

Trial-to-trial, uncertainty-based adjustment of decision boundaries in visual categorization

Ahmad T. Qamar^{a,1}, R. James Cotton^{a,1}, Ryan G. George^{a,1}, Jeffrey M. Beck^b, Eugenia Prezhdo^a, Allison Laudano^a, Andreas S. Tolias^a, and Wei Ji Ma^{a,2,3}

^aDepartment of Neuroscience, Baylor College of Medicine, Houston, TX 77030; and ^bDepartment of Brain and Cognitive Science, University of Rochester, Rochester, NY 14627

Edited by Barbara Anne Doshier, University of California, Irvine, CA, and approved October 15, 2013 (received for review November 13, 2012)

Categorization is a cornerstone of perception and cognition. Computationally, categorization amounts to applying decision boundaries in the space of stimulus features. We designed a visual categorization task in which optimal performance requires observers to incorporate trial-to-trial knowledge of the level of sensory uncertainty when setting their decision boundaries. We found that humans and monkeys did adjust their decision boundaries from trial to trial as the level of sensory noise varied, with some subjects performing near optimally. We constructed a neural network that implements uncertainty-based, near-optimal adjustment of decision boundaries. Divisive normalization emerges automatically as a key neural operation in this network. Our results offer an integrated computational and mechanistic framework for categorization under uncertainty.

Bayesian inference | vision | decision-making | optimality

Imagine a woman is approaching you from a distance and you are trying to determine whether or not she is the friend you are waiting for. Because of various sources of noise, your observations of her facial features, hair color, etc. will be uncertain. A sensible strategy would be to be more tolerant to deviations between your observations and your knowledge of your friend's looks when she is far away than when she is close by and your observations are less uncertain. In this categorization problem, you are determining whether the image of the approaching woman falls into the narrow category of images of your friend or the wide category of images of all other people. Categorization can be modeled as a process of applying one or more decision boundaries to a noisy measurement in a space of stimulus features (1–7). The example suggests that adjusting such decision boundaries based on the current level of sensory uncertainty might be a better strategy than using uncertainty-independent decision boundaries.

Previous studies have not addressed whether organisms adjust their decision boundaries from trial to trial according to the level of sensory uncertainty. Perceptual studies of categorization under sensory uncertainty have typically used category distributions for which the level of uncertainty was irrelevant for optimal behavior (2, 3, 6, 8). For example, in a classic task, observers categorize the direction of motion of a set of dots coherently moving to the left or to the right, in the presence of distractor dots moving in random directions (8). Regardless of the level of sensory noise corrupting the brain's measurement of the net motion direction, the optimal decision is simply to report whether this measurement was to the right or to the left. In other words, applying a fixed decision boundary to a scalar estimate is optimal in this task; no knowledge of uncertainty about motion direction is needed. In cognitive models of categorization, dynamic decision boundaries have been invoked to explain a broad range of phenomena, including sequential effects (9, 10), context effects (11), and generalization (12). However, these studies limited themselves to fixed levels of sensory noise and were not able to demonstrate optimality of behavior. Thus, a dichotomy exists: perceptual models are often normative and describe behavior in tasks with variable sensory uncertainty but trivial category distributions, whereas cognitive models examine more complex

forms of categorization but are typically nonnormative and ignore the role of sensory uncertainty.

Here, we attempt to connect these domains using a visual categorization task in which sensory noise is varied unpredictably from trial to trial. Our simple experimental design allows us to determine how observers should adjust their decision boundaries to achieve optimal performance; thus, our approach is normative. We found that humans and monkeys do adjust their decision boundaries from trial to trial according to sensory uncertainty. We also constructed a biologically inspired neural network model that can perform near-optimal, uncertainty-based adjustment of decision boundaries. Thus, we offer both a computational and a mechanistic account of brain function in a task in which trial-to-trial sensory uncertainty drives decision boundary dynamics.

Results

Task. Human and monkey observers categorized the orientation of a drifting grating. The two categories ($C = 1, 2$) were defined by Gaussian probability distributions with the same mean (defined as 0°) and different SDs, denoted by σ_1 and σ_2 for categories 1 and 2, respectively (Fig. 1*A* and *B*). As in related tasks (13, 14), the overlap of these distributions introduces ambiguity: a given orientation can come from either category, and therefore, categorization performance cannot be perfect even in the absence of sensory noise. Observers were trained using high-contrast stimuli and trial-to-trial feedback. During testing, contrast was varied from trial to trial to manipulate sensory uncertainty. Human observers did not receive trial-to-trial feedback during testing.

