

Perceptual Decision-Making as Probabilistic Inference by Neural Sampling

Highlights

- Task structure constrains feedback signals in a probabilistic inference framework
- These feedback signals cause choice probabilities and noise correlations
- Task-dependent noise correlations are predicted to increase with learning
- In our neural sampling-based representation, feedback signals cause a confirmation bias

Authors

Ralf M. Haefner, Pietro Berkes, József Fiser

Correspondence

ralf.haefner@gmail.com

In Brief

Feedback signals are ubiquitous in cortex yet underconstrained by empirical data. Haefner et al. derive predictions for their effect on sensory representations as a function of the behavioral task, with implications for the role of correlated variability in sensory coding.

Perceptual Decision-Making as Probabilistic Inference by Neural Sampling

Ralf M. Haefner,^{1,*} Pietro Berkes,² and József Fiser³

¹Brain & Cognitive Sciences, University of Rochester, Rochester, NY 14627, USA

²Sloan-Swartz Center for Theoretical Neurobiology, Brandeis University, Waltham, MA 02454, USA

³Department of Cognitive Science, Central European University, Budapest 1051, Hungary

*Correspondence: ralf.haefner@gmail.com

<http://dx.doi.org/10.1016/j.neuron.2016.03.020>

SUMMARY

We address two main challenges facing systems neuroscience today: understanding the nature and function of cortical feedback between sensory areas and of correlated variability. Starting from the old idea of perception as probabilistic inference, we show how to use knowledge of the psychophysical task to make testable predictions for the influence of feedback signals on early sensory representations. Applying our framework to a two-alternative forced choice task paradigm, we can explain multiple empirical findings that have been hard to account for by the traditional feedforward model of sensory processing, including the task dependence of neural response correlations and the diverging time courses of choice probabilities and psychophysical kernels. Our model makes new predictions and characterizes a component of correlated variability that represents task-related information rather than performance-degrading noise. It demonstrates a normative way to integrate sensory and cognitive components into physiologically testable models of perceptual decision-making.

INTRODUCTION

Almost 150 years ago, Helmholtz proposed that visual perception is an inference process (Helmholtz, 1867). He suggested that at any point in time the brain combines its prior knowledge about the external world with the incoming sensory information to compute the best explanation for its inputs. Substantial evidence from both psychophysics (Kersten et al., 2004) and recent physiological findings (Berkes et al., 2011) suggests that the responses of visual neurons are influenced by such prior information, also called an internal model. Since sensory information is usually uncertain and often incomplete, this proposed inference and the assumed internal model both need to be probabilistic in nature (Fiser et al., 2010; Pouget et al., 2013). Furthermore, it has been noted that probabilistic inference in a hierarchical model requires a flow of information remarkably similar to the one within the visual system: feedforward from the retina and feedforward,

recurrent, and feedback within the cortex (Mumford, 1992; Lee and Mumford, 2003). However, the major challenge to generating neurophysiologically testable predictions based on probabilistic inference, and for incorporating feedback in this framework, has been that the nature of the internal model that the brain uses for general vision is currently unknown.

We overcame this challenge by studying the consequences of probabilistic inference in a well-controlled task in which the generative model of the sensory inputs is under the control of the experimenter. For the subject to perform probabilistic inference in such a task, the subject's brain needs to learn this experimenter-defined generative model. Learning such a model can be interpreted as a perturbation on the unknown internal model for general vision. We can therefore use our knowledge of this experimenter-defined perturbation to make predictions about how the neural responses should change as a result of it. We applied these ideas to the well-studied two-alternative forced choice (2AFC) task paradigm and, assuming that the brain performs inference by neural sampling (Hoyer and Hyvärinen, 2003; Fiser et al., 2010), we generated predictions for both the responses of sensory neurons and for psychophysical measurements. We found that our probabilistic inference model correctly reproduces key experimental observations in 2AFC tasks: the task dependency of noise correlations (Cohen and Newsome, 2008), the temporal increase of the correlation between sensory responses and behavior (called choice probability [CP]; Cohen and Newsome, 2009; Nienborg and Cumming, 2009; Nienborg et al., 2012), and the decrease in the correlation between stimulus and behavior (called psychophysical kernel [PK]; Neri et al., 1999; Ahumada, 2002) in some tasks (Nienborg and Cumming, 2009), but not others (Brunton et al., 2013) (Figure 1).

We also derived a number of specific predictions based on our model that can be readily tested in empirical studies. First, we predict that the noise correlations for sensory responses in a 2AFC task have two maxima and two minima whose locations are defined by the task-relevant stimuli. We further predict an increase in the amplitude as defined by those maxima and minima on the timescale of perceptual learning. Psychophysically, we predict that the relative weighting of evidence throughout the trial should depend on stimulus strength. Since our model reflects the structure of the task and was not designed to fit any of the existing observations made in specific experiments, we predict that our findings generalize to other stimuli, modalities, and sensory areas in the brain. In general, the fact that our framework's predictions directly reflect the spatial and temporal structure of

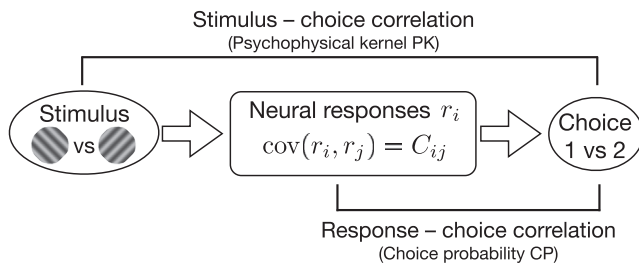


Figure 1. Schematic of the Perceptual Decision-Making System Modeled in This Paper

Information about a visual stimulus is represented by the responses r_i of a population of sensory neurons whose activity influences the behavioral choice of the subject. This paper focuses on the following three observables in this system: (1) the structure of the correlations between neural responses (noise correlation) believed to be important for the amount of information that can be represented by neural responses (Zohary et al., 1994; Abbott and Dayan, 1999; Ecker et al., 2011); (2) the time course of the magnitude of the choice-triggered stimulus average, also called psychophysical kernel (PK), which quantifies the influence of the stimulus on the final decision as a function of the time when the stimulus is presented during the trial (Neri et al., 1999; Ahumada, 2002; Nienborg and Cumming, 2009; Brunton et al., 2013); and (3) the correlation between a single neuron's response and the final choice, usually quantified as choice probability (CP; Britten et al., 1996; Nienborg et al., 2012; Haefner et al., 2013).

the experimenter-defined task makes them easy to test. Importantly, since its predictions concern the full statistical structure of neural responses, including higher-level correlations, our framework provides theory-driven guidance on how to analyze the high-dimensional data obtained from increasingly common population-recording techniques.

RESULTS

A basic assumption in our framework is that sensory neurons in the cortex represent the brain's belief about some aspect of the outside world (Lee and Mumford, 2003). In the case of primary visual cortex (V1), we assumed this to be how well a Gabor-shaped feature describes the image on the retina at the location of the neuron's receptive field (Olshausen and Field, 1997). From that follows that the activity of V1 neurons will depend *both* on the image on the retina *and* on any prior knowledge about the outside world in the rest of the brain. The former, called likelihood in probabilistic inference, is communicated to V1 via feedforward connections, while the latter, the prior, arrives via feedback connections. Training a subject on a given psychophysical task induces such prior knowledge in the subject's brain about the stimulus that is being presented. In this paper, we show how to use knowledge about the task structure to make predictions about its influence on V1 responses.

Qualitative Intuitions

Let's assume a simple coarse orientation discrimination task where the subject has to report whether either a vertical or a horizontal grating is embedded in a noisy image. Then, on any given trial, if the subject believes the vertical grating to be present, neurons representing vertical orientations will be more active than baseline, and neurons representing horizontal orientations

will be less active, *even if there is no signal in the stimulus at all*. This means that those neurons will be correlated with the choice (commonly quantified as CPs, Figure 6B). Furthermore, as a side effect of this correlation between neural response and perceptual decision, the correlations between the responses of neurons supporting the same choice will be higher than between the responses of neurons supporting opposing choices (Figure 4). Finally, since the internal belief about the correct choice develops gradually by integrating the sensory information over the course of a trial, the magnitude of those correlations also should increase over the course of a trial (Figure 6). As a side effect, the dynamic feedback about the current belief about the correct choice on the sensory responses may, in the case of weak external inputs, lead to a self-reinforcing loop with the result that early evidence has a bigger influence on the final decision than late evidence (Figure 7).

