Neural Coding of Uncertainty and Probability

Wei Ji Ma¹ and Mehrdad Jazayeri²

¹Center for Neural Science and Department of Psychology, New York University, New York, New York 10003; email: weijima@nyu.edu

²McGovern Institute for Brain Research and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139; email: mjaz@mit.edu

Annu. Rev. Neurosci. 2014. 37:205-20

The Annual Review of Neuroscience is online at neuro.annualreviews.org

This article's doi: 10.1146/annurev-neuro-071013-014017

Copyright © 2014 by Annual Reviews. All rights reserved

Keywords

Bayesian inference, decision making, perception, population encoding

Abstract

Organisms must act in the face of sensory, motor, and reward uncertainty stemming from a pandemonium of stochasticity and missing information. In many tasks, organisms can make better decisions if they have at their disposal a representation of the uncertainty associated with task-relevant variables. We formalize this problem using Bayesian decision theory and review recent behavioral and neural evidence that the brain may use knowledge of uncertainty, confidence, and probability.

PROBABILITY, BELIEFS, AND UNCERTAINTY	206
THEORY OF PROBABILISTIC COMPUTATION	207
BEHAVIORAL EVIDENCE FOR PROBABILISTIC COMPUTATION	209
THE FLEXIBILITY OF PROBABILISTIC COMPUTATION	210
NEURAL REALIZATIONS OF PROBABILISTIC COMPUTATION	212
NEURAL SIGNATURES OF UNCERTAINTY	214
FUTURE DIRECTIONS	215

PROBABILITY, BELIEFS, AND UNCERTAINTY

The information that organisms have available to make decisions is remarkably limited and impoverished. For example, visual signals are degraded in the dark, other individuals' internal states are not directly accessible, and the amount of food available in food sources may vary depending on many unknown factors. Moreover, our nervous system is not perfect. For example, threedimensional visual objects are projected onto two-dimensional sensors, causing ambiguity (Kersten et al. 2004), and responses to the same sensory event are inherently variable (Faisal et al. 2008).

Knowing the nature of internal and external stochastic processes will generally help organisms to make better decisions. In sensory processing, to infer the state of the world from variable and ambiguous sensory inputs, the brain would benefit from knowing the probabilistic relationships between stimuli and the sensory responses they evoke. On the motor side, knowing the statistics of one's motor variability can enhance the effectiveness of movements. Knowledge of the variability of rewards and costs could also be highly informative. Higher faculties such as anticipation, planning, decision making, and thinking can all benefit from knowledge about the probabilistic contingencies in the environment and stochasticity within the nervous system.

"Belief" is a term used to describe an agent's knowledge of probabilistic information about variables that describe the state of the world, the state of the body, or a mental state. Mathematically, belief can be viewed as a "subjective probability": The stronger one's belief in a particular proposition, the higher the corresponding subjective probability. For example, when trying to cross a road, you could maintain a belief distribution over the speed of an approaching car. Some beliefs may be hardwired from birth, whereas others could be acquired flexibly through experience or vary from trial to trial as the sensory input varies. For an optimal observer, beliefs are based on the actual probabilistic relationships between variables, but in general they need not be. Uncertainty is typically specified by some measure of the width of the belief distribution. In the road-crossing example, sensory uncertainty could be defined as the standard deviation of the belief distribution over the car speed (e.g., 30 ± 2 km/h): In fog, the uncertainty may be higher (e.g., 30 ± 10 km/h).

The primary alternative to using belief distributions and uncertainty is to use point estimates of variables, e.g., "The car's speed is 30 km/h." However, in most realistic conditions, uncertainty information is relevant for decision making: You might cross the road if a car's speed is 30 ± 2 km/h but not if it is 30 ± 10 km/h. The same holds in higher cognition: When you hear an outrageous statement but are uncertain whether the speaker is joking, you may seek clarification instead of getting angry. In this paper, we review recent studies that have shown that, under suitable conditions, humans and animals behave as if they do make use of belief distributions and uncertainty. We also provide a critical view of what it means to use belief distributions and discuss outstanding questions related to the neural basis of such probabilistic computation.



Figure 1

A probabilistic formalization of different classes of decision-making tasks. To simplify notation, we ignore temporal dynamics. Each node contains a random variable, which is described by a probability distribution that is conditioned on the variables from which arrows point. For example, the statistics of the outcome o depend on the action taken, a (we ignore potential other dependencies for simplicity). The brain's goal is to realize an appropriate mapping from neural activity r to motor commands a. This can involve taking into account uncertainty/probabilistic information.

THEORY OF PROBABILISTIC COMPUTATION

Decision-making studies have advanced along three relatively independent tracks: (a) Perceptual decisions are concerned with judgments about a stimulus; (b) sensorimotor decisions incorporate sensorimotor contingencies to specify when, where, and how to make movements; and (c) rewardbased decisions focus on behavioral responses that are based on an assessment of the utility of choice options. Figure 1 shows a schematic that encompasses these different classes of decision making. The so-called forward or generative model characterizes the true probabilistic relationships between different variables. Suppose there is a particular aspect of the world (the world includes one's own body), denoted C, that is of interest, for example, whether just-expired milk is still good. Relevant to C are certain sensory features, denoted s, for example, the odorants of the milk. Each value of C and each value of s have a certain frequency of occurring together in the world, which is represented by a joint probability distribution p(C,s). Upon interrogating the environment, the sensory feature s evokes a neural population activity, denoted r, for example, the olfactory response to the odorants. Because of the inherent variability in neural responses, the relationship between r and s is stochastic and is best described by a probability distribution p(r|s). Decision making or acting is the process of mapping the neural activity r to a decision or action (motor command), denoted a, for example, to drink or not to drink the milk. In perceptual tasks, a is simply the observer's estimate of the world state C. The action a leads to an outcome o in the external world, for example, the milk being ingested. This mapping may be subject to motor noise and can be formulated by $p(\rho \mid a)$. Finally, the state of the world and the outcome may lead to reward or punishment, denoted R, for example, getting sick from drinking bad milk. In general,

reward may be stochastic for a given world state C and outcome v; for example, there might be 50% chance of getting sick after drinking bad milk.

We can use **Figure 1** to define more rigorously different types of belief and uncertainty. Let us first consider an optimal (ideal) observer, for whom beliefs follow directly from the true probability distributions. Sensory or perceptual beliefs are over *C* and *s* and are captured by the distribution p(C,s|r); in our examples, these would be distributions over car speed or milk quality. Outcome or motor beliefs are expressed by p(o|a), which represents the agent's knowledge of the nature of the noise that will influence a given motor command. Finally, the belief distribution over reward is p(R|o,C): For example, one option might give you a 50% chance of receiving \$20, while the other option yields \$8 for sure. In perceptual and sensorimotor tasks, the contingency between outcome and reward is typically deterministic, which reduces p(R|o,C) to a deterministic reward or cost function; for example, when *o* is an estimate of *C* and both are continuous, the cost function can be defined as the squared estimation error $(o - C)^2$. These three types of belief distributions and their corresponding measure of uncertainty are distinct and may have different neural underpinnings.

The fundamental question we ask is whether and how the brain incorporates belief distributions or uncertainty in the mapping from response *r* to action *a*. At one extreme, the brain would not use belief distributions at all, but instead use point estimates of stimuli, even when all the dependencies shown in **Figure 1** are stochastic. At the other extreme, the brain acts as an optimal observer that knows everything about the generative model shown in **Figure 1** and uses this knowledge in the best possible way. In between lie many possibilities, including cases in which the brain utilizes either belief distributions over a subset of variables or belief distributions that deviate from the true distributions.