Theory. The statistical structure of the task, also called the generative model, contains three variables (Fig. 1*C*): category

Significance

Categorization is an important part of perception and cognition. For example, an animal must successfully categorize a disturbing sound as being due to the wind or to a predator. Computationally, categorization amounts to applying decision boundaries to noisy stimulus measurements. Here, we examine how these decision boundaries change as the quality of the sensory evidence varies unpredictably from trial to trial. We show that both humans and monkeys adjust their decision boundaries from trial to trial, often near-optimally. We further show how a neural network can perform this computation near-optimally. Our results might lead to a better understanding of categorization.

Author contributions: R.J.C., A.S.T., and W.J.M. designed research; A.T.Q., R.J.C., R.G.G., J.M.B., E.P., A.L., A.S.T., and W.J.M. performed research; A.T.Q., R.G.G., E.P., and W.J.M. analyzed data; and W.J.M. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹A.Q., R.J.C., and R.G. contributed equally to this work.

²Present address: Center for Neural Science and Department of Psychology, New York University, New York, NY 10003.

³To whom correspondence should be addressed. E-mail: weijima@nyu.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1219756110/-DCSupplemental.

C , stimulus orientation s , and measurement x . On each trial, C is drawn randomly and determines whether s is drawn from the narrow Gaussian, $p(s | C = 1)$ with SD σ_1 or the wide Gaussian, $p(s | C = 2)$ with SD σ_2 . We assume that on each trial, the true orientation is corrupted by sensory noise to give rise to the observer's measurement of orientation, x . We denote by $p(x | s)$ the probability distribution over x for a given stimulus orientation s . We assume it to be Gaussian with mean s and SD σ (6). This SD is experimentally manipulated through contrast.

On a given trial, the observer uses the measurement x to infer category C . An optimal observer would do this by computing the posterior probability distribution over category, denoted $p(C | x)$, which indicates the degree of belief that the category was C , based on x . It is convenient to express the posterior in terms of the log posterior ratio, d , which is, using Bayes' rule, the sum of the log likelihood ratio and the log prior ratio

$$d = \log \frac{p(C=1 | x)}{p(C=2 | x)} = \log \frac{p(x | C=1)}{p(x | C=2)} + \log \frac{p_1}{1-p_1}, \quad [1]$$

where p_1 is the observer's prior belief that $C = 1$. The absolute value of d is one possible measure of decision confidence. The difficulty in computing the category likelihood $p(x | C)$ is that the stimulus s that caused x is unknown. The optimal observer deals with this by multiplying, for every possible s separately, the probability that the observed x came from this s with the probability of this s under the hypothesized category, and finally summing over all s

$$p(x | C) = \int p(x | s)p(s | C)ds = \int L_x(s)p(s | C)ds. \quad [2]$$

Here, we use the notation $L_x(s)$ to denote the likelihood of a stimulus value s based on a measurement x , $L_x(s) = p(x | s)$. The width of this function, also σ , measures sensory uncertainty (Fig. 1E). Thus, Eq. 2 reflects the combination of sensory

information, $L_x(s)$, with category information, $p(s | C)$. An integral over a hidden variable, as in Eq. 2, is known as marginalization and is a central operation in Bayesian computation (15–18). A straightforward calculation gives (SI Text)

$$d = k_1 - k_2 x^2, \quad [3]$$

where $k_1 = \frac{1}{2} \log \frac{\sigma^2 + \sigma_2^2}{\sigma^2 + \sigma_1^2} + \log \frac{p_1}{1-p_1}$ and $k_2 = \frac{\sigma_2^2 - \sigma_1^2}{2(\sigma^2 + \sigma_1^2)(\sigma^2 + \sigma_2^2)}$. The decision strategy that maximizes accuracy is the maximum-a-posteriori (MAP) read-out, i.e., to report the value of C for which $p(C | x)$ is larger. This rule is equivalent to reporting category 1 when d is positive, or in other words, when $|x| < \sqrt{\frac{k_1}{k_2}} \equiv k$. Thus, the optimal observer reports category 1 when the measurement lies within the interval from $-k$ to k and category 2 otherwise. Critically, the optimal boundary or criterion k depends on the sensory uncertainty σ : when evidence is uncertain, the optimal observer is more willing to attribute measurements far away from zero to category 1 (Fig. 1D). This effect, which reflects the intuition of the friend recognition example, is a direct consequence of the shape of the category distributions. We consider two variants of the optimal model: one in which $p_1 = 0.5$, reflecting the experimental statistics (which we call the Opt model), and one in which p_1 is a free parameter (Opt-P model).

The main alternative to the optimal model is one in which the observer uses a fixed decision boundary (Fixed model). Then, the decision rule is $|x| < k_0$, with k_0 a constant.