The Model

To make the above intuitions quantitative, we proceed in three steps. First, we need to make an assumption about the generative model that the brain has learned and in which it performs inference. Second, we need to make an assumption about how the variables in this model are related to sensory neurons. Third, we need to specify how the brain's beliefs are related to the responses of those neurons.

The generative model of a psychophysical task consists of the relationship between the available behavioral choices and the stimulus, as well as the prior probabilities of each choice being the correct one. In 2AFC tasks, the behaviorally relevant variable is binary and related to one of two possible choices. In visual tasks, the sensory input is the luminance pattern on the retina. For concreteness, we present our results in the context of an orientation discrimination task (Figure 2A). However, our results are also applicable to 2AFC tasks based on other stimuli (e.g., motion and disparity) or modalities (e.g., auditory and vestibular). We compared our model predictions with empirical data from a coarse motion direction discrimination task (Cohen and Newsome, 2008, 2009), a coarse disparity discrimination task (Nienborg and Cumming, 2009), and an auditory task (Brunton et al., 2013).

First, the correct decision, D , is drawn from a prior distribution (here we assumed $p(D=1) = p(D=2) = 0.5$; Figure 2B). D determines which one of two target images (here grating images) will be embedded in Gaussian noise. We represent the presence or absence of each target by g_1 and g_2 being 0 or 1, respectively. Probabilistic inference in this task entails computing the probability distribution over the relevant unobserved variables (here D), given the observed ones (here the image I) $p(D|I)$. While the experimenters control the image I , they learn about the subject's belief about D in any one trial by recording the subject's choice.

The subject achieves optimal performance if the internal model that its brain has learned agrees with the actual (external) model that generates the stimulus defined by the experimenter. However, while the experimenter's model defines the theoretical optimum, the model learned by the brain is likely to deviate from it (Ma, 2012). Here we model two such deviations. The first deviation is due to the fact that the brain has to learn the task by

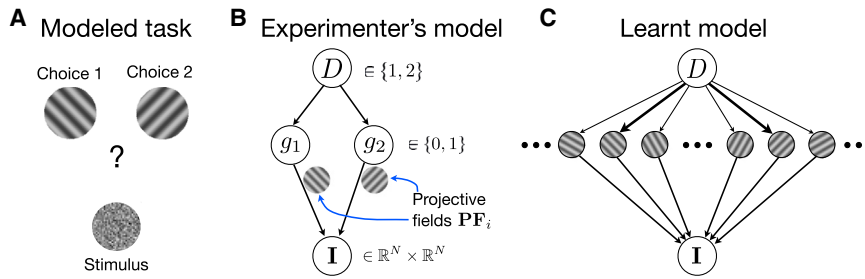


Figure 2. The Structure of the Task

(A) The task is to decide whether a noisy display was caused by a -45° - or $+45^\circ$ -oriented grating. (B) Experimenter's model. First the rewarded decision D is chosen from some prior distribution, $p(D)$. Depending on the value of D , 1 or 2, the corresponding grating variable g_1 or g_2 is set to 1, while the value of the other grating variable is set to zero. The image I , consisting of $N \times N$ pixels on the screen, is drawn from some noisy distribution around the signal/target images, $g_1 T_1 + g_2 T_2$. (C) Model learned by the brain. Due to the noise in the stimulus, the correct decisions corresponding

to the two values of D may be caused not just by the gratings with exactly -45° or exactly $+45^\circ$ orientation but also, with smaller probability, by gratings of similar orientation (indicated by arrows of different width). The relationships between D and the g_i values are not deterministic but probabilistic and learned by the subject by correlating the rewarded responses with the stimuli preceding them. Note that arrows here and in all following figures other than Figure 3 indicate statistical dependencies in a Bayesian network, *not* information flow (which is shown in Figures 3B and 7C).

correlating rewarded choices with preceding noisy stimuli (Qamar et al., 2013). As a result, there will be some uncertainty about the exact task-relevant orientations (Figure 2C). Furthermore, while the relationships between D and g_i are deterministic in the experimenter's model, such that $p(g_1 = 1 | D = 1) = 1$ and $p(g_2 = 1 | D = 1) = 0$, for instance, allowing for other explanations for the image on the retina will change that probability to intermediate values between 0 and 1. Our assumptions about these deviations manifest themselves as two free parameters in our model (Experimental Procedures).

To make experimentally testable predictions for the neuronal responses in V1, we combine the generative model of the task with the probabilistic version of an established sparse-coding model for V1 (Olshausen and Field, 1997; Hoyer and Hyvärinen, 2003). In this model, the responses of V1 simple cells, $\{x_i\}$, are assumed to represent the intensity of an oriented Gabor-shaped feature in the image at a particular location (Figure 3; Experimental Procedures). The ideas underlying our approach are that (1) the brain will perform the task based on the responses of its existing V1 neurons; and (2) the relationship between retina and V1 is not substantially altered by learning a task during adulthood, long after the critical period. As a result, each variable g_i in our task model now represents the task-relevant responses of neurons in our V1 model $\{x_i\}$ rather than a pixel pattern on the screen directly (Figure 3; Experimental Procedures). If the brain indeed solves the task by performing probabilistic inference over all unobserved variables given the observed one (I), the task structure (which defines the upper part of our model) acts as a prior on the activity of the V1 neurons represented by \mathbf{x} . Hence, task-related knowledge, i.e., expectations about which stimuli are more or less likely in the sensory inputs, is thereby incorporated in the sensory neurons' responses. This top-down influence induces correlations between the V1 neurons (noise correlations), between the neurons and the decision (CPs), and between the stimulus and the decision (PK). Since these three types of correlations are directly observable, we use them to link our model predictions to empirical data.

Finally, we need to specify the algorithm by which the brain performs inference in the generative model defined by Figure 3 and how its beliefs about the variables in this model are represented by neural responses. While both parametric (Ma et al., 2006) and sampling-based implementations have been pro-

posed, we assumed a neural sampling-based representation supported by many studies in cognition (both for inference, reviewed in Vul et al., 2014, and learning [Bonawitz et al., 2014]) and increasingly also sensory processing (Hoyer and Hyvärinen, 2003; Fiser et al., 2010; Berkes et al., 2011; Buesing et al., 2011; Pecevski et al., 2011; also see Moreno-Bote et al., 2011 for an approach combining both parametric and sampling-based representations). In short, the neural sampling hypothesis proposes that the brain performs inference, i.e., computes $p(D, \mathbf{g}, \mathbf{x} | I)$, by generating a sequence of samples from this probability distribution based on the generative model that it has learned. In our case, a single sample, indexed by k , from this probability distribution $p(D, \mathbf{g}, \mathbf{x} | I)$ is a vector $(D^{(k)}, g_1^{(k)}, g_2^{(k)}, \dots, x_1^{(k)}, x_2^{(k)}, \dots)$. In a sampling-based representation, the marginal probability $p(D | I)$ can be deduced from the sequence of samples $D^{(k)}$. This means that a decision-making area, computing the probability that the display on the screen was caused by one orientation rather than the other one, can simply count the number of occurrences of $D^{(k)} = 1$ and $D^{(k)} = 2$ over time, for instance, by increasing or decreasing the activity of a pool of neurons, as has been proposed previously (Gold and Shadlen, 2007).

The generative model (Experimental Procedures) specifies the sampling equations (Supplemental Experimental Procedures) that constitute a mechanistic model for how each V1 neuron in our model updates its firing rate (represented by x_i), depending on its inputs from the retina (feedforward), the other V1 neurons (lateral), and the rest of the brain (feedback, specified by the task model). In our simulations, we modeled a trial of 1-s duration by generating 80 samples from the full model as described in Figure 3 (see Experimental Procedures). The 80 samples for a single simulated V1 neuron represent both the neuron's firing rate changing throughout the trial and the brain's evolving belief about the intensity of the corresponding Gabor-shaped feature in the image. The sum of all 80 samples for a particular x_i represents the spike count over the duration of the entire trial used to compute noise correlations for Figures 4 and 5. Individual samples are used to compute the instantaneous correlation between a neuron's response and the decision for Figure 6C.

