. This produced a cue effect time course for each task condition. (By construction, the cue effect time courses for the conditions with neutral cues were zero). Averaging across motion direction, coherence and choice yielded an average cue-related activity modulation for each neuron for a given cue direction. Ninety-five percent confidence intervals were defined by the cross-cell mean cue effect ± 1.96 SDs. This parametric formulation for the confidence interval was justified by highly linear quantile–quantile plots (data not shown) of cue effect sizes for T_{in} and T_{out} cues.

We also estimated the time at which the cue effect begins to differ between trials in which the monkey chose T_{in} versus T_{out} . Separate time courses of the cue effect were computed as above using only T_{in} -choice trials or T_{out} -choice trials. Subtracting these time courses yielded a choice-conditioned difference in the cue effect. We defined a difference threshold as the mean difference plus three SDs, based on activity from a 275 ms period before motion onset. The latency was defined as the first time point at which the choice-conditioned difference exceeded this threshold.

To explore whether the neural cue effect was related to the behavioral effect, we computed correlations between metrics of these two effects. The neural cue effect was quantified, for each cue direction, by first computing firing rates averaged over the 100 ms period before motion onset and then subtracting this average rate on T_{out} -cue trials from the average rate on T_{in} -cue trials. The behavioral effect was computed by subtracting the cue-related bias (as described in the previous section) for T_{in} -cue trials from the cue-related bias for T_{out} -cue trials, yielding positive values when T_{in} cues lead to more T_{in} choices than do T_{out} cues. The overall association between these two metrics was quantified by Spearman correlation, as the behavioral metric was not normally distributed.

Neural data analysis: MT. To examine effects of the cue on MT responses to motion, we constructed neural response curves for each cue direction, where the normalized firing rate during the 1000 ms motion stimulus is plotted as a function of the signed coherence (see Fig. 7B). The neuron's firing rate for each combination of signed coherence and cue direction was normalized to that neuron's response on zero-coherence, neutral-cue trials. To test for a significant effect of cue direction on MT responses to motion, we analyzed the raw (i.e., non-normalized) mean firing rates with a two-factor repeated-measures ANOVA, with cue direction and signed coherence as factors.

Figure 7B suggests that the presence of a conflict between the directions of the cue and the motion may influence MT responses. To test this, we applied a two-factor repeated-measures ANOVA to the raw firing rates during the motion stimulus with factors of signed coherence and conflict. Conflict conditions were defined as situations in which the cue indicated one direction of motion while the random-dot stimulus contained robust motion (2c, 4c, and 8c%) in the opposite direction. Non-conflict conditions were defined as the corresponding neutral-cue trials.

As the above analyses yielded no significant effect of cue direction on MT responses (see Results), we deemed it important to assess the sensitivity of our analysis. We performed simulations to estimate the statistical power of the ANOVA described above. In brief, we generated data sets powered identically to ours, but with known cue-related effects on firing rates. We analyzed these data sets with the same two-factor repeated-

measures ANOVA and compared the sensitivity of our analyses to the size of the imposed firing rate effect.

More specifically, the baseline responses for each simulated neuron, at each coherence, were taken as the neutral-cue firing rates of a real neuron. To be conservative, we assumed that null-direction cues had no effect on the mean response and that preferred-direction cues provided a simple additive offset to the mean firing rate on neutral-cue trials. We then generated trial-by-trial firing rates by sampling from Gaussian distributions with the appropriate means and variances. Assuming Poisson-like firing statistics, rather than Gaussian, produced similar results. We matched the number of trials for each simulated neuron to the corresponding number in our measured data set. Given the sampled firing rates, we computed means and repeated the ANOVA as in our real data set. We repeated this procedure 1000 times and tallied the number of significant main effects of cue direction. Finally, we searched for values of the firing rate offset due to a preferred-direction cue that would be detected as significant with 90% and 50% reliability by our analysis.

We examined whether receptive field eccentricity was related to the magnitude of the neural cue effect. For each neuron, we estimated the cue effect by subtracting the firing rate during the motion stimulus on neutral-cue trials from the corresponding firing rate on preferred-direction-cue trials. This was done separately for each motion direction and coherence and averaged across these conditions to yield a metric reflecting the cue effect for each neuron. The data set was split according to the median eccentricity of the receptive fields, and the average cue effect was computed for neurons with receptive fields nearer and farther from the fovea.

Choice probabilities (CPs) were computed according to the method of Britten et al. (1996). For each distinct combination of motion direction and coherence, firing rates during the motion display were normalized by computing z-scores. Within each distinct stimulus condition, normalized responses were then divided into two groups based on whether the monkey chose the preferred or null direction of the neuron on each trial. Normalized response distributions for preferred-choice trials were then combined across all distinct stimulus conditions for which the monkey made at least one choice in each direction, and this was repeated for response distributions from null-choice trials. Finally, we calculated a grand choice probability for the neutral-cue trials by applying receiver operating characteristic analysis to the distributions of normalized responses sorted by choice. This entire procedure was then repeated for valid-cue trials to generate a grand choice probability under value cue conditions. Permutation tests were used to evaluate the significance of CPs (Britten et al., 1996; Uka and DeAngelis, 2004).

Integration-to-bound model. We simulated the monkey's behavior using an integration-to-bound framework (Mazurek et al., 2003; Gold and Shadlen, 2007) to estimate the expected time course of the cue effect. In our model, each possible choice had a separate integrator associated with it. When the accumulated evidence in one integrator exceeded a bound, the corresponding choice was recorded for that trial. The bound was the same for both integrators. As the motion stimulus in our task was of a fixed duration, motion evidence was presented to the integrators for a fixed number of time steps. If neither integrator reached its bound by the end of the motion presentation, the decision for that trial was generated randomly, weighted by the relative amount of accumulated evidence at that time. For example, if the first integrator had accumulated half of the required evidence and the second integrator had accumulated three-fourths of the required evidence, then there was a 40% chance $[0.5/(0.5 + 0.75) * 100]$ of choosing the first target and a 60% chance of choosing the second target. An alternative implementation, which, if neither bound is reached, selects the integrator with the most evidence, changes the fit parameters but does not change the final pattern of the cue effect time course.

We modeled the moment-to-moment evidence about motion as samples from a normal distribution, with a mean that was linearly related to the signed motion coherence ($\mu = kC + \mu_0$) and a variance (σ^2) that was constant across trials. For each integrator, if a sample of evidence brought the integrator closer to its bound, the value was added in full. If that sample supported the opposite choice, the value was scaled by a constant (ϕ) and subtracted from the running total. This approach reflects the




Figure 1. Task design for the standard version of the cued direction discrimination task. Top panels illustrate the display seen by the animal during cue presentation (left), motion stimulus presentation (center), and target selection (right). The timing of the task events is depicted at the bottom. All numbers are durations given in milliseconds. Red arrows mark times of the three screen snapshots.

observation that LIP activity does not change symmetrically for motion toward T_{in} versus that toward T_{out} (Churchland et al., 2008), as also observed in our data (see Fig. 3). Finally, prior expectations regarding choice were modeled by adding an initial positive offset to the integrator associated with the cued direction.