Of course, the optimal model is not the only possible model in which the observer takes into account uncertainty on a trial-to-trial basis, even when we restrict ourselves to decision rules of the form $|x| < \text{function of } \sigma$. As a first step in exploring this model space, we test all linear functions of σ (we call this model Lin- σ): the observer uses the rule $|x| < k_0 \left(1 + \frac{\sigma}{\sigma_p}\right)$. We also test a model in which the observer applies a fixed boundary not to the measurement x but to the MAP estimate of the stimulus, $\frac{\sigma_p^2}{\sigma_p^2 + \sigma^2} x$, obtained under a Gaussian prior with mean 0 and SD σ_p . The decision rule is then $|x| < k_0 \left(1 + \frac{\sigma^2}{\sigma_p^2}\right)$. We call this the Quad- σ model (Fig. 1E). In all models except for the Fixed model, the observer takes into account the trial-to-trial level of sensory uncertainty; these models therefore describe probabilistic, but not necessarily optimal, computation (19).

In each model, we describe the relationship between noise variance σ^2 and contrast c (expressed as a proportion, not as a percentage) as a power law with a baseline, $\sigma^2(c) = (\alpha c)^{-\beta} + \gamma$, and include a lapse rate λ to account for random guesses and unintended responses (20). The Opt model has four free parameters (α , β , γ , and λ), Opt-P and Fixed each have one more (p_1 and k , respectively), and Lin- σ and Quad- σ each have two more (k and σ_p). Thus, Lin- σ and Quad- σ can be considered more flexible models than Opt and Opt-P. Models are summarized in Table S1.

Behavioral Results. We first obtained maximum-likelihood estimates of the model parameters (Table S2). To compare models, we then computed both the marginal log likelihood (using the Laplace approximation) and the Akaike information criterion (AIC) for each model and each subject (Materials and Methods).

Our goal was to determine whether observers take into account sensory uncertainty in setting their decision boundaries. Thus, we are interested in whether or not the Fixed model accounts better for the data than all probabilistic models. It did not for any human subject or either monkey, according to either measure (Table 1 and Table S3). Moreover, the differences between the best probabilistic model and the Fixed model were large; to illustrate, Jeffreys (21) considered a marginal log likelihood difference of more than $\log(30) = 3.4$ very strong evidence. Subjects differed in which probabilistic model described their data best, with Opt, Opt-P, Lin- σ , and Quad- σ each winning for at least

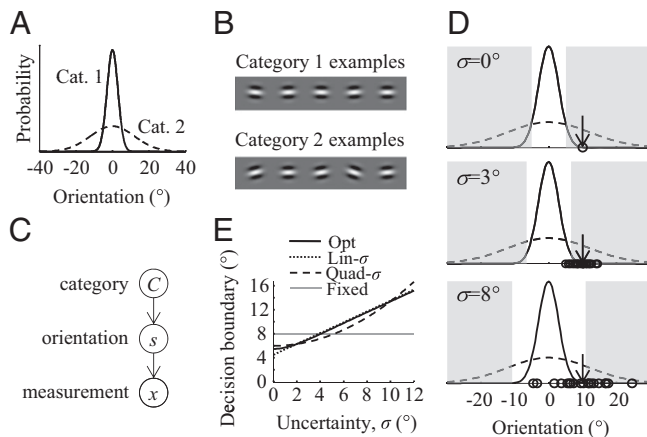


Fig. 1. Task and models. (A) Probability distributions over orientation for Categories 1 and 2. The distributions have the same mean (0°) but different SDs ($\sigma_1 = 3^\circ$ and $\sigma_2 = 12^\circ$; for Monkey L, $\sigma_2 = 15^\circ$). (B) Sample stimuli drawn from each category. (C) Generative model of the task (see text). (D) A completely certain ($\sigma = 0^\circ$) optimal observer would set the decision boundaries $\pm k$ at the intersection points of the two category distributions (black curves). The shaded areas indicate where the optimal observer would respond Category 2. When sensory noise is larger, not only will the measurements (open circles) be more variable for a given true orientation (arrow), but the optimal observer will also move the decision boundaries to larger values. In the experiment, noise levels are interleaved. (E) Decision boundary as a function of uncertainty level under four models: optimal, linear (example with $k_0 = 4.5^\circ$ and $\sigma_p = 5^\circ$), quadratic (example with $k_0 = 6^\circ$ and $\sigma_p = 9^\circ$), and Fixed (example with $k_0 = 8^\circ$).