Task-Induced Correlations between Sensory Neurons

Figure 4 shows the trial-to-trial correlation coefficients between the sensory neurons in our model while the system is performing

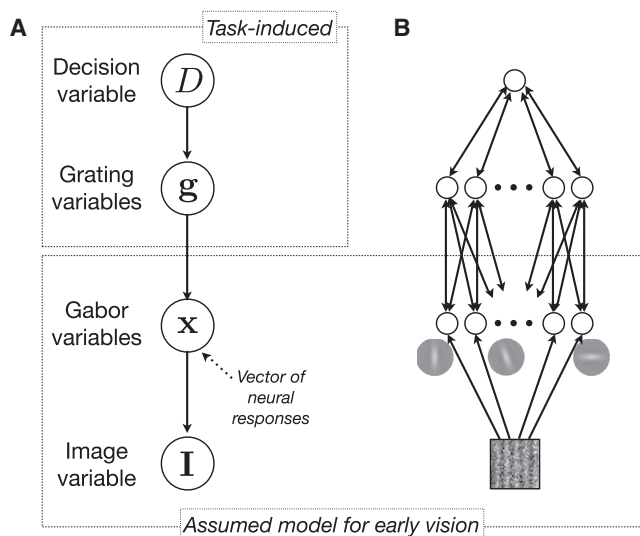


Figure 3. Full Probabilistic Model Consisting of a Sparse-Coding Model for Early Vision and the Generative Model for the Task

(A) Graphical representation of the combined generative model is shown. (B) Correspondence to biology. We assume \mathbf{x} to be represented by early visual neurons, for instance in V1. While our results do not depend on the details of the implementation of grating and decision variables in the brain, we imagine them being represented in mid- or high-level visual areas and decision-making areas, respectively.

inference on a stimulus without signal (no grating present), commonly called noise correlations (Cohen and Kohn, 2011; Nienborg et al., 2012). The most important feature to note are the two maxima and the two minima in the correlation structure occurring at the task-relevant orientations: 45° versus 135° (Figure 4A) and vertical versus horizontal (Figure 4B). The correlation peaks are a direct result of the fact that neurons with stimulus preferences close to the same task-relevant stimulus will be increased and decreased *together* as the subject's belief about the correct choice varies from trial to trial. At the same time, neurons that support different choices will have reduced correlations since their responses are increased and decreased on different trials. This deviation from rotational symmetry is task dependent in that the locations of peaks and troughs are entirely determined by the task-relevant orientations. Furthermore, our framework can be used easily to predict the noise correlation structure for other psychophysical tasks. For example, for a three-alternative forced choice (3AFC) task, our framework predicts the noise correlation structure to have three peaks along the diagonal and three corresponding troughs on each side of the diagonal (Figure 5A). Similarly, a detection task is predicted to induce a single peak at the task-relevant stimulus (Figure 5B).

A signature of this task dependence of the correlation structure has been observed previously in area MT in a coarse motion direction task (Cohen and Newsome, 2008). To facilitate comparison with the data from that experiment, we collapsed our correlation matrix and plotted the correlation coefficient as a function of the difference in the preferred orientations of the two neurons (Figure 4C). Following Cohen and Newsome (2008), we plotted this dependency separately for neuron pairs

where both neurons prefer the same task-relevant orientation (red) and for neuron pairs supporting opposing choices (blue). As in the empirical data (Figure 4D), we found a reduction in correlations between neurons supporting different choices compared to pairs of neurons supporting the same choice. Our model prediction deviates from this dataset at the largest differences in orientation. However, it fully agrees with more recent data (Bondy and Cumming, 2013) that does not find a crossover in the correlation profiles at large differences between preferred orientation. Future studies will be needed to understand this difference between datasets, since both 2AFC tasks—one in the motion direction domain (Cohen and Newsome, 2008) and one in the orientation domain (Bondy and Cumming, 2013)—are identical from the perspective of our framework and, hence, model.

The second feature to note in the correlation structure (Figures 4A and 4B) is the elongated shape of the peaks and troughs. This is due to the task uncertainty included in our model. The more precisely the task-relevant orientations are known by the subject, the closer to rotationally symmetric they are predicted to be. This prediction is consistent with the observation that a substantial uncertainty is required in our model (parameter κ ; Experimental Procedures) to find a reasonable agreement with the data of Cohen and Newsome (2009), where the monkeys are cued to the task-relevant directions on a trial-by-trial basis, while there is very little sign of any elongated shape in their correlation matrix and, hence, task uncertainty in the data of Bondy and Cumming (2013), who train their monkeys for several days on every new pair of task-relevant orientations before measuring the correlation matrix.

Our model does not predict the absolute magnitude of the correlation matrix and, hence, the overall level of correlations. Any non-task-specific changes in the overall level of excitability (e.g., due to changes in alertness; Ecker et al., 2014) would add a positive offset to the correlations in Figures 4 and 5. However, since the amplitude of the correlation matrix (measured as peak minus trough) directly reflects the degree of task knowledge, another new prediction of our framework is an *increase* in the amplitude with learning (Figure 8). As the brain learns a better approximation of the experimenter's task, its prior becomes stronger and the top-down influence increases in our model, a process that may be related to perceptual learning (Goldstone, 1998).

Correlations between Sensory Responses and Behavior

Another much-studied empirical quantity in the context of perceptual decision-making is CP (Britten et al., 1996; Shadlen et al., 1996; Haefner et al., 2013; Nienborg et al., 2012), essentially a measure for the correlation between the variability in the response of a sensory neuron and the behavioral choice in each trial (Experimental Procedures). Our model makes three predictions in agreement with existing empirical data. First, it predicts CPs to grow as a function of time (Figure 6C, compare with data redrawn in Figure 6D). Whether CPs reflect the causal influence of a sensory neuron on the decision or whether they reflect information received by the neuron about the decision formed in higher cortical areas has been a matter of debate (Shadlen et al., 1996; Nienborg and Cumming, 2009, 2010). In

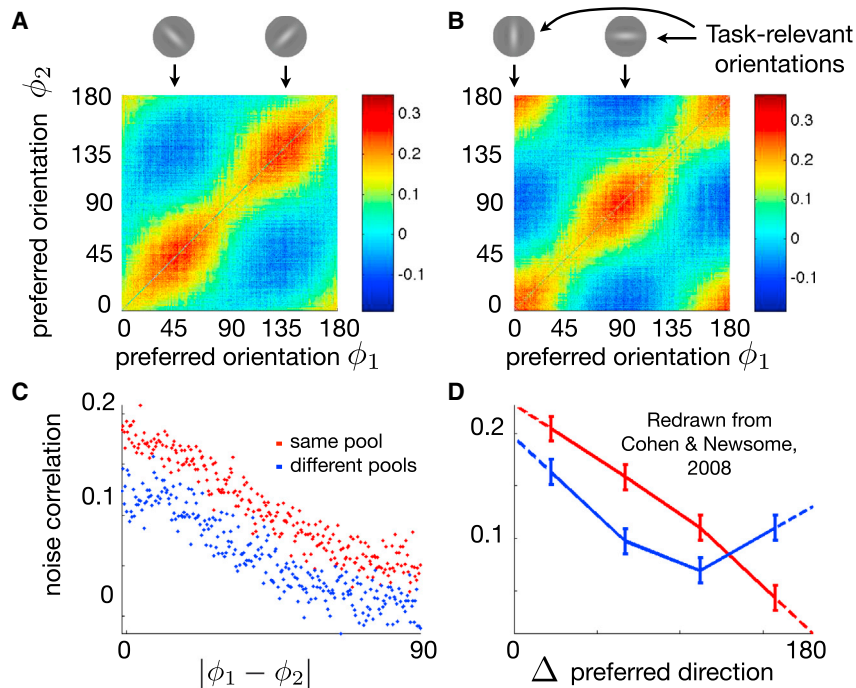


Figure 4. Model Predictions for Noise Correlations between the Sensory Neurons in Our Model

(A) Full correlation matrix with neurons sorted by their preferred orientation. The task-relevant orientations are indicated at 45° and 135°. Only the qualitative shape is predicted by our model; the overall magnitude depends on the particular choice of model parameters (and can be fit to data).

(B) Same as (A) is shown, but for a cardinal orientations task.

(C) Correlation coefficients as a function of difference in preferred orientation between the two neurons in a pair. Red represents all pairs where each neuron's preferred orientation is closest to the same task-relevant orientation (i.e., supports the same choice). Blue represents all pairs in which the two neurons are aligned with different choices. The variability around the means reflects the measurement noise from 1,000 simulated trials.

(D) Noise correlations recorded between neurons in area MT for a motion direction task (Cohen and Newsome, 2008; data are replotted from Figure 4 therein). Since motion direction has a period of 360° compared to a period of 180° for orientation, differences in preferred orientation between 0° and 90° in (C) are comparable to differences in preferred direction between 0° and 180° in (D).

our model, CPs are the result of both feedforward (bottom-up) and feedback (top-down) processing. A stochastic increase in firing of a sensory neuron does increase the probability of a decision consistent with the activity of this neuron (feedforward). In the absence of noise correlations caused by the decision variable (Ecker et al., 2010), the strength of this effect depends inversely on the number of sensory neurons contributing to the decision (Shadlen et al., 1996; Haefner et al., 2013).

