

Cracking the neural code for sensory perception by combining statistics, intervention and behavior

Stefano Panzeri^{1,2*+}, Christopher D. Harvey^{3*}, Eugenio Piasini¹, Peter E. Latham⁴, & Tommaso Fellin^{2,5*}

¹ Neural Computation Laboratory, Istituto Italiano di Tecnologia, Rovereto, Italy

² Neural Coding Laboratory, Istituto Italiano di Tecnologia, Rovereto, Italy

³ Department of Neurobiology, Harvard Medical School, Boston, MA, USA

⁴ Gatsby Computational Neuroscience Unit, University College London, London, UK

⁵ Optical Approaches to Brain Function Laboratory, Istituto Italiano di Tecnologia, Genoa, Italy

* Correspondence should be addressed to: stefano.panzeri@iit.it, Christopher_Harvey@hms.harvard.edu, tommaso.fellin@iit.it

+ Lead contact: stefano.panzeri@iit.it

Summary

The two basic processes underlying perceptual decisions – how neural responses encode stimuli, and how they inform behavioral choices – have mainly been studied separately. Thus, although many spatiotemporal features of neural population activity, or “neural codes,” have been shown to carry sensory information, it is often unknown whether the brain uses these features for perception. To address this issue, we propose a new framework centered on redefining the neural code as the neural features that carry sensory information used by the animal to drive appropriate behavior; that is, the features that have an intersection between sensory and choice information. We show how this framework leads to a new statistical analysis of neural activity recorded during behavior that can identify such neural codes, and we discuss how to combine intersection-based analysis of neural recordings with intervention on neural activity to determine definitively whether specific neural activity features are involved in a task.

Keywords: neural coding, choice, information, optogenetics, population coding, behavior

Introduction

To survive, organisms must both accurately represent stimuli in the outside world, and use that representation to generate beneficial behavioral actions. Historically, these two processes – the mapping from stimuli to neural responses, and the mapping from neural activity to behavior – have mainly been treated separately. Of the two, the former has received the most attention. Often referred to as the “neural coding problem,” its goal is to determine what features of neural activity carry information about external stimuli. This approach has led to many empirical and theoretical proposals about the spatial and temporal features of neural population activity, or “neural codes,” that represent sensory information (Buonomano and Maass, 2009; Harvey et al., 2012; Harvey et al., 2013; Kayser et al., 2009; Luczak et al., 2015; Panzeri et al., 2010; Shamir, 2014). However, there is still no consensus about the neural code for most sensory stimuli in most areas of the nervous system.

The lack of consensus arises in part because, while it is established that certain features of neural population responses carry information about specific stimuli, it is unclear whether the brain uses the information in these features to perform sensory perception (Engineer et al., 2008; Jacobs et al., 2009; Luna et al., 2005; Victor and Nirenberg, 2008). In principle, the link between sensory information that is present and sensory information that is read out to inform choices can be probed using the animal’s behavioral report of sensory stimuli. In addition, improvements in techniques to perturb activity of neural populations during behavior (Boyden et al., 2005; Deisseroth and Schnitzer, 2013; Emiliani et al., 2015; Tehovnik et al., 2006) now make it possible to test causally hypotheses about the neural code, by “writing” on the neural tissue putative information, and then measuring the behavior elicited by this manipulation. However, progress in cracking the neural code has been limited by the lack of a conceptual framework that fully integrates the advantages offered by behavioral, neurophysiological, statistical, and interventional techniques.

Here we elaborate such a conceptual framework, which at its core is based on a change in how a neural code should be defined. We propose that a neural code should be defined as the set of neural response features carrying sensory information that, crucially, is used by the animal to drive appropriate behavior; that is, the set of neural response features that have an intersection between sensory and choice information. In the following we discuss this framework and its implications for designing and interpreting experiments aimed at cracking the neural code, as well as some theoretical and experimental challenges that arise from it.

What it takes to crack the neural code underlying a sensory percept

To illustrate our new framework, we consider a perceptual discrimination task in which an animal has to extract information present in the sensory environment and, based on that information, choose an appropriate action. For definiteness, we assume (Fig. 1A) a two-alternatives forced-choice discrimination task: the animal has to extract color information from a visual stimulus (that is decide whether a green ($s=1$) or a blue ($s=2$) stimulus was presented) and decide whether to choose to move left (choice $c=1$) or right ($c=2$), with the correct choice resulting in a reward (we numbered choices so that $c=1$ is the correct rewarded choice for $s=1$ and $c=2$ is the correct choice for $s=2$). We suppose that an experimenter is recording the activity of a population of sensory neurons (visual neurons in this example) while the animal performs the task. We would like to determine whether the activity of these neurons contributes causally to the animal’s perception and behavioral choice.