Table 1. Model comparison using log marginal likelihood for main experiment

Subject	Opt	Opt-P	Lin- σ	Quad- σ	Fixed	DIFF
Human 1	-1,479.0	-1,429.1	-1,422.9	-1,425.9	-1,452.0	-29.1
Human 2	-1,642.7	-1,646.2	-1,648.6	-1,660.6	-1,718.6	-75.9
Human 3	-891.7	-893.0	-897.6	-897.7	-919.6	-27.9
Human 4	-1,396.6	-1,383.5	-1,353.7	-1,364.8	-1,500.6	-146.8
Human 5	-1,280.6	-1,272.4	-1,241.3	-1,245.4	-1,318.0	-76.7
Human 6	-1,152.0	-1,131.0	-1,134.9	-1,135.6	-1,162.5	-31.5
Monkey A	-33,875.2	-32,862.9	-32,954.4	-32,762.6	-35,354.8	-2,592.2
Monkey L	-71,973.6	-71,787.5	-71,464.0	-71,501.7	-73,167.5	-1,703.5

Numbers are marginal log likelihoods obtained using the Laplace approximation. Shaded in green are the models whose values fall within $\log(30)$ of the value of the best model. The Fixed model is never among them. DIFF, difference between the Fixed model and the best probabilistic model.

one human subject, according to either method. The data of monkey A were best described by Quad- σ and those of monkey L best by Lin- σ . These results suggest that observers use different and sometimes suboptimal strategies to incorporate sensory uncertainty information on a trial-by-trial basis, but no observer ignores this information.

Fig. 2 shows three types of psychometric curves with corresponding model fits: proportion correct as a function of contrast (Fig. 2A), proportion of category 1 reports as a function of contrast and true category (Fig. 2B), and as a function of contrast and orientation (Fig. 2C). For each curve, the Fixed model provides a much worse fit than each of the probabilistic models (measured by RMSE), providing further indication that observers compute their decision boundaries in an uncertainty-dependent manner. In Fig. 2C, we observe that the curve widens as contrast decreases; because the Fixed model has a fixed decision boundary, it cannot account for this widening. Among the probabilistic models, the Lin- σ model fits best overall, in accordance with Table 1 and Table S3, but the differences are small. Although the Opt model provides a slightly worse fit, it should be kept in mind that in this model, the decision boundary is rigidly specified as a function of uncertainty σ : in contrast to the Lin- σ and Quad- σ models, the Opt model does not introduce any free parameters in the decision boundary function. In this light, the good fit of the Opt model is remarkable.

To agnostically estimate the uncertainty-dependent decision boundary humans use, we fitted an additional, flexible

model, in which a separate boundary is fitted at each contrast level (SI Text). The decision rule is then $|x| < k_c$, where k_c is the boundary at contrast c . The resulting estimates of k_c are plotted in Fig. 2D. We find a significant effect of contrast on k_c [repeated-measures ANOVA: $F(5,25) = 12.4$, $P < 10^{-5}$]. The probabilistic models account for the trend in k_c much better than the Fixed model (in RMSE).

To confirm that the fitted values of sensory uncertainty σ in the probabilistic models are meaningful, we measured them in an independent experiment. The same six human observers performed a left/right orientation discrimination task (Materials and Methods) under the same contrasts and other experimental settings as in the categorization experiment. The estimates of σ obtained from the categorization task were strongly correlated with those obtained from the discrimination task (Fig. 2E; Pearson $r = 0.89$, $P < 10^{-3}$). Using the estimates of σ from the discrimination task instead of those obtained from the estimates of α , β , and γ in the flexible model, although worse, still produces reasonable probabilistic model fits to the decision boundary function (Fig. 2F). Together, these results are evidence that the fitted values of σ in the main experiment are meaningful.

To ascertain that our results did not depend on the choice of vertical as the central orientation, we repeated the experiment using 45° clockwise from vertical instead. Results were consistent with the main experiment (SI Text and Fig. S1).

Monkey summary statistics are shown in Fig. 3 and Fig. S2. Again, the Fixed model did not account well for the data. For

