Decision-Related Activity in Sensory Neurons: Correlations Among Neurons and with Behavior

Hendrikje Nienborg,[1,*] Marlene R. Cohen,[2,*]
and Bruce G. Cumming[3]

[1]Werner Reichardt Center for Integrative Neuroscience, 72076 Tuebingen, Germany; email: hnienb@gmail.com

[2]Department of Neuroscience and Center for the Neural Basis of Cognition, University of Pittsburgh, Pittsburgh, Pennsylvania 15213; email: cohenm@pitt.edu

[3]Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Bethesda, Maryland 20892; email: bgc@lsr.nei.nih.gov

Annu. Rev. Neurosci. 2012. 35:463–83

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

This article's doi:
10.1146/annurev-neuro-062111-150403

0147-006X/12/0721-0463$20.00

*These authors contributed equally to this work.

Keywords

choice probability, decision making, noise correlation

Abstract

Neurons in early sensory cortex show weak but systematic correlations with perceptual decisions when trained animals perform at psychophysical threshold. These correlations are observed across repeated presentations of identical stimuli and cannot be explained by variation in external factors. The relationship between the activity of individual sensory neurons and the animal's behavioral choice means that even neurons in early sensory cortex carry information about an upcoming decision. This relationship, termed choice probability, may reflect the effect of fluctuations in neuronal firing rate on the animal's decision, but it can also reflect modulation of sensory responses by cognitive factors, or network properties such as variability that is shared among populations of neurons. Here, we review recent work clarifying the relationship among fluctuations in the responses of individual neurons, correlated variability, and behavior in a variety of tasks and cortical areas. We also discuss the possibility that choice probability may in part reflect the influence of cognitive factors on sensory neurons and explore the situations in which choice probability can be used to make inferences about the role of particular sensory neurons in the decision-making process.

Contents

INTRODUCTION

How the responses of individual neurons and their interactions with nearby neurons relate to perception is critical to understanding how the brain generates our mental world. Until the late 1980s, this question was addressed predominantly by studying a subject's perceptual performance and the response properties of individual sensory neurons in separate experiments (reviewed in Parker & Newsome 1998). Since then, numerous studies have simultaneously measured the activity of individual neurons and subjects' perceptual judgments.

Simultaneous neuronal recordings and behavioral measurements provide the opportunity to determine the relationship between the neuron's responses and the animal's behavioral choices. Cortical neurons respond variably to repeated presentations of the same stimulus (Tolhurst et al. 1983) (see **Figure 1**). For near-threshold perceptual judgments, observers make variable choices as well, making errors on some fraction of behavioral trials. Measuring the extent to which the fluctuations in the responses of an individual neuron predict perceptual judgments may reveal important information about the role played by single neurons in a behavioral task.

The relationship between the trial-to-trial fluctuations in the activity of individual sensory neurons and perceptual decisions is typically quantified using a measure called choice probability (CP). The idea of CP was first introduced by Britten and colleagues (1996), who recorded the activity of direction-selective neurons in the middle temporal area (MT) while monkeys performed a two-alternative forced-choice motion-direction-discrimination task (**Figure 1*a***). The authors noticed that they could predict, from the responses of single MT neurons, whether, for example, a monkey would report that a random dot stimulus containing no net motion was moving upward or downward (**Figure 1*b*,*c***).

The authors used CP to quantify the extent to which they could use the neuron's responses to a given stimulus to predict the animal's perceptual decisions. If, for example, the authors recorded from a neuron whose preferred direction was upward, they hypothesized that the animal would be more likely to report upward motion on trials in which that neuron fired more than its average. Conversely, they hypothesized that the animal would be more likely to report downward motion on trials in which the neuron fired less than its average. CP quantifies the accuracy of this prediction by representing the proportion of trials on which an ideal observer could predict the animal's choices given the firing rate of the neuron (**Figure 1*c***) (Shadlen et al. 1996). A CP of 1 would mean that the neuron always fired more on trials when the monkey reported upward motion than on trials

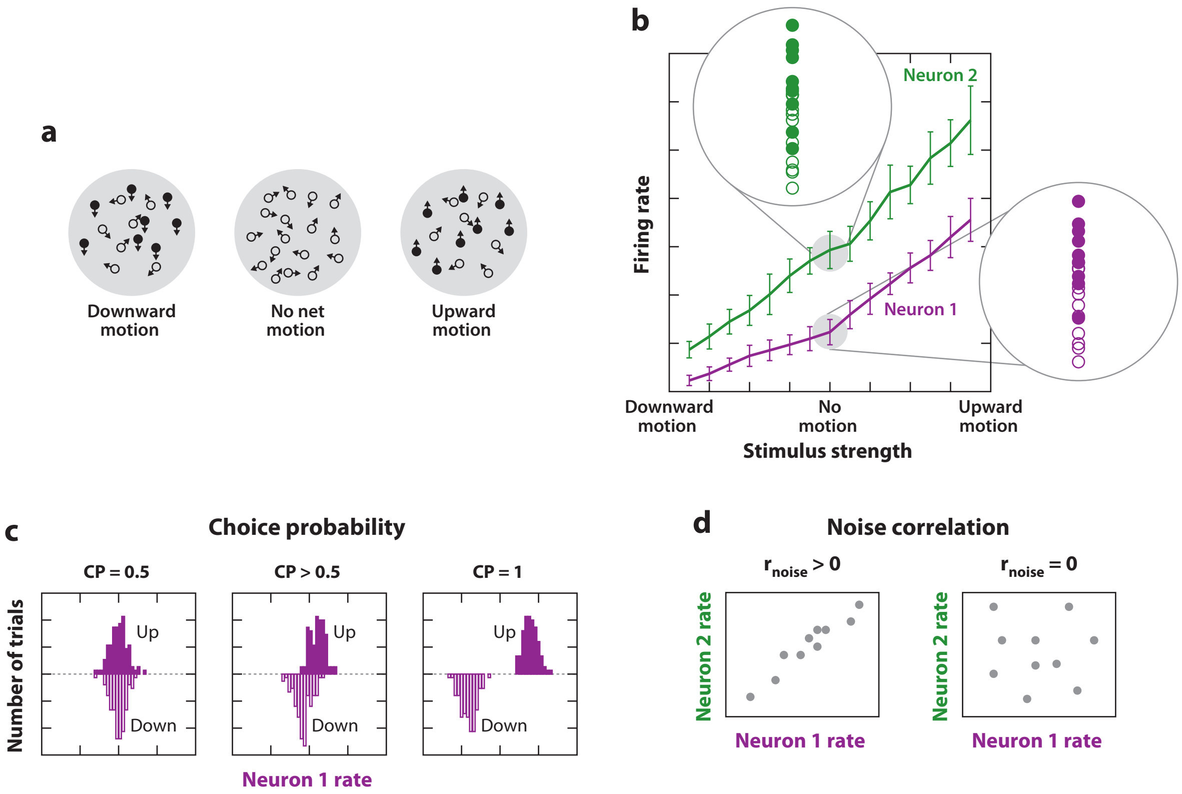

Figure 1

Measuring and calculating choice probability and noise correlation. (*a*) The original choice probability (CP) study (Britten et al. 1996) used a task in which monkeys reported whether the coherent motion in a random dot display moved in one of two opposite directions, illustrated here for an up-down discrimination. Signal strength was adjusted by varying the fraction of dots moving coherently in one direction (*filled circles*), whereas the remaining dots moved at random (*noise dots*, *open circles*). CP was usually calculated from responses to stimuli with no coherent motion in which all the dots moved randomly (*center panel*). (*b*) Average firing rate responses of two hypothetical middle temporal area (MT) neurons with upward preferred directions to random dot stimuli of various motion strengths. The circles zoom in on responses to stimuli with no net motion to illustrate that the responses of each neuron will be different on each stimulus presentation. These trial-to-trial fluctuations in neuronal response are the basis for CP (which measures the correlation between these fluctuations and behavior) and noise correlation (which measures the correlation between these fluctuations in a pair of neurons). In this illustration, filled circles represent responses on trials in which the animal reported upward motion, and open circles represent trials in which the animal reported downward motion. (*c*) Frequency histograms of responses of the hypothetical magenta neuron to stimuli with no net motion separated by whether the animal reported upward motion (*top*, *filled histograms*) or downward motion (*downward*, *open histograms*). CP measures the probability with which an ideal observer could predict the animal's choices from the responses of this neuron. CP is defined such that if the distributions of responses are identical for up and down choices (*left histograms*), CP = 0.5 because an ideal observer would be at chance, or 50% correct performance. If the histograms for the two choices did not overlap (*right histograms*), CP = 1 because the observer would be able to use this neuron to predict choices with 100% accuracy. Note that if the spike counts were lower when the animal reported the neuron's preferred direction (rates on up trials were less than on down trials in this case), CP would be less than 0.5. (*d*) Noise correlation is calculated using the same trial-to-trial fluctuations as used in CP but collapsing across choices. The correlation between two neurons is defined as the Pearson's correlation coefficient between the responses of a pair of neurons to repeated presentations of the same stimulus. This panel illustrates two potential noise correlations for the pair of hypothetical neurons in panel *b*, one in which fluctuations in the responses are correlated (*left panel*; each point represents responses on one presentation of a stimulus with no net motion) and one in which the responses are uncorrelated. In visual cortex, noise correlations tend to be weak but positive (for review, see Cohen & Kohn 2011).

when it reported downward motion, whereas a CP of 0 would indicate that the neuron always fired more when the monkey reported downward motion. A CP of 0.5 would indicate no relationship between the neuron's responses and the animal's decisions. Empirically, the authors found a mean CP of 0.56, meaning that the activity of MT neurons is weakly but consistently related to perceptual choices in this task (Britten et al. 1996).