In our model, CPs at the beginning of the trial are primarily due to this feedforward pathway, since no coherent top-down belief has formed yet, i.e., the current belief over D is only weakly correlated with the final choice. While we assume an unbiased prior over the correct decision, biased expectations about the upcoming correct choice generally increase CPs at stimulus onset, albeit through the feedback pathway. At any point throughout the trial, the brain's belief about the correct decision reflects the accumulated sensory information presented earlier in the trial. This constitutes prior information about the likely retinal image at that point, which influences the brain's belief about the content of the image and, hence, the responses of sensory neurons representing this belief. In general, as the trial progresses and the top-down belief about the correct decision becomes stronger, CPs are enhanced by an increasing top-down component and increasingly reflect the accumulation of evidence about the decision variable (Figures 6A and 6C). The crucial point here is that the posterior belief over x at some time t within the trial depends on both the current observation I_t , but also all previous observations $I_{1..t}$. Information about these previous observations is communicated to a neuron representing x_t via the posterior belief over the correct decision, $p_t(D)$ at time t (Figure 6E, i.e., $p_t(D) \equiv p(D | I_1, \dots, I_t)$).

Second, CPs in our model are largest for those neurons whose preferred orientation is closest to the task-relevant orientations (Figure 6B), a relationship in agreement with empirical data (Cohen and Newsome, 2009; Bosking and Maunsell, 2011). Since neurons whose preferred orientation is aligned with the task axes are also the most informative about the correct decision, the relationship in Figure 6B implies a correlation between neuro-metric thresholds and CPs, consistent with empirical findings (Nienborg et al., 2012). As for noise correlations described above, our model predictions concern the qualitative shape of the CP dependence on time and preferred orientation, not their magnitude. Regardless of the particular parameter values in our model, CPs increase over time and are largest for neurons most modulated by the task-relevant stimulus dimension. This is in agreement with empirical evidence not just for coarse but also for fine discrimination tasks (Purushothaman and Bradley, 2005) where neurons with the steepest tuning curve slope have the highest CPs. Third, as for the amplitude of the noise correlations, the magnitude of CPs is related to the degree to which the brain has learned the task model. This predicts that the CP for task-relevant neurons should increase with learning, as has been observed empirically (Law and Gold, 2008).

Correlations between Stimulus and Behavior

The strength of the correlation between stimulus and behavior in 2AFC tasks is typically measured by the PK (Neri et al., 1999; Ahumada, 2002; Nienborg and Cumming, 2009). The PK quantifies how strongly the evidence in the stimulus is weighted in the decision-making as a function of the time at which the evidence is presented during the trial. Our model predicts that the weighting decreases over time (Figure 7A) so that evidence presented early in the trial has a larger influence on the final decision

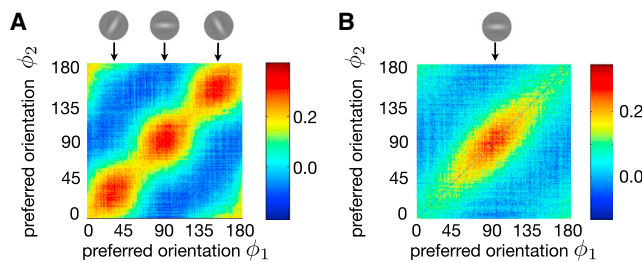


Figure 5. Predicted Noise Correlations for Alternative Tasks Analogous to Figure 4

(A) Three-alternative choice task (3AFC) is shown.
(B) Detection task is shown.

than evidence presented near the end. This order effect is not inherent to probabilistic inference—exact inference leads to a constant PK as would be expected from the optimal solution—but is a consequence of performing approximate online inference (sequential sampling in our model, also see [Supplemental Experimental Procedures](#)). In our model, the reason for this decrease is a self-reinforcing feedback loop involving sensory and decision neurons ([Figure 7C](#)). In the absence of any stimulus information, when the responses of the sensory neurons are entirely determined by their prior, small deviations from an exactly even (50–50) belief over D will bias the brain’s belief about x by top-down belief propagation. This bias in the sensory representation in turn reinforces the initial deviation in the brain’s belief over D in a feedforward fashion. Even if there is information in the stimulus, as long as it does not dominate the top-down prior (as is usually the case in threshold psychophysics), this mechanism leads to early evidence having a larger impact on the final choice than evidence presented later in the trial. Furthermore, the strength of this order effect should depend on how strong the influence of the bottom-up stimulus evidence is on the sensory representation compared to the top-down prior.

Both predictions are in agreement with empirical data. A decreasing PK was found in the same coarse discrimination task ([Figure 7B](#)) that reported CPs to be increasing over time ([Figure 6C](#)) ([Nienborg and Cumming, 2009](#)). Here the stimulus consisted of weak disparity signals embedded in noise. A recent study in rats using a Poisson-clicks stimulus, on the other hand, found a constant PK when the task was to report whether more clicks were played to the left or the right ear ([Brunton et al., 2013](#)). Here the individual clicks were highly super-threshold such that the influence of the top-down prior compared to the bottom-up likelihood was presumably negligible. Therefore, we predict that as the volume of the clicks is reduced to detection threshold level, or if the clicks are embedded in distractor noise, the PK for the Poisson-clicks task should become downward sloping, too. We emphasize that the model was not designed to contain this self-reinforcement of beliefs but that this feature emerged as a side effect of performing probabilistic inference in an online fashion using a neural sampling-based implementation in a hierarchical model. Its mathematical form is determined by the structure of the task and it shares the same parameters that determine magnitude and time course of noise correlations and CPs.

DISCUSSION

In this paper we have demonstrated how the structure of a psychophysical task can be used to make empirically testable predictions for neural responses from the old idea of visual perception as probabilistic inference, even though the brain’s internal model for general vision is unknown. We have applied these ideas to a simple 2AFC task and found that our model naturally accounts for several findings that are hard to explain in the traditional feedforward framework and makes new predictions for neurophysiology and psychophysics. The crucial difference to previous models of this task, probabilistic ([Beck et al., 2008](#)) or not ([Gold and Shadlen, 2007](#)), was our assumption that the brain performs inference not only about the task-relevant variables but also the variables represented by sensory neurons. Strictly speaking, to perform a given task, one only has to compute the posterior over the task-relevant variable(s). However, in general, these variables could be (and in real-life usually are) many, including low-level as well as high-level features represented at different levels of the visual hierarchy, and, furthermore, these variables are usually not even pre-specified. Decision-making in such a context requires inference over all latent/unobserved variables, not just a subset pre-defined by the experimenter. This assumption implies that the responses of sensory neurons represent posterior beliefs, i.e., not just certain stimulus features but also incorporating prior information from the rest of the brain ([Lee and Mumford, 2003](#)). This provides a functional constraint on the top-down connectivity that is heavily under-constrained by empirical data and, therefore, often ignored (but see [Rao and Ballard, 1999](#) and [Wimmer et al., 2015](#)). Since the experimenter controls the task that determines the functional form of this top-down influence in our normative model, the validity of this normative constraint can be tested very directly. The first new and specific prediction of our model, about the structure of the full correlation matrix in primary visual cortex during a coarse orientation discrimination task ([Figure 4A](#)), has since been confirmed by independent preliminary evidence ([Bondy and Cumming, 2013](#)).