Because directional motion signals did not arrive instantly at LIP, we enforced a fixed latency on when motion signals began to be accumulated in our simulation. The latency, 290 ms, was based upon our measured neural data (see above, Neural data analysis: LIP).

In total, five free parameters specify the model: (1) the proportionality constant, k , between coherence and the mean value of the evidence; (2) the mean evidence offset, μ_0 , which accommodates small choice biases (independent of the cue) that are present without coherent motion; (3) the variance, σ^2 , of the evidence; (4) the scaling parameter, ϕ , for evidence supporting the opposite choice (which ranged from 0 to 1); and (5) the initial offset associated with the cue. We simulated 400 trials for each distinct combination of direction and coherence to generate psychometric curves.

Behavioral data from neutral-cue and T_{in} -cue trials with 1000 ms motion durations were combined across the two monkeys and used to fit the model. Model fits were quantified by applying an inverse logistic transform to the measured and simulated psychometric curves to linearize the data and then computing the correlation coefficient between the two data sets. This procedure was performed separately for neutral-cue and T_{in} -cue trials, and the resulting correlation coefficients were averaged to estimate the overall model fit. Using the best-fitting parameters, we reran the simulation for 1000 trials per stimulus condition and examined the evolution of the accumulated evidence over time. Averaging across motion direction and coherence yielded time courses for neutral-cue and T_{in} -cue trials, and the time course of the cue effect was simply the difference between these time courses.

A control model was implemented that did not include a bound on the integration. In this “unbounded” model, the integrator with the highest value at the end of the motion was selected as the decision. This yielded a model with only four free parameters, as the variance of the evidence is irrelevant in this scale-free model. Again, the model parameters were fit to the behavior, and the time course of the cue effect was computed as before.

Results

Monkeys were trained to perform a cued direction discrimination task (Fig. 1) (see Materials and Methods). In this task, the monkey earned rewards for reporting the direction of motion in a random dot display. An arrow cue preceded the motion stimulus and was partially predictive of the motion direction. In particular, arrow cues pointing toward the T_{in} target predicted that the motion was twice as likely to be directed toward T_{in} than T_{out} , and vice versa for trials with a T_{out} -cue. Neutral (bidirectional) cues indicated that the upcoming motion was equally likely to be in either direction. We demonstrate that these cues biased monkeys’ behavioral responses, and then we explore how the cues altered neural responses.

Effect of prior cue on motion discrimination behavior

Figure 2A shows the fraction of T_{in} choices as a function of coherence, sorted by the direction of the arrow cue, for a single recording session with 1000-ms-duration trials. In general, as more random dots moved coherently toward T_{in} (i.e., greater positive coherence), the monkey more frequently chose T_{in} . Importantly, when motion was preceded by a T_{in} cue (green), the monkey more frequently chose T_{in} than on neutral-cue trials (red). Likewise, the monkey made fewer T_{in} choices when cued toward T_{out} (blue). To quantify these effects, we fit a logistic function (Eq. 1) to the data and estimated the behavioral bias associated with each cue direction (see Materials and Methods). Although the monkey had a slight overall bias toward choosing T_{out} in this session, there was also a clear effect of cue direction on choices. Presentation of a T_{in} cue yielded a positive bias (leftward shift) equivalent to 7.3% coherence, and presentation of a T_{out} cue yielded a bias equivalent to -2.3% coherence. These cue effects were significant by logistic regression analysis ($p < 0.0001$).

Averaging each monkey’s behavior across sessions indicates that the monkeys chose a target more frequently when the cue predicted motion in that direction than on corresponding neutral-cue trials (Fig. 2B, C). The effect of the cue across sessions was significant by logistic regression analysis (Monkey 1, $p < 0.0001$; Monkey 2, $p < 0.005$). The cue effect was particularly strong when the motion stimulus contained little to no coherent motion. On these trials, the stimulus provided poor quality information about motion direction, and the monkey could maximize its expected reward by weighting the cue more heavily.

A separate set of experiments further established that arrow cues influenced choices. Monkey 1 performed a shortened version of the cued direction discrimination task in which the motion display persisted for only 250 ms on each trial. With less motion information available, the monkey relied more heavily on the prior information provided by the cue, as seen by the larger separation of the curves in Figure 2D relative to those in Figure 2B. The bias induced by the cue was significantly greater in the sessions with only short-duration trials ($p < 10^{-18}$, Mann-Whitney U test).

LIP responses to motion stimuli

Previous studies involving choice tasks have shown that LIP neurons increase their firing rates leading up to a saccade to a target in the cell's response field. Furthermore, these studies suggest that LIP activity builds up faster when stronger sensory evidence guides the choice (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Churchland et al., 2008; Kiani et al., 2008).

We evaluated this hypothesis in our data set, focusing on correct, long-duration, neutral-cue trials. SDFs in Figure 3 show the population average time course of LIP responses to motion, sorted by coherence and choice. LIP firing rates on T_{in} -choice trials (solid) transiently decreased, rose quickly for about 250 ms, remained steady or rose gently through the end of the stimulus period, and increased sharply in anticipation of a saccade into the cell's response field (Shadlen and Newsome, 2001). Interestingly, the slopes of the LIP responses during the initial rise did not vary with coherence. To quantify this, we computed slopes from 150 to 300 ms after motion onset (Fig. 3, gray shaded region). There was no consistent ordering of the slopes according to the coherence, during motion either toward T_{in} ($p = 0.08$, t test; see Materials and Methods) or toward T_{out} ($p = 0.71$). However, greater motion coherence was associated with higher average firing rates ($p < 10^{-4}$; linear regression) in the plateau or gently rising response period (500–1000 ms after motion onset; Fig. 3, solid curves). Population SDFs for T_{out} -choice trials (Fig. 3, dashed lines) show an analogous pattern in the opposite direction: firing rates declined starting about 200 ms after onset of motion, and the firing rate waned faster on trials with stronger T_{out} motion.

These differences in the SDFs related to stimulus strength are subtle relative to some previous studies (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Churchland et al., 2008; Kiani et al., 2008), in part because our range of motion coherences was constrained so that the monkeys' choices would be effectively biased by the prior (see Discussion).