Although CPs for sensory neurons are typically small (usually <0.6), it is remarkable that they are detectable at all, given that many thousands of cortical neurons are activated by a typical sensory stimulus. If the responses of all these thousands of neurons contribute independently to a decision, the relationship between the one that an experimenter happens to find with an electrode and the animal's behavior should be negligible. If, however, very few neurons are involved in the decision, the chances of finding one of them during a recording session should be very low.

Quantitative simulations have been invaluable for understanding how it is possible that the activity of any one neuron could correlate with behavior when so many neurons are potentially involved in the decision. The dominant theoretical framework has been a simple pooling model, developed to describe the data from the original study by Britten and colleagues (Shadlen et al. 1996). This model necessarily represents a considerable simplification of the cortical circuitry, but because of its simplicity, it has provided a useful and influential framework in which to think about CPs. In the motion direction–discrimination task, the model discriminates upward from downward motion by comparing the activity of two pools of direction-selective neurons, one preferring upward motion, the other preferring downward (**Figure 2*a***). The activity of the neurons in each pool is summed, and the model reports the direction of motion corresponding to the pool with greater activity. When the authors assumed that the responses of each neuron fluctuated independently of all other neurons, the model indicated that a mean CP as high as 0.56 (Britten et al. 1996) implies that the decision was based on a very small number of neurons (<10).

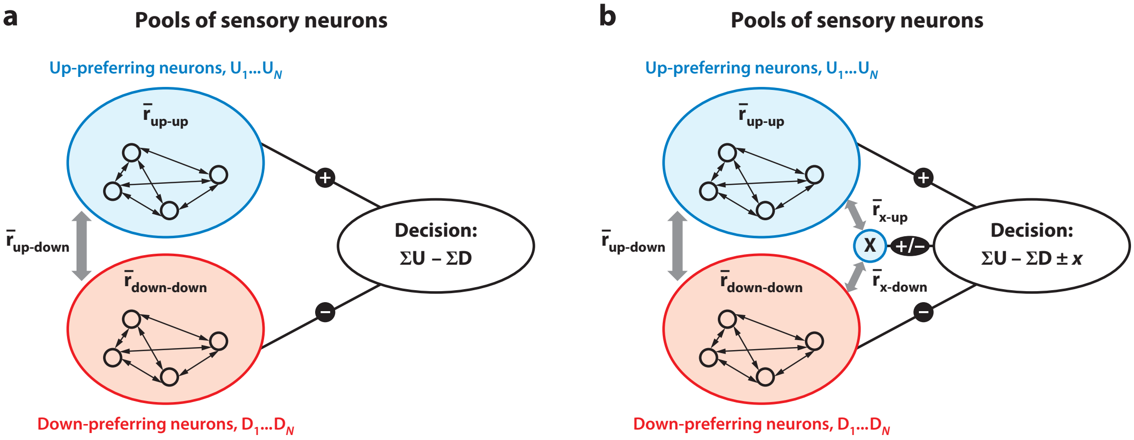

Figure 2

(*a*) Schematic of the pooling model. The spikes from one group of neurons, preferring upward motion (*blue*), are summed, as are the responses of a group preferring the opposite direction (*red*). The decision is based on the difference of these sums: $\sum U - \sum D$. Double-headed arrows indicate interneuronal noise correlations. (*b*) A single neuron (x) is added. Whether this neuron is part of the "up" pool (positive weight) the down pool (negative weight) or neither (0 weight) has negligible effect on the decision if the pools are large. If fluctuations in spike count of neuron *X* are correlated with those of neurons in one decision pool, then neuron *X* will be correlated with the decision. If neuron *X* is correlated with neurons in both pools to equal extents, then there will be no correlation with the decision.

The resolution to the apparent paradox between large numbers of neurons and measurable CPs lies in the observation that the responses of nearby cortical neurons tend to fluctuate in a correlated manner (**Figure 1*d***). CP measures the relationship between fluctuations in the response of one neuron to repeated presentations of the same stimulus and the animal's choices. The term noise correlation refers to the relationship between those same response fluctuations in pairs of nearby neurons (Bair et al. 2001). Noise correlation (also called spike count correlation or interneuronal correlation) is simply the Pearson's correlation between the trial-to-trial fluctuations in the responses of a pair of neurons. In sensory cortex, noise correlations among pairs of nearby neurons with similar tuning are typically weak but positive (for review, see Cohen & Kohn 2011).

The existence of positive noise correlations explains the existence of CPs in large pools of neurons. The pooling model hypothesizes that the animal's decision is related to the summed activity of many neurons. If the responses of all neurons fluctuated independently, the summed activity of all the neurons that contribute to a decision would be nearly constant because the noise would be averaged away. If, however, the responses of all neurons fluctuate up and down together (as is the case when noise correlations are positive), the summed activity of even thousands of neurons will fluctuate as well. In this scenario, the activity of any one neuron will be correlated with behavior simply because it is correlated with all the rest of the neurons and therefore with the monkey's percept. Measured CPs therefore depend critically on the presence of noise correlations within a population of neurons.

Since the original study by Britten and colleagues, our knowledge of the conditions under which CPs are observed, and our understanding of what they can tell us about the process of perceptual decision making, has greatly expanded. Our goals here are to review studies of CP in different systems and tasks, to explore the implications of CPs for population coding and decision making, and to discuss how recent and future experimental and theoretical work will shed light on the neuronal mechanisms underlying perception. We discuss experimental methods for determining the number of neurons involved in a perceptual decision, the role of noise correlations in determining observed CPs, evidence that CPs may at least partially reflect the impact of top-down signals on the activity of sensory neurons, and cases in which CPs can and cannot be used to infer the way that the responses of many neurons are combined to guide decisions.

DETERMINING THE NUMBER OF NEURONS INVOLVED IN A DECISION

Above, we describe two important factors that interact to determine CP: the number of neurons that contribute to the decision (pool size) and the extent to which fluctuations in activity are correlated between neurons. At one extreme is the possibility that fluctuations in activity are independent in different neurons. Under this assumption, the pooling model shows that typical CPs imply that very few neurons are involved in the decision (pool sizes of fewer than 10 neurons). Larger pool sizes require the existence of noise correlations to explain CP (Shadlen et al. 1996). Most studies reporting CP have not simultaneously measured interneuronal correlations from the same trials. In such studies, since noise correlations might have been very small, small pool sizes remain a viable theoretical explanation.

One difficulty with proposing small decision pools is explaining how it is possible to record significant CP on average in a large group of neurons sampled at random from the population. If all these neurons were recorded using the same physical stimulus, the observation of CP in many neurons would rule out the possibility of very small pools; if only very few neurons contribute to a decision, the probability of encountering such a neuron would be tiny. However, in most studies, a number of stimulus parameters are adjusted to make signals from a given cell especially relevant. If the stimulus is sufficiently tailored in this way, it becomes possible that for each of the specific stimuli used, the pool size underlying psychophysical responses is indeed small. Some studies that tried to ensure very specific matching of stimuli to neuronal preferences have argued in favor of such a scheme (Ghose & Harrison 2009), and such data are certainly compatible with small pool sizes. However, because noise correlations were not measured, the data are equally compatible with large pool sizes and correlated activity.

One experimental observation argues strongly against the idea of small pool sizes. In the widely used random dot direction task, CP changes quite gradually as the stimulus orientation moves away from a neuron's preferred direction (Bosking & Maunsell 2011, Cohen & Newsome 2009). In contrast, if the careful tailoring of stimulus parameters leads to CP because of small pool sizes, changing the stimulus should lead to substantial reductions in CP. To our knowledge, no study has demonstrated a rapid decrease in CP with changes in the stimulus, making it unlikely that small pools explain CP.

In the absence of clear evidence that small pools of cortical neurons underlie sensory decisions, the remainder of the review makes the assumption that large pools of neurons are responsible. As mentioned above, the existence of CPs when pool sizes are large implies that fluctuations in activity are correlated between neurons. We now discuss these correlations in more detail.