An important feature of our model prediction for the noise correlation matrix between sensory neurons is that its principal structure (two maxima and two minima and their location) is the result of a normative approach and was not fit to the data we aimed to explain. We assumed that the brain performs the task using probabilistic inference. We did *not* assume that the brain does so optimally and our approach makes explicit the way in which the brain deviates from optimality. Suboptimality in our model is the consequence of three features. First, sampling-based approximations of probabilities converge to the exact solution only in the limit of infinitely many samples. Since the brain can only generate a finite number of samples, its solution to any problem will deviate from the optimal solution. Second, in conjunction with an online processing of the evidence, we show this to lead to an overweighting of early evidence, another deviation from optimality. Third, the internal model that the brain has learned about the task will generally deviate from the true external, experimenter-defined one. These deviations from optimality are made explicit by our model and, in fact, gave rise to its free parameters: the number of samples

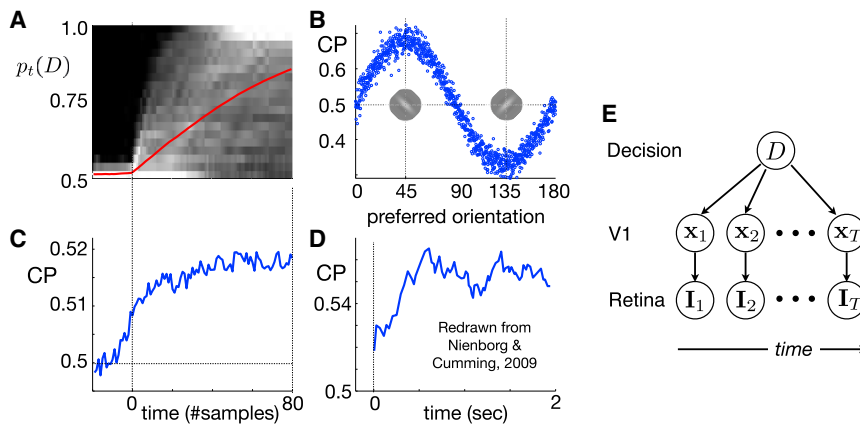


Figure 6. Model Prediction for CPs

(A) Temporal evolution of the belief about the correct decision D . Red curve shows the average posterior $p_t(D=1)$ as a function of time for those trials for which the final choice was $D=1$. It is fixed to 0.5 until trial onset and accumulates evidence throughout the trial. The grayscale map shows the distribution of beliefs as a function of time (high probability in white and zero probability in black). Each column represents a histogram of $p_t(D=1)$ at time t . (B) CP as a function of preferred orientation. Here CPs were computed across the entire trial rather than in small time bins as in (C), which is why their magnitude is greater than in (C). (C) Magnitude of CP as a function of time is shown. (D) CP magnitude as a function of time based on data recorded in area V2 for a disparity direction task is shown (Nienborg and Cumming, 2009). (E) Generative model used for the task extended into the time domain is shown (Figure 3A after integrating out the grating variables).

generated per trial, the strength of the brain's belief that each rewarded decision is preceded by a grating of a particular orientation in the stimulus, and the strength of the belief that a particular grating is characterized by a particular response of sensory neurons.

At one extreme of the parameter range, the brain has not learned the task at all and the task, therefore, has no impact on sensory responses: task-dependent noise correlations, CPs and PK would then be zero. The reason for this is that, in the absence of a task, our model only consists of a fixed sparse-coding model for V1. However, we remain agnostic about the non-task-related correlations in sensory neurons and, strictly speaking, our model only makes predictions for the difference between before and after task learning. At the other extreme of the parameter range, the internal model of the brain becomes identical to the experimenter's model, at which point the psychophysical performance of the model becomes optimal given the constraints of the V1 part of the model (which we assume to be fixed) and the online processing constraint. A direct prediction from our modeling is, therefore, that task-induced CPs and correlations should increase during task learning (Figure 8). We emphasize that our qualitative predictions hold across the entire model parameter regime: independent of the precise values for the parameters, CPs increase over time, PKs decrease, and the noise correlation structure has two maxima and two minima whose locations are defined by the task-relevant stimuli. At first sight, our prediction that the amplitude of the predicted correlation structure should increase with perceptual learning appears to be at odds with the finding that perceptual learning decreases noise correlations in visual and vestibular tasks in area MSTd (Gu et al., 2011). However, that reduction was only seen for the average correlations (for which our model was not designed to make a prediction) and not for the slope of the relationship between noise correlations and signal correlations, which is only loosely related to the amplitude of the task-dependent correlation matrix.

The role of correlations between sensory neurons in the probabilistic inference framework is very different from that in the traditional feedforward information-processing view. While noise

correlations typically limit the information about the external stimulus that can be represented by populations of sensory neurons (Zohary et al., 1994; Averbeck et al., 2006; Moreno-Bote et al., 2014), or at least complicate the readout (Shamir and Sompolinsky, 2006; Ecker et al., 2011), in the probabilistic inference model task-dependent correlations reflect prior information about the structure of the outside world (Berkes et al., 2011). This information is communicated to the relevant sensory neurons and, hence, modifies their responses. This belief propagation is best seen in the increasing time course of CPs, which mirrors the formation of the brain's cognitive belief about the correct decision (Figures 6A and 6C).

While attention and belief propagation have been argued to employ the same biological mechanism (Krug, 2004), they make different neurophysiological predictions in a 2AFC task. To increase psychophysical performance in the traditional feedforward framework, attention needs to increase the responses of neurons supporting both choices *on the same trial*. Otherwise, enhancing the responses of neurons supporting one choice only will, in the absence of any external attention cue that is valid more often than not (a cue that is not present in the experiments we model), simply lead to biased decision-making decreasing psychophysical performance. Belief propagation, on the other hand, increases the responses of only those neurons that support the choice believed to be more likely by the rest of the brain. Only the latter mechanism, but not the performance-improving kind of attention, leads to CPs. Furthermore, a gain change that varies from trial to trial but acts equally on neurons supporting either choice (as required in order to improve performance) implies equal noise correlations independent of which choice the neurons support, in contrast to empirical findings (Cohen and Newsome, 2008). (Also, the implied correlation matrix would have peaks where our prediction shows troughs.) Finally, since an alternation of attention would reduce performance and lead to an inverse relationship between correlation strength and performance (Ecker et al., 2016) (in contrast to Figure 8), and since attention recently was found to be associated primarily with decreased variability in input common to sensory neurons independent of their task-related stimulus preference (Rabinowitz

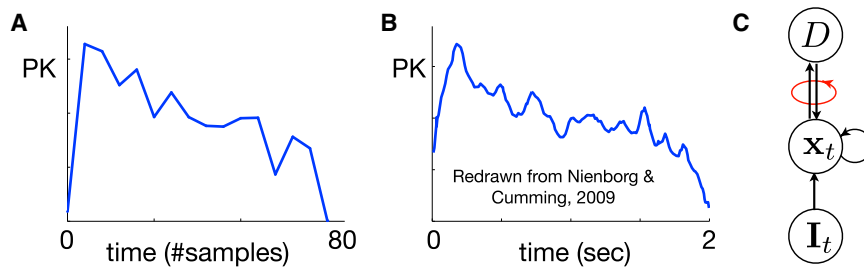


Figure 7. Time Dependency of PK

(A) Model prediction. Evidence early in the trial is weighted more strongly than evidence presented late in the trial.

(B) Empirical PK in a disparity discrimination task is shown (Nienborg and Cumming, 2009).

(C) Information flow in our model: feedforward from the retinal image I_t to the sensory representation x_t , and from the sensory representation to the decision D . Feedback is from decision to sensory representation and lateral within the sensory representation (for equations see Supplemental Experimental Procedures). Feedforward and feedback between sensory representation and decision area form a self-reinforcing loop entailing a decreasing PK.

et al., 2015), we argue that belief propagation is the more parsimonious explanation for the observed task-dependent correlations.

The observation that the time courses of CPs and PKs differ during the stimulus presentation has been a challenge to feedforward models of sensory processing and has been taken as evidence for a feedback component in CPs (Nienborg and Cumming, 2009). These findings are fully consistent with our model where CPs contain both a feedforward and a feedback component. The existence of a feedforward component can most easily be seen by the fact that CPs are larger than 0.5 at stimulus onset when the model's belief about the correct decision is still 50-50. The magnitude of this bottom-up component follows the same decreasing time course as the PK (Figure 7A) (Nienborg and Cumming, 2009), with the difference accounted for by top-down belief propagation. Furthermore, we have demonstrated that the very same feedback signal that causes the increasing CPs also causes a decreasing PK due to a positive feedback loop between decision-making neurons and sensory neurons. We stress that we did not hand-craft the model structure in order to fit these observations but that this feedback signal is a direct result of performing online inference by sampling in this task.

There exist alternative explanations for a decreasing PK suggesting that the brain inappropriately uses strategies optimal for reaction-time tasks (integration-to-bound, Gold and Shadlen, 2007) or hypothesizing costs intrinsic to accumulating evidence (Drugowitsch et al., 2012). However, our explanation demonstrates that even if a neural decision circuit *itself* equally integrates the signals it receives from sensory neurons over time, as has been suggested recently (Brunton et al., 2013), its top-down influence on the very same sensory neurons can lead to early evidence being weighted more strongly than evidence presented at the end of the trial if that evidence is weak. Future work on extending our model to a reaction-time paradigm and comparing its predictions with reaction-time data will be able to assess the relative importance and explanatory power of these different hypotheses.