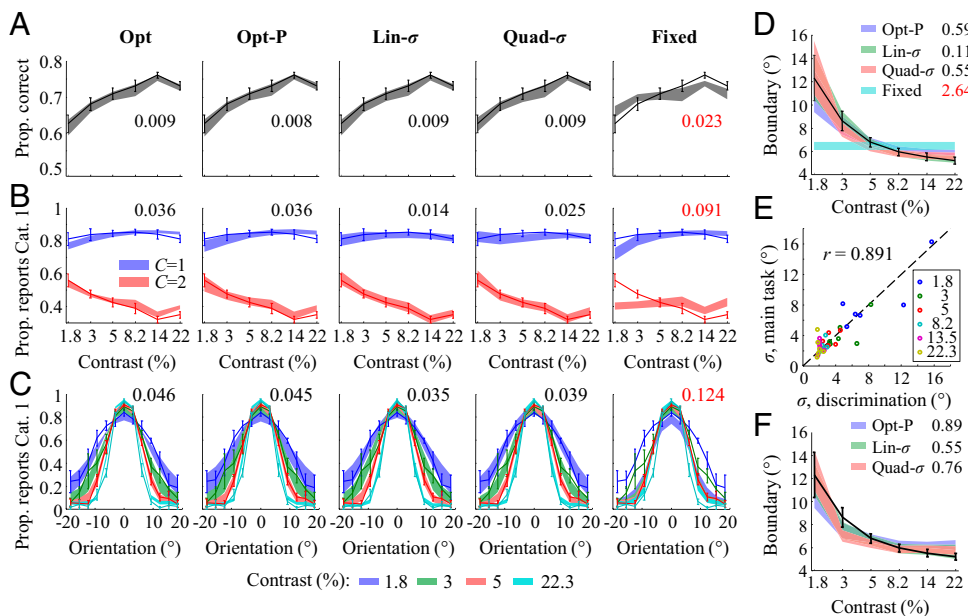


Fig. 2. (A) Proportion correct versus contrast, with model fits. Model fits often do not look smooth because they were computed based on the orientations actually presented in the experiment, and those were drawn randomly (see SI Text). Here and elsewhere, error bars and shaded areas indicate 1 SEM, and numbers indicate root mean squared error between data and model means (red is worst). (B) Proportion "Category 1" responses versus contrast, separated by true category, with model fits. (C) Proportion "Category 1" responses versus orientation and contrast, with model fits. For visibility, not all contrasts are plotted. (D) Decision boundaries fitted separately at each contrast level (error bars: data; shaded areas: models). (E) Estimates of sensory uncertainty σ estimated from the categorization task against ones obtained from an independent orientation discrimination task; color labels contrast. (F) Decision boundaries predicted by three probabilistic models based on the uncertainty estimates from the discrimination task.

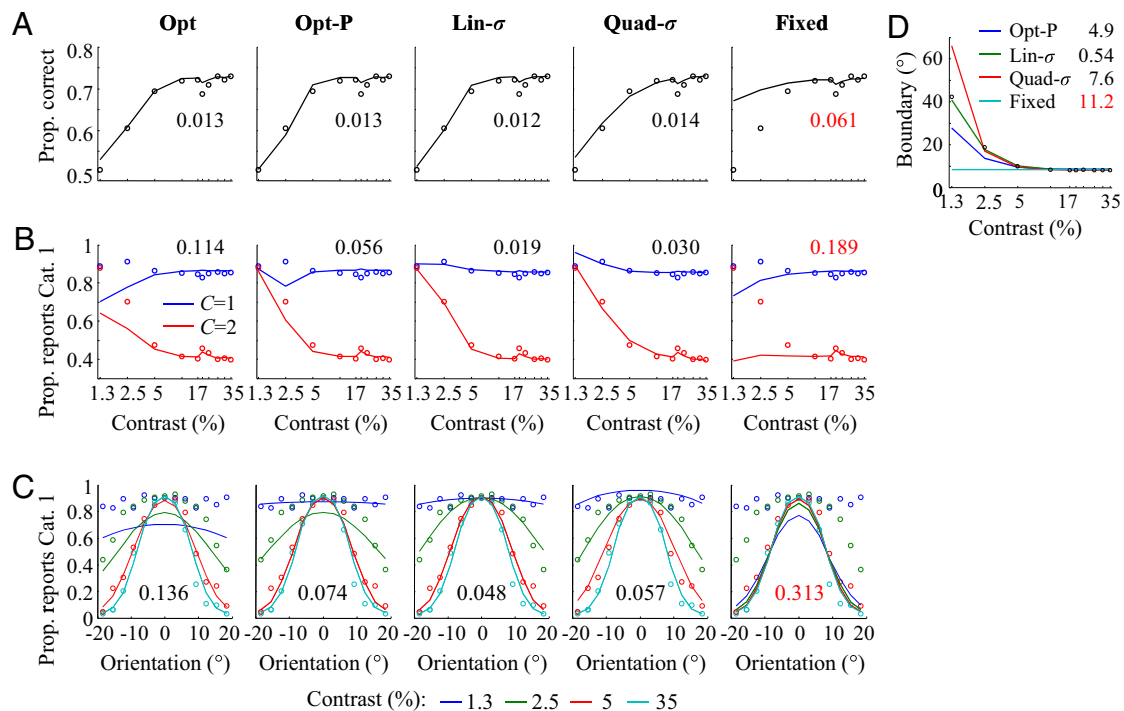


Fig. 3. As in Fig. 2, but for Monkey L.

monkey L, Lin- σ provided the best fit (Fig. 3), and for monkey A, Quad- σ (Fig. S2). Taken together, our results indicate that both humans and monkeys adjust their decision boundaries from trial to trial based on sensory uncertainty, but use diverse and sometimes suboptimal strategies to do so.