INTERNEURONAL NOISE CORRELATIONS AND CHOICE PROBABILITY

The Importance of Noise Correlation Structure

Based on the limited data available at the time, the original pooling model used a correlation structure (a relationship between noise correlation and tuning or pool identity) consisting of weak positive noise correlations between neurons within a pool (which we term $r_{up\text{-}up}$ to refer to correlations between pairs of neurons preferring upward motion) and no noise correlations between neurons in opposite pools (e.g., between an upward- and a downward-preferring neuron, $r_{up\text{-}down} = 0$). The CPs produced by the pooling model in this scheme depend on both the correlations between neurons and on how the neuronal signals are pooled (or read out by downstream neurons). To understand the relationship between readout, correlation, and CP, it is useful to consider what determines the CP found in any one model neuron (neuron "X" in **Figure 2*b***). A recent study (Haefner et al. 2012) derived an expression for this quantity and showed that

$$CP_i - 0.5 \propto \sum_{j=1}^{N} w_j r_{ij} / \beta N, \qquad 1.$$

where CPi is the CP observed in the i^{th} neuron, w_j is the contribution of neuron j to the decision (pooling weight), and r_{ij} is the noise correlation between neuron i and neuron j. The term β depends on the weighted mean of all correlations but is the same for all neurons in a given model.

In fact, the relationship is not exactly proportional, but over the range of observed CP, Equation 1 is an excellent approximation. When the pool of neurons is large, the contribution of neuron i to the sum on the right-hand side of Equation 1 becomes negligible. Therefore, the pooling weight assigned to any one neuron has no impact on the CP recorded from that neuron. Rather, CP is determined by the mean correlation between that neuron and all other neurons, weighted according to pooling weight of each of the other neurons.

The idea that an individual neuron's pooling weight does not affect its CP can be illustrated in some informative scenarios. Consider a situation in which there are just two opposing sets of neurons ("up" and "down" pools). In this case, the two pools are given opposite weights. For example, all up-preferring neurons might have a weight of +1, and down-preferring neurons might have a weight of −1 (meaning that they contribute evidence against, rather than in support of, an up decision). In the simplest model used by Shadlen et al. (1996), correlations between all pairs of up-preferring neurons take the same value ($r_{up\text{-}up}$), whereas correlations between pairs involving one up-preferring neurons and one down-preferring neuron take a different value ($r_{up\text{-}down}$). If $r_{up\text{-}down}$ is zero, Equation 1 shows that CP depends only on $r_{up\text{-}up}$, the correlation between neurons belonging to a given pool, as found in Shadlen et al. (1996).

However, several studies have revealed that even neurons with different preferred directions tend to have weakly positive noise correlations (Cohen & Newsome 2008, Gutnisky & Dragoi 2008, Huang & Lisberger 2009, Jermakowicz et al. 2009, Kohn & Smith 2005, Smith & Kohn 2008, Zohary et al. 1994), so it is more realistic to consider the case where $r_{up\text{-}down}$ is also positive. In this case, Equation 1 shows that CP depends on $r_{up\text{-}up} - r_{up\text{-}down}$, as found in numerical simulations (Nienborg & Cumming 2010). CP does not simply depend on the overall level of correlations between all pairs of neurons in the decision pools. It depends on the structure of these correlations; there must be a specific relationship between pooling weight and noise correlation. The simplest case is that correlations between neurons within a single pool (neurons with weights of the same sign) are higher than correlations between neurons that belong to different pools (weights of opposite signs).

Because the pooling weight of any individual neuron does not affect that neuron's CP, Equation 1 accounts for simulations that included a population of model neurons making no contribution to a decision [i.e., have a weight of 0 (Cohen & Newsome 2009)]. These "irrelevant" neurons can have CP just as large as neurons that do contribute to the decision. As long as both types of neurons have the same correlations with the other neurons in the decision pool, they will have the same CP regardless of their pooling weight. In **Figure 2*b***, if the set of correlations between neuron "X" and all other neurons is fixed and N is large, then the CP observed in neuron "X" is the same whether the weight (w) is -1, 0, or 1. This lack of dependence on pooling weight occurs because the contribution of this neuron to the decision, $w_x x$ (where x is the response of neuron "X"), is negligible when N is large. Consequently any correlation between x and the decision arises only because of correlations between neuron "X" and the pools of upward- or downward-preferring neurons (U_i and D_i). For the same reasons, neurons that do contribute to a decision will not show CP if they lack the appropriate correlations. If the mean correlation between neuron "X" and all up neurons equals the mean correlation between "X" and all down neuron correlations, then CP will be 0.5, regardless of the weight w_x given to that neuron (Equation 1).

Neurons that carry the most reliable signals for discrimination (quantified using a neurometric threshold; for review, see Parker & Newsome 1998) tend to show larger CPs (Britten et al. 1996, Celebrini & Newsome 1994, Gu et al. 2008, Parker et al. 2002, Purushothaman & Bradley 2005, Romo et al. 2002, Uka & DeAngelis 2004). Because the weight assigned to a neuron does not influence its CP, this result cannot be explained by suggesting that the animals assign greater weight to these more reliable neurons. Rather, the relationship between neuronal sensitivity and CP has been explained by suggesting that neurons with weak signals also have weaker noise correlations with neurons carrying strong signals (Shadlen et al. 1996). Although this correlation between sensitivity and CP is commonly found, 20-fold changes in neuronal sensitivity have been reported across studies with little difference in CP (e.g., Britten et al. 1992, 1996; Cook & Maunsell 2002; Purushothaman & Bradley 2005). Conversely, studies finding higher CPs than typical (e.g., Dodd et al. 2001) do not show unusually high neuronal sensitivity (Parker et al. 2002). Thus, across studies, there is not a consistent relationship between CP and the ratio of neurometric threshold to psychometric threshold. This variation between studies can nonetheless be explained by differences in how neurons are weighted: Whereas psychometric thresholds depend on the relative weights of sensitive and less sensitive neurons in a pool, CPs depend on the correlation structure and only indirectly on the weights of the neurons within a pool (Equation 1). A study that looked at CP over several months of training found that as behavioral thresholds improved, CP for the more sensitive neurons gradually increased, whereas neurometric thresholds were unchanged (Law & Gold 2008). This finding was explained by suggesting that the animals learn to adjust pooling weights to favor more sensitive neurons. Equation 1 shows that this explanation requires that the less sensitive neurons that have lower weights also have weak noise correlations with more sensitive neurons. If this were not true, the terms $w_j r_{ij}$ contributed by the insensitive neurons would still produce high CP early in training.

At first sight, this dissociation between pooling weight and CP for individual neurons suggests that it is not possible to infer anything about the contributions of neurons to decisions based on CP. However, considering the entire population, CP does depend on both correlations and pooling weights. For this reason, it may still be possible to use measures of CP combined with measures of noise correlation to say something about the contribution of identified subpopulations of neurons to a decision (see Using Choice Probability to Infer Readout). Because this process still makes assumptions about what gives rise to a structured pattern of noise correlations, it is important first to consider how this pattern of correlations might arise.

Sources of Correlation Leading to Choice Probability

Equation 1 shows that in discrimination tasks, CP will arise when correlations within a pool are higher than those between pools. This finding suggests that neurons within a pool supporting one decision receive a common signal that is not shared between pools supporting opposite decisions. Because correlations are likely the largest determinant of CP, understanding the origin of this common signal is critical for understanding the origin of CPs.

Three potential (not mutually exclusive) sources of common signals could lead to CPs.

1. The common signal could reflect noise in shared, feed-forward sensory afferents. This includes sensory noise at the level of the sensory receptors and noise arising from action-potential propagation and synaptic transmission. As neurons with similar tuning preference share more inputs than do neurons with different tuning preferences the noise in such feed-forward afferents could cause higher correlations in similarly tuned neurons than in neurons with different tuning, as is required for CPs.

   The spatial structure of correlations in cortex suggests that sensory afferents are not the only inputs producing noise correlation. If they were, then noise correlations should be restricted to distances over which neurons receive common feed-forward input. In contrast, measurements in V1 have found that correlations are present over large distances of cortex (up to 10 mm; e.g., Smith & Kohn 2008), whereas correlations in retinal ganglion cells are spatially very localized and restricted largely to directly neighboring pairs (Greschner et al. 2011). This broad spatial structure suggests that, at least in area V1, shared noise in feed-forward inputs is not the only source of variability contributing to noise correlation. These spatially extended correlations in V1 mean that subsequent processing stages (e.g., MT) might show correlations over large distances that reflect correlations in the afferent input from V1.
2. The common signal could be generated within a sensory area (and perhaps transmitted in a feed-forward way to downstream areas). This signal could arise from mutual connectivity between neurons with similar tuning preference such that they share locally generated noise or from horizontal connections among functionally similar neurons (Ahmed et al. 2012, Malach et al. 1997). This source of noise likely contributes to CP, given that the majority of inputs in

cortex originate from intrinsic connections (Salin & Bullier 1995) and synaptic transmission is a substantial source of noise (Faisal et al. 2008).