Our study complements a growing body of literature suggesting that sampling-based probabilistic inference may underlie higher perceptual and cognitive processes (Griffiths et al., 2012; Vul et al., 2014), and it suggests that probabilities also might be represented by samples in lower sensory processing (given the compatibility with the data considered here). It is particularly intriguing that basic visual perception appears to

be susceptible to the same confirmation bias that is ubiquitous in the context of higher cognitive reasoning (Nickerson, 1998). If the same two mechanisms that contribute to a decreasing PK in our model (a sequential sampling approximation and evidence accumulation on the basis of inferred beliefs instead of directly observed information) also underlie higher-order cognition, then this suggests a new candidate explanation for the confirmation bias, different from previous accounts (Lieder et al., 2012). More generally, our framework bridges cognitive science and systems neuroscience by constructing a rational process model (Griffiths et al., 2012) combining the generative model for the psychophysical task with knowledge about the biological architecture and the neural sampling hypothesis.

Our work differs from a previous probabilistic model (Beck et al., 2008) of the 2AFC task in two important aspects. First, we assume that individual cortical sensory neurons represent marginal posteriors over unobserved variables, while Beck et al. (2008) assume that the population represents the likelihood of a single unobserved variable. This aspect requires recurrent connections within the sensory area as well as feedback connections from higher areas that are missing in their model (see Beck et al., 2011, 2012 and Grabska-Barwinska et al., 2013 for probabilistic models incorporating recurrent connections). Without the feedback connections in our model there would be no task-dependent noise correlations, nor would there be CPs that increase while the PK decreases. Second, our model differs from that of Beck et al. in that the representation of probabilities is sampling based (Fiser et al., 2010; Hoyer and Hyvärinen, 2003) rather than based on a probabilistic population code (PPC) (Ma et al., 2006). Preliminary evidence suggests, however, that our results can be generalized to other neural representations of probabilities, including PPC-based ones (Haefner, 2014; A. Pouget, personal communication).

It has been suggested previously that priors also may be encoded in feedforward weights (Ganguli and Simoncelli, 2010; Wei and Stocker, 2012). While this is likely for long-term priors reflecting permanent statistics of the natural environment, it appears impossible for the task-dependent context information considered in our study that can vary on a trial-by-trial basis (Cohen and Newsome, 2008) or indeed the influence of the stimulus information at the beginning of a trial on later sensory beliefs within the same trial (Figure 6E).

A recent paper (Wimmer et al., 2015) has presented a mechanistic model that reproduces an increasing CP and a decreasing

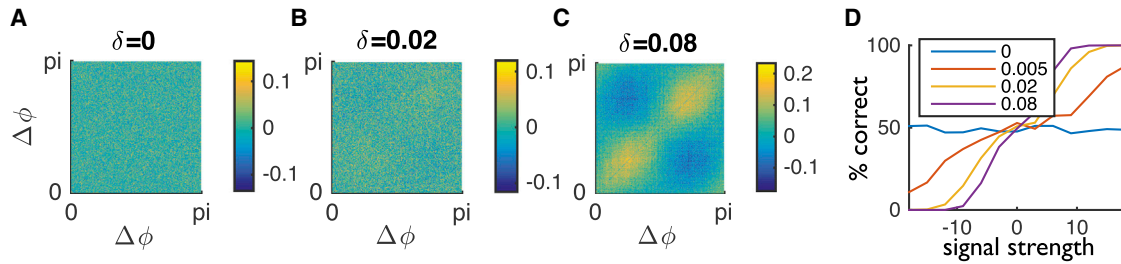


Figure 8. Relationship between Amplitude of Noise Correlation Structure and Psychophysical Performance: Stronger Correlations Reflect Higher Performance

(A–C) Correlation structure for three models that differ in the strength of the bottom-up and top-down connections (quantified by parameter δ [see [Experimental Procedures](#)], indicated on top of each panel). As δ increases from 0 (no learning), the correlation structure reflecting the 2AFC task emerges. Note the different color scales in the three panels. $\delta = 0$ (A); $\delta = 0.02$ (B); $\delta = 0.08$ (C).

(D) Psychometric curves for four models. At $\delta = 0$ performance is at chance and as δ increases the curves become steeper, i.e., performance increases. The legend indicates the value of δ for each of the simulated models. Note that we kept κ constant for simplicity; during actual learning both δ and κ will change. For description of model ([Schwartz and Simoncelli, 2001](#)) and corresponding population responses, see [Figure S2](#).

PK. There are similarities between the computational motifs of mutual inhibition and a positive feedback loop between sensory and decision neurons in their model and in ours. However, unlike our model theirs is constructed in a bottom-up/data-driven fashion and therefore cannot address the question of why the brain is structured the way it is, especially since mutual inhibition (explaining away in our model) and feedback connections (belief propagation) decrease the psychophysical performance of their model. Furthermore, their model is designed specifically for the 2AFC task under consideration, without relating it to the general probabilistic inference problems that the brain has to solve during natural vision.

A major open question in systems neuroscience concerns the origin and role of correlated neuronal variability. It has long been known that neuronal responses are correlated, and the development of population-recording techniques is finally making detailed measurements of these correlations feasible, in particular their magnitude, shape, and dependency on both stimulus and cognitive variables. In the framework of probabilistic inference, the generative model for the stimulus in the task acts as a prior on the variables represented by the sensory responses. Since the generative model for the task is entirely under the control of the experimenter, this allows for designing experiments with very direct predictions. We therefore hope our results can inspire future experiments and provide a framework for analyzing multi-neuron recordings in the awake brain and for linking these recordings to otherwise hard-to-measure cognitive factors like beliefs and task strategies.

EXPERIMENTAL PROCEDURES

Generative Model Learned by the Brain

We assume that the brain has learned the two crucial features of the coarse orientation discrimination task that we are modeling: that there are exactly two possible decisions, $D = 1$ and $D = 2$, and that they are related to the orientation of the visual stimulus presented on the screen. We further assume that the subject has learned that both choices are equally likely a priori, $p_D(D = 1) = p_D(D = 2) = 0.5$. From correlating correct choices with the preceding stimulus, the brain has to learn the task-relevant orientations ψ_1 and ψ_2 . Inevitably, there will be some level of uncertainty (tolerance) about the precise orientations indicating choice 1 and choice 2. We model this uncertainty by

a circular Gaussian (von Mises) function around the correct orientation as follows:

$$g_i | D \sim \text{Bernoulli} \left\{ \frac{1}{n_g / \kappa} \exp \left[\kappa \cos 2 \left(\phi_i^{(g)} - \psi_D \right) \right] \right\} \quad (\text{Equation 1})$$

where g_i is one of n_g binary variables indicating the presence of a grating with orientation $\phi_i^{(g)}$ in the stimulus, and ψ_D corresponds to the target orientation associated with decision D . The normalization is chosen to be $\sum_i^{n_g} p(g_i = 1 | D) = 1$, such that on average one of the grating variables is on. κ is a free parameter in our model that can range from 0 to ∞ , where $\kappa = 0$ represents no knowledge about which orientations are task relevant and ∞ represents the limit of perfect learning, where the internal model is identical to the generative model used by the experimenter. We assume that the task is performed on the basis of an early sensory representation (\mathbf{x} , here primary visual cortex), such that each orientation variable g_i corresponds to a particular expected value for \mathbf{x} . Again, we make a canonical circular Gaussian assumption about the dependency between a grating variable g_k representing orientation $\phi_k^{(g)}$ and a Gabor-shaped feature x_i with orientation $\phi_i^{(x)}$ in the image as follows:

$$E[x_i | \mathbf{g}] = 1 + \delta \sum_{k=1}^{n_g} g_k \exp \left[\lambda \cos 2 \left(\phi_i^{(x)} - \phi_k^{(g)} \right) \right]. \quad (\text{Equation 2})$$

Here, $0 < \lambda$ and $0 < \delta$ are free parameters (discussed below). Since the x_i are meant to model neurons in V1, we model their relationship to the image on the retina by a standard sparse-coding model ([Olshausen and Field, 1997](#); [Hoyer and Hyvriinen, 2003](#)) as follows:

$$p(x_i | \mathbf{g}) = \frac{1}{\tau_i} \exp(-x_i / \tau_i) H(x_i) \quad (\text{Equation 3})$$

$$p(\mathbf{l} | \mathbf{x}) = \mathcal{N} \left(\mathbf{l} : \frac{1}{n_x} \sum_{i=1}^{n_x} \mathbf{P}\mathbf{F}_i x_i, \mathbf{1} \right) \quad (\text{Equation 4})$$

with $\tau_i \equiv E[x_i | \mathbf{g}]$ as defined in Equation 2, H is the Heaviside function restricting \mathbf{x} to positive values, and $\mathbf{1}$ is the identity matrix. (The x_i can be interpreted as firing rates here with no specific assumption about the spike production mechanism.) The matrices $\mathbf{P}\mathbf{F}_i$ contain the projective fields for each x_i and the actual image on the screen is a noisy version of their linear superposition weighted by \mathbf{x} . Since the only stimulus dimension relevant for our results is orientation, we simply assume a bank of Gabor-shaped filters (aspect ratio 2:1), which only differ in their orientation, uniformly spaced between 0 and π . (While the properties of sensory neurons are highly heterogeneous [e.g., [Ecker et al., 2011](#) and [Goris et al., 2015](#)], this heterogeneity is inconsequential for the results presented here.) For $\delta = 0$, the model for the sensory responses becomes independent of the task and the prior over \mathbf{x} is sparse and independent. As δ increases, the expected intensity of the pattern associated with the presence of

a grating in the stimulus increases. The generative model described here determines the sampling equations that are provided in the [Supplemental Experimental Procedures](#).