LIP responses to arrow-cue onset

Next, we examined how LIP neurons responded to the onset of the arrow cue. Note that this cue was placed at the fovea, outside the neurons' response fields. If prior knowledge afforded by the cue affects the decision-making process by shifting the decision variable toward a bound, and if LIP activity represents this decision variable, then we expect the cue to modify LIP firing rates before onset of the motion stimulus. Alternatively, if the cue alters sensory representations of the relevant stimuli (i.e., motion signals in area MT), we expect no cue effects before the motion begins. The left panel in Figure 4A shows population SDFs aligned to cue onset for T_{in} -cue trials (green), neutral-cue trials (red), and T_{out} -cue trials (blue). Relative to neutral-cue responses, LIP neurons fired more strongly when the arrow cue specified T_{in} . In addition, there was a slight suppression in activity when the cue indicated T_{out} . The transient dip in activity after cue onset is reminiscent of the dip noted at motion onset during

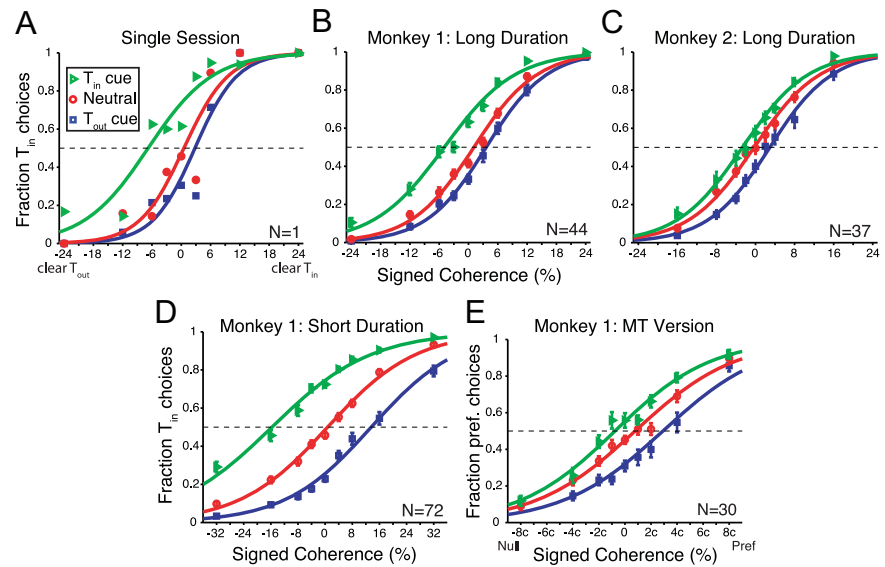


Figure 2. Prior expectations bias monkeys' perceptual reports. Psychometric functions are plotted as fraction T_{in} choices as a function of signed coherence, where large positive coherences indicate clear motion toward T_{in} , and large negative coherence indicates clear motion toward T_{out} . Data are plotted separately based on the direction of the cue. Data in **A–C** are from 1000 ms motion duration trials of the standard (LIP) version of the cued direction discrimination task. **A**, Single-session psychometric functions for Monkey 1. **B**, Average psychometric functions from 44 recording sessions for Monkey 1. **C**, Average psychometric functions from 37 recording sessions for Monkey 2. **D**, Average psychometric functions from 72 recording sessions from the short motion duration (250 ms) version of the task. Note the broader range of motion coherences in this version. **E**, Average psychometric functions from 30 recording sessions from the MT version of the task. Because the coherence range varied across sessions in this task, data are combined by ordinal coherence and plotted in terms of the lowest coherence for that session, denoted c (see Materials and Methods). Error bars indicate SEs. Curves are best-fitting logistic functions (Eq. 1). Dashed lines indicate chance performance.

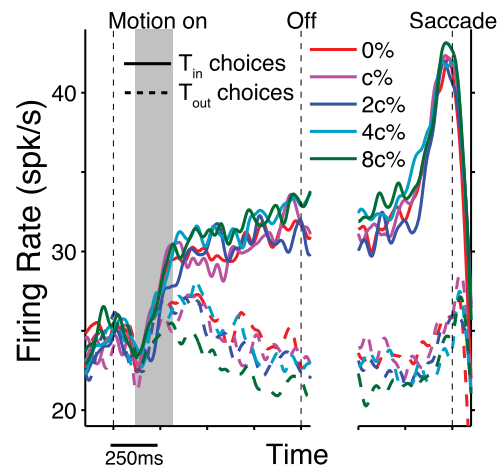


Figure 3. LIP responses to motion reflect choice and stimulus strength. Population SDFs grouped by choice and motion strength, aligned to motion onset (left) and saccade initiation (right). Solid and dashed curves represent trials in which the monkey chose T_{in} and T_{out} , respectively. Color indicates ordinal coherence, such that 0% denotes ambiguous motion and 8c% denotes the strongest motion coherence for each monkey (see Materials and Methods). Only correct, neutral-cue, 1000 ms motion trials are included here (all neutral-cue, 1000 ms motion trials were included for the 0% coherence curves). Vertical dashed lines mark the times of motion onset, motion offset, and saccade initiation. $n = 81$ cells from both monkeys performing the mixed-duration task. The gray vertical band denotes the time interval over which response slopes were computed.

previous studies (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Churchland et al., 2008; Kiani et al., 2008; Kiani and Shadlen, 2009; Rorie et al., 2010).

The effect of the cue on LIP responses before motion onset was very unlikely driven by a visual response to the cue itself. There

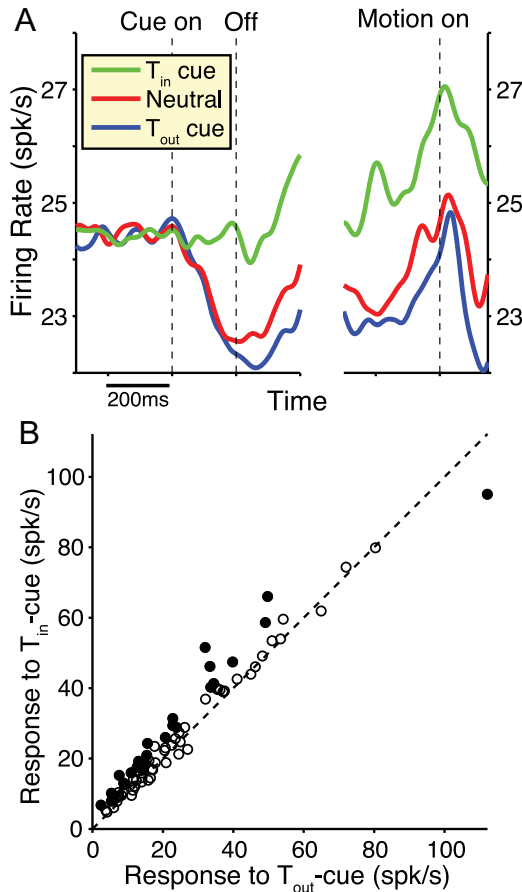


Figure 4. LIP responses diverge upon delivery of the arrow cue. **A**, Populations SDFs are shown, grouped by cue direction and aligned to either cue onset (left) or motion onset (right). Given the variable cue-motion ISI, the earliest motion stimuli appeared 150 ms after cue offset. All trials contributed to the SDFs regardless of the motion parameters. **B**, Firing rates over the final 100 ms before motion onset are plotted for T_{in} -cued trials versus T_{out} -cued trials, with each datum representing one neuron. Filled circles represent a significant difference between the firing rates ($p < 0.05$, t test). The dashed line represents the unity-slope diagonal. $n = 81$ cells from both monkeys performing the mixed-duration task.

was no systematic difference in the effect of cue direction on LIP responses during the 200 ms cue display as a function of response field location: the effect of cue direction for neurons with response fields close to the fovea ($< 10^\circ$ from the fovea) was only 0.81 spikes per second greater than that for neurons with response fields farther from the fovea. Moreover, this difference was not significant ($p = 0.32$, Mann–Whitney U test).