These first two sources of correlation are compatible with the idea that noise correlations reflect hard-wired mechanisms that are a fixed property of the network. Such inputs likely make an important contribution to CPs, and many CP results can be well explained by these sources of correlations. More recent data, reviewed below, however, suggest that the structure of the correlation is more flexible and can depend, e.g., on the task an animal performs. Therefore, a third, more flexible, common signal likely contributes to CP.

3. The flexible common signal underlying noise correlations and CPs could reflect a top-down signal (i.e., a signal originating anywhere but the ascending pathway of the sensory-processing hierarchy preceding the sensory area from which CPs are recorded). This common signal could reflect a cognitive process such as feature attention and change with task instructions or any other downstream signal related to the decision. Such a top-down origin of the common signal implies that the correlation structure is not fixed but can change dynamically with the animal's cognitive state, e.g., with task instructions. Recent evidence suggests that top-down signals contribute to noise correlations, changes in the structure of the noise correlations, and CPs.

Several recent studies have found that noise correlations can be changed dynamically by the task instructions given to the animal. Two studies manipulating spatial attention found that noise correlation decreased when spatial attention was directed to the receptive field (Cohen & Maunsell 2009, Mitchell et al. 2009). However, although the changes in the mean noise correlation with spatial attention were dramatic, the structure of the correlation (the relationship between noise correlation and tuning curve similarity) was unchanged (Cohen & Maunsell 2009). Therefore, if spatial attention decreases overall noise correlation without changing the structure of the correlation in all tasks, it would not lead to CP changes. (Note that given that Cohen & Maunsell used a detection task, the observed changes in mean correlation have consequences for measurements of "detect probabilities" in their task, as addressed below in Detection Tasks: A Different Relationship Between Noise Correlations, Neuronal Activity, and Behavior.)

A different study looked explicitly at correlation in two discrimination tasks (Cohen & Newsome 2008) and found changes in the correlation structure that would lead to changes in CP. In interleaved trials, the animals were required to discriminate the direction of motion along orthogonal axes of motion direction, e.g., left versus right in one trial, then up versus down in a subsequent trial. The discrimination axes were carefully chosen for any given pair of neurons to ensure that for one axis the two MT neurons contributed to the same decision pool, whereas for the orthogonal axis they contributed to opposite decision pools. Cohen & Newsome found systematic changes in correlation as a function of the task. For neurons with preferred directions that differed by less than 135°, they found higher correlations when both neurons contributed to the same decision pool, compared with when they contributed to opposite pools. These correlation measures were all made using responses to a single visual stimulus, dots moving with zero coherence. Only the task context in which they were presented had changed, suggesting that these changes in correlation reflect some top-down input. The authors' simulations demonstrate that correlation change could result from fluctuations in feature-selective attention to the two choice directions. For example, while discriminating upward from downward motion, the animal may attend more strongly to upward motion on some trials and downward motion on others. These fluctuations would add a positively correlated signal to neurons belonging to the same decision pool and a negatively correlated signal to neurons in opposite pools. Using the basic pooling model, Nienborg & Cumming (2010) estimated that the signal coming from these fluctuations accounts for approximately half the observed CP in MT.

However, using the pooling model in this way makes an important assumption: that the changes in spike count produced by feature attention still contribute to the animal's decision. Alternately, the animal could know its attentional state on each trial and discount it when making a decision. If attentional changes are not discounted, and hence do contribute to decisions, there should be a systematic relationship between the attentional state and choice. Evidence in favor of this was found by Nienborg & Cumming (2009), who used a variant of reverse-correlation analysis to estimate changes in the neuronal response function with choice. They found that choices to a neuron's preferred stimulus feature were associated with an increase in neuronal response gain, similar to effects of attention.

If such a feature-selective signal is present before the decision is formed (possibly reflecting a bias or expectation), it may bias the decision via its effect on the sensory neurons. Nienborg & Cumming (2009) found indirect evidence to support this notion. On trials where small rewards were available, animals made less use of visual information in the stimulus, presumably relying more on biases. The trials with small reward were associated with slightly higher CP, suggesting that fluctuations in bias contributed to CP.

Alternatively, a feature-selective top-down signal could occur after the decision is formed (postdecision). The role of such a feature-selective feedback signal may be to serve perceptual stability, in particular when the sensory signals are weak or ambiguous, the situation in which CPs are typically measured. Nienborg & Cumming (2009) also found evidence that postdecision top-down signals may contribute to CPs. The study used a reverse-correlation approach ("psychophysical reverse correlation") to quantify how the monkeys weighted the relevant information in the visual stimulus and simultaneously measured CPs in V2. Although the weight that the animals gave to the visual stimulus decreased over the course of the trial, CPs did not decrease. As the animals give less weight to the visual inputs, stochastic variation in the neural representation of those inputs should also have less impact on choice, and hence CP should fall. Nienborg & Cumming proposed that a feature-selective feedback signal (reflecting the decision) supported CP at the end of the trial. If these changes in noise correlation during stimulus presentation were a fixed property of feedforward correlations, this might also explain the result. These changes in noise correlation during stimulus presentation being a property of fixed, feedforward noise correlations might also explain the result. Data on the structure of the correlation over time from untrained fixating or anesthetized animals are sparse and differ between studies (Samonds et al. 2009; M.A. Smith, M.A. Sommer, A. Kohn, unpublished observations), making it unclear whether the structure of feed-forward noise correlations changes in a way to account for the time course of CPs.

Further experiments will be required to understand the relationship between the effect of top-down signals on the structure of noise correlations, on the sensory representation, and on the readout of that representation. Without new data, it will be difficult to provide quantitative estimates of the contribution that top-down signals make to the structure of noise correlations, CP, and the extent to which these precede or follow the decision.

Choice Probability for More than One Stimulus Attribute Indicates Highly Structured or Flexible Correlations

Neurons in MT show CP for direction-discrimination tasks (Britten et al. 1996), disparity-discrimination tasks (Uka & DeAngelis 2004), speed-discrimination tasks (Liu & Newsome 2005), and a task that required subjects to identify the conjunctions of disparity and motion (differentiating "near and left" from "near and right"; Dodd et al. 2001). As laid out above (see The Importance of Noise Correlation Structure), the observation of CPs in each of these tasks implies that the correlation between neurons supporting the same decision is higher than the correlation between neurons supporting different decisions.

The CP observed in the disparity task implies that correlations between neurons with similar disparity preferences (e.g., both preferring "near" disparities) are higher than correlations between neurons with dissimilar disparity preferences (e.g., one preferring near disparity, one preferring far disparity) (**Figure 3*b***). For the direction-discrimination task, a mean CP in right-preferring neurons is observed because they have larger noise correlations with one another than with left-preferring neurons. But this population of right-preferring neurons includes both near- and far-preferring neurons (**Figure 3*a***), and correlations between these pairs need to be low to explain CP in the disparity task. In the