The neural code in tasks such as this one involves two crucial stages in the information processing chain (Fig. 1A). The first stage is *sensory coding*: the mapping, on each trial, of the sensory stimulus to neural population activity. The second stage is *information readout*: the mapping from neural population activity to behavioral choice.

An important observation is that sensory coding and information readout can be based on distinct features: some features of neural activity used to encode sensory information may be ignored when information is read out, and vice versa. For example, suppose that, as is often found (Panzeri et al., 2010; Shriki et al., 2012; Shusterman et al., 2011; Victor, 2000; Zuo et al., 2015), the precise timing of spikes and the spike rate are both informative about the stimulus. Although there is information in spike timing, the downstream neural circuit may be sensitive only to the rate, and thus unable to use information contained in spike timing. Or, the downstream circuit may be sensitive to spike timing, but if extracting spike timing information requires an independent knowledge of the stimulus presentation time to which the downstream circuit does not have access, the readout will not be able to use spike timing information. In a less extreme case, the same features of neural activity may be used for both sensory coding and information readout, but they may be weighted differently by different sets of neurons. For example, the timing of spikes may carry more information about the stimulus than the number of spikes, but the spike rate may weigh more than spike timing in reaching a behavioral choice.

Although characterizing separately the features of neural activity used for sensory coding and for information readout can provide insight into neural information transmission, we argue that to crack the neural code it is essential to consider the *intersection* between sensory coding and information readout (Fig. 1A), defined as the set of neural features carrying sensory information that is read out to inform a behavioral choice. The only information that matters for task performance is the information at this intersection. In fact, only features that lie at this intersection can be used to convert sensory perception into appropriate behavioral actions and can help the animal perform a perceptual discrimination task. We therefore define the neural code that allows the animal to do the task to be the “intersection” features of neural activity carrying sensory information that is read out for behavioral choice.

In the following, we propose a framework for identifying the information at the intersection of sensory coding and information readout. We propose a combination of statistical approaches, behavior, and interventional manipulations (Fig. 1B). Statistical approaches can be used on single trials to identify the neural activity features that covary with the sensory stimuli and behavioral choices; they are, therefore, critical for forming hypotheses about the features of the neural activity that both contain sensory information and are used by the information readout. These hypotheses can be tested using experiments in which sensory stimuli are replaced with (or accompanied by) direct manipulation of neural population activity (Fig. 1B). The manipulation of the specific features of neural population activity that take part in sensory coding and the examination of how these manipulations affect the animal’s behavioral choices probe causally the intersection between sensory information and readout.

Examples of candidate neural codes

Before detailing the concepts behind this proposed framework, we first provide examples to illustrate the types of neural codes and questions that could be addressed. In all these

examples, we suppose that we record (either from the same brain location or from multiple locations) neural population activity. That activity consists of n neural features, denoted r_1, \dots, r_n . We would like to determine which of these features, either individually or jointly, carry sensory information that is essential for performing the perceptual discrimination task. For simplicity, hereafter we focus on two features, r_1 and r_2 , but our framework is general enough to deal with an arbitrary number of features (see Supplemental Information).

A common example of studies of population coding (Fig. 2A) considers as candidate neural codes two features, r_1 and r_2 , defined respectively as the total spike count of two populations of neurons. This has been the focus of many recent studies designed to test whether activity in specific neural populations is essential for accurate performance in sensory discrimination tasks (Chen et al., 2011; Guo et al., 2014; Hernandez et al., 2010; Peng et al., 2015). Those spike counts could be from spatially separated populations in two different brain regions, as shown in Fig. 2A, or they could be from different genetically or functionally defined cell classes in the same brain region (Baden et al., 2016; Chen et al., 2013; Li et al., 2015; Wilson et al., 2012). More sophisticated examples of features of neural population activity may involve low-dimensional projections of the activity of large neuronal populations (Cunningham and Yu, 2014). These could be, for example, the first two principal components of neural population activity, and would consist of weighted sums of spike counts of the recorded population (Fig. 2B). Major open questions that follow from these studies are (Otchy et al., 2015): which populations are instructive for the task (provide sensory information used for perceptual discrimination), which populations are permissive for the task (modulate task performance without directly contributing any specific sensory information), and which populations have no causal role in the sensory discrimination task despite having sensory information? The neural code is expected to be present in instructive populations. In contrast, permissive areas could provide task-relevant modulation that is not related to the sensory stimuli, such as attention or saliency signals. Populations with no causal role may still contain task-related information if it is inherited from instructive regions.

Other questions relevant for population coding regard which neurons are required for sensory information coding and perception (Houweling and Brecht, 2008; Huber et al., 2008; Reich et al., 2001). For example, often only a relatively small fraction of neurons in a population have sharp tuning profiles to the stimuli, whereas the majority of neurons have weak and/or mixed tuning to many different variables (Meister et al., 2013; Rigotti et al., 2013). Information about stimuli can be decoded from both types of neurons, but it remains a major open question whether only the sharply tuned neurons or other neurons as well can contribute to behavioral discrimination (Morcos and Harvey, 2016). A related question is: how many neurons are required for sensory perception? This question can be investigated by determining the smallest subpopulation of neurons that carries all information used for perception.