The right panel of Figure 4A shows population activity aligned to the onset of motion and indicates that response modulation due to the cue was present at the time of motion onset. To quantify the observation that the cue effect was present before motion onset, we averaged firing rates in the 100 ms period preceding motion onset. This time window was chosen to avoid any overlap with the cue display. With respect to neutral-cue trials, LIP firing rates during this interval were significantly amplified on T_{in} -cue trials relative to neutral cues ($p < 10^{-5}$, paired t test), with no significant effect on T_{out} -cue trials ($p = 0.22$). The majority of individual neurons (65 of 81) showed greater activity in this 100 ms window for T_{in} cues than for T_{out} cues (Fig. 4B). Among neurons with a significant difference in activity between the two conditions ($p < 0.05$, t test; Fig. 4B, filled circles), 27 of 28 neurons showed stronger activity on T_{in} -cue trials.

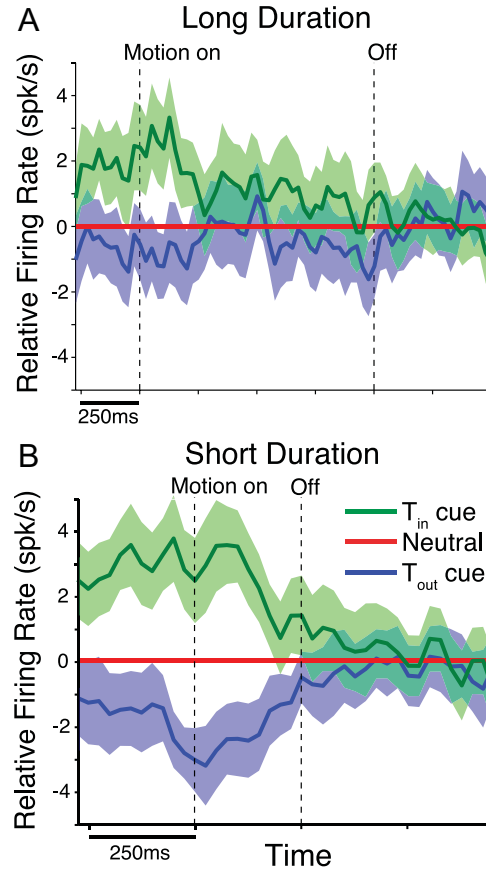


Figure 5. Time course of LIP responses reveals an effect of cue direction that wanes with time. Time courses reflect the difference in activity between trials of each cue type and the neutral-cue trials, averaged across neurons. Time courses were constructed using 50 ms windows slid across the data in 25 ms steps. The shaded areas reflect 95% confidence intervals across neurons. **A**, Time course of the cue effect constructed from the 1000 ms duration motion trials from the mixed-duration sessions ($n = 81$ cells from two monkeys). **B**, Time course constructed from separate sessions containing only 250 ms motion stimuli ($n = 72$ cells from one monkey). Note the different time scales on the abscissa.

Effect of cue on LIP responses during motion integration

The previous analyses demonstrate that both motion stimuli and arrow cues modulate LIP firing rates, and that these modulations were in the expected directions. However, if LIP activity truly reflects the integration of incoming sensory evidence with the prior, then cue-related modulations should be detected while the motion evidence is being evaluated.

To visualize the evolution of the cue-related response modulation during presentation of the motion stimulus, we computed a time course of the cue effect. The time course was quantified as the response (firing rate) on cued trials minus the response on neutral-cue trials, averaged across motion directions, coherences, and the monkey’s choices. To estimate how long the neural effect of the cue persisted, we generated 95% confidence intervals and asked when the effect was no longer distinguishable from zero. An effect of T_{in} cues is clearly present early on, but the effect fades over the course of the motion presentation (Fig. 5A, green). Comparing directly, the average effect of T_{in} cues decreased significantly between the initial and final 100 ms of the motion stimulus (T_{in} cue, $p < 0.003$; T_{out} cue, $p = 0.61$, two-tailed paired t tests).

Similar cue-related effects were observed for each monkey individually: T_{in} cues amplified LIP responses (relative to neutral-cue responses) over the first 100 ms of the motion stim-

ulus (Monkey 1, $p < 0.001$; Monkey 2, $p < 0.03$, paired t tests), and effects were greater in the first 100 ms than in the last 100 ms (Monkey 1, $p < 0.02$; Monkey 2, $p < 0.05$, paired t tests).

In contrast, T_{out} cues were associated with a slight suppression of firing rate relative to neutral-cue trials (Fig. 5A, blue). When averaged across the entire motion period, the effect approached significance ($p = 0.051$, paired t test).

If the cue effect were implemented by changing the threshold of accumulated evidence, the integration-to-bound model would predict different plateau firing rates late in the motion presentation. This would manifest as a persistent difference in response between cue types. However, cue direction has no significant effect on LIP firing rates during the final 250 ms of the motion stimulus ($p = 0.15$, one-factor repeated-measures ANOVA), arguing against a change in the neuronal decision threshold. Visually, this is supported by Figure 5A, in which the firing rates for T_{in} and T_{out} cues converge toward the firing rate for the neutral cue late in the trial, once the decision has been made.

We confirmed these cue effects in a separate data set collected from Monkey 1. In these sessions, the monkey was only exposed to trials with 250 ms motion durations. Data from these trials were analyzed in an identical manner to the long-duration data set, yielding another time course of the cue effect (Fig. 5B). The same general pattern held: T_{in} cues increased responses before motion onset (relative to neutral cues), whereas T_{out} cues decreased responses. These effects peaked near the onset of motion and decayed during the stimulus period. Compared to the long-duration stimuli, the effect of the cue on LIP responses during the 100 ms period before motion onset was significantly greater for both T_{in} cues ($p < 0.05$, Mann–Whitney U test) and T_{out} cues ($p < 0.04$). The more pronounced effects of the cue on short-duration LIP responses match nicely with the more pronounced behavioral effects in the short duration task (Fig. 2D). Hence, when less motion information is available, the cue exerts a greater effect on both LIP responses and behavior.