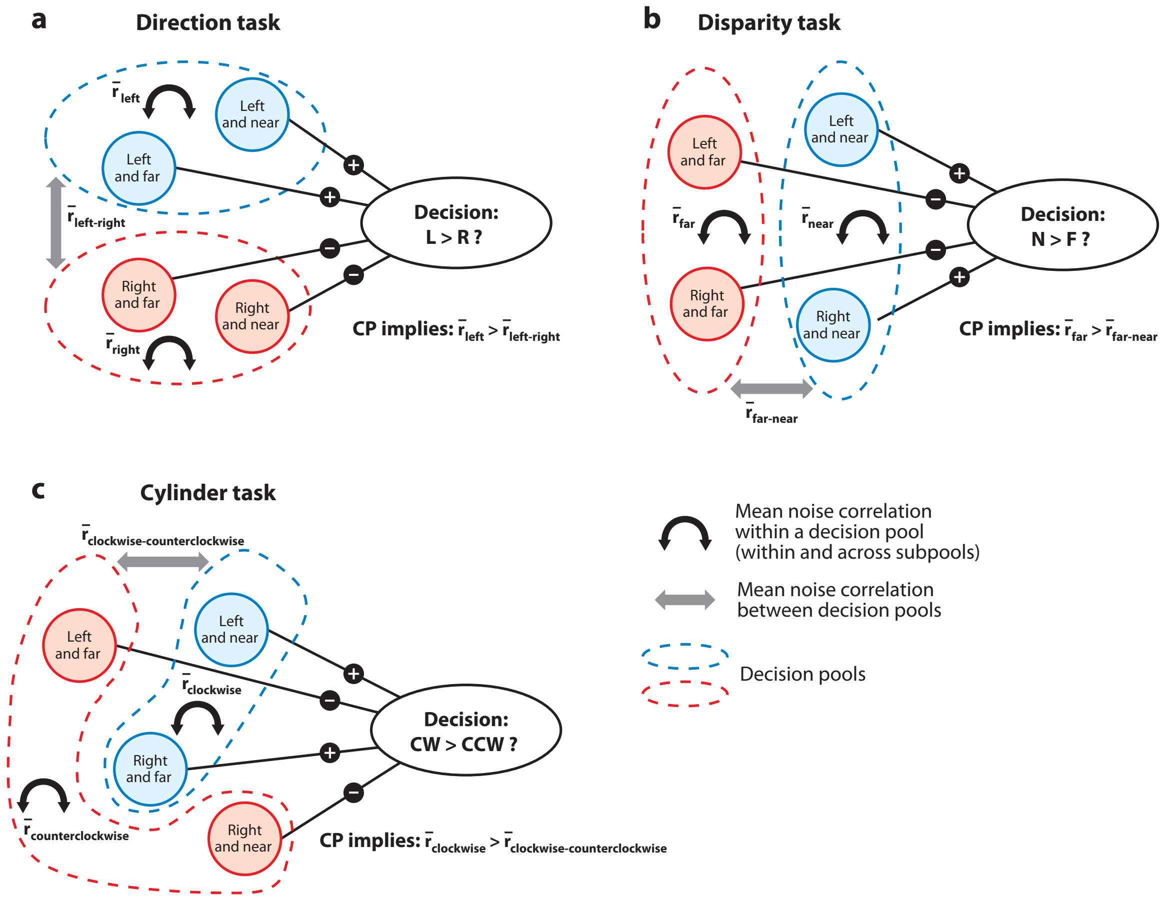

Figure 3

Using the same neurons for different tasks (motion and disparity). Four pools of neurons are shown with preferences for each conjunction: left and far, left and near, right and far, and right and near. Depending on the task being performed, these groups must be pooled differently. Dashed ellipses group together neurons in a decision pool in each case. (*a*) Direction discrimination. (*b*) Disparity discrimination. (*c*) A task involving two transparent surfaces forming the appearance of a rotating cylinder. Animals reported the conjunction of disparity and motion, corresponding to identifying whether the cylinder rotation was clockwise or counterclockwise. In each case, the requirements for interneuronal correlation to produce choice probability (CP) are different. A single pattern of interneuronal noise correlations that can explain CP in all three tasks needs to be highly structured.

conjunction task (**Figure 3*c***) near and right neurons are pooled with far and left neurons. These two groups belong to opposite pools in the direction task and in the disparity task. It is not clear that a single fixed set of interneuronal correlations can explain CP in all three tasks, and no quantitative studies have explored this.

This argument compares different studies, so it is important to confirm that CP for different tasks can be measured in the same neurons. Sasaki & Uka (2009) recorded from MT neurons in animals trained to perform both disparity- and depth-discrimination tasks and found significant CP for both tasks. They found that the direction-discrimination task produced substantially larger CPs when the animals' behavioral report involved the same eye movement used to report preferred disparity. This may be explained in principle if the animals were not reliably switching tasks according to their instructions, although the authors present evidence against this interpretation. If the effect is not explained by poor behavior, it implies that the strongest correlations exist among neurons with preferences for disparity and motion that follow the (arbitrary) convention chosen by the experimenters. It is difficult to imagine that this represents an innate property of MT. The alternative is to suggest that training has established a particular structure of interneuronal correlations in these animals. If this does happen, there would be no need to find a fixed set of noise correlations that can explain CPs for multiple tasks. Rather, the correlation structure changes with the task because of training. More work with animals trained on multiple tasks will be required to address this point. Of course, a significant top-down component that changes with task also removes the need for a single, fixed, correlation structure.

Detection Tasks: A Different Relationship Between Noise Correlations, Neuronal Activity, and Behavior

The impact of noise correlations on the relationship between the activity of any given neuron and behavior depends on the way that information from multiple neurons is combined to drive a perceptual decision. So far, our discussion has focused on discrimination tasks, in which the logical decision rule is to compute the difference in activity between two pools of neurons (each pool favoring one of the two possible choices). Other tasks, however, have different ideal algorithms for combining the activity of multiple cells.

One commonly studied task with a fundamentally different relationship among noise correlations, activity in single neurons, and behavior is a change-detection task. Rather than asking the subject to differentiate between, for example, upward from downward motion, these tasks require a subject to notice a change in a single stimulus (Bosking & Maunsell 2011; Cohen & Maunsell 2010, 2011; Cook & Maunsell 2002) or the onset of a weak stimulus (Palmer et al. 2007).

In certain detection tasks, ideal decoding often involves simply taking the (perhaps weighted) average response of a single pool of neurons. For example, the vast majority of cells respond more to a stimulus than to a uniform background, so the onset of a low-contrast stimulus could be detected simply by comparing the average response to some criterion level (Palmer et al. 2007). In other tasks, prolonged exposure to a base stimulus may lead to adaptation of neuronal responses, so most neurons will show increased responses to the changed stimulus when that adaptation is released (Cohen & Maunsell 2010, 2011).

Even when a decision is based on the average of a single neuronal pool, one may quantify the relationship between the activity of individual neurons and behavior in the same way as for a discrimination task. The detect probability (DP) is the probability that an ideal observer could predict whether an animal will detect a near-threshold change in a stimulus on the basis of the fluctuations in the responses of an individual neuron. The term detect probability was chosen to distinguish it from CP in discrimination tasks.

Unlike CPs, DPs in tasks in which the decision is based on the average activity in a single group of neurons are monotonically related to the average amount of noise correlation in the pool (rather than the difference in correlation between two pools). This result can be seen from Equation 1 if all the weights are positive. When noise correlations are low, the average response upon which the decision is based does not vary much, so the activity of one neuron is not very predictive of the animal's choices. When correlations are high, fluctuations in each neuron's responses will be correlated with the decision because they are correlated with the activity of the pool as a whole.

One important consequence of this different relationship between noise correlations and DPs is that cognitive factors and experimental artifacts can have a greater effect on DPs than CPs. Because CPs are affected only by the difference in the correlations within and between two pools, they can be affected only by a cognitive factor that affects the two pools differently. Suppose that an animal fails to direct spatial attention to the stimulus during one trial. This will reduce neuronal firing rate and reduce the probability that an animal detects any change, producing a correlation between firing rate and choice. In the direction-discrimination task, spatial attention should affect "up" and "down" pools equally and hence will not produce systematic CP. In this way, fluctuations in global factors such as arousal, alertness, or motivation or in many experimental artifacts such as fixational eye movements or blinks can introduce correlations in firing rate across the entire population and affect DP. Because CP does not depend on the overall correlation level, but on the differences in correlation between pools, these global factors should not affect CP. A study of the effects of fixational eye movements found that these did indeed affect DP, but not CP (Herrington et al. 2009). These cognitive and experimental factors therefore pose some difficulty for interpreting DPs in tasks in which the decision might be based on the average response of all neurons.

Detection tasks using stimulus changes that elicit increases in the responses of some neurons and decreases in others can be useful for determining the impact of global factors on DP. In this situation, the animal cannot detect the change simply by summing the response of all neurons, so must compare two pools. Because global factors will affect both pools, while also affecting psychophysical performance, they will produce different DP in the two groups. For example, changes in speed (Price & Born 2010) or the onset of coherent motion (Bosking & Maunsell 2011) elicit both increases and decreases in MT neurons, depending on the relationship between the particular stimulus and tuning of the neuron under study. Bosking & Maunsell (2011) found that neurons whose preferred direction matched the direction of the onset of a coherent motion stimulus had DP that was significantly greater than chance (0.58), whereas neurons with opposite tuning had a DP of 0.46 (significantly less than 0.5). Because the responses of neurons whose preferred direction matched or opposed the direction of the coherent motion are equally informative for solving this task, the responses of the two groups of neurons should ideally be given equal and opposite weight in the decision. If this were true and there were no influence of global factors on DPs, the DPs for the two groups should deviate from 0.5 in equal and opposite amounts. The observed difference in the magnitude of DP for two groups (0.58–0.5 = 0.08 for neurons whose preferred direction matched the stimulus direction and 0.5–0.46 = 0.04 for neurons whose preferred direction was opposite the stimulus direction), therefore, places an upper bound on the effects of global factors on DP.

USING CHOICE PROBABILITY TO INFER READOUT

The existence and pattern of CP are often used to make inferences about the nature of the mechanisms that read out the responses of sensory neurons to drive decisions, but CP's dependence on correlations makes these inferences difficult (Nienborg & Cumming 2007). As Equation 1 shows, for any given neuron, it is impossible to determine whether CP reflects a causal contribution to a decision or whether it reflects responses that are simply correlated with other neurons that do.