Another set of questions considers the role of spike timing in sensory coding and perception (Fig. 2C-D). Spike timing could be measured with respect to the stimulus presentation time, an internal brain rhythm (Kayser et al., 2009; O'Keefe and Recce, 1993), or a rhythmic active sampling process such as sniffing (Shusterman et al., 2011). In many cases both spike timing and spike count carry sensory information (in the example of Fig. 2C stimulus $s=1$ elicits responses with fewer and earlier spikes than does $s=2$). Although it is accepted that spike timing carries sensory information, whether or not timing is used for behavior has been vigorously debated (Engineer et al., 2008; Harvey et al., 2013; Jacobs et al., 2009; Luna et

al., 2005; Zuo et al., 2015). For example, it is still debated whether the sensory information carried by millisecond-scale spike timing is redundant with that provided by the total spike count in a longer response window of hundreds of milliseconds, whether the information in spike timing measured with respect to stimulus onset can be accessed by a downstream neural decoder, and whether recurrent circuits in higher cortical areas can extract millisecond-scale information.

Also of interest is whether the complex aspects of the temporal structure of spike trains could be part of the neural code. One possibility is that the regularity of spike timing of single neurons or the coordination of spike timing across cells carries information about the stimulus (as in the example of Fig. 2D, where stimulus 1 elicits more regular spike trains than stimulus 2). The regularity or temporal coordination across cells may also have a large effect on the readout (Doron et al., 2014; Jia et al., 2013; Nikolic et al., 2013): for example, spikes closer in time may elicit a larger post-synaptic response and so may have a crucial impact on task performance. However, some studies have suggested that temporal coordination does not have a behavioral effect, but instead all spikes are weighted the same by the readout (Histed and Maunsell, 2014).

In what follows we will consider, for simplicity, two features, and we will generically refer to these features as r_1 and r_2 . These features could refer to spike timing and spike counts, the mean firing rate in two different brain regions, the activity of two different cell types, and so on.

Determining the single-trial intersection between sensory information coding and information readout using statistical analysis of neural recordings

We now consider how to identify, using statistical measures applied to recordings of neural activity during behavior experiments, three conceptually important domains of interest in the neural information space (Fig. 1A): the “sensory information” domain (the features of neural population activity that carry stimulus information), the “readout” domain (the features that influence the computation of choice), and the “intersection” between the two domains (the features that carry sensory information used to compute choice).

Throughout this article, to illustrate neural coding and stimulus and choice domains, we use scatterplots of simulated responses characterized by two features; each dot in the two-dimensional feature plane (r_1, r_2) represents a single-trial response color-coded for that trial’s stimulus ($s=1$: green; $s=2$, blue) (see also Fig. 2A, right panel for a schematization of this representation). Each dot therefore shows the simulated neural response of feature r_1 and r_2 on each individual trial.

A simple way to visualize how neural response features encode sensory stimuli is to compute a sensory decoding boundary (Quiñ Quiroga and Panzeri, 2009) – shortened to “*sensory boundary*” hereafter – that best separates trials by stimulus (i.e., that best separates the blue and green dots in the plots in Fig. 3). This boundary (black dashed line in the r_1, r_2 plane in Fig. 3A₁, B₁, C₁) can be used as a rule to decide which stimulus most likely caused a given single-trial neural response. Similarly, we can visualize how neural response features are used to produce a choice with a “*decision boundary*” (Haefner et al., 2013), visualized as red dashed line in Fig. 3A₁, B₁, C₁. In these figures, this decision boundary coincides with the line that best separates trials by choice. Responses that lead to correct choices are shown as filled

dots; those leading to incorrect choice are shown as open circles. The orientation of the boundaries determines the relative importance of each feature in sensory coding or choice: a diagonal boundary gives weight to both features, whereas a horizontal or vertical line gives weight only to r_2 or only to r_1 , respectively.

To quantify how well each feature or set of features carries information about stimulus or choice, we use the fraction correct. In terms of the illustrations of Fig. 3, the fraction of correctly decoded stimuli is the fraction of green or blue dots that fall on the correct side of the sensory boundary (below or above the sensory boundary for the green, $s=1$, and blue, $s=2$, stimulus, respectively). Other measures, such as those based on signal detection theory (Britten et al., 1996; Shadlen et al., 1996) or information theory (Quian Quiroga and Panzeri, 2009) can be used instead, and are discussed in Supplemental Information. We use fraction correct primarily because it is simple and intuitive, but we could use any of the other measures without changing the basic framework. To emphasize the generality of our reasoning, hereafter we often refer to fraction correct as “information”. If the fraction correct refers to the decoded stimulus we call it “sensory information” or “stimulus information”; if it refers to decoded choice we call it “choice information”.