This result prompted us to examine whether there is a correlation between the strength of neural and behavioral cue effects across sessions. The mixed-duration data set revealed a near-significant association between behavioral bias and the neural effect of the cue (Spearman $r = 0.21$, $p = 0.06$) (see Materials and Methods). Overall, this analysis suggests a functional link between the neuronal and behavioral data.

Diffusion model of cue effects

Our neural data argue that prior knowledge that varies from trial to trial can indeed modulate LIP activity, and that these effects persist well into presentation of the motion stimulus. One might argue that the gradual disappearance of cue effects during the motion stimulus suggests that LIP might not reflect how priors are combined with sensory evidence to form decisions. However, the absence of cue-related activity late in the trial does not necessarily imply that the early representation of priors in LIP is discarded. Rather, as demonstrated in the following simulations, waning of the cue effect with time is expected if the decision process is compatible with an integration-to-bound framework.

We simulated the monkey's choices using an integration-to-bound model to investigate the effect of the cue over time (see Materials and Methods). Although our LIP data did not reveal strong sorting of responses by coherence (Fig. 3), as predicted by an integration-to-bound model, this framework has proven useful in modeling behavior and LIP responses in similar tasks (Kiani et al., 2008). It is possible that our use of a restricted range of coherences limited our ability to observe the coherence depen-

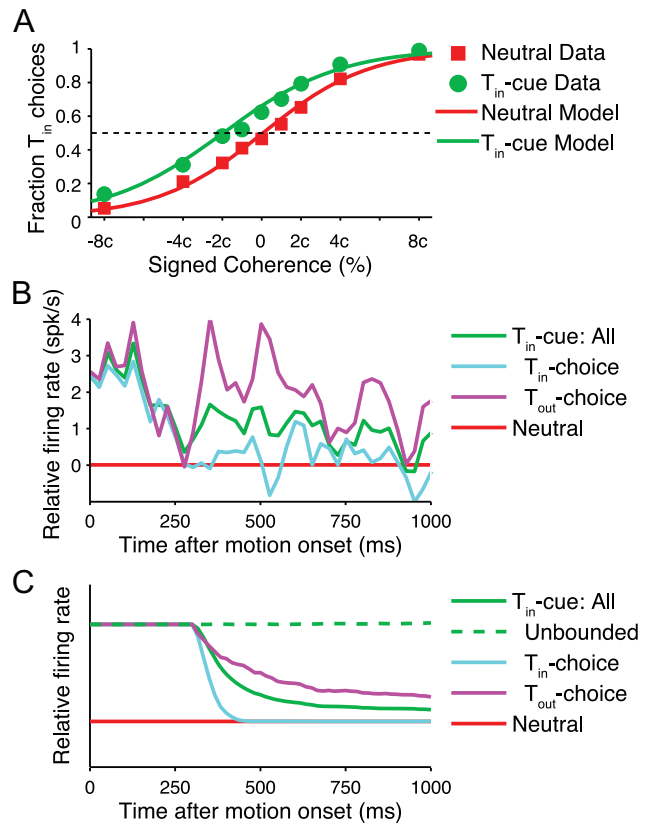


Figure 6. Integration-to-bound simulation predicts waning cue effect. **A**, Red and green symbols depict measured psychometric discrimination performance as a function of signed coherence for neutral-cue and T_{in} -cue trials, respectively. Data include only 1000 ms duration motion trials from the mixed-duration task, averaged across both monkeys. Smooth curves show the model performance after it is fit to the data. The dashed line indicates chance performance. **B**, Empirical time courses of the cue effect, expressed as the difference in SDF between each directional-cue condition and the corresponding neutral-cue condition. The time courses were computed across all trials with T_{in} cues (green), as well as for only those T_{in} -cued trials in which the monkey chose T_{in} (cyan) or T_{out} (magenta). **C**, Time courses for the model predictions plotted in an analogous fashion, constructed using average decision-variable as a proxy for average firing rate. The green dashed curve is constructed from an alternative integration model that does not include a bound (see Materials and Methods).

dence of responses. In our simulations, the effect of the cue was modeled as an initial offset in the decision variable associated with the cued direction, analogous to the effect observed in LIP. We did not implement the cue effect as a sensory gain modulation because our data from area MT do not support this possibility, as shown in the following section. For simplicity, this simulation only modeled neutral and T_{in} -cued trials, as the effect of T_{out} cues on LIP activity was modest (Fig. 5A, blue). Note that the simulated integration process was not leaky, implying that any effect of the cue in the model should influence the decisions. We fit the simulation parameters to match the monkeys' averaged behavioral data for neutral-cue and T_{in} -cue trials, and Figure 6A shows that the model fits (smooth curves) accurately characterize the fraction of T_{in} choices made by the monkeys (symbols) for both neutral-cue (red) and T_{in} -cue (green) trials.

Having found simulation parameters that reproduce the behavior, we looked inside the simulation at the accumulated evidence in favor of T_{in} —an observation akin to measuring peristimulus time histograms (PSTHs) from an LIP neuron. A time course of the cue effect was derived from these simulated PSTHs using the same methods used to analyze the neural data (Fig. 5A) (see Materials and Methods). The resultant time course

(Fig. 6C, solid green) shows that the cue effect is strong early in the trial but fades toward zero as motion information arrives, paralleling the decline of the cue effect seen in the neural data (Fig. 6B, green, reproduced from Fig. 5A). Although the exact temporal profile of the simulated cue effect does not match the decline measured in neurons, we made no attempt to tune the model to fit the time course, except for including a delay parameter that reflects motion signal latencies in LIP.

The basic intuition for this result is as follows. According to the integration-to-bound model, after activity reaches a bound, it remains there until the monkey indicates its choice (Shadlen and Newsome, 2001; Roitman and Shadlen, 2002; Kiani et al., 2008; Rorie et al., 2010). Although T_{in} -cue trials generally reach the bound earlier, firing rates on both neutral-cue and T_{in} -cue trials eventually reach the same bound, indicating that the separation between the firing rate trajectories—the cue effect—vanishes late in the trial (Huk and Shadlen, 2005). In Bayesian terms, the prior information is not separately represented in late-trial LIP activity, but is instead merged with the sensory evidence to form the posterior probability.

This argument leads to two predictions. First, if the fading of the neural effect of the cue is dependent on a bound, then a decision-making model that does not implement a bound should not show this decline. To test this prediction, we fit our behavioral data with an unbounded model in which a choice is rendered on each trial based on which direction has more accumulated evidence at the end of the motion stimulus (see Materials and Methods). As expected, the resultant time course of the cue effect indeed shows no decay in the absence of a bound (Fig. 6C, green dashed line). To evaluate how well the unbounded and bounded integration models fit the behavioral data, we linearized the data with an inverse logistic transform and computed the correlation coefficients between the data and predictions of each model. The behavioral data were well captured by both the bounded model ($r = 0.950$) and the unbounded model ($r = 0.954$). Thus, the different time courses of the cue effect between the two models cannot be attributed to model fits of differing quality.