Nonetheless, CPs can place some constraints on possible readout mechanisms. For example, in MT, neurons with opposite preferred directions but similar spatial receptive fields tend to be weakly but positively correlated (Cohen & Newsome 2008, Huang & Lisberger 2009, Zohary et al. 1994). In Bosking & Maunsell's (2011) detection task, if the neurons whose preferred directions were opposite the stimulus direction were not involved in the decision, their positive correlation with neurons whose tuning matched the stimulus direction would have led to DPs greater than 0.5. That the authors observed DPs less than 0.5 suggests that the readout mechanism accounts for the fact that decreases, rather than increases, in the responses of these neurons signal motion onset. As described above (see The Importance of Noise Correlation Structure), this conclusion depends on both the observed DP and the interneuronal noise correlation. Because the latter was not measured, the authors' conclusion requires an assumption that the correlation structure during their task was the same as that reported by other investigators.

CP has also been used to assess the aspects of the neural code that are relevant for guiding behavior. Salinas et al. (2000) showed CP for spike counts (or rates) in primary sensory cortex in a tactile frequency-discrimination task. They did not observe significant CP for measures of the periodicity of the response (their estimate of a spike-timing signal). Comparing these results suggests that firing rates, rather than spike timing, affect decisions. However, this result may also reflect differences in how these different properties are correlated between neurons. We point out above (see The Importance of Noise Correlation Structure) that if a neuron's firing rate does not show appropriate correlations with other neurons, then that neuron will not show CP even if it contributes to the decision. The same argument applies to spiking periodicity: If trial-to-trial fluctuations in this periodicity are not correlated between neurons, or noise correlations in periodicity do not show the appropriate structure, then CP will be absent even if the information is being used to guide decisions. A study of choice-related responses in the rat olfactory bulb (Cury & Uchida 2011) found the opposite result: Fine temporal responses were correlated with the animal's reaction time in an olfactory discrimination task, but the mean rate was not correlated. In the future, we must identify the correlation structures for different measures of neuronal activity to determine whether rate and temporal codes play different roles in two tasks, in different brain areas, or across different species.

It Is Difficult to Infer the Time Window in Which Decisions Are Made

Ideally, CPs (and all analyses of neuronal responses contributing to perceptual decisions) would be calculated using spikes recorded over the same period during which the animal makes the perceptual decision. The timing and duration of this window likely depend on the stimulus: Tasks involving noisy stimuli or stimuli that evolve over time (Britten et al. 1996, Nienborg & Cumming 2006, 2009, Uka & DeAngelis 2006) may benefit from longer viewing durations than would studies that use high-contrast, noise-free stimuli (Ghose & Harrison 2009, Price & Born 2010). The decision window also likely depends on the task: Object recognition or discrimination may require longer viewing durations than simple onset detection tasks. In general, more difficult tasks require longer durations than easier ones do.

The decision window used by the animal can be short, variable, and difficult to determine. Many studies use stimulus presentations that last greater than one second, but evidence suggests that animals typically avail themselves of only a few hundred ms of the available stimulus duration, even in tasks in which integrating evidence over time is beneficial to performance (Cohen & Newsome 2009, Kiani et al. 2008, Roitman & Shadlen 2002). When longer stimuli are available, animals may ignore a stimulus during, for example, the period immediately after stimulus onset or after the decision has been made (Kiani et al. 2008, Nienborg & Cumming 2009). This issue is made more complicated for the experimenter because the dynamics of the decision may be different on every behavioral trial.

Reaction time tasks (in which subjects are free to respond as soon as they are ready) can be useful for identifying the decision period used by the subject because the response time provides an upper bound for the time used and because most animals are motivated to make quick decisions to receive a reward (or an easier trial) sooner. Viewing durations that more closely match the decision window provide more accurate measurements of the spikes that may be involved in the decision. In a reaction-time version of the motion-direction discrimination task, substantially less noise in the decision process is necessary to explain observed neuronal sensitivities and CP (Cohen & Newsome 2009). However, even reaction-time tasks inevitably include nondecision time, and the stimulus-presentation time may include motor-preparation delays (which often occur after stimulus offset in fixed-duration tasks). Behavioral and neuronal data can be useful for estimating these nondecision periods (Huk & Shadlen 2005, Janssen & Shadlen 2005, Mazurek et al. 2003, Palmer et al. 2005, Stanford et al. 2010). In practice, however, the amount of data required to estimate CP accurately and the fact that the animal's decision-making process likely varies from trial to trial make it difficult to estimate the decision window precisely.

The timescale of interneuronal noise correlations and autocorrelations in a single neuron's responses can affect how much misidentifying the decision period will affect CP measurements. If the timescale of interneuronal correlations is substantially shorter than the decision period and if fluctuations in neuronal responses are independent across time (i.e., the responses show little autocorrelation), then CP should be high during the decision period and at chance during other times. Therefore, counting spikes during nondecision periods should lead to lower-than-accurate measurements of CP. If, however, correlations have a long timescale or there are substantial autocorrelations, then firing rates outside the decision period should predict choices. In this case, longer measurement windows will not bias measurements of CP and may lead to more accurate measures because task-independent noise will be averaged out.

Several lines of evidence suggest that CPs can remain high during nondecision times. Most directly, several studies have examined the time course of CP by computing CP in sliding windows throughout the stimulus-presentation period. Although the time course varies somewhat across studies (and even across subjects within a study) and these calculations suffer from the weakness that the decision period may not be the same on every trial, most studies find that in fixed-duration tasks, CPs are high throughout the stimulus-presentation period. CP tends to rise after stimulus onset and remains above chance through the stimulus period (Britten et al. 1996, Celebrini & Newsome 1994, Nienborg & Cumming 2006, 2009, Uka & DeAngelis 2004). In reaction-time tasks, CP tends to decay before the saccade, which could be because the decision must be expedited to allow for motor preparation (Cohen & Newsome 2009, Cook & Maunsell 2002, Price & Born 2010).

Measurements of interneuronal correlations, autocorrelations, and fluctuations in top-down or cognitive factors are consistent with the idea that CPs could be high outside the decision period. Interneuronal correlations are typically dominated by fluctuations on the timescale of tens to hundreds of milliseconds (Bair et al. 2001, Ecker et al. 2010, Kohn & Smith 2005, Mitchell et al. 2009). MT neurons show a small amount of autocorrelation on a similar timescale (Bair et al. 2001), suggesting that CPs will remain high beyond the decision period. Two studies using briefly presented stimuli have demonstrated CPs that lasted less than 100 ms (Ghose & Harrison 2009, Price & Born 2010), challenging the above description. Such rapid decisions may be dominated by the activity of a very few neurons (Ghose &

Harrison 2009), in which case a component of CP no longer depends on correlations (see Determining the Number of Neurons Involved in a Decision). If the pools are small enough, timing of activity in single neurons could be reflected in CP. Alternatively, the temporal correlations in firing under these circumstances (not reported) may have been short enough to explain the results.

The timescale of fluctuations in top-down or cognitive factors is typically longer than that of a perceptual decision. Early attempts to measure fluctuations in spatial (Cohen & Maunsell 2010) and feature attention (Cohen & Maunsell 2011) based on the activity of a few dozen sensory neurons suggest that attention fluctuates on timescales ranging from a few hundred milliseconds to tens of seconds. Even tasks that push animals to shift attention as quickly as possible report relatively slow changes. In the absence of salient changes in a visual stimulus (exogenous attention), animals can only shift spatial attention approximately every 400 ms (Cheal & Lyon 1991, Muller et al. 1998, Muller & Rabbitt 1989). Even exogenously cued shifts of attention take at least 100–200 ms (Bisley & Goldberg 2003, Cheal & Lyon 1991, Herrington & Assad 2009, Krose & Julesz 1989, Muller & Rabbitt 1989, Nakayama & Mackeben 1989).

Together, these studies suggest that CP will remain high during nondecision times, which indicates that measuring CP is not an effective way to measure a decision period or estimate the algorithm by which an animal integrates evidence. On the flip side, the relative constancy of CP throughout a stimulus presentation indicates that it may be possible to measure CPs fairly accurately even though the precise decision window is difficult to determine.

Combining Measures of Noise Correlation and Choice Probability to Infer Readout

Observing CP in a single neuron does not imply that the neuron contributes to the decision. Therefore, it is difficult to use CP to infer the algorithms by which sensory information is read out to drive decisions. However, in cases when both CPs and noise correlations are measured within and between identified groups of neurons, one can make reliable inferences concerning how a population of neurons is read out.

A powerful example of this notion comes from a series of studies investigating how neurons in the dorsal subdivision of the medial superior temporal area (MSTd) are related to judgments of the direction of self-motion. Neurons in MSTd carry signals encoding the direction of self-motion, even when an animal is moved in total darkness (information believed to be derived from the vestibular system). These neurons have CPs during vestibular stimulation (Gu et al. 2007). MSTd also contains many cells that encode the direction of self-motion simulated by a visual stimulus when the animal is stationary. These cells also show CPs during the visual task (Gu et al. 2008).