We say that a neural response feature, r_i , carries sensory or choice information if the value of the presented stimulus or the animal’s choice can be predicted from the single-trial values of this feature. Stimulus information and sensory boundaries are typically computed presenting two or more different stimuli, and quantifying how well the stimulus-specific distributions of neural response features are separated by sensory boundaries (Quian Quiroga and Panzeri, 2009). Choice information has been typically computed separately from stimulus information (Britten et al., 1996), by evaluating decision boundaries from distributions of responses with no sensory signal or at fixed sensory stimulus (to eliminate spurious choice variations of neural response arising from their stimulus-related variations).

To understand the neural code associated with a particular task, it is relatively obvious that we need to consider both stimulus and choice information. If a response feature carries stimulus but not choice information, then the sensory information it carries isn’t used for the task. If a response feature carries choice but not stimulus information, then although it may contribute to choice, or relay or execute the result of the decision making, it still cannot be used *per se* to increase task performance because it does not carry information about the sensory variable to be discriminated (Koulakov et al., 2005). However, a fact that has been underappreciated so far is that a neural feature can carry both sensory and choice information, but still not contribute to task performance. This could happen, for example, when features carry both stimulus and choice information, but the rule used to encode sensory information is incompatible with the rule used to read them out.

We illustrate this in Fig. 3A. Suppose that in this figure, r_1 and r_2 are the times of the first spike of two different neurons. These features are signal-correlated (Averbeck et al., 2006); that is, both neurons spike earlier (corresponding to smaller values of both r_1 and r_2 in the scatterplot in Fig. 3A) to the green stimulus ($s=1$) and later (corresponding to larger values of r_1 and r_2 in the scatterplot) to the blue stimulus ($s=2$), with no “noise” correlations (Averbeck et al., 2006) between the activity of these neurons at fixed stimulus. For this encoding scheme, higher values of r_1+r_2 indicate that the blue stimulus is more likely, and so the sensory boundary is anti-diagonal: it is the line $r_1+r_2 = \text{constant}$. Suppose, though, that the readout does not have access to the stimulus time. In such a case, the only information the readout can use is the relative time of firing between the two neurons. This is the difference,

$r_1 - r_2$, and so the decision boundary is $r_1 - r_2 = \text{constant}$. In this case the responses carry information about both stimulus and the choice, but the responses cannot be used to perform the task – the orthogonal sensory and decision boundaries mean the animal’s choice is unrelated to the stimulus.

The case illustrated in Fig. 3A could happen also in studying neural population coding rather than spike timing. For example, r_1 and r_2 could be weighted sums of activity of neurons within a large population (as in Fig. 2B), with stimulus encoded by the sum of the two neural features and choice by the most active of the two features (which feature is most active is revealed by the sign of $r_1 - r_2$). Also in this “population” interpretation of Fig. 3A, none of the stimulus information in population activity could be used to perform the task.

Investigating whether neural stimulus information is usefully read out for task performance requires quantifying whether neural discrimination predicts behavioral discrimination. This has traditionally been addressed by evaluating the similarity between neurometric functions (quantifying the trial-averaged performance in discriminating various pairs of stimuli using one or more neural response features) and psychometric functions (quantifying the animal’s trial-averaged performance in discriminating the same set of pairs of stimuli). If a set of response features contributes to task performance, psychometric and neurometric functions should be similar (stimulus pairs discriminations that are easier for the animal should also be easier for the considered response features, and so on). This approach has provided numerous insights in sensory coding across several modalities (Engineer et al., 2008; Newsome et al., 1989; Romo and Salinas, 2003). For example, it was used to study the role of spike rates and spike times of somatosensory neurons for tactile perception of low-frequency (8-16 Hz) skin vibrations (Romo and Salinas, 2003). Although most such neurons encoded vibration frequency by spike rate, some neurons encoded it by spike times fired in phase with skin deflections. However, the neurometric performance of spike rates correlated better to the psychometric one than that of spike times, suggesting that spike rates produce this sensation (Romo and Salinas, 2003). A similar approach applied to high frequency vibrations (>100 Hz) suggested that discriminating high frequency vibrations relies on both spike times and rates (Harvey et al., 2013).