Bayesian decision theory specifies how an optimal observer would utilize belief distributions. The agent first computes the probability of receiving reward R when sensory activity is r and the planned action is a. This probability is

$$p(R|r,a) = \iint p(R|o,C)p(o|a)p(C|r)dodC.$$
1.

The three factors in the integrand correspond to the three belief distributions: reward, outcome, and sensory beliefs, respectively. The optimal observer uses these belief distributions to average over the world state *C* and outcome *o* that are not known (such integration is also called marginalization). Optimality is defined as executing actions *a* that maximize utility under the distribution p(R|r,a). Utility could simply be the expected value of *R*, or it could be a complicated nonlinear function of the distribution p(R|r,a), for example in order to account for risk aversion (Kahneman & Tversky 1979, Glimcher et al. 2008).

The distribution p(C|r) is known as the posterior distribution over *C* and represents knowledge about the world state *C* after combining sensory information with prior expectations. The posterior can be further broken down (using Bayes' rule) as the normalized product of a likelihood p(r|C) and a prior p(C). The likelihood p(r|C) is computed by integrating the product of p(r|s) (the likelihood over *s*) and p(s|C) over *s* (another marginalization); this operation transforms beliefs over *s* into beliefs over *C*. Let us consider a simple example of a sensory decision in which an observer receives a unit reward (R = 1) for correctly estimating the state of the world *C*. In this case, because the outcome *o* and action *a* are determined by the observer's estimate of *C*, denoted \hat{C} , Equation 1 can be simplified to $p(R = 1|r, \hat{C}) = p(C = \hat{C}|r)$. Then, maximizing reward reduces to maximizing the posterior, i.e., the strength of the observer's belief that \hat{C} is the true state of the world.

Bayesian decision theory provides a rigorous definition for confidence, namely the belief associated with the proposition that the observer has chosen or intends to choose. For example, in the sensory decision task described above, confidence would be the posterior probability of the observer's estimate of the state of the world, $p(C = \hat{C} | r)$ (or a monotonic function of it). More generally, confidence can be defined as the observer's belief that the chosen action maximizes utility (de Martino et al. 2013). A basic way to measure confidence in humans is using a discrete rating scale (Peirce & Jastrow 1884), but we review several other methods below (see Neural Signatures of Uncertainty). Confidence generally correlates with task performance, and the strength of this correlation—measuring the accuracy of one's knowledge of the quality of one's decisions—is itself typically correlated with, although sometimes dissociable from, performance (Fleming & Dolan 2012).

Using belief distributions in decision making, also called probabilistic computation, is neither necessary nor sufficient for optimal performance (Ma 2012). This is because the Bayesian decision strategy ultimately amounts to a specific deterministic mapping from neural activity r to an action a: a = F(r). For example, in simple two-alternative detection or discrimination tasks, this mapping reduces to a comparison of a point estimate of the stimulus with a fixed criterion, without any need to represent uncertainty (Green & Swets 1966). Indeed, observers in signal detection theory models do not commonly utilize an internal measure of uncertainty. However, in most tasks of realistic complexity, optimal performance does require keeping track of entire belief distributions, and this fact is typically used to connect optimal performance to probabilistic computation. Conversely, using belief distributions or uncertainty does not guarantee optimality. The behavior of an agent with incorrect beliefs q(R | o, C), q(o | a), q(r | C), or q(C) or of an agent who does not compute Equation 1 correctly may be suboptimal even though the agent performs probabilistic computation. Here, our focus is not on characterizing the extent to which observers behave optimally, but rather on the ways in which belief distributions or uncertainty information may be used to guide behavior.

BEHAVIORAL EVIDENCE FOR PROBABILISTIC COMPUTATION

A typical behavioral experiment specifies the task contingencies (distributions over C, s, and R) and analyzes behavioral outcomes (o) to infer the mapping from r to a. To know whether the brain uses uncertainty information, tasks have been designed in which maximization of accuracy or reward requires the observer to take trial-to-trial variations in uncertainty into account. For example, we can vary the likelihood over C by changing the stimulus s or the reliability of the information rprovides about s, the prior over C by changing the statistics of C in the world, or reward beliefs by changing the probabilities of reward.

We first consider a study in which trial-to-trial feedback was provided and sensory reliability was not varied. When trial-to-trial feedback is provided, it may be possible for the brain to gradually learn to implement an optimal stimulus-to-response strategy without performing true probabilistic computation, i.e., without using belief distributions on every trial. Yang & Shadlen (2007) trained monkeys to choose between two options using evidence provided by four visual shapes. Each shape was associated with a log likelihood ratio (log LR), and the cumulative log LR specified the odds with which the option would be rewarded. They found that monkeys combined the evidence from individual shapes and made decisions based on the cumulative log LR. Although it is possible that the animals used uncertainty information to perform the task, it is also possible that the animals used the extensive training period to establish a stimulus-response mapping that assigned a suitable weight to each shape for or against the two options.

One line of evidence in support of probabilistic computation and against stimulus-response mapping strategy comes from studies in which the reliability of a stimulus was varied on a

200

trial-by-trial basis. These studies have shown that subjects adjust their behavior based on the reliabilities of stimuli even when the point estimates of the stimuli remain unchanged. For example, in a set of cue combination experiments, Angelaki and colleagues showed that monkeys optimally combine an optic flow (visual) cue of varying reliability with a self-motion (vestibular) cue to make judgments about heading direction (Gu et al. 2008, Morgan et al. 2008, Fetsch et al. 2009), suggesting that the animal's decisions are based on a trial-to-trial estimate of stimulus reliability. Similarly, subjects take stimulus reliability into account in gaze direction perception (Mareschal et al. 2013), coincidence detection (Miyazaki et al. 2005), time interval reproduction (Jazayeri & Shadlen 2010, Acerbi et al. 2012), speeded reaching movements (Tassinari et al. 2006, Landy et al. 2012), and dynamic sensorimotor tasks (Faisal & Wolpert 2009, Turnham et al. 2011, O'Reilly et al. 2013), as well as when the number of reliable stimuli (Van den Berg et al. 2012) or the reward contingencies (Feng et al. 2009, Kiani & Shadlen 2009) are varied.

Another line of evidence in support of the probabilistic computation comes from studies in which subjects were shown to take sensory uncertainty into account even when trial-to-trial feedback was randomized or completely withheld. One study had human subjects perform a discrimination task with asymmetric rewards but without feedback until after many trials (Whiteley & Sahani 2008). Subjects approximately maximized reward, suggesting that they used trial-to-trial sensory uncertainty information. Analogous conclusions were reached in a comparable auditory task (Maiworm et al. 2011), a spatial reasoning task in natural scenes (D'Antona et al. 2013), an auditory-visual causal inference task (Kording et al. 2007), and speech perception studies (Ma et al. 2009, Bejjanki et al. 2011). A Bayesian integration model of sound localization by the barn owl (Fischer & Pena 2011) was based on data obtained using random rewards (Hausmann et al. 2009). These results provide stronger evidence for the hypothesis that the brain computes with sensory uncertainty.