The generative model described above will be normative, i.e., correct, only for stimuli that actually have been generated from it. Despite its deviations from the models actually used by the particular experiments to which we compared our data (Cohen and Newsome, 2008; Nienborg and Cumming, 2009; Brunton et al., 2013), it captures the two crucial features of a 2AFC orientation discrimination task: (1) the expectation of an oriented stimulus of some orientation, and (2) the expectation that only one of two possible orientations (or ranges of orientations) will be shown. κ and λ together with the Gabor-shaped filters determine the orientation bandwidth of the expected oriented stimulus, and δ can be interpreted as the expected saliency or contrast of the signal in the noise.

In tasks where the stimulus is dynamically changing within a trial, evidence about the correct decision has to be accumulated over time. In this case, the brain will learn that, while there is only one correct choice per trial, the image and latent variables other than D can change from stimulus frame to stimulus frame. As a result, the posterior over D evolves over time starting with the initial prior $p_0(D)$, with $p_t(D) \sim p_{t-1}(D)p(\mathbf{l}_t | D)$ for independent \mathbf{l}_{t-1} and \mathbf{l}_t (Beck et al., 2008). Such an accumulation of evidence is compatible with existing neurophysiological observations in putative decision-making regions of the brain (Gold and Shadlen, 2007), providing a graded representation of the instantaneous belief over the correct decision D . In our hierarchical model, the posterior over D is not updated based on the observed variable \mathbf{l} directly but based on the samples drawn from the posterior over \mathbf{g} , $p(\mathbf{g} | \mathbf{l}_t)$. We denote the number of samples generated on the timescale on which the brains assumes the \mathbf{l}_t to be statistically independent of each other by n_s . In our model, we update the belief over D in an online fashion (Gold and Shadlen, 2007) based on the individual samples as follows:

$$\log \frac{p_t^{(k)}(D=1)}{p_t^{(k)}(D=2)} \approx \log \frac{p_{t-1}^{(k)}(D=1)}{p_{t-1}^{(k)}(D=2)} + \frac{1}{n_s} \log \frac{p(\mathbf{g}_t^{(k)} | D=1)}{p(\mathbf{g}_t^{(k)} | D=2)}$$

with $k = 1 \dots n_s$ and $p_t^{(1)}(D) \equiv p_{t-1}^{(n_s)}(D)$, and where $\mathbf{g}_t^{(k)}$ is the corresponding samples drawn from $p(\mathbf{g}_t | \mathbf{l}_t, D)$. n_s constitutes a free parameter in our model and can be constrained experimentally by the time course of the CP: the smaller the n_s the faster the temporal increase in CP.

In our simulations we explicitly model \mathbf{x} and \mathbf{g} . Furthermore, we assume that an area exists in the brain (e.g., LIP; Gold and Shadlen, 2007; Beck et al., 2008) that accumulates the sensory evidence and represents the brain's current belief over the correct decision. This belief over D acts as a top-down prior for the sensory representation, as required by full probabilistic inference over all latent variables, i.e., including \mathbf{x} and \mathbf{g} , and not just the decision variable D . We note that two features responsible for the decreasing PK are easily overcome in typical machine learning applications. First, if the updating of the posterior over D is not performed after every sensory sample but in batches, e.g., after every ten samples, then evidence is weighted more equally over time. We do not think this is biologically plausible since it would require the brain to cache incoming samples and effectively ignore them during the caching time, before updating the posterior, without explicit knowledge when would be the right time for an update of its belief over D . Instead, we believe that it is more plausible that evidence is continuously integrated. Second, generating more than one chain of samples in parallel (e.g., particle filtering; Doucet, 2001) also would prevent the PK from decreasing and has been suggested previously (Lee and Mumford, 2003). Investigating the source of the decreasing PK in more detail as suggested in the [Discussion](#) will therefore give further insights into what kind of sampling scheme the brain is implementing.

Numerical Details

We performed Gibbs sampling in the generative model (see the [Supplemental Experimental Procedures](#)). Trials start after the burn-in period and last 80 complete samples, i.e., 80 samples are generated from each variable in the model. For the simulations underlying [Figures 4, 5, 6](#), and [7](#) we used the following: $\kappa = 1$, $\lambda = 3$, $\delta = 0.08$, $n_x = 1,024$ and $n_g = 256$. The only exception is the 3AFC [Figure 5A](#) where we used $\kappa = 3$, i.e., a narrower range of orientations related

to each choice. We further assumed $n_s = 20$. Since these are free parameters in our model, we concentrate on qualitative rather than quantitative predictions. However, when modeling a particular experiment, they can be constrained by fitting them to observations, in particular CPs, noise correlations, and classification images.

The noise correlations and CP time course were computed using a gray screen as stimulus in order to eliminate stimulus-induced influences and to avoid having to correct CPs for them (Nienborg and Cumming, 2009). The PK time course was computed using a dynamic stimulus in which randomly switching gratings were embedded in Gaussian noise; 16 independent stimulus frames were presented per trial.

CP and PK

The CP of a particular neuron with respect to choice 1 can be defined as the probability that a random sample from the neuron's response distribution preceding choice 1 is larger than a sample from the same neuron's response distribution preceding choice 2 (Britten et al., 1996).

The PK can be defined as the amplitude of the classification image as a function of time (Nienborg and Cumming, 2009). The classification image is the difference between the mean stimulus preceding choice 1 and the mean stimulus preceding choice 2 (Neri et al., 1999; Ahumada, 2002). To compute the time course of the PK, one computes the classification image (essentially the choice-triggered average of the stimulus) as a function of time, t , within a trial, $PK_t = \langle s_t \rangle_{D=1} - \langle s_t \rangle_{D=2}$, where $\langle s_t \rangle_{D=d}$ represents the average over the external stimuli presented at time t , for all trials that led to decision $D = d$.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures, two figures, and one table can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2016.03.020>.

AUTHOR CONTRIBUTIONS

R.M.H., P.B., and J.F. conceived the project and designed the model. R.M.H. implemented the model and performed the analyses. R.M.H. and J.F. wrote the manuscript.

ACKNOWLEDGMENTS

We sincerely thank the many colleagues with whom we have discussed this work and who have provided us with valuable advice on content and presentation. This work was supported by the Swartz Foundation and by National Science Foundation (NSF) grant IOS-1120938.

Received: November 9, 2014

Revised: October 5, 2015

Accepted: March 12, 2016

Published: April 14, 2016

REFERENCES

- Abbott, L.F., and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Comput.* *11*, 91–101.
- Ahumada, A.J., Jr. (2002). Classification image weights and internal noise level estimation. *J. Vis.* *2*, 121–131.
- Averbeck, B.B., Latham, P.E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* *7*, 358–366.
- Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E., and Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron* *60*, 1142–1152.
- Beck, J.M., Latham, P.E., and Pouget, A. (2011). Marginalization in neural circuits with divisive normalization. *J. Neurosci.* *31*, 15310–15319.