Although many cells respond to both visual and vestibular signals, they do not necessarily signal the same motion direction for both cues. So-called incongruent cells are activated most strongly by visual stimuli indicating one direction but by vestibular stimulation indicating the opposite direction (Gu et al. 2008). A similar number of congruent cells prefer the same direction of motion regardless of the cue used. This heterogeneous population offers some unique insights into how the sensory information is used because congruent and incongruent neurons make distinctive contributions to CP in different scenarios.

Figure 4 illustrates the contribution of four groups of neurons (congruent and incongruent for each direction) to decisions in a self-motion direction-discrimination task. CP in these groups depends on both the correlations between the groups and on how the population is read out. Provided noise correlations are highest for pairs of neurons with similar tuning for both parameters (e.g., matched incongruent pairs have higher correlation than do congruent-incongruent pairs), then CP in incongruent neurons reflects the way they are read out. Gu and colleagues (2011) measured interneuronal correlations for these neurons

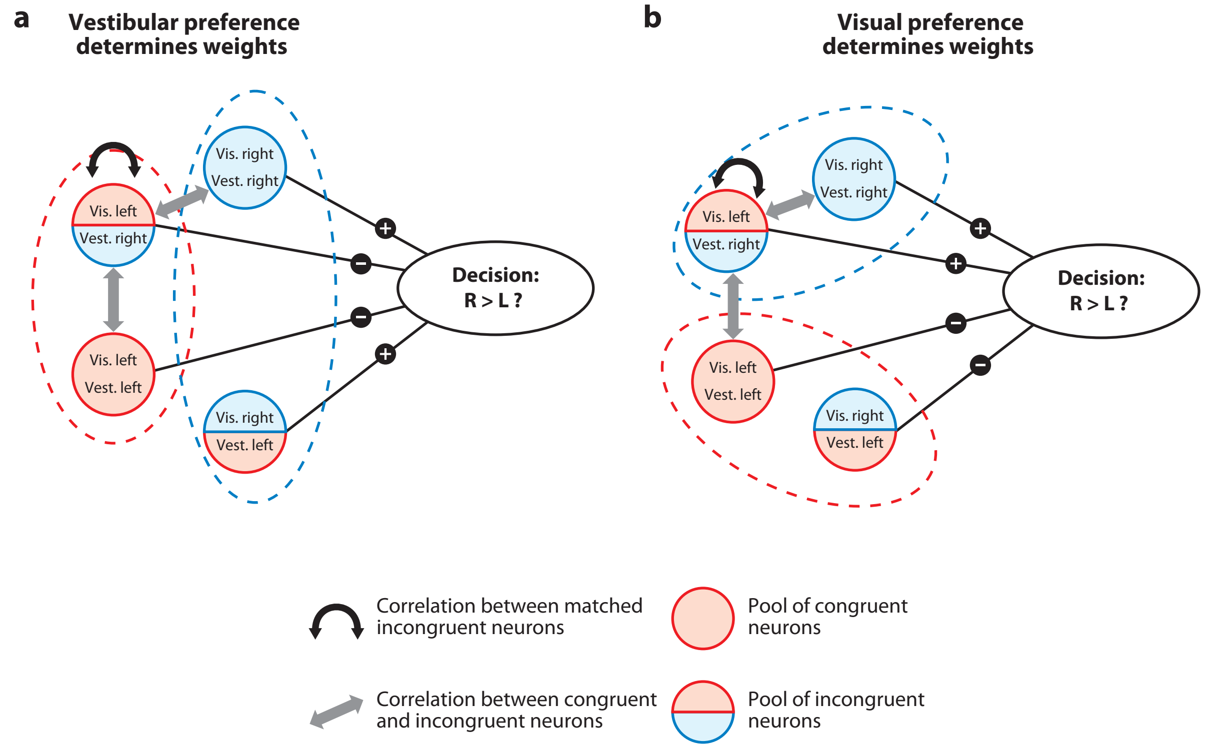

Figure 4

Combining measures of choice probability (CP) with measures of interneuronal correlation can test models of readout. Area MST (medial superior temporal) contains neurons that signal self-motion to visual and vestibular cues. Direction preference is shown by color (*Blue* = *right*, *Red* = *left*). Congruent neurons show the same preferred directions for both cues, whereas incongruent neurons show opposite preferences (*hence two colors*). Interneuronal noise correlations are shown with double-headed arrows. The darker color for correlations between neurons within a single pool indicates that these are stronger than correlations between congruent and incongruent neurons. For simplicity, only three arrows are shown, illustrating how the readout rule affects CP in visual left/vestibular right incongruent neurons. (*a*) Neurons are pooled according to vestibular preference (*dashed ellipses show pooling rule*). (*b*) Neurons are pooled according to visual preference. In both cases, there are weak correlations with congruent neurons in both pools. These correlations have opposite effects on the CP of neurons in the visual left/vestibular right group, which tend to cancel. For that reason, correlations with the other neurons in the same incongruent group largely determine the sign of CP in incongruent neurons. The result is that the sign of the CP in this group reflects how it is read out.

and found the typical relationship between tuning curve similarity and noise correlation in which correlations are highest for neurons with similar tuning. Gu and colleagues (2008) found that CP for incongruent cells was greater than 0.5 when calculated according to the vestibular preference and was therefore systematically less than 0.5 when calculated with reference to the visual preference. This observation suggests that these animals read out the activity of incongruent cells according to the direction of their vestibular preference. Of course, the results do not prove that the responses of these neurons causally contribute to the animal's decision. It is always possible that some other population of neurons forms the decision pool, and neurons in MSTd simply have appropriate noise correlations with that population. However, these experiments do demonstrate that the incongruent neurons are not read out according to their visual preference. Without measures of interneuronal correlation, even this statement would not be possible. This example illustrates that measures of CP combined with measures of interneuronal correlation do allow strong tests of hypotheses describing the way that neuronal populations are read out to guide decisions.

CONCLUSIONS

The most straightforward interpretation of the widely observed relationship between the activity of neurons in sensory cortex and animals' behavioral choices is that random fluctuations in the activity of sensory neurons influence perceptual decisions. If the decision is supported by large pools of neurons (more than ∼100 neurons), these random fluctuations must be correlated between members of a pool. Understanding which signals give rise to these correlations is therefore central to the interpretation of CPs. The results of multiple studies suggest that the correlation structure is not fixed but depends on the task an animal performs. Recent evidence suggests that, at least for some tasks, part of this signal reflects the influence of cognitive factors on sensory neurons, but there is currently no agreed upon method that allows the relative magnitude of flexible top-down and hard-wired bottom-up components to be quantified.

The difficulty of quantifying the relative influence of various sources of correlations on CP illustrates a wider problem: that any attempt to answer a question such as this depends on the model used. Cognitive factors are not part of the pooling model, which to date has dominated thinking about CP. More complex models of the relationship between sensory neurons and perceptual decisions will probably be required in the future. Nonlinear summation of neuronal responses (even simple nonlinearities such as thresholding) could lead to a situation in which correlated activity in some pairs of neurons has more impact on choices than it does in other pairs. Exploring correlations between identified neuronal subtypes may also help to identify populations with distinctive contributions to a decision. It may be important to consider recurrent interactions between neurons in a pool, neurons in different pools, and even neurons in different cortical areas, especially as these may have a profound effect on interneuronal correlations. Using different, more sophisticated models to study decision making may also change the questions that seem most relevant. For example, in models that include recurrent interactions, asking whether CPs and noise correlations arise from common feed-forward or top-down inputs may not be possible or even sensible.

Clearly, there is room for much more sophisticated models. The development of such models, however, must depend on empirical data that are not explained by simple pooling models. The paucity of such data illustrates the power of the pooling model. Simultaneous recordings from large groups of neurons in animals performing threshold psychophysics may be able to provide such data. New technologies for manipulating the activity of identified subgroups of neurons are also likely to be useful. Combining these new physiological methods with more refined ways of measuring animal behavior may provide a basis for models that go beyond simple pooling and clarify how activity in sensory neurons supports perceptual decisions.