Some studies combined both solutions: They varied the reliability of the sensory information from trial to trial and withheld trial-to-trial feedback. This approach has been used to study cue combination (reviewed in Trommershauser et al. 2011) as well as the integration of sensory inputs with a prior distribution in domains as diverse as speed perception (Stocker & Simoncelli 2006), orientation perception (Girshick et al. 2011), duration estimation (Ahrens & Sahani 2011, Cicchini et al. 2012), and reaching movements (Kording & Wolpert 2004, Battaglia & Schrater 2007). An additional reason for withholding feedback in some of these studies was to make subjects use prior distributions derived from natural statistics rather than from experimental statistics. The strategy of withholding feedback and varying reliability has been extended to multiple-item categorical tasks, namely visual search (Ma et al. 2011), change detection (Keshvari et al. 2012), oddity detection (Hillis et al. 2002; as modeled by Hospedales & Vijayakumar 2009), and simultaneity judgment (Magnotti et al. 2013). In an orientation categorization task, it was shown that observers used stimulus reliability to adjust category boundaries from trial to trial in a near-optimal manner (Qamar et al. 2013). Finally, in a novel design, subjects selected one of two images on which they wanted to perform an orientation identification task (Barthelme & Mamassian 2009). Subjects chose the most informative image most of the time, and their performance was higher than when they were not allowed to choose. Together, these studies provide converging evidence that in many behavioral settings the brain has trial-to-trial access to sensory uncertainty information.

THE FLEXIBILITY OF PROBABILISTIC COMPUTATION

The concern about feedback highlights a more general question regarding probabilistic computation: To what extent do subjects utilize various types of belief distributions flexibly across tasks, modalities, and motor effectors? The degree of flexibility has implications for how the brain



Figure 2

Three possible levels of flexibility of probabilistic computation. The boxes marked F indicate the mapping from neural activity r to action a, as implemented in the network architecture. Feeding into these is the information extracted from r that is used by the brain on a trial-to-trial basis. Information that is not available or used has to be incorporated through the parameters of the mapping function F (shown inside the box) that cannot be modified flexibly. By "point estimates" we refer to point estimates of stimulus features of interest (but not of stimulus features that correlate with sensory uncertainty), of outcomes, and of rewards.

may implement probabilistic computation in neural circuits. At the most flexible level (Figure 2, top row), the sensory likelihoods (say, over s and over C), prior, outcome belief distribution, and reward belief distribution (or cost function) are all independently accessible, presumably each through the activity of a population of neurons, and are used on each trial through Equation 1. In this case, the mapping function F from r to a is realized through a transformation that does not depend on beliefs and uncertainties. This strategy requires a complicated neural machinery but has the advantage of being able to instantly accommodate any change in sensory and motor variability and to integrate any known priors and cost functions. At the other extreme (Figure 2, *bottom row*), on each trial the subject has access only to point estimates of the sensory stimuli. As such, optimal behavior is possible only if the width of the likelihood function, prior, outcome beliefs, and reward beliefs are learned as parameters of the mapping function F. Therefore, any change in any of these beliefs requires a modification of the mapping function. When feedback is provided, the subject can learn F through trial and error (Law & Gold 2008), but this may take many trials. Withholding feedback makes it easier to assess the flexibility of probabilistic computation, although, in simple tasks, belief distributions can also be learned in an unsupervised fashion (Raphan & Simoncelli 2011). In principle, an observer's behavior may show intermediate levels of flexibility where some-but not all-belief distributions are readily accessible. For example, the brain may be able to flexibly accommodate changes in sensory or motor uncertainty, but require continuous learning to accommodate priors and reward beliefs (Figure 2, middle row).

One way to measure the flexibility of probabilistic computations is to test whether previously learned belief distributions would transfer to other conditions (Seydell et al. 2008, Maloney & Mamassian 2009). The most convincing case for probabilistic computation could be made if subjects utilize belief distributions flexibly across tasks, environmental statistics, sensory modalities, motor effectors, and cost functions. Varying sensory uncertainty from trial to trial, as described above, addresses one aspect of this flexibility (**Figure 2**, *middle row*), but the question of flexibility with respect to other belief distributions (**Figure 2**, *top row*) remains unanswered. Two studies that had subjects take into account prior information that varied from trial to trial found probabilistic but suboptimal behavior (Hudson et al. 2007, Acerbi et al. 2014). One human fMRI study found that changes in the likelihood and prior can be attributed to different brain areas (Vilares et al. 2012), but these signals could also be correlates of other factors such as attention or task difficulty.

Trommershäuser and colleagues conducted a series of studies (reviewed in Trommershauser et al. 2008) in which human subjects made speeded hand movements to a green disc on a screen while trying to avoid a red disc. The cost function was varied by changing the point value of a disc or the degree of overlap. In these experiments, subjects obtained near-maximal reward given their level of motor noise. The investigators found no evidence of learning, suggesting that subjects instantly incorporated previously acquired implicit knowledge of their motor uncertainty. Seydell et al. (2008) tested transfer in a similar task by comparing performance under a particular cost function between observers trained on the same cost function and observers trained on a different cost function; no difference was found, suggesting that outcome uncertainty was encoded and used. Fleming et al. (2013) asked whether humans could optimally combine sensory uncertainty information with knowledge of motor uncertainty and a cost function that changed from trial to trial. Subjects overweighted sensory uncertainty relative to their motor uncertainty, suggesting that sensory beliefs were not properly taken into account.

NEURAL REALIZATIONS OF PROBABILISTIC COMPUTATION

If organisms use information about sensory uncertainty in decision making on a trial-by-trial basis, then the question arises as to how sensory uncertainty is represented in neural populations. Theoretical schemes proposed to answer this question have been reviewed elsewhere (Jazaveri 2008, Ma et al. 2008, Vilares & Kording 2011, Pouget et al. 2013), and we only briefly summarize them here. In probabilistic population codes (Foldiak 1993, Sanger 1996, Pouget et al. 2003, Jazaveri & Movshon 2006, Ma et al. 2006), the brain has knowledge of the generative process p(r|s) and thus automatically possesses a likelihood function over s when given a pattern of activity r. Uncertainty is implicitly represented in the population activity; for example, more uncertainty may correspond to a lower total spike count. Neural operations performed on input populations correspond to manipulations of the corresponding likelihood functions. To give this framework predictive power, one needs to assume a specific form of neural variability, p(r|s). One family of distributions that has been proposed is "Poisson-like" (Ma et al. 2006). Poisson-like variability is well-established at the level of single neurons (Dean 1981, Tolhurst et al. 1983, Britten et al. 1993, Softky & Koch 1993), but recent developments have also found evidence in support of this functional description at the population level (Graf et al. 2011, Berens et al. 2012). Under the Poisson-like assumption, the logarithm of the likelihood function is linear in population activity (Jazayeri & Movshon 2006, Ma et al. 2006). As a corollary, the log LR in a two-alternative discrimination task can be straightforwardly written as a linear combination of neural activity (Jazayeri & Movshon 2006, Ma et al. 2006) or, in the reduced case of two neurons, as a function of the difference between those two neurons (Beck et al. 2008). Assuming Poisson-like probabilistic

population codes, Fetsch et al. (2012) obtained likelihoods from neural activity in cortical area MSTd, where visual and vestibular cues for self-motion are integrated, and accurately predicted the monkey's cue integration behavior, including a moderate deviation from optimality.

Another class of models is based on the assumption that the activity of a neuron at a given point in time is a sample from the belief distribution that is to be represented (Hoyer & Hyvarinen 2003, Paulin 2005, Fiser et al. 2010, Shi et al. 2010, Griffiths et al. 2012). In these so-called sampling codes, the entire probability distribution and the corresponding uncertainty are represented across time or across a neural population. These schemes hold that the probability of a variable of interest is directly mapped onto firing rate. In support of this proposal, the spontaneous activity in the ferret visual cortex may reflect the statistics of the environment (Berkes et al. 2011). Moreover, there is some behavioral evidence that observers sample from the posterior distribution (Moreno-Bote et al. 2011, Gershman et al. 2012). Sampling codes of this kind have not been thoroughly formalized, but certain versions of its formulation may be implausible. For example, if singleneuron firing rate were a sample of a sensory belief distribution, then firing rate variability should increase with uncertainty, which is inconsistent with the variance-reducing effect that decreasing contrast has on visual cortical neurons (Tolhurst et al. 1983).