Neural Network Results. Our next goal is to provide a proof of concept that a biologically plausible neural network can adjust the decision boundary from trial to trial based on sensory uncertainty. We do this for the optimal model, because in this model, on each trial, the observer does not simply evaluate the decision rule, but also has a representation of the probability that the decision is correct. Although we did not explore it experimentally, knowing this probability is important for combining category information with information about the consequences of actions. For example, an animal might categorize a sound as being caused by the wind, but if the posterior probability of wind is 55% and that of predator is 45%, the best course of action would still be to run. Thus, for this animal, it is important to know whether the posterior probability of wind is 55% or 99%. In many forms of decision-making, observers must possess knowledge of posterior probability to maximize reward (3, 22–24). Although such a posterior probability is naturally given by the optimal model, it might also be possible to associate (suboptimal) posterior probabilities with the decision rules in the Lin- σ and Quad- σ models. However, this is not obvious and we leave it for future work.

To obtain a neural code, we replace the scalar observation x by the vector of spike counts in a population of orientation-tuned neurons, denoted \mathbf{r} (Fig. 4A). We assume variability in the exponential family with linear sufficient statistics (Poisson-like variability) (25), which is consistent with the physiology of primary visual cortex (26, 27). Then, the distribution of \mathbf{r} across trials for a given stimulus s can be described by

$$p(\mathbf{r} | s, g) = \varphi(\mathbf{r}, g) e^{\mathbf{h}(s) \cdot \mathbf{r}}, \quad [4]$$

where g denotes the gain (mean amplitude) of the population, which is affected by contrast, and φ is an arbitrary function. Using the framework of probabilistic population coding (25), all available

information about the stimulus is contained in the neural likelihood function, $L_{\mathbf{r}}(s) = p(\mathbf{r} | s) = \int p(\mathbf{r} | s, g) p(g) dg = \Phi(\mathbf{r}) e^{\mathbf{h}(s) \cdot \mathbf{r}}$ (Fig. 4B), where Φ is easily expressed in terms of φ . We assume that $\mathbf{h}(s)$ is a quadratic function of s and thus can be written as $\mathbf{h}(s) = -\frac{1}{2}s^2 \mathbf{a} + s \mathbf{b}$, where \mathbf{a} and \mathbf{b} are constant vectors, so that the likelihood $L_{\mathbf{r}}(s)$ is an (unnormalized) Gaussian. The mean $\frac{\mathbf{b} \cdot \mathbf{r}}{\mathbf{a} \cdot \mathbf{r}}$ and variance $\frac{1}{\mathbf{a} \cdot \mathbf{r}}$ of this likelihood function correspond to x and σ^2 in the behavioral model, respectively. In the special case of independent Poisson variability with Gaussian tuning curves, \mathbf{a} is a vector whose entries are equal to the inverse squared widths of the neurons' tuning curves, and $\frac{\mathbf{b} \cdot \mathbf{r}}{\mathbf{a} \cdot \mathbf{r}}$ is the center-of-mass decoder (28), the real-line analog of the population vector decoder (29) (SI Text). The log likelihood ratio over category can then be found from Eq. 3 by substituting the neural quantities for x and σ^2

$$d = \log \frac{p(\mathbf{r} | C=1)}{p(\mathbf{r} | C=2)} = \frac{1}{2} \log \frac{1 + \sigma_2^2 \mathbf{a} \cdot \mathbf{r}}{1 + \sigma_1^2 \mathbf{a} \cdot \mathbf{r}} - \frac{(\sigma_2^2 - \sigma_1^2) (\mathbf{b} \cdot \mathbf{r})^2}{2(1 + \sigma_1^2 \mathbf{a} \cdot \mathbf{r})(1 + \sigma_2^2 \mathbf{a} \cdot \mathbf{r})}. \quad [5]$$

Compared with Eq. 3, we have left out the log prior ratio; this is a constant shift and therefore easily implemented. A key aspect of Eq. 5 is that the log posterior ratio is a nonlinear function of \mathbf{r} , in line with probabilistic population code implementations of other Bayesian computations that require marginalization (17, 18).