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

Ahmed B, Cordery PM, McLelland D, Bair W, Krug K. 2012. Long-range clustered connections within extrastriate visual area V5/MT of the rhesus macaque. *Cereb. Cortex* 22:60–73

Bair W, Zohary E, Newsome WT. 2001. Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J. Neurosci.* 21:1676–97

Bisley JW, Goldberg ME. 2003. Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299:81–86

Bosking WH, Maunsell JH. 2011. Effects of stimulus direction on the correlation between behavior and single units in area MT during a motion detection task. *J. Neurosci.* 31:8230–38

Britten KH, Shadlen MN, Celebrini S, Newsome WT, Movshon JA. 1996. A relationship between behavioural choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* 13:87–100

Britten KH, Shadlen MN, Newsome WT, Movshon JA. 1992. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* 12:4745–65

Celebrini S, Newsome WT. 1994. Neuronal and psychophysical sensitivity to motion signals in extrastriate area MST of the macaque monkey. *J. Neurosci.* 14:4109–24

Cheal M, Lyon DR. 1991. Central and peripheral precuing of forced-choice discrimination. *Q. J. Exp. Psychol. A* 43:859–80

Cohen MR, Kohn A. 2011. Measuring and interpreting neuronal correlations. *Nat. Neurosci.* 14:811–19

Cohen MR, Maunsell JH. 2009. Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* 12:1594–600

Cohen MR, Maunsell JH. 2010. A neuronal population measure of attention predicts behavioral performance on individual trials. *J. Neurosci.* 30:15241–53

Cohen MR, Maunsell JH. 2011. Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron* 70:1192–204

Cohen MR, Newsome WT. 2008. Context-dependent changes in functional circuitry in visual area MT. *Neuron* 60:162–73

Cohen MR, Newsome WT. 2009. Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. *J. Neurosci.* 29:6635–48

Cook EP, Maunsell JH. 2002. Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. *J. Neurosci.* 22:1994–2004

Cury KM, Uchida N. 2011. Robust odor coding via inhalation-coupled transient activity in the mammalian olfactory bulb. *Neuron* 68:570–85

Dodd JV, Krug K, Cumming BG, Parker AJ. 2001. Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. *J. Neurosci.* 21:4809–21

Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS. 2010. Decorrelated neuronal firing in cortical microcircuits. *Science* 327:584–87

Faisal AA, Selen LP, Wolpert DM. 2008. Noise in the nervous system. *Nat. Rev. Neurosci.* 9:292–303

Ghose GM, Harrison IT. 2009. Temporal precision of neuronal information in a rapid perceptual judgment. *J. Neurophysiol.* 101:1480–93

Greschner M, Shlens J, Bakolitsa C, Field GD, Gauthier JL, et al. 2011. Correlated firing among major ganglion cell types in primate retina. *J. Physiol.* 589:75–86

Gu Y, Angelaki DE, Deangelis GC. 2008. Neural correlates of multisensory cue integration in macaque MSTd. *Nat. Neurosci.* 11:1201–10

Gu Y, DeAngelis GC, Angelaki DE. 2007. A functional link between area MSTd and heading perception based on vestibular signals. *Nat. Neurosci.* 10:1038–47

Gu Y, Liu S, Fetsch CR, Yang Y, Fok S, et al. 2011. Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* 71:750–61

Gutnisky DA, Dragoi V. 2008. Adaptive coding of visual information in neural populations. *Nature* 452:220–24

Haefner RM, Gerwinn S, Macke JH, Bethge M. 2012. Inferring decoding strategy from choice probabilities in the presence of noise correlations. *Nat. Precedings.* **http://hdl.handle.net/10101/npre.2012.7014.1**

Herrington TM, Assad JA. 2009. Neural activity in the middle temporal area and lateral intraparietal area during endogenously cued shifts of attention. *J. Neurosci.* 29:14160–76

Herrington TM, Masse NY, Hachmeh KJ, Smith JE, Assad JA, Cook EP. 2009. The effect of microsaccades on the correlation between neural activity and behavior in middle temporal, ventral intraparietal, and lateral intraparietal areas. *J. Neurosci.* 29:5793–805

Huang X, Lisberger SG. 2009. Noise correlations in cortical area MT and their potential impact on trial-by-trial variation in the direction and speed of smooth-pursuit eye movements. *J. Neurophysiol.* 101:3012–30

Huk AC, Shadlen MN. 2005. Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* 25:10420–36

Janssen P, Shadlen MN. 2005. A representation of the hazard rate of elapsed time in macaque area LIP. *Nat. Neurosci.* 8:234–41

Jermakowicz WJ, Chen X, Khaytin I, Bonds AB, Casagrande VA. 2009. Relationship between spontaneous and evoked spike-time correlations in primate visual cortex. *J. Neurophysiol.* 101:2279–89

Kiani R, Hanks TD, Shadlen MN. 2008. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.* 28:3017–29

Kohn A, Smith MA. 2005. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci*. 25:3661–73

Krose BJ, Julesz B. 1989. The control and speed of shifts of attention. *Vis. Res*. 29:1607–19

Law CT, Gold JI. 2008. Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nat. Neurosci*. 11:505–13

Liu J, Newsome WT. 2005. Correlation between speed perception and neural activity in the middle temporal visual area. *J. Neurosci*. 25:711–22

Malach R, Schirman TD, Harel M, Tootell RB, Malonek D. 1997. Organization of intrinsic connections in owl monkey area MT. *Cereb. Cortex* 7:386–93

Mazurek ME, Roitman JD, Ditterich J, Shadlen MN. 2003. A role for neural integrators in perceptual decision making. *Cereb. Cortex* 13:1257–69

Mitchell JF, Sundberg KA, Reynolds JH. 2009. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63:879–88

Muller HJ, Rabbitt PM. 1989. Reflexive and voluntary orienting of visual attention: time course of activation and resistance to interruption. *J. Exp. Psychol. Hum. Percept. Perform*. 15:315–30

Muller MM, Teder-Salejarvi W, Hillyard SA. 1998. The time course of cortical facilitation during cued shifts of spatial attention. *Nat. Neurosci*. 1:631–34

Nakayama K, Mackeben M. 1989. Sustained and transient components of focal visual attention. *Vis. Res*. 29:1631–47

Nienborg H, Cumming BG. 2006. Macaque V2 neurons, but not V1 neurons, show choice-related activity. *J. Neurosci.* 26:9567–78

Nienborg H, Cumming BG. 2007. Psychophysically measured task strategy for disparity discrimination is reflected in V2 neurons. *Nat. Neurosci.* 10:1608–14

Nienborg H, Cumming BG. 2009. Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* 459:89–92

Nienborg H, Cumming B. 2010. Correlations between the activity of sensory neurons and behavior: How much do they tell us about a neuron's causality? *Curr. Opin. Neurobiol*. 20:376–81

Palmer C, Cheng SY, Seidemann E. 2007. Linking neuronal and behavioral performance in a reaction-time visual detection task. *J. Neurosci*. 27:8122–37

Palmer J, Huk AC, Shadlen MN. 2005. The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis*. 5:376–404

Parker AJ, Krug K, Cumming BG. 2002. Neuronal activity and its links with the perception of multi-stable figures. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 357:1053–62

Parker AJ, Newsome WT. 1998. Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci*. 21:227–77

Price NS, Born RT. 2010. Timescales of sensory- and decision-related activity in the middle temporal and medial superior temporal areas. *J. Neurosci*. 30:14036–45

Purushothaman G, Bradley DC. 2005. Neural population code for fine perceptual decisions in area MT. *Nat. Neurosci*. 8:99–106

Roitman JD, Shadlen MN. 2002. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci*. 22:9475–89

Romo R, Hernandez A, Zainos A, Lemus L, Brody CD. 2002. Neuronal correlates of decision-making in secondary somatosensory cortex. *Nat. Neurosci*. 5:1217–25

Salin PA, Bullier J. 1995. Corticocortical connections in the visual system: structure and function. *Physiol. Rev.* 75:107–54

Salinas E, Hernandez A, Zainos A, Romo R. 2000. Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *J. Neurosci.* 20:5503–15

Samonds JM, Potetz BR, Lee TS. 2009. Cooperative and competitive interactions facilitate stereo computations in macaque primary visual cortex. *J. Neurosci.* 29:15780–95

Sasaki R, Uka T. 2009. Dynamic readout of behaviorally relevant signals from area MT during task switching. *Neuron* 62:147–57

Shadlen MN, Britten KH, Newsome WT, Movshon JA. 1996. A computational analysis of the relationship between neuronal and behavioural responses to visual motion. *J. Neurosci.* 16:1486–510

Smith MA, Kohn A. 2008. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J. Neurosci.* 28:12591–603

Stanford TR, Shankar S, Massoglia DP, Costello MG, Salinas E. 2010. Perceptual decision making in less than 30 milliseconds. *Nat. Neurosci.* 13:379–85

Tolhurst DJ, Movshon JA, Dean AF. 1983. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vis. Res.* 23:775–85

Uka T, DeAngelis GC. 2004. Contribution of area MT to stereoscopic depth perception: choice-related response modulations reflect task strategy. *Neuron* 42:297–310

Uka T, DeAngelis GC. 2006. Linking neural representation to function in stereoscopic depth perception: roles of the middle temporal area in coarse versus fine disparity discrimination. *J. Neurosci.* 26:6791–802

Zohary E, Shadlen MN, Newsome WT. 1994. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–43

Annual Review of Neuroscience

Volume 35, 2012

Contents

Indexes

Errata

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