A potential problem with comparing neurometric and psychometric functions is that these functions may be similar even when the sensory and choice information do not intersect at all. The reason why this may happen is that it is based on comparing trial-averaged quantities, rather than comparing sensory information and animal’s choice in single trials. To understand the possible problems of only comparing neurometric and psychometric functions, consider a new scenario (Fig. S1A). The scenario is in part similar to that of Fig. 3A: r_1, r_2 are again the first spike times of two different neurons, and they are tuned to the stimuli and contribute to choice; and, as in Fig 3A, in this new example of Fig S1A both neurons spike earlier to the green stimulus ($s=1$) and later to the blue one ($s=2$), leading to an anti-diagonal sensory boundary (line $r_1 + r_2 = \text{constant}$), and the readout uses the relative time of firing $r_1 - r_2$ between the two neurons (the decision boundary projected on the r_1, r_2 plane is $r_1 - r_2 = \text{constant}$). However, suppose that in the example of Fig. S1A the actual choice (unlike in Fig. 3A) depends also on a third neural feature, r_3 , which we’ll take to be the sum of the firing rate of the two neurons. Assume also that, crucially, the stimulus dependence of the firing rate r_3 is similar to that of both r_1 and r_2 , so that stimulus $s=1$ elicits both earlier spike and lower rates than stimulus $s=2$ does. Suppose finally that the experimenter now tunes the task difficulty by varying some “stimulus signal intensity” parameter whose effect on neural firing is to change the separation between the clouds of the $s=1$ (green) and $s=2$ (blue) stimulus-specific

responses (Supplementary Fig. S1A₂). As the task becomes more difficult the animal's psychometric performance decreases, as does the decoding neurometric performance (because the blue and green stimulus-specific distributions of points get closer). We can plot neurometric and psychometric performance as a function of signal intensity, and they will have similar shape: both will be near chance when signal intensity is small and the stimulus-specific distributions of r_1, r_2 largely overlap (Fig S1A₃), and will be nearly perfect when signal intensity is high and the stimulus-specific distributions of r_1, r_2 are far apart (Fig. S1A₅). Thus, in this example (Fig S1A), statistical analysis will show that spike timing features r_1, r_2 have sensory information (because r_1+r_2 is stimulus-dependent), have choice information (even at fixed stimulus, reflecting that r_1-r_2 impacts on choice), and the neurometric function of r_1, r_2 is similar to the psychometric function (because the stimulus dependence of both r_1 and r_2 is similar to that of the firing rate r_3 which is the only contributor to task performance). Yet r_1, r_2 do not contribute to task performance because none of the sensory information they carry is read due to the orthogonality of the sensory and decision boundaries. That r_1, r_2 do not contribute to task performance can only be discovered observing that the trial-to-trial fluctuations of the accuracy of sensory information in r_1, r_2 , encoded only by r_1+r_2 , does not influence at all behavior, as the decision depends on r_1-r_2 . For example, in trials when r_1+r_2 indicates the presence of a stimulus different from that presented, behavior is not less (or more) likely to be correct because of this stimulus coding error (Fig. S1A₃₋₅).

These examples illustrate a general fact: it is not possible to determine if sensory information is transmitted to the readout using the trial-averaged stimulus and choice information, either separately or in combination. It is, instead, necessary to investigate the effect of sensory coding on information readout within a single trial. We therefore propose the use of a measure we call intersection information, denoted I . Conceptually, intersection information is large only if the neural features carry a large amount of information about the stimulus, and that information is used to inform choice – so that, based on these features, the animal is correct most of the time.

A quantitative description of I was derived in (Zuo et al., 2015). The authors reasoned that, if feature r_i contributes to task performance, there should be an association on each trial between the accuracy of sensory information provided through that feature and behavioral choice. In other words, on trials in which r_i provides accurate sensory evidence (stimulus is decoded correctly from r_i), then the likelihood of correct choice should increase. Thus, the simplest operational definition of the intersection information, I , for a particular feature is the probability that on a single trial the stimulus is decoded correctly from r_i and the animal makes the correct choice (see Supplemental Information for additional details, in particular Eq. (S7)).

Intersection information can be used to rank features according to their potential importance for task performance. Importantly, it is high if there is a large amount of stimulus information and readout is near-optimal. It is low, on the other hand, if a neural response feature has only sensory information but very little choice information, or vice versa, or if the rule used for sensory coding is incompatible with the rule used by the readout.

We illustrate intersection information using three examples (Figs. 3A₁, 3B₁, and 3C₁), with null (chance-level), intermediate and high values of intersection information, respectively. In these plots, we divide the r_1, r_2 feature space into four possible areas based on the sensory and decision boundaries: $\hat{s} = 1, c = 1$; $\hat{s} = 1, c = 2$; $\hat{s} = 2, c = 1$; $\hat{s} = 2, c = 2$ (\hat{s} is the decoded

stimulus, which can be different from the stimulus, s , presented to the animal). The intersection information is the fraction of trials that are decoded correctly and result in a correct behavioral choice; these trials correspond to the filled dots, indicating a trial with correct choice, shown in the regions in Figs. 3A₂, B₂, C₂ colored with the decoded stimulus color code. The larger are the colored areas, the larger is the intersection information. Chance level for the intersection measure is when there is no relationship between the stimulus decoded by neural activity and the choice taken by the animal at fixed stimulus (the chance level of intersection equals the product of the probability of a correct behavioral choice and the probability of correctly decoding the stimulus; see Supplemental Information for details). This is the case in Fig. 3A, where the sensory and decision boundaries are orthogonal. Because trials that provide faithful stimulus information are just as likely to result in correct as incorrect choices, there is chance-level intersection information (see Fig S1D for the intersection information values in these examples).