Explicit probability codes compose a third class of neural codes for uncertainty (Barlow & Levick 1969, Anderson 1994, Anastasio et al. 2000, Barber et al. 2003, Lee & Mumford 2003, Rao 2004, Deneve 2008). In this class, the activity of a neuron tuned to a stimulus feature is monotonically related to the probability density of that feature (typically through a linear or logarithmic transformation). Higher uncertainty is then represented by a wider activation pattern across the population.

More physiological evidence is needed to distinguish between these schemes. However, any scheme must address how basic Bayesian computations can be implemented using biologically plausible neural operations. This is still a work in progress. An important computation is to combine a likelihood function with a prior. Several recent studies have proposed that priors over sensory variables are encoded through the organization and distribution of tuning curves in the sensory representation. For example, more neurons may be dedicated to stimuli that have higher probability (Fischer & Pena 2011, Ganguli & Simoncelli 2011, Girshick et al. 2011). If the density of neurons encodes the prior and the sensory tuning function is proportional to the likelihood. a simple population vector decoder can estimate the most probable stimulus (Fischer & Pena 2011). More sophisticated representations of the prior are also possible, if the decoder properly weights the sensory responses (Ganguli & Simoncelli 2011). Finally, if the logarithm of the prior is encoded by the tuning functions (Simoncelli 2009), then an optimal linear decoder could integrate the prior and likelihood information (Jazaveri & Movshon 2006). In a different view, integration of the prior and likelihood information is mediated by interactions between spontaneous activity and stimulus-evoked sensory responses (Berkes et al. 2011). Priors may also be encoded by neurons downstream of sensory representations (Basso & Wurtz 1997, Platt & Glimcher 1999, Janssen & Shadlen 2005, Churchland et al. 2008). In this scenario, the prior could exert its effect on behavior either through feedback mechanisms akin to attentional modulation (Ghose & Maunsell 2002) or through linear operations that exploit Poisson-like sensory variability (Ma et al. 2006).

Computation can also consist of combining multiple likelihoods over the same variable, as in cue combination, or of transforming a likelihood over one variable (say, *s* in **Figure 1**) into a likelihood over another variable (say, *C* in **Figure 1**), as in categorization tasks. Most work on the implementation of these computations has been done within a Poisson-like probabilistic population code framework; herein, cue combination and evidence accumulation are implemented through linear operations on neural activity (Jazayeri & Movshon 2006, Ma et al. 2006, Beck et al. 2008), but Kalman filtering (as used in motor control and visual tracking) and many forms of

categorization require quadratic operations and divisive normalization (Beck et al. 2011, Ma et al. 2011, Qamar et al. 2013). Virtually no work has been done on how mid-level and high-level visual computations, such as inferring Gestalt or obtaining viewpoint invariance, are performed probabilistically by a neural network. Finally, any probabilistic representation is bound to change at some point downstream, because the organism must make an estimate or execute an action. Depending on task demands, training regimen, and species, this process may take place in different brain areas, including the premotor and association areas in cortex (Jazayeri & Movshon 2006), in subcortical areas (Beck et al. 2008), or at the level of the muscles (Simoncelli 2009).

NEURAL SIGNATURES OF UNCERTAINTY

Recent studies have begun to examine the neural representation of uncertainty in animal models. Kepecs et al. (2008) trained rats to categorize a mixture of two odors and controlled sensory uncertainty by varying the proportion of each odor and the category boundary. Recording from neurons in the orbitofrontal cortex established a neural correlate of the animal's confidence (decision certainty), which was consistent with the predictions of a computational model. The authors also used a separate task to ensure that rats were able to use the uncertainty information: When given the option of initiating a new trial instead of waiting for a reward, rats resorted to this option more often when uncertainty was high. However, because orbitofrontal activity was not recorded during the task in which animals had to use uncertainty information, whether the recorded orbitofrontal signals contributed to the animal's measure of confidence remains unclear. In another study, Kiani & Shadlen (2009) trained monkeys to perform a motion discrimination task in which the monkeys could opt out of the decision on some trials by choosing a small but certain reward. The animals did so when sensory evidence was weak, and firing rates of LIP neurons on those trials were in between the activity levels of when the monkey chose either category. These findings confirm earlier work that LIP responses in this task vary monotonically with the log-LR (Gold & Shadlen 2007). Whether animals use this signal to make their trial-by-trial confidence judgments remains undetermined. In another study, monkeys used saccades to make a bet on whether the decision they made on a prior visual search task was correct (Middlebrooks & Sommer 2012). The animal's confidence, which was inferred from the magnitude of the bet, was reflected in a simultaneous recording of neural activity in the supplementary eye field. Most recently, Komura et al. (2013) trained monkeys in a direction discrimination task in which the animal either could discriminate the direction of a cloud of moving dots and receive a large reward for correct judgments or could opt out of the decision task altogether to receive a smaller reward. They found that neural activity in the pulvinar nucleus of the thalamus decreased with sensory uncertainty. Moreover, reversible inactivation of the region of interest in the dorsal pulvinar increased the proportion of opt-out responses, suggesting that signals in this region of the pulvinar contribute to the animal's assessment of uncertainty.

In reward-based decision making, recent neural studies have begun to examine the neural correlates of expected reward (mean of R given a) and risk (variance of R given a). A thorough discussion of this literature can be found elsewhere (Rangel et al. 2008, Rushworth & Behrens 2008, Lee et al. 2012); we highlight only two studies. Preuschoff et al. (2006) asked subjects to bet which of two randomly drawn integers m and n between 1 and 10 was greater. After placing the bet, subjects were informed of the value of m. This manipulation dissociated expected reward from risk: Expected reward was monotonically related to m, whereas risk was highest at intermediate values of m. Using fMRI in humans, this study and similar ones found correlates of risk in the striatum, insula, and lateral orbitofrontal cortex. Behrens et al. (2007) used a task in which humans chose between two options and gained information about the probability of each leading to reward

by observing either a stable or a volatile history of rewards. The authors found activation in the anterior cingulate cortex correlated with the subject's estimate of volatility.

Animal studies have also investigated the topic of expected reward and reward uncertainty (Schultz 2000, Hikosaka et al. 2008, Rangel et al. 2008). Early studies found a crucial role for dopamine in reward probability and risk (Fiorillo et al. 2003). Others have examined the ways in which the probability or utility of reward influences neural activity and choice behavior (Platt & Glimcher 1999, Sugrue et al. 2004, Lau & Glimcher 2008, So & Stuphorn 2010, Chen et al. 2013) as well as how trial-to-trial variations of expected reward control behavior in stochastic environments (Barraclough et al. 2004, Dorris & Glimcher 2004).

Taken together, these studies provide evidence for widespread representation of different kinds of uncertainty and open the door to several novel lines of inquiry. For example, what are the differences and similarities between the neural codes for the three different types of beliefs? What are the algorithms and neural mechanisms by which neurons measure or represent uncertainty? How do the computational principles for estimating confidence generalize across brain areas and behavioral contexts? Do the observed neural correlates of confidence causally contribute to an animal's ability to assess sensory uncertainty and decision confidence?