A second prediction from the model relates to an asymmetry between the two choices that is apparent in our data as well as in the existing literature: the bound in the neural activity manifests clearly only for choices toward T_{in} (Roitman and Shadlen, 2002; Churchland et al., 2008; Kiani et al., 2008). Therefore, the model predicts the decline of the neural cue effect to be more pronounced during trials in which the monkey chooses T_{in} compared to trials for which the animal chooses T_{out} . Figure 6C shows the predicted time course of the cue effect, computed separately for trials in which the simulation chose T_{in} (cyan) or T_{out} (magenta). The cue effect fades faster when T_{in} was chosen. Returning to the neural data from LIP, we constructed separate neuronal time courses for each chosen target and discovered the same pattern: the cue effect fades faster when the monkeys choose T_{in} (Fig. 6B, cyan) than when they choose T_{out} (magenta). Thus, if one assumes that the bounded integration framework is correct, these data strongly support the notion that prior expectations and incoming sensory evidence are merged in the firing rates of LIP neurons, and that this merger is associated with a bias in the monkey's choices.

Effect of cues on sensory gain

Although the LIP recordings suggest that cueing the probable direction of motion shifts the decision variable toward the appropriate bound, it remains possible that prior knowledge might also

amplify sensory responses related to more likely stimuli. This would manifest as a change in the slope of the rising phase of LIP activity, increasing the slope when motion toward T_{in} is more likely and decreasing the slope when motion toward T_{out} is more likely.

We tested this idea in the main data set, looking at the effect of cue direction on the slope of the steeply rising phase of neural responses (150–350 ms after motion onset) (Fig. 3, gray band) for trials in which the monkey correctly chooses T_{in} . A repeated-measures ANOVA with coherence and cue direction as factors showed that motion coherence did not significantly modify the slope of the response ($p = 0.56$), but cue direction did indeed significantly alter the slope ($p = 0.001$). However, the rank ordering of the mean slopes was opposite of what would be required to explain the behavioral bias: T_{in} cues decreased the slope of the response relative to neutral cues and T_{out} cues in our data set. This unexpected effect may have been related to the fact that a smaller firing rate change is required to reach the bound on T_{in} cues. Repeating the analysis using slightly different time windows never yielded firing rate slopes that were greater on T_{in} -cued trials than on other trials. During the later phase of gently rising responses (500–1000 ms after motion onset), there was no significant effect of motion coherence ($p = 0.97$) or cue direction ($p = 0.13$, repeated-measures ANOVA) on the slope of the neural response.

The presence of the shifted baseline response hampers our ability to detect slope changes in LIP responses. Therefore, to test more directly whether priors modify sensory representations, we recorded from neurons in area MT, which represent the direction of motion in fixation or discrimination tasks (Maunsell and Van Essen, 1983; Britten et al., 1992). During MT recordings, the monkey performed a version of the cued direction discrimination task in which the motion stimulus was tailored to maximally drive MT responses (Fig. 7A). Importantly, the axis of discrimination was aligned to the neuron's preferred direction. Although there were some differences between the LIP and MT tasks in terms of timing and spatial layout, the essential elements were the same with respect to the partial validity of the arrow cue (see Materials and Methods).

The monkey's behavior in the MT version of the task (Fig. 2E) was similar to the behavior seen in the LIP version (Fig. 2B). The monkey clearly based his choices on the motion stimulus, since stronger motion in either direction elicited more choices in that direction. Furthermore, the arrow cues biased the monkey's choices in the direction of the cue ($p < 10^{-5}$, logistic regression analysis of pooled data; $n = 30$).

Given a clear behavioral effect in this version of the task, we tested whether the direction of the cue modulated MT responses, as might be expected under the sensory gain hypothesis. For each cue direction, MT firing rates during the motion stimulus were normalized and plotted as a function of signed coherence. Figure 7B shows population response functions for the three cue conditions. MT activity increased sharply with the strength of coherent motion in the preferred direction, consistent with previous findings (Britten et al., 1993). Strikingly, the neural response functions for the different cue directions are almost completely overlapping. A repeated-measures ANOVA indicates no significant effect of cue direction on MT responses ($p = 0.29$, main effect of cue direction; $p < 0.0001$, main effect of signed coherence). The absence of cue effects was apparent at the single neuron level as well, with a significant main effect of cue direction found in only 1 of 30 MT neurons (3.3%)—a rate indistinguishable from chance (5%).

Although the MT responses did not match the predictions of the sensory-gain hypothesis, Figure 7B leaves open the possibility that cues may alter MT activity when the direction of the cue is in conflict with the direction of the motion stimulus: responses appear to be increased on preferred-cue trials with null-direction motion, and may also be slightly reduced on null-cue trials with preferred-direction motion. However, statistical testing failed to reveal a significant main effect of conflict ($p = 0.83$, two-factor repeated-measures ANOVA with factors of signed coherence and conflict; see Materials and Methods). Even looking only at conditions that appear to be most influenced by the cue failed to reveal a significant effect: using only trials with strong motion in the null direction (i.e., $-4c\%$ and $-8c\%$), a two-factor repeated-measures ANOVA with factors of cue direction and coherence showed no significant effect of cue direction ($p = 0.25$). In sum, our analyses show no convincing evidence that cues modify MT responses during the motion presentation.

One difference between the two versions of the task was that the visual stimulus was located peripherally in the MT version of the task (Fig. 7A) but was placed at the fovea in the LIP version (Fig. 1A). If foveal stimuli were more conducive to measurable neural cue effects, we would expect a more pronounced effect for MT neurons with receptive fields that lay closer to the fovea. We tested this hypothesis in our data, dividing our MT neurons according to the median receptive field eccentricity (10.1° from the fovea). The average cue effect during the motion period (see Materials and Methods) did not differ significantly with the eccentricity of the receptive field. Specifically, this analysis yielded a 0.03 spikes per second greater cue effect for neurons located further from the fovea, a nonsignificant effect ($p = 0.51$, Mann–Whitney U test).

During the delay period before motion onset, cue type did not modulate MT activity. To demonstrate this, we first divided the delay-period response of each cell on cued trials by that cell's response to neutral-cue trials. Figure 7C shows that these normalized responses do not differ significantly with the direction of the cue ($p = 0.09$ main effect of cue direction by one-factor repeated-measures ANOVA), with the greatest responses seen on neutral-cue trials. Therefore, the MT recordings support the conclusion derived from the LIP experiments, namely, that prior information does not influence behavior by amplifying sensory responses to motion.