We propose that categorization is performed by a feedforward neural network that is Poisson-like both in input and output (Fig. 4C). For concreteness, one could think of the input layer as primary visual cortex, encoding orientation, and of the output layer as prefrontal cortex, encoding category; however, the computation does not depend on these labels. The assumption that category is encoded by a Poisson-like output population \mathbf{z} is supported by extant physiological findings in decision-making areas (30). The problem is then to find the mapping from \mathbf{r} to \mathbf{z} such that \mathbf{z} encodes not just category, but the optimal likelihood of category. Poisson-like variability in the output can be described by a probability distribution analogous to Eq. 4, namely $p(\mathbf{z} | C, g_z) = \varphi_{\mathbf{z}}(\mathbf{z}, g_z) e^{\mathbf{H}(C) \cdot \mathbf{z}}$.

probabilistic population code, the amplitude of the population pattern of activity encodes certainty. As a result, the stimulus estimate, x , must be computed using an amplitude-invariant operation. A representative example is the center-of-mass (or population vector) decoder, which computes the weighted sum of the neurons' preferred orientations, with weights equal to the neurons' spike counts; because the sum of the spike counts appears in the denominator, the center-of-mass decoder automatically contains a divisive normalization. Bayesian decision-making is performed by applying an operation to x and σ^2 . In the present task, the decision variable is a quadratic function of x (Eq. 3), explaining why quadratic operations and quadratic divisive normalization appear in the neural implementation (Eq. 5).

Quadratic operations and divisive normalization have proven crucial for implementing a number of seemingly disparate forms of Bayesian inference (19, 20). At first sight, these frequent appearances might be surprising, because in the probabilistic population code implementation of cue combination, neither quadratic operations nor divisive normalization are needed (27). However, this is because cue combination is special in that optimal inference only involves the ratio x/σ^2 , in which the divisive normalization is cancelled out to produce an operation linear in neural activity. However, this type of cancellation is the exception rather than the rule. In virtually every other form of Bayesian decision-making, the optimal decision variable will involve quadratic operations and divisive normalization when implemented with probabilistic population codes.

We make several predictions for physiology. First, we predict that populations of orientation-sensitive neurons in early visual cortex not only encode an estimate of orientation but also trial-to-trial information about sensory uncertainty and that this encoded sensory uncertainty correlates with the animal's decision boundary as obtained from its behavior in the categorization task. Second, we predict that the activity of category-sensitive neurons, which might be found in prefrontal cortex (2) or lateral intraparietal cortex (LIP)

(1, 5), is linearly related to the logarithm of the posterior probability ratio over category. This prediction is consistent with patterns found in LIP, where category probability can be reconstructed from a logistic mapping on neural activity (33). Third, we predict that sensory uncertainty decoded on each trial from early visual areas is propagated (along with the best estimate of orientation) to categorization neurons, so that including this decoded uncertainty as an explanatory factor should help to predict the activity of those neurons. Finally, we predict that if divisive normalization is selectively removed from a neural circuit involved in computing category, then the observer will become severely suboptimal in a way predicted by the QUAD network (Fig. 5). These predictions illustrate how a probabilistic, computationally [in the sense of Marr (40)] driven approach to categorization can guide the generation of hypotheses about neural mechanisms.

Materials and Methods

Details of all methods are provided in *SI Text*. Categories were defined by normal distributions with means 0° and SDs $\sigma_1 = 3^\circ$ and $\sigma_2 = 12^\circ$ (except for monkey L, for whom $\sigma_2 = 15^\circ$). On each trial, a category was selected with equal probability. The stimulus was a drifting Gabor (50 ms for humans; 500 ms for monkeys) whose orientation s was drawn from the distribution of the selected category. Contrast was varied randomly from trial to trial. For monkeys, through training, the narrow distribution was associated with a red target and the wide distribution with a green target; the monkey reported category through a saccade to the red or the green target. Humans responded using a key press. Stimuli were delivered using Psychophysics Toolbox for Matlab (Mathworks). Models were fitted using maximum-likelihood estimation, implemented through a conjugate gradient method. Bayes' factors were computed using the Laplace approximation. Neural networks were trained using a stochastic gradient descent on the Kullback-Leibler divergence between the true and the approximated posterior.

ACKNOWLEDGMENTS. This work was supported by National Science Foundation Grant IIS-1132009 (Collaborative Research in Computational Neuroscience) (to W.J.M. and A.S.T.).