FUTURE DIRECTIONS

The study of probabilistic computation by the brain is still in its early stages, and many open questions remain. At the behavioral level, experiments must establish the extent to which belief distributions are encoded and utilized in a flexible manner, as schematized in **Figure 2**. One strategy is to assess how organisms act in the face of multiple types of uncertainty, as was done in a recent study by Fleming et al. (2013). Such tasks are also ecologically relevant: As the example of spoiled milk illustrates, natural behavior can often improve by combining knowledge about multiple types of belief. It is often fruitful to study sensory, outcome, or reward uncertainty in isolation, but we believe progress can be made by combining established paradigms from these individual domains. In any experimental work, particular care should be taken to address the impact of feedback on claims of probabilistic computation.

Recent work has explored the possibility that neural computation is probabilistic but suboptimal (Beck et al. 2012, Whiteley & Sahani 2012, Orhan et al. 2014, Acerbi et al. 2014). Suboptimality can arise from at least two sources: (*a*) wrong or incompletely learned beliefs and (*b*) neural networks implementing approximate computations. It will be important to develop behavioral paradigms to tease apart and characterize the factors that contribute to suboptimal performance. This line of work may shed light on disorders of higher brain function that are associated with faulty probabilistic computations.

A direction for physiological investigation is to move beyond correlation and establish a causal link between behavior and neural activity associated with belief and uncertainty. Also important are understanding the learning mechanisms at multiple levels of analysis—from synapses to circuits that enable the brain to encode beliefs as well as assessing if different types of belief are used and represented differently. Mechanisms of integration are also poorly understood: For example, we do not know how neurons combine various types of belief. Finally, it will be valuable to design experiments to test the theoretical schemes for the neural representation of probability.

On the theoretical side, a challenge is to propose a flexible and general framework for the neural implementation of Equation 1. Theoretical schemes have focused predominantly on relatively simple sensory problems, for example, the probabilistic representation of a single, one-dimensional stimulus feature. Future studies should expand the reach of probabilistic models to more complex and naturalistic stimuli. They should address how a particular representational scheme can be

used to perform complex computations, for example, categorization in a high-dimensional space or realization of the Gestalt "principles" of perception. An intriguing idea is that in dynamic natural scenes, such as when trying to predict whether a stack of blocks will topple, people build beliefs over possible futures by mentally simulating the dynamics, in this example, the laws of physics (Battaglia et al. 2013). The application of probabilistic models to the problem of object recognition (Kersten et al. 2004), which traditional models seek to explain without reference to beliefs (DiCarlo et al. 2012), needs to be explored. A higher-order form of inference that has barely been studied at the neural level is structure learning (Tenenbaum & Griffiths 2001, Kemp & Tenenbaum 2008, Braun et al. 2010, Pouget et al. 2013), the process of learning generalizable rules for categorizing stimuli or performing actions. Study of structure learning, generalization, and model selection could help to bridge the gap between simple psychophysical tasks and more cognitive domains. Finally, capacity limitations in the encoding stage have largely been ignored in probabilistic models of decision making but deserve attention (Palmer et al. 1990, Keshvari et al. 2013, Mazyar et al. 2013).

The idea that the brain computes with belief distributions has already had a profound impact on neuroscience, psychology, and cognitive science. In the coming years, we foresee greater convergence among these fields and progress in applying the concepts of belief and uncertainty to understand the computational and neural underpinnings of more complex and more natural behaviors.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

LITERATURE CITED

- Acerbi L, Vijayakumar S, Wolpert DM. 2014. On the origins of suboptimality in human probabilistic inference. PLOS Comp. Biol. In press
- Acerbi L, Wolpert DM, Vijayakumar S. 2012. Internal representations of temporal statistics and feedback calibrate motor-sensory interval timing. PLoS Comput. Biol. 8(11):e1002771
- Ahrens M, Sahani M. 2011. Observers exploit stochastic models of sensory change to help judge the passage of time. Curr. Biol. 21:1–7
- Anastasio TJ, Patton PE, Belkacem-Boussaid K. 2000. Using Bayes' rule to model multisensory enhancement in the superior colliculus. *Neural Comput.* 12(5):1165–87
- Anderson C. 1994. Neurobiological computational systems. In Computational Intelligence Imitating Life, pp. 213–22. New York: IEEE Press
- Barber MJ, Clark JW, Anderson CH. 2003. Neural representation of probabilistic information. Neural Comput. 15(8):1843–64
- Barlow HB, Levick WR. 1969. Three factors limiting the reliable detection of light by retinal ganglion cells of the cat. J. Physiol. 200:1–24
- Barraclough DJ, Conroy ML, et al. 2004. Prefrontal cortex and decision making in a mixed-strategy game. Nat. Neurosci. 7:404–10
- Barthelme S, Mamassian P. 2009. Evaluation of objective uncertainty in the visual system. PLoS Comput. Biol. 5(9):e1000504
- Basso MA, Wurtz RH. 1997. Modulation of neuronal activity by target uncertainty. Nature 389:66-69
- Battaglia PW, Hamrick JB, Tenenbaum JB. 2013. Simulation as an engine of physical scene understanding. Proc. Natl. Acad. Sci. USA 110(45):18327–32
- Battaglia PW, Schrater PR. 2007. Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. J. Neurosci. 27(26):6984–94

- Beck J, Ma WJ, Kiani R, Hanks T, Churchland AK, et al. 2008. Bayesian decision making with probabilistic population codes. *Neuron* 60(6):1142–52
- Beck JM, Latham PE, Pouget A. 2011. Marginalization in neural circuits with divisive normalization. J. Neurosci. 31(43):15310–19
- Beck JM, Ma WJ, Pitkow X, Latham PE, Pouget A. 2012. Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron* 74(1):30–39
- Behrens TEJ, Woolrich MW, Walton ME, Rushworth MFS. 2007. Learning the value of information in an uncertain world. Nat. Neurosci. 10(9):1214–21
- Bejjanki VR, Clayards M, Knill DC, Aslin RN. 2011. Cue integration in categorical tasks: insights from audio-visual speech perception. PLoS ONE 6(5):e19812
- Berens P, Ecker AS, Cotton RJ, Ma WJ, Bethge M, Tolias AS. 2012. A fast and simple population code for orientation in primate V1. J. Neurosci. 32(31):10618–26
- Berkes P, Orban G, Lengyel M, Fiser J. 2011. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331(6013):83–87
- Braun DA, Mehring C, Wolpert DM. 2010. Structure learning in action. Behav. Brain Res. 206(2):157-65
- Britten KH, Shadlen MN, Newsome WT, Movshon JA. 1993. Responses of neurons in macaque MT to stochastic motion signals. Vis. Neurosci. 10(6):1157–69
- Chen X, Mihalas S, Niebur E, Stuphorn V. 2013. Mechanisms underlying the influence of saliency on valuebased decisions. *J. Vis.* 13(12):18
- Churchland AK, Kiani R, Shadlen MN. 2008. Decision-making with multiple alternatives. *Nat. Neurosci.* 11(6):693–702
- Cicchini GM, Arrighi R, Cecchetti L, Giusti M, Burr DC. 2012. Optimal encoding of interval timing in expert percussionists. J. Neurosci. 32(3):1056–60
- D'Antona AD, Perry JS, Geisler WS. 2013. Humans make efficient use of natural image statistics when performing spatial interpolation. *J. Vis.* 13(14):11
- de Martino B, Fleming SM, Garrett N, Dolan RJ. 2013. Confidence in value-based choice. *Nat. Neurosci.* 16:105–10
- Dean AF. 1981. The variability of discharge of simple cells in the cat striate cortex. Exp. Brain Res. 44:437-40
- Deneve S. 2008. Bayesian spiking neurons I: inference. Neural Comput. 20(1):91-117
- DiCarlo JJ, Zoccolan D, Rust NC. 2012. How does the brain solve visual object recognition? *Neuron* 73(3):415–34
- Dorris MC, Glimcher PW. 2004. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron* 44(2):365–78
- Faisal A, Selen LPJ, Wolpert DM. 2008. Noise in the nervous system. Nat. Rev. Neurosci. 9(4):292-303
- Faisal AA, Wolpert DM. 2009. Near optimal combination of sensory and motor uncertainty in time during a naturalistic perception-action task. J. Neurophysiol. 101(4):1901–12
- Feng S, Holmes P, Rorie A, Newsome WT. 2009. Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Comput. Biol.* 5(2):e1000284
- Fetsch CR, Pouget A, DeAngelis GC, Angelaki DE. 2012. Neural correlates of reliability-based cue weighting during multisensory integration. Nat. Neurosci. 15:146–54
- Fetsch CR, Turner AH, DeAngelis GC, Angelaki DE. 2009. Dynamic reweighting of visual and vestibular cues during self-motion perception. J. Neurosci. 29(49):15601–12
- Fiorillo CD, Tobler PN, Schultz W. 2003. Discrete coding of reward probability and uncertainty by dopamine neurons. Science 299(5614):1898–902
- Fischer BJ, Pena JL. 2011. Owl's behavior and neural representation predicted by Bayesian inference. Nat. Neurosci. 14(8):1061–66
- Fiser J, Berkes P, Orbán G, Lengyel M. 2010. Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn. Sci.* 14(3):119–30
- Fleming SM, Dolan RJ. 2012. The neural basis of metacognitive ability. *Phil. Trans. R. Soc. B* 367(1594):1338–49
- Fleming SM, Maloney LT, Daw ND. 2013. The irrationality of categorical perception. J. Neurosci. 33(49):19060–70