Intersection information is intermediate when only some of the features of neural activity carrying sensory information are read out while the information of other is lost before the readout stage. This is the case in Fig. 3B, when both r_1 and r_2 carry sensory information but only r_2 is read out. This may correspond, for example, to a case when both spike rate r_1 and spike timing r_2 of a neuron carry information, but only rate r_1 is used for behavior (similarly to the case of (O'Connor et al., 2013)), for example because the readout mechanism is not sensitive enough to precise spike timing.

Intersection information is largest when the optimal sensory boundary and the decision boundary coincide, as in Fig. 3C, so that all sensory information is optimally used to perform the task. This is the case when all measured features of neural activity that carry sensory information directly contribute to the animal's choice. In this example (Fig. 3C), trials that lead to correct stimulus decoding from the joint features, r_1 and r_2 (those below the diagonal for the green stimulus, $s=1$, and those above the diagonal for the blue stimulus, $s=2$), always lead to correct behavioral choices. Trials leading to incorrect stimulus decoding from r_1 and r_2 (above the diagonal for $s=1$ and below it for $s=2$) always lead to incorrect behavioral choices. This situation is reminiscent of texture encoding by somatosensory cortical neurons (Zuo et al., 2015), in which both spike rate and timing seem to carry sensory information that is used for behavioral discrimination.

The above simple reasoning can be extended to provide more refined measurements of the relationship between sensory information in neural activity and behavioral choice. For example, one could also measure (Zuo et al., 2015) what we call the fraction of intersection information fII , defined as the fraction of trials with correct stimulus decoding that have correct behavior. Unlike II , fII is not sensitive to the amount of sensory information (the fraction of trials the stimulus is decoded correctly from neural feature r), but only to the proportion of these correctly decoded trials that lead to correct behavior. Thus, fII is an indicator of the optimality of the readout – in the linear case, the alignment between the sensory and decision boundaries – rather than the total impact of the code on task performance. Measuring both II and fII could be useful to determine whether a moderate amount of intersection information, II , is because the feature has a moderate amount of information but is efficiently read or because the feature has high information but not read out very efficiently. Moreover, given that if a feature r_i contributes to task performance, then, in trials when r_i provides inaccurate evidence (stimulus is decoded incorrectly), the likelihood of correct choice should decrease, an additional separate quantification of the

agreement of stimulus information and behavioral choice in incorrect trials would complement intersection measures (see (Zuo et al., 2015) and Supplemental Information).

The purely statistical approach to measure intersection information is most straightforward if all response features are statistically independent, because in that case the intersection information approach applied to set of features would unambiguously identify the contribution of those features to task performance. However, often features are not independent. For example, if the features are the activity of neurons in different brain regions, these features might be partly correlated if there are connections between the two regions. Alternatively, if the features are spike timing and spike count, both involve the same spikes and so may be dependent. The presence of dependencies among features complicates the interpretation of intersection information. In particular, it raises two critical questions: First, does a set of features with intersection information contribute to task performance, or instead reflect only a correlation with other features that truly contributes to task performance? Second, does each neural response feature provide unique intersection information that is not provided by other features?

To illustrate the complications induced by correlations among features, we also consider the intersection information from one feature rather than two. We return to Fig. 3, for which responses are signal-correlated in all panels (that is, responses to $s=1$ are on average lower than those to $s=2$ for both features (Averbeck et al., 2006)). We first consider a case (as in Fig. 3B) for which both features carry information about the stimulus and are partly correlated (because of signal correlations) but only feature r_1 is read out. Suppose that we apply our statistical analysis to feature r_2 ; that is, we decode the stimulus using only r_2 , which can be done optimally by decoding responses in the lower and upper half of the r_1, r_2 space as $\hat{s} = 1$ and $\hat{s} = 2$ respectively. We will find higher-than-chance intersection and choice information (as shown by the fact that lower values of r_2 are found in trials with choice $c=1$ than in trials with $c=2$, see Fig. S1B) even though r_2 is not read out. That's only because r_2 is correlated with r_1 , which is the feature that is truly read out.

If we record from both features, we can differentiate, just from statistical analyses of neural recordings, between the case when only one feature is read out (Fig. 3B) and the case when both are read out (Fig. 3C). If, as in Fig. 3C, r_1 and r_2 carry complementary sensory information (the diagonal sensory boundary implies that both features should be used for optimal decoding) and if the readout uses both features (the decision boundary is also diagonal), then intersection information obtained when decoding the stimulus using two features will be larger than the intersection information obtained when decoding the stimulus with either feature alone. This is because decoding the stimulus with only one feature will lose the complementary task information present in the other feature and so the task performance will suffer (Fig. S1C-D). Thus, a statistical signature that task performance benefits from both features is that using both feature increases the intersection information (Fig. S1D). However, if we cannot record both features (and, more generally, all features that carry intersection information), the only way to fully prove which features contribute to task performance is to use interventional methods. That's the subject of the next section.