Although this data set provides no evidence for an effect of arrow cues on MT responses, it is possible that our sample was underpowered to detect an effect. To address this concern, we took a simulation approach (see Materials and Methods). We measured the means and variances of the neutral-cue firing rates and added a perturbation to the means to mimic small effects of cue direction. We then generated artificial MT responses from




Figure 7. MT responses are not affected by the cue. **A**, Task design for the MT version of the cued direction discrimination task. The format is similar to that of Figure 1. Note that motion stimuli were placed eccentrically within the recorded MT neuron's receptive field. "Pref" denotes the MT neuron's preferred direction of motion, and "Null" denotes the opposite direction. **B**, Neural response functions constructed as population average firing rates as a function of signed coherence, sorted by cue direction. Each neuron's firing rates are normalized to the firing rates on neutral-cue, zero-coherence trials before averaging. Signed coherences are ordinal and are expressed in units of c , the lowest nonzero coherence. **C**, Population delay-period firing rates as a function of cue direction. Firing rates are normalized to the neutral-cue firing rate. All error bars indicate SE. $n = 30$ cells from one monkey performing the MT version of the task.

these modified distributions and performed the repeated-measures ANOVAs on these simulated data. Given our collection of 30 neurons, our analysis would have a 90% chance of detecting effect sizes as small as a difference of 0.65 spikes between preferred-cue trials and neutral-cue trials. If preferred cues added merely 0.38 spikes to MT responses, then our data set would have had a 50% chance to detect the difference. These analyses argue that our experiment had sufficient power to detect even very weak effects of the cue on MT firing rates.

Finally, although the arrow cue did not appear to directly modulate MT responses, it might have modulated the contribution of MT responses to perceptual decisions. Previous studies have demonstrated weak trial-by-trial correlations between firing rate fluctuations and choices, a phenomenon quantified using choice probabilities (Britten et al., 1996; Uka and DeAngelis, 2004; Purushothaman and Bradley, 2005; Cohen and Newsome, 2009). If choice probabilities reflect a causal contribution of MT activity to the animal's decision, then factors that influence choices downstream of MT, such as prior probabilities in our experiment, ought to degrade choice probabilities. In our data set, the average choice probability was not significantly different

from chance for valid cues (CP, 0.51; $p = 0.11$, t test). In contrast, the grand choice probability was significantly greater than chance for neutral-cue trials (CP, 0.53; $p < 0.03$), consistent with previous literature (Britten et al., 1996). This difference in choice probability between valid and neutral cues, although modest, is consistent with the cue acting at a stage of processing downstream from area MT (such as LIP).

Discussion

We examined neural computations that allow prior knowledge to influence perceptual decisions. Two monkeys performed a cued direction discrimination task, in which judgments of motion direction were biased by a partially predictive arrow cue that changed from trial to trial. LIP recordings showed that altering the prior probability for a particular direction of motion produces shifts in firing rates that persist into the motion presentation. In the context of integration-to-bound models, this finding suggests that priors influence choices by bringing the decision variable closer to the likely bound.

This implementation is parsimonious because it allows priors and sensory evidence to be treated equivalently as probabilities. Just as it has been argued that the overall decision variable evolves over time by adding the log probabilities of the sensory evidence (Gold and Shadlen, 2007), this study supports the notion that a representation of the prior contributes to the decision variable by an additive process. Such an implementation could be plausibly implemented by realistic neuronal populations (Ma et al., 2006; Beck et al., 2008), and empirical data have suggested that LIP activity reflects log probabilities (Yang and Shadlen, 2007).

If LIP reflects a true decision variable, then any manipulation that biases decisions should affect LIP responses. Rorie et al. (2010) used asymmetric rewards to bias decisions in a similar motion discrimination task. Like our result, they observed that LIP responses were shifted before motion. Furthermore, these shifts faded during the motion display, although the timing of their effects did not exactly match ours. Overall, the similarity in the results, despite very different manipulations of prior probability, further supports the idea that decision circuitry in LIP may generally combine bias signals with sensory evidence. The neural instantiation of the bias signal itself is an important question meriting further study.

We also examined the possibility that prior knowledge influences choices by scaling sensory representations of visual motion. In the integration-to-bound framework, this would manifest as a change in the slope of LIP firing rates, but our LIP recordings did not support this prediction. We also took a less model-dependent approach to evaluating the sensory gain hypothesis by recording directly from MT neurons in one monkey. Consistent with the LIP results, MT responses were unchanged by the arrow cue. To the extent that MT provides the critical sensory evidence for this task, our findings suggest that the dynamic prior in our task was mainly manifest at the level of decision circuitry (e.g., LIP).

In apparent contrast with our findings, recent studies have suggested that priors influence low-level sensory neural representations (Fischer and Pena, 2011; Girshick et al., 2011). Fischer and Pena (2011) reported that the distribution of preferred interaural time differences (ITDs) in owl tectum neurons reflects the prior probability of finding prey at the spatial location corresponding to that ITD. One possible explanation is that the implementation of priors depends on the time-scale of the prior knowledge base. It may be sensible to implement priors in the distribution of tuning properties of sensory neurons when the prior reflects statistics of the environment that are stable on long time scales, as in

the owl tectum experiments. In contrast, having the prior implemented downstream of the sensory representation may allow decisions to accommodate rapidly changing priors, as in our cued discrimination task. Thus, we suggest that static or slowly varying priors might be implemented at the level of sensory representations, whereas rapidly varying priors (or rapid shifts between a small set of static priors) should act at the level of decision circuitry.

Trial-by-trial correlations between MT responses and choices (choice probabilities) have been described previously (Britten et al., 1996; Dodd et al., 2001; Uka and DeAngelis, 2004; Purushothaman and Bradley, 2005) and were seen in our neutral-cue condition. Choice probabilities may imply that MT responses contribute causally to perceptual decisions, a view referred to as the “bottom-up” interpretation of choice probability. An alternative, “top-down” interpretation posits that after the decision has been made, MT receives inputs that reflect the choice on that trial (Krug, 2004). Our finding, that the average choice probability is no longer significantly greater than chance when the monkey is cued to expect motion is consistent with the bottom-up view, as any influence of the cue on decisions downstream of the MT would be expected to weaken choice probabilities. However, our results certainly do not preclude a top-down contribution to choice probabilities, for which there is evidence (Nienborg and Cumming, 2009). A purely top-down scheme could explain our results if the presence of a cue gates the feedback signal that is sent to the MT, although we consider this explanation less parsimonious.

The lack of effect of the arrow cue on MT responses in our task is notable given that a previous study showed selective responses to static shapes in MT when animals had associated those shapes with particular directions of motion (Schlack and Albright, 2007). Although MT activity was unchanged by priors, the MT may still veridically reflect the percept of the stimulus, without reflecting the animal's choice. For example, the monkey might perceive the direction of a weakly coherent stimulus as incrementally favoring the uncued direction, but may nevertheless choose the direction that is more likely based on the prior.