- Foldiak P. 1993. The 'ideal homunculus': statistical inference from neural population responses. In Computation and Neural Systems, ed. F Eeckman, J Bower, pp. 55–60. Norwell, MA: Kluwer Acad.
- Ganguli D, Simoncelli EP. 2011. Implicit encoding of prior probabilities in optimal neural populations. In Advances in Neural Information Processing Systems, ed. J Shawe-Taylor, RS Zemel, P Bartlett, F Pereira, KW Weinberger, pp. 658–66. Cambridge, MA: MIT Press
- Gershman SJ, Vul E, Tenenbaum JB. 2012. Multistability and perceptual inference. Neural Comput. 24:1-24
- Ghose GM, Maunsell JH. 2002. Attentional modulation in visual cortex depends on task timing. Nature 419(6907):616–20
- Girshick AR, Landy MS, Simoncelli EP. 2011. Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nat. Neurosci.* 14:926–32
- Glimcher PW, Fehr E, Camerer C, Poldrack RA, eds. 2008. Neuroeconomics: Decision Making and the Brain. New York: Academic
- Gold JI, Shadlen MN. 2007. The neural basis of decision making. Annu. Rev. Neurosci. 30:535-74
- Graf AB, Kohn A, Jazayeri M, Movshon JA. 2011. Decoding the activity of neuronal populations in macaque primary visual cortex. Nat. Neurosci. 14(2):239–45
- Green DM, Swets JA. 1966. Signal Detection Theory and Psychophysics. Los Altos, CA: Wiley
- Griffiths TL, Vul E, Sanborn AN. 2012. Bridging levels of analysis for probabilistic models of cognition. Curr. Dir. Psychol. Sci. 21(4):263–68
- Gu Y, Angelaki DE, DeAngelis GC. 2008. Neural correlates of multisensory cue integration in macaque MSTd. Nat. Neurosci. 11(10):1201–10
- Hausmann L, von Campenhausen M, Endler F, Singheiser M, Wagner H. 2009. Improvements of sound localization abilities by the facial ruff of the barn owl (*Tyto alba*) as demonstrated by virtual ruff removal. *PLoS ONE* 4:e7721
- Hikosaka O, Bromberg-Martin E, Hong S, Matsumoto M. 2008. New insights on the subcortical representation of reward. *Curr. Opin. Neurobiol.* 18(2):203–8
- Hillis JM, Ernst MO, Banks MS, Landy MS. 2002. Combining sensory information: mandatory fusion within, but not between, senses. Science 298(5598):1627–30
- Hospedales T, Vijayakumar S. 2009. Multisensory oddity detection as Bayesian inference. *PLoS ONE* 4(1):e4205
- Hoyer PO, Hyvärinen A. 2003. Interpreting Neural Response Variability as Monte Carlo Sampling of the Posterior. Cambridge, MA: MIT Press
- Hudson TE, Maloney LT, Landy MS. 2007. Movement planning with probabilistic target information. *J. Neurophysiol.* 98(5):3034–46
- Janssen P, Shadlen MN. 2005. A representation of the hazard rate of elapsed time in macaque area LIP. Nat. Neurosci. 8(2):234–41
- Jazayeri M. 2008. Probabilistic sensory recoding. Curr. Opin. Neurobiol. 18(4):431-37
- Jazayeri M, Movshon JA. 2006. Optimal representation of sensory information by neural populations. *Nat. Neurosci.* 9(5):690–96
- Jazayeri M, Shadlen MN. 2010. Temporal context calibrates interval timing. Nat. Neurosci. 13(8):1020-26
- Kahneman D, Tversky A. 1979. Prospect theory: an analysis of decision under risk. Econometrica 47:263–91
- Kemp C, Tenenbaum JB. 2008. The discovery of structural form. Proc. Natl. Acad. Sci. USA 105(31):10687-92
- Kepecs A, Uchida N, Zariwala HA, Mainen ZF. 2008. Neural correlates, computation and behavioural impact of decision confidence. *Nature* 455:227–33
- Kersten D, Mamassian P, Yuille A. 2004. Object perception as Bayesian inference. Annu. Rev. Psychol. 55:271– 304
- Keshvari S, Van den Berg R, Ma WJ. 2012. Probabilistic computation in human perception under variability in encoding precision. PLoS ONE 7(6):e40216
- Kiani R, Shadlen MN. 2009. Representation of confidence associated with a decision by neurons in the parietal cortex. Science 324:759–64
- Komura Y, Nikkuni A, Hirashima N, Uetake T, Miyamoto A. 2013. Responses of pulvinar neurons reflect a subject's confidence in visual categorization. *Nat. Neurosci.* 16(6):749–55
- Körding KP, Beierholm U, Ma WJ, Quartz S, Tenenbaum JB, Shams L. 2007. Causal inference in multisensory perception. PLoS ONE 2(9):e943

Kording KP, Wolpert DM. 2004. Bayesian integration in sensorimotor learning. Nature 427(6971):244-47