Causal interventional testing of the neural code

Why do we need intervention?

The statistical methods described above for determining sensory, choice and intersection information are useful for identifying potential neural coding mechanisms, and for forming hypotheses about information coding and transmission. However, as just discussed, because response features are often correlated and because we do not usually have experimental access to all of them, whether a neural feature carries information in the intersection between sensory coding and readout can ultimately only be proved with intervention. Before discussing how to design interventional experiments that test intersection information, it is useful to consider why interventional manipulations of neural activity are so crucial to prove hypotheses. (In the following, we refer to “statistical” information measures as shorthand for measures of information obtained from recorded natural unperturbed neural activity, and “interventional” information measures to indicate information estimates from neural activity imposed by intervention).

Suppose that statistical measures like those described in the previous section found that a neural feature, r_1 , carries stimulus, choice and intersection information. An interpretation of this result is that r_1 provides essential stimulus information to the decision readout (this is indicated in Fig. 4A by the arrow from r_1 to the choice, c). However, another interpretation (sketched in Fig. 4B), one that is still compatible with these statistical measures, is that r_1 *does not* transmit information to c (not even indirectly). Instead, it only receives a copy of the information that other neural features (such as r_2 in Fig. 4B) do transmit to the choice. In this case, the sensory information in r_1 is not causally involved in the decision (as indicated by the lack of arrow from r_1 to c in Fig. 4B), but r_1 correlates with the decision because it correlates with r_2 , the decision’s cause.

Intervention can disambiguate these two scenarios by imposing a chosen value on r_1 , one that is decided by the experimenter and so is independent of r_2 . By doing that, we break any possible effect of r_2 , or of any other possible variable, on r_1 (Fig. 4E). In this case, any observed relationship between r_1 and choice must be due to the causal effect of r_1 on choice (Pearl, 2009).

In the following we discuss how to design a causal intervention experiment that tests if neural features carry intersection information. We are interested in an intervention design that can tell whether r_1 transmits stimulus information used for decision (as in Fig. 4F, indicated by the arrow between r_1 and choice c being colored like a stimulus) or r_1 informs c but does not transmit stimulus information contributing to task performance (as in Fig. 4E, indicated by the arrow between r_1 and c not being colored).

Intervention on neural activity and intersection between sensory information and readout

Here we examine cases in which we can both record and manipulate (in the same animal, but not necessarily at the same time) neural features r_1 and r_2 during a perceptual discrimination task.

Let us first consider a causal intervention on the neural features. Suppose that we impose a number of different values of r_1, r_2 in a series of intervention trials (“lightning bolt” symbols in Fig. 5, colored by the behavioral choice they elicit) and we measure the choice taken by the animal. In our examples, choice is determined by the red dashed decision boundary in the r_1, r_2 space. Observing the correspondence between the value imposed on r_1, r_2 and the

animal's choice would easily determine the orientation of the decision boundary (Fig 5A-C). From this interventionally determined decision boundary, choice information can be obtained exactly as in the statistical case.

Applying the same reasoning used for the statistical case, an interventional measure of intersection information is the fraction of trials on which the animal's choice reports the stimulus that would be decoded (using the sensory boundary acquired with statistical analysis of neural responses) from the imposed neural activity pattern (as above, this can be assessed against chance level). Application of this interventional measure of intersection information to our examples in Fig. 5 shows that interventional intersection information captures the alignment of the sensory and decision boundary. It is high when, as in Fig. 5C, the animal choice ($c=1$, pink; $c=2$, dark red) always corresponds to the stimulus decoded from neural activity (in Fig. 5C, the case of maximal intersection, all patterns in the $\hat{s}=1$ "green" decoding region lead to $c=1$, and the same applies to the $\hat{s}=2, c=2$ region); it is null (chance-level) when sensory and decision boundaries are mismatched (as in Fig. 5A, where half of imposed patterns in either stimulus decoding region lead to choice $c=1$ and half to $c=2$).

A critical observation is that the intersection information computed via intervention may be different from that computed using purely statistical analysis. That can happen, as discussed above, if the neural features are correlated with variables that did carry intersection information, but did not themselves provide any information about choice. For example, if a statistically determined non-null choice or intersection information in one feature just reflects a top down choice signal and not a causal contribution of the feature to choice, this feature will show null (chance-level) intersection information with intervention. Thus, of particular interest are cases in which the decision boundary is orthogonal to the sensory boundary under intervention, but not in trials without an intervention (O'Connor et al., 2013). When that's the case, the neural features under investigation carry no intersection information and the intersection or choice information determined statistically mean that the considered features only correlate with the true factors that are instructive for task performance.