- Freedman DJ, Assad JA (2006) Experience-dependent representation of visual categories in parietal cortex. *Nature* 443(7107):85–88.
- Freedman DJ, Riesenhuber M, Poggio T, Miller EK (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291(5502):312–316.
- Kepecs A, Uchida N, Zariwala HA, Mainen ZF (2008) Neural correlates, computation and behavioural impact of decision confidence. *Nature* 455(7210):227–231.
- Ashby FG, Lee WW (1991) Predicting similarity and categorization from identification. *J Exp Psychol Gen* 120(2):150–172.
- Swaminathan SK, Freedman DJ (2012) Preferential encoding of visual categories in parietal cortex compared with prefrontal cortex. *Nat Neurosci* 15(2):315–320.
- Green DM, Swets JA (1966) *Signal Detection Theory and Psychophysics* (John Wiley & Sons, Los Altos, CA).
- Ashby FG, Maddox WT (2005) Human category learning. *Annu Rev Psychol* 56:149–178.
- Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. *Nature* 341(6237):52–54.
- Stewart N, Brown GD, Chater N (2002) Sequence effects in categorization of simple perceptual stimuli. *J Exp Psychol Learn Mem Cogn* 28(1):3–11.
- Petrov AA, Anderson JR (2005) The dynamics of scaling: A memory-based anchor model of category rating and absolute identification. *Psychol Rev* 112(2):383–416.
- Lewandowsky S, Kalish M, Ngang SK (2002) Simplified learning in complex situations: Knowledge partitioning in function learning. *J Exp Psychol Gen* 131(2):163–193.
- Lamberts K (1994) Flexible tuning of similarity in exemplar-based categorization. *J Exp Psychol Learn Mem Cogn* 20(5):1003–1021.
- Sanborn AN, Griffiths TL, Shiffrin RM (2010) Uncovering mental representations with Markov chain Monte Carlo. *Cognit Psychol* 60(2):63–106.
- Liu Z, Knill DC, Kersten D (1995) Object classification for human and ideal observers. *Vision Res* 35(4):549–568.
- Kersten D, Mamassian P, Yuille A (2004) Object perception as Bayesian inference. *Annu Rev Psychol* 55:271–304.
- Knill DC (2003) Mixture models and the probabilistic structure of depth cues. *Vision Res* 43(7):831–854.
- Ma WJ, Navalpakkam V, Beck JM, Berg Rv, Pouget A (2011) Behavior and neural basis of near-optimal visual search. *Nat Neurosci* 14(6):783–790.
- Beck JM, Latham PE, Pouget A (2011) Marginalization in neural circuits with divisive normalization. *J Neurosci* 31(43):15310–15319.
- Ma WJ (2012) Organizing probabilistic models of perception. *Trends Cogn Sci* 16(10):511–518.
- Wichmann FA, Hill NJ (2001) The psychometric function: I. Fitting, sampling, and goodness of fit. *Percept Psychophys* 63(8):1293–1313.
- Jeffreys H (1961) *The Theory of Probability* (Oxford Univ Press, Oxford), 3rd Ed, p 470.
- Whiteley L, Sahani M (2008) Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. *J Vis* 8(3):1–15.
- Maiworm M, König P, Röder B (2011) Integrative processing of perception and reward in an auditory localization paradigm. *Exp Psychol* 58(3):217–226.
- Kiani R, Shadlen MN (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* 324(5928):759–764.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9(11):1432–1438.
- Graf ABA, Kohn A, Jazayeri M, Movshon JA (2011) Decoding the activity of neuronal populations in macaque primary visual cortex. *Nat Neurosci* 14(2):239–245.
- Berens P, et al. (2012) A fast and simple population code for orientation in primate V1. *J Neurosci* 32(31):10618–10626.
- Ma WJ, Pouget A (2009) *Population Coding: Theoretic Aspects. Encyclopedia of Neuroscience* (Elsevier, New York), Vol 7, pp 749–755.
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2(11):1527–1537.
- Yang T, Shadlen MN (2007) Probabilistic reasoning by neurons. *Nature* 447(7148):1075–1080.
- Trotter Y, Celebrini S (1999) Gaze direction controls response gain in primary visual-cortex neurons. *Nature* 398(6724):239–242.
- Galletti C, Battaglini PP (1989) Gaze-dependent visual neurons in area V3A of monkey prestriate cortex. *J Neurosci* 9(4):1112–1125.
- Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons. *Science* 230(4724):456–458.
- Heeger DJ (1992) Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9(2):181–197.
- Carandini M, Heeger DJ (2012) Normalization as a canonical neural computation. *Nat Rev Neurosci* 13(1):51–62.
- Knill DC, Richards W, eds (1996) *Perception as Bayesian Inference* (Cambridge Univ Press, New York).
- Trommershauser J, Kording K, Landy MS, eds (2011) *Sensory Cue Integration* (Oxford Univ Press, New York).
- van den Berg R, Vogel M, Josic K, Ma WJ (2012) Optimal inference of sameness. *Proc Natl Acad Sci USA* 109(8):3178–3183.
- Keshvari S, van den Berg R, Ma WJ (2012) Probabilistic computation in human perception under variability in encoding precision. *PLoS ONE* 7(6):e40216.
- Marr D (1982) *Vision* (MIT Press, Cambridge, MA).