- Landy MS, Trommershauser J, Daw ND. 2012. Dynamic estimation of task-relevant variance in movement under risk. 7. Neurosci. 32(37):12702–11
- Lau B, Glimcher PW. 2008. Value representations in the primate striatum during matching behavior. *Neuron* 58(3):451–63
- Law CT, Gold JI. 2008. Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat. Neurosci. 11(4):505–13
- Lee D, Seo H, Jung MW. 2012. Neural basis of reinforcement learning and decision making. Annu. Rev. Neurosci. 35:287–308
- Lee TS, Mumford D. 2003. Hierarchical Bayesian inference in the visual cortex. J. Opt. Soc. Am. 20(7):1434-48
- Ma WJ. 2012. Organizing probabilistic models of perception. Trends Cogn. Sci. 16(10):511-18
- Ma WJ, Beck JM, Latham PE, Pouget A. 2006. Bayesian inference with probabilistic population codes. Nat. Neurosci. 9(11):1432–38
- Ma WJ, Beck JM, Pouget A. 2008. Spiking networks for Bayesian inference and choice. Curr. Opin. Neurobiol. 18:217–22
- Ma WJ, Navalpakkam V, Beck JM, van den Berg R, Pouget A. 2011. Behavior and neural basis of near-optimal visual search. Nat. Neurosci. 14(6):783–90
- Ma WJ, Zhou X, Ross LA, Foxe JJ, Parra LC. 2009. Lip-reading aids word recognition most in moderate noise: a Bayesian explanation using high-dimensional feature space. PLoS ONE 4(3):e4638
- Magnotti JF, Ma WJ, Beauchamp MS. 2013. Causal inference of asynchronous audiovisual speech. Front. Psychol. 4:798
- Maiworm M, König R, Röder B. 2011. Integrative processing of perception and reward in an auditory localization paradigm. *Exp. Psychol.* 58(3):217–26
- Maloney LT, Mamassian P. 2009. Bayesian decision theory as a model of human visual perception: testing Bayesian transfer. Vis. Neurosci. 26(1):147–55
- Mareschal I, Calder AJ, Clifford CWG. 2013. Humans have an expectation that gaze is directed toward them. *Curr. Bio.* 23(8):717–21
- Mazyar H, van den Berg R, Seilheimer RL, Ma WJ. 2013. Independence is elusive: set size effects on encoding precision in visual search. 7. Vis. 13(5). doi: 10.1167/13.5.8
- Middlebrooks PG, Sommer MA. 2012. Neuronal correlates of metacognition in primate frontal cortex. Neuron 75(3):517–30
- Miyazaki M, Nozaki D, Nakajima Y. 2005. Testing Bayesian models of human coincidence timing. J. Neurophysiol. 94(1):395–99
- Moreno-Bote R, Knill DC, Pouget A. 2011. Bayesian sampling in visual perception. Proc. Natl. Acad. Sci. USA 108(30):12491–96
- Morgan ML, DeAngelis GC, Angelaki DE. 2008. Multisensory integration in macaque visual cortex depends on cue reliability. *Neuron* 59:662–73
- Najemnik J, Geisler WS. 2005. Optimal eye movement strategies in visual search. Nature 434(7031):387-91
- O'Reilly JX, Jbabdi S, Rushworth MFS, Behrens TEJ. 2013. Brain systems for probabilistic and dynamic prediction: computational specificity and integration. *PLoS Biol.* 11(9):e1001662
- Orhan AE, Sims CR, Jacobs RA, Knill DC. 2014. The adaptive nature of visual working memory. *Curr. Dir. Psychol. Sci.* In press
- Palmer J. 1990. Attentional limits on the perception and memory of visual information. J. Exp. Psychol. Hum. Percept. Perform. 16:332–50
- Paulin MG. 2005. Evolution of the cerebellum as a neuronal machine for Bayesian state estimation. J. Neural Eng. 2(3):S219–34
- Peirce CS, Jastrow J. 1884. On small differences of sensation. Mem. Natl. Acad. Sci. 3:73-83
- Platt ML, Glimcher PW. 1999. Neural correlates of decision variables in parietal cortex. *Nature* 400(6741):233–38
- Pouget A, Beck JM, Ma WJ, Latham PE. 2013. Probabilistic brains: knowns and unknowns. Nat. Neurosci. 16:1170–78
- Pouget A, Dayan P, Zemel RS. 2003. Inference and computation with population codes. Annu. Rev. Neurosci. 26:381–410

- Preuschoff K, Bossaerts P, Quartz SR. 2006. Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* 51(3):381–90
- Qamar AT, Cotton RJ, George RG, Beck JM, Prezhdo E, et al. 2013. Trial-to-trial, uncertainty-based adjustment of decision boundaries in visual categorization. Proc. Natl. Acad. Sci. USA 110(50):20332–37
- Rangel A, Camerer C, Montague PR. 2008. A framework for studying the neurobiology of value-based decision making. Nat. Rev. Neurosci. 9(7):545–56
- Rao RP. 2004. Bayesian computation in recurrent neural circuits. Neural Comput. 16(1):1-38
- Raphan M, Simoncelli EP. 2011. Least squares estimation without priors or supervision. Neural Comput. 23(2):374–420
- Rushworth MFS, Behrens TEJ. 2008. Choice, uncertainty and value in prefrontal and cingulate cortex. Nat. Neurosci. 11(4):389–97
- Sanger T. 1996. Probability density estimation for the interpretation of neural population codes. J. Neurophysiol. 76(4):2790–93
- Schultz W. 2000. Multiple reward signals in the brain. Nat. Rev. Neurosci. 1(3):199-207
- Seydell A, McCann BC, Trommershäuser J, Knill DC. 2008. Learning stochastic reward distributions in a speeded pointing task. J. Neurosci. 28(17):4356–67
- Shi L, Griffiths TL, Feldman NH, Sanborn AN. 2010. Exemplar models as a mechanism for performing Bayesian inference. *Psychon. Bull. Rev.* 17(4):443–64
- Simoncelli EP. 2009. Optimal estimation in sensory systems. In *The Cognitive Neurosciences*, Vol. 4, ed. M Gazzanig, pp. 525–35. Cambridge, MA: MIT Press
- So NY, Stuphorn V. 2010. Supplementary eye field encodes option and action value for saccades with variable reward. J. Neurophysiol. 104(5):2634–53
- Softky WR, Koch C. 1993. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. J. Neurosci. 13(1):334–50
- Stocker AA, Simoncelli EP. 2006. Noise characteristics and prior expectations in human visual speed perception. Nat. Neurosci. 9(4):578–85
- Sugrue LP, Corrado GS, Newsome WT. 2004. Matching behavior and the representation of value in the parietal cortex. Science 304(5678):1782–87
- Tassinari H, Hudson TE, Landy MS. 2006. Combining priors and noisy visual cues in a rapid pointing task. *J. Neurosci.* 26:10154–63
- Tenenbaum JB, Griffiths TL. 2001. Generalization, similarity, and Bayesian inference. *Behav. Brain Sci.* 24(4):629–40; discuss. 652–791
- Tolhurst D, Movshon J, Dean AF. 1983. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vis. Res.* 23:775–85
- Trommershauser J, Kording K, Landy MS, eds. 2011. Sensory Cue Integration. New York: Oxford Univ. Press
- Trommershäuser J, Maloney L, Landy MS. 2008. Decision making, movement planning and statistical decision theory. Trends Cogn. Sci. 12(8):291–97
- Turnham EJ, Braun DA, Wolpert DM. 2011. Inferring visuomotor priors for sensorimotor learning. PLoS Comput. Biol. 7(3):e1001112
- van den Berg R, Vogel M, Josić K, Ma WJ. 2012. Optimal inference of sameness. Proc. Natl. Acad. Sci. USA 109(8):3178–83
- Vilares I, Howard JD, Fernandes HL, Gottfried JA, Kording KP. 2012. Differential representations of prior and likelihood uncertainty in the human brain. *Curr. Biol.* 22(18):1641–48
- Vilares I, Kording KP. 2011. Bayesian models: the structure of the world, uncertainty, behavior, and the brain. Ann. NY Acad. Sci. 1224:22–39
- Whiteley L, Sahani M. 2008. Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. J. Vis. 8(3):1–15