Relation to previous studies of priors and LIP

A small number of previous studies have examined how priors influence neural representations of a decision in LIP, and there are both similarities and important differences with our findings. Platt and Glimcher (1999) manipulated the prior probability of the target location in a saccade task and found that LIP firing rates increased with the probability that the saccade target was located in the response field. Like in our task, this initial increase was maximal before target onset and decayed leading up to the actual saccade. However, unlike in our task, knowledge of the prior in that study did not contribute to successful performance. Rather, the influence of the prior reflected the temporary statistics of the environment.

In contrast, Gold et al. (2008) failed to find an effect of priors on LIP activity (or MT responses). In monkeys learning to discriminate motion direction, the monkey's recent history acted as a prior distribution that biased choices. Gold et al. (2008) did not find a significant neural correlate of this dependence in LIP responses, but they acknowledge that this lack of effect could reflect a lack of statistical power.

Most recently, a markedly different result was reported by Hanks et al. (2011) in the context of a reaction-time version of the direction discrimination task. By manipulating the prior probabilities of the two opposite motion directions over blocks of 200–1000 trials, they too found a small shift in baseline LIP responses that depended on prior probability. However, they found that

this bias signal grows, not fades, over the course of the trial. This striking difference in the time course of the prior, compared with our finding, may derive from key differences in task design. In the reaction-time task of Hanks et al. (2011), a longer elapsed time implies a weaker motion stimulus. In this case, it makes sense to put greater weight on the prior as time proceeds. In our fixed-duration task, there is no such relationship between elapsed time and motion strength, and the termination of the decision process diminishes the measured effect of priors.

Together, these studies suggest that neural manifestations of priors may be sensitive to task demands and the time scale over which priors are stable, and further research is needed to dissect these possible influences. Our study is unique in demonstrating a single-unit correlate of stimulus priors that vary rapidly from trial to trial and appear to be manifest only at the level of decision circuitry and not lower-level sensory representations.

LIP responses and integration-to-bound models

Much of our analysis and simulation has used the integration-to-bound framework. This model has elegantly explained choice behavior in numerous studies, including situations in which an integration-to-bound framework does not maximize rewards (Kiani et al., 2008; Rorie et al., 2010). However, it remains unclear to what extent LIP activity represents the integration process. LIP responses in prior studies have not always reflected perfect integration of sensory evidence. For example, studies have shown only modest sorting of slopes by coherence (Kiani et al., 2008), or nonmonotonic firing rates (Roitman et al. 2002). Furthermore, these same studies have argued that LIP activity also represents other signals in these tasks, including urgency signals (Churchland et al., 2008; Hanks et al., 2011) and decision confidence (Kiani et al., 2009). Details of task design may therefore determine the extent to which LIP activity clearly reflects an integration process. Our study may have been limited by the modest range of motion coherences used in the task, which was necessary to ensure a balanced use of the cue and motion signals. Even if our study suggests future refinements in our understanding of LIP as an integrator, the fading of the cue-related signal supports the notion that LIP activity reflects a bounded process.

References

- Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, Shadlen MN, Latham PE, Pouget A (2008) Probabilistic population codes for Bayesian decision making. *Neuron* 60:1142–1152.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10:1157–1169.
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13:87–100.
- Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. *Nat Neurosci* 11:693–702.
- Cisek P (2007) A parallel framework for interactive behavior. *Prog Brain Res* 165:475–492.
- Cohen MR, Newsome WT (2009) Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. *J Neurosci* 29:6635–6648.
- DeAngelis GC, Newsome WT (2004) Perceptual “read-out” of conjoined direction and disparity maps in extrastriate area MT. *PLoS Biol* 2:E77.
- Dodd JV, Krug K, Cumming BG, Parker AJ (2001) Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. *J Neurosci* 21:4809–4821.
- Fischer BJ, Peña JL (2011) Owl’s behavior and neural representation predicted by Bayesian inference. *Nat Neurosci* 14:1061–1066.
- Girshick AR, Landy MS, Simoncelli EP (2011) Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nat Neurosci* 14:926–932.
- Gold JL, Shadlen MN (2007) The neural basis of decision making. *Annu Rev Neurosci* 30:535–574.
- Gold JL, Law CT, Connolly P, Benuer S (2008) The relative influences of priors and sensory evidence on an oculomotor decision variable during perceptual learning. *J Neurophysiol* 100:2653–2668.
- Hanks TD, Ditterich J, Shadlen MN (2006) Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat Neurosci* 9:682–689.
- Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN (2011) Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J Neurosci* 31:6339–6352.
- Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci* 25:10420–10436.
- Kiani R, Shadlen MN (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* 324:759–764.
- Kiani R, Hanks TD, Shadlen MN (2008) Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J Neurosci* 28:3017–3029.
- Krug K (2004) A common neuronal code for perceptual processes in visual cortex? Comparing choice and attentional correlates in V5/MT. *Philos Trans R Soc Lond B Biol Sci* 359:929–941.
- Link SW, Heath RA (1975) A sequential theory of psychological discrimination. *Psychometrika* 40:77–105.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438.
- Maunsell JH, Van Essen DC (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *J Neurophysiol* 49:1127–1147.
- Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. *Cereb Cortex* 13:1257–1269.
- Nienborg H, Cumming BG (2009) Decision-related activity in sensory neurons reflects more than a neuron’s causal effect. *Nature* 459:89–92.
- Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238.
- Premereur E, Vanduffel W, Janssen P (2011) Functional heterogeneity of macaque lateral intraparietal neurons. *J Neurosci* 31:12307–12317.
- Purushothaman G, Bradley DC (2005) Neural population code for fine perceptual decisions in area MT. *Nat Neurosci* 8:99–106.
- Reynolds JH, Chelazzi L (2004) Attentional modulation of visual processing. *Annu Rev Neurosci* 27:611–647.
- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475–9489.
- Rorie AE, Gao J, McClelland JL, Newsome WT (2010) Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS ONE* 5:e9308.
- Schlack A, Albright TD (2007) Remembering visual motion: neural correlates of associative plasticity and motion recall in cortical area MT. *Neuron* 53:881–890.
- Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86:1916–1936.
- Smith PL, Ratcliff R (2004) Psychology and neurobiology of simple decisions. *Trends Neurosci* 27:161–168.
- Treue S (2001) Neural correlates of attention in primate visual cortex. *Trends Neurosci* 24:295–300.
- Uka T, DeAngelis GC (2003) Contribution of middle temporal area to coarse depth discrimination: comparison of neuronal and psychophysical sensitivity. *J Neurosci* 23:3515–3530.
- Uka T, DeAngelis GC (2004) Contribution of area MT to stereoscopic depth perception: choice-related response modulations reflect task strategy. *Neuron* 42:297–310.
- Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* 447:1075–1080.