- Beck, J.M., Ma, W.J., Pitkow, X., Latham, P.E., and Pouget, A. (2012). Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* 74, 30–39.
- Berkes, P., Orbán, G., Lengyel, M., and Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331, 83–87.
- Bonawitz, E., Denison, S., Griffiths, T.L., and Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends Cogn. Sci.* 18, 497–500.
- Bondy, A., and Cumming, B.G. (2013). Top down signals influence the distribution of noise correlations amongst sensory neurons. In 2013 Neuroscience Meeting Planner (Society for Neuroscience).
- Bosking, W.H., and Maunsell, J.H.R. (2011). Effects of stimulus direction on the correlation between behavior and single units in area MT during a motion detection task. *J. Neurosci.* 31, 8230–8238.
- Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* 13, 87–100.
- Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science* 340, 95–98.
- Buesing, L., Bill, J., Nessler, B., and Maass, W. (2011). Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS Comput. Biol.* 7, e1002211.
- Cohen, M.R., and Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nat. Neurosci.* 14, 811–819.
- Cohen, M.R., and Newsome, W.T. (2008). Context-dependent changes in functional circuitry in visual area MT. *Neuron* 60, 162–173.
- Cohen, M.R., and Newsome, W.T. (2009). Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. *J. Neurosci.* 29, 6635–6648.
- Doucet, A. (2001). *Sequential Monte Carlo Methods* (Wiley Online Library).
- Drugowitsch, J., Moreno-Bote, R., Churchland, A.K., Shadlen, M.N., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *J. Neurosci.* 32, 3612–3628.
- Ecker, A.S., Berens, P., Keliris, G.A., Bethge, M., Logothetis, N.K., and Tolias, A.S. (2010). Decorrelated neuronal firing in cortical microcircuits. *Science* 327, 584–587.
- Ecker, A.S., Berens, P., Tolias, A.S., and Bethge, M. (2011). The effect of noise correlations in populations of diversely tuned neurons. *J. Neurosci.* 31, 14272–14283.
- Ecker, A.S., Berens, P., Cotton, R.J., Subramanian, M., Denfield, G.H., Cadwell, C.R., Smirnakis, S.M., Bethge, M., and Tolias, A.S. (2014). State dependence of noise correlations in macaque primary visual cortex. *Neuron* 82, 235–248.
- Ecker, A.S., Denfield, G.H., Bethge, M., and Tolias, A.S. (2016). On the structure of neuronal population activity under fluctuations in attentional state. *J. Neurosci.* 36, 1775–1789.
- Fiser, J., Berkes, P., Orbán, G., and Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn. Sci.* 14, 119–130.
- Ganguli, D., and Simoncelli, E.P. (2010). Implicit encoding of prior probabilities in optimal neural populations. In *Advances in Neural Information Processing Systems* 23, J.D. Lafferty, C.K.I. Williams, J. Shawe-Taylor, R.S. Zemel, and A. Culotta, eds. (Neural Information Processing Systems Foundation), pp. 6–9.
- Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574.
- Goldstone, R.L. (1998). Perceptual learning. *Annu. Rev. Psychol.* 49, 585–612.
- Goris, R.L., Simoncelli, E.P., and Movshon, J.A. (2015). Origin and function of tuning diversity in macaque visual cortex. *Neuron* 88, 819–831.
- Grabska-Barwinska, A., Beck, J., Pouget, A., and Latham, P.E. (2013). Demixing odors—fast inference in olfaction. In *Advances in Neural Information Processing Systems* 26, C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, eds. (Neural Information Processing Systems Foundation), pp. 1–9.
- Griffiths, T.L., Vul, E., and Sanborn, A.N. (2012). Bridging levels of analysis for probabilistic models of cognition. *Curr. Dir. Psychol. Sci.* 21, 263–268.
- Gu, Y., Liu, S., Fetsch, C.R., Yang, Y., Fok, S., Sunkara, A., DeAngelis, G.C., and Angelaki, D.E. (2011). Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* 71, 750–761.
- Haefner, R.M. (2014). On the relationship between stimulus-evoked and choice-related responses and correlations during perceptual decision-making in a probabilistic inference framework. In 2014 Neuroscience Meeting Planner (Society for Neuroscience).
- Haefner, R.M., Gerwinn, S., Macke, J.H., and Bethge, M. (2013). Inferring decoding strategies from choice probabilities in the presence of correlated variability. *Nat. Neurosci.* 16, 235–242.
- Hoyer, P.O.P., and Hyvärinen, A. (2003). Interpreting neural response variability as Monte Carlo sampling of the posterior. In *Advances in Neural Information Processing Systems* 15, S. Becker, S. Thrun, and K. Obermayer, eds. (Neural Information Processing Systems Foundation), pp. 293–300.
- Kersten, D., Mamassian, P., and Yuille, A. (2004). Object perception as Bayesian inference. *Annu. Rev. Psychol.* 55, 271–304.
- 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.
- Law, C.-T.T., and Gold, J.I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat. Neurosci.* 11, 505–513.
- Lee, T.S., and Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 20, 1434–1448.
- Lieder, F., Griffiths, T.L., and Goodman, N.D. (2012). Burn-in, bias, and the rationality of anchoring. In *Advances in Neural Information Processing Systems* 25, F. Pereira, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, eds. (Neural Information Processing Systems Foundation), pp. 2699–2707.
- Ma, W.J. (2012). Organizing probabilistic models of perception. *Trends Cogn. Sci.* 16, 511–518.
- Ma, W.J., Beck, J.M., Latham, P.E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438.
- Moreno-Bote, R., Knill, D.C., and Pouget, A. (2011). Bayesian sampling in visual perception. *Proc. Natl. Acad. Sci. USA* 108, 12491–12496.
- Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., and Pouget, A. (2014). Information-limiting correlations. *Nat. Neurosci.* 17, 1410–1417.
- Mumford, D. (1992). On the computational architecture of the neocortex. II. The role of cortico-cortical loops. *Biol. Cybern.* 66, 241–251.
- Neri, P., Parker, A.J., and Blakemore, C. (1999). Probing the human stereoscopic system with reverse correlation. *Nature* 401, 695–698.
- Nickerson, R.S. (1998). Confirmation bias: a ubiquitous phenomenon in many guises. *Rev. Gen. Psychol.* 2, 175–220.
- Nienborg, H., and Cumming, B.G. (2009). Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* 459, 89–92.
- Nienborg, H., and Cumming, B. (2010). Correlations between the activity of sensory neurons and behavior: how much do they tell us about a neuron's causality? *Curr. Opin. Neurobiol.* 20, 376–381.
- Nienborg, H., Cohen, M.R., and Cumming, B.G. (2012). Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Annu. Rev. Neurosci.* 35, 463–483.
- Olshausen, B.A., and Field, D.J. (1997). Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision Res.* 37, 3311–3325.
- Pecevski, D., Buesing, L., and Maass, W. (2011). Probabilistic inference in general graphical models through sampling in stochastic networks of spiking neurons. *PLoS Comput. Biol.* 7, e1002294.
- Pouget, A., Beck, J.M., Ma, W.J., and Latham, P.E. (2013). Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* 16, 1170–1178.

- Purushothaman, G., and Bradley, D.C. (2005). Neural population code for fine perceptual decisions in area MT. *Nat. Neurosci.* *8*, 99–106.
- Qamar, A.T., Cotton, R.J., George, R.G., Beck, J.M., Prezhdo, E., Laudano, A., Tolia, A.S., and Ma, W.J. (2013). Trial-to-trial, uncertainty-based adjustment of decision boundaries in visual categorization. *Proc. Natl. Acad. Sci. USA* *110*, 20332–20337.
- Rabinowitz, N.C., Goris, R.L., Cohen, M., and Simoncelli, E.P. (2015). Attention stabilizes the shared gain of V4 populations. *eLife* *4*, e08998.
- Rao, R.P., and Ballard, D.H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* *2*, 79–87.
- Schwartz, O., and Simoncelli, E.P. (2001). Natural signal statistics and sensory gain control. *Nat. Neurosci.* *4*, 819–825.
- Shadlen, M.N., Britten, K.H., Newsome, W.T., and Movshon, J.A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* *16*, 1486–1510.
- Shamir, M., and Sompolinsky, H. (2006). Implications of neuronal diversity on population coding. *Neural Comput.* *18*, 1951–1986.
- Helmholtz, H. (1867). *Handbuch der physiologischen Optik* (Leipzig: Leopold Voss).
- Vul, E., Goodman, N., Griffiths, T.L., and Tenenbaum, J.B. (2014). One and done? Optimal decisions from very few samples. *Cogn. Sci.* *38*, 599–637.
- Wei, X.-X., and Stocker, A. (2012). Efficient coding provides a direct link between prior and likelihood in perceptual Bayesian inference. In *Advances in Neural Information Processing Systems 25*, F. Pereira, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, eds. (Neural Information Processing Systems Foundation), pp. 1–9.
- Wimmer, K., Compte, A., Roxin, A., Peixoto, D., Renart, A., and de la Rocha, J. (2015). Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. *Nat. Commun.* *6*, 6177.
- Zohary, E., Shadlen, M.N., and Newsome, W.T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* *370*, 140–143.