Whiteley L, Sahani M. 2012. Attention in a Bayesian framework. Front. Hum. Neurosci. 6:100

Yang T, Shadlen MN. 2007. Probabilistic reasoning by neurons. Nature 447:1075-80

Contents

v

Annual Review of Neuroscience

Embodied Cognition and Mirror Neurons: A Critical Assessment Alfonso Caramazza, Stefano Anzellotti, Lukas Strnad, and Angelika Lingnau	. 1
Translational Control in Synaptic Plasticity and Cognitive Dysfunction Shelly A. Buffington, Wei Huang, and Mauro Costa-Mattioli	17
The Perirhinal Cortex Wendy A. Suzuki and Yuji Naya	39
Autophagy and Its Normal and Pathogenic States in the Brain Ai Yamamoto and Zhenyu Yue	55
Apolipoprotein E in Alzheimer's Disease: An Update Jin-Tai Yu, Lan Tan, and John Hardy	79
Function and Dysfunction of Hypocretin/Orexin: An Energetics Point of View <i>Xiao-Bing Gao and Tamas Horvath</i>	01
Reassessing Models of Basal Ganglia Function and Dysfunction Alexandra B. Nelson and Anatol C. Kreitzer	17
A Mitocentric View of Parkinson's Disease Nele A. Haelterman, Wan Hee Yoon, Hector Sandoval, Manish Jaiswal, Joshua M. Shulman, and Hugo J. Bellen1	37
Coupling Mechanism and Significance of the BOLD Signal: A Status Report <i>Elizabeth M.C. Hillman</i>	61
Cortical Control of Whisker Movement Carl C.H. Petersen	83
Neural Coding of Uncertainty and Probability Wei Ji Ma and Mehrdad Jazayeri	05
Neural Tube Defects Nicholas D.E. Greene and Andrew J. Copp	21
Functions and Dysfunctions of Adult Hippocampal Neurogenesis Kimberly M. Christian, Hongjun Song, and Guo-li Ming	43
Emotion and Decision Making: Multiple Modulatory Neural Circuits Elizabeth A. Phelps, Karolina M. Lempert, and Peter Sokol-Hessner	63

Basal Ganglia Circuits for Reward Value–Guided Behavior Okihide Hikosaka, Hyoung F. Kim, Masaharu Yasuda, and Shinya Yamamoto 289
Motion-Detecting Circuits in Flies: Coming into View Marion Silies, Daryl M. Gohl, and Thomas R. Clandinin
Neuromodulation of Circuits with Variable Parameters: Single Neurons and Small Circuits Reveal Principles of State-Dependent and Robust Neuromodulation <i>Eve Marder, Timothy O'Leary, and Sonal Shruti</i>
The Neurobiology of Language Beyond Single Words Peter Hagoort and Peter Indefrey
Coding and Transformations in the Olfactory System Naoshige Uchida, Cindy Poo, and Rafi Haddad
Chemogenetic Tools to Interrogate Brain Functions Scott M. Sternson and Bryan L. Roth
Meta-Analysis in Human Neuroimaging: Computational Modeling of Large-Scale Databases Peter T. Fox, Jack L. Lancaster, Angela R. Laird, and Simon B. Eickhoff
Decoding Neural Representational Spaces Using Multivariate Pattern Analysis James V. Haxby, Andrew C. Connolly, and J. Swaroop Guntupalli
Measuring Consciousness in Severely Damaged Brains Olivia Gosseries, Haibo Di, Steven Laureys, and Mélanie Boly
Generating Human Neurons In Vitro and Using Them to Understand Neuropsychiatric Disease Sergiu P. Paşca, Georgia Panagiotakos, and Ricardo E. Dolmetsch
Neuropeptidergic Control of Sleep and Wakefulness Constance Richter, Ian G. Woods, and Alexander F. Schier

Indexes

Cumulative Index of Contributing Authors, Volumes 28-37	533
Cumulative Index of Article Titles, Volumes 28–37	537

Errata

An online log of corrections to *Annual Review of Neuroscience* articles may be found at http://www.annualreviews.org/errata/neuro



It's about time. Your time. It's time well spent.

New From Annual Reviews:

Annual Review of Statistics and Its Application

Volume 1 • Online January 2014 • http://statistics.annualreviews.org

Editor: Stephen E. Fienberg, Carnegie Mellon University Associate Editors: Nancy Reid, University of Toronto Stephen M. Stigler, University of Chicago

The Annual Review of Statistics and Its Application aims to inform statisticians and quantitative methodologists, as well as all scientists and users of statistics about major methodological advances and the computational tools that allow for their implementation. It will include developments in the field of statistics, including theoretical statistical underpinnings of new methodology, as well as developments in specific application domains such as biostatistics and bioinformatics, economics, machine learning, psychology, sociology, and aspects of the physical sciences.

Complimentary online access to the first volume will be available until January 2015.

TABLE OF CONTENTS:

- What Is Statistics? Stephen E. Fienberg
- A Systematic Statistical Approach to Evaluating Evidence from Observational Studies, David Madigan, Paul E. Stang, Jesse A. Berlin, Martijn Schuemie, J. Marc Overhage, Marc A. Suchard, Bill Dumouchel, Abraham G. Hartzema, Patrick B. Ryan
- The Role of Statistics in the Discovery of a Higgs Boson, David A. van Dyk
- Brain Imaging Analysis, F. DuBois Bowman
- Statistics and Climate, Peter Guttorp
- Climate Simulators and Climate Projections, Jonathan Rougier, Michael Goldstein
- Probabilistic Forecasting, Tilmann Gneiting, Matthias Katzfuss
- Bayesian Computational Tools, Christian P. Robert
- Bayesian Computation Via Markov Chain Monte Carlo, Radu V. Craiu, Jeffrey S. Rosenthal
- Build, Compute, Critique, Repeat: Data Analysis with Latent Variable Models, David M. Blei
- Structured Regularizers for High-Dimensional Problems: Statistical and Computational Issues, Martin J. Wainwright

- High-Dimensional Statistics with a View Toward Applications in Biology, Peter Bühlmann, Markus Kalisch, Lukas Meier
- Next-Generation Statistical Genetics: Modeling, Penalization, and Optimization in High-Dimensional Data, Kenneth Lange, Jeanette C. Papp, Janet S. Sinsheimer, Eric M. Sobel
- Breaking Bad: Two Decades of Life-Course Data Analysis in Criminology, Developmental Psychology, and Beyond, Elena A. Erosheva, Ross L. Matsueda, Donatello Telesca
- Event History Analysis, Niels Keiding
- Statistical Evaluation of Forensic DNA Profile Evidence, Christopher D. Steele, David J. Balding
- Using League Table Rankings in Public Policy Formation: Statistical Issues, Harvey Goldstein
- Statistical Ecology, Ruth King
- Estimating the Number of Species in Microbial Diversity Studies, John Bunge, Amy Willis, Fiona Walsh
- *Dynamic Treatment Regimes,* Bibhas Chakraborty, Susan A. Murphy
- Statistics and Related Topics in Single-Molecule Biophysics, Hong Qian, S.C. Kou
- Statistics and Quantitative Risk Management for Banking and Insurance, Paul Embrechts, Marius Hofert

Access this and all other Annual Reviews journals via your institution at www.annualreviews.org.

ANNUAL REVIEWS | Connect With Our Experts

Tel: 800.523.8635 (US/CAN) | Tel: 650.493.4400 | Fax: 650.424.0910 | Email: service@annualreviews.org