The values of the interventionally-evoked neural features in an experiment are arbitrarily determined by the experimenter. The chosen evoked neural features may be designed to drive behavior robustly, but may occur very rarely during perception in natural conditions. This may lead to an over-estimation of their importance for task performance. To correct for this problem, when computing interventional intersection information, we should weigh intervention results with the probability distribution over stimuli and responses that occur under natural conditions (see Supplemental Information). Thus, evaluating the causal impact of a neural code with intervention experiments ultimately demands a statistical analysis of the probability of naturally occurring patterns during the presentation of each stimulus while performing the task.

By analogy to what we proposed for the statistical measures, we can design intervention experiments that address whether two neural response features, r_1 , r_2 , that are correlated during measures of natural neural activity both contribute causally to choice and to task performance. Suppose that (as in in Fig. 5B,C) we recorded two correlated features in unperturbed (i.e., no intervention) conditions and that we would like to determine interventionally if the readout uses both such sources of sensory information to perform the task. Designing such experiment requires manipulating both features at the same time, and then comparing interventional intersection information of the joint features and of the individual ones. If the experimenter designs a set of intervention patterns that generate

uncorrelated feature values, then only the features that carry sensory information and are read out will show higher-than-chance intersection information. If the experimental design cannot fully decorrelate the features evoked by intervention, then a complementary contribution to task performance of the two features will still be revealed interventionally when finding that adding a feature increases the interventional intersection information, exactly as in the statistical analysis case.

Above we argued that statistical analysis is not sufficient to determine whether there is intersection information in a set of neural features. In the following, we argue that causal manipulations of neural activity alone are also not sufficient to determine whether there is intersection information in a set of neural features. Experiments are frequently designed such that an animal is trained to discriminate neural activity patterns that are created artificially using interventional approaches, such as microstimulation or optogenetics, without direct regard to how these patterns may encode sensory stimuli. This approach is extremely powerful for testing the capabilities of the readout. For example, this approach has been used to test the sensitivity of the readout to precisely timed neural activity (Doron et al., 2014; Yang et al., 2008; Yang and Zador, 2012) and also to test the minimal number of neurons or spikes to which a readout could be sensitive (Houweling and Brecht, 2008; Huber et al., 2008). Thus this approach can be used to infer intersection information indirectly (by comparing the features that carry sensory information with those that can be detected by the readout). However, this approach is insufficient to determine directly if a given neural feature is used for performing specific sensory discrimination tasks, and to evaluate how much this feature contributes to sensory discrimination.

Because both statistical and interventional approaches by themselves are not sufficient to test neural codes, we propose a scheme in which statistical analyses must be used to generate hypotheses about neural codes, and interventional experiments must be used to test them.

Neurophysiological examples of the potential advantages of application of the statistical and interventional concept of intersection

The foundations underlying intersection between sensory information and readout can be traced back to the work of Newsome and colleagues on visual motion perception in primates (Britten et al., 1996; Newsome et al., 1989). These studies showed that visual area MT encoded visual motion information in its firing rate: higher firing rates indicate motion along the neuron's preferred direction. They established a statistical relationship between the animal's choice in a visual motion discrimination task and the firing rate of MT neurons in the same trial (Britten et al., 1996). The causal role of the firing rate of MT neurons in motion perception was interventionally demonstrated showing that microstimulation of this region biases perception of motion direction (Newsome et al., 1989). Such studies continue today, taking also advantage of modern genetic, optogenetic and recording techniques. One example is the study of the neural coding of sweet and bitter taste in mice. The authors first established that anatomically separate populations of neurons responded to sweet and bitter taste, and thus carried stimulus information (Chen et al., 2011). An optogenetic intervention was then used to activate the spatially separated 'sweet' and 'bitter' populations (Peng et al., 2015). These intervention experiments elicited behavioral responses as expected for a mouse's response to sweet and bitter tastes. These studies therefore reveal intersection information in neural codes as spatially segregated response patterns using a combination of stimulus information, statistical analysis, and intervention.

These studies investigated simple properties of an individual neural feature (firing rate of classes of neurons) and followed implicitly part of the logic of the framework proposed here, although they did not measure a single-trial statistical intersection that we propose. Measuring the intersection information becomes, however, crucial in more complex scenarios in which (unlike the cases considered above) either a clear hypothesis about the neural code does not exist *a priori* (as may happen when analyzing coding of complex natural stimuli, rather than simpler laboratory stimuli) or when there are multiple, perhaps partly correlated, candidate features for the neural code that all seem, statistically, to contribute to choice or stimulus. In these cases, it is necessary to evaluate quantitatively the contribution of each feature to behavior. Below we discuss how the full or partial application of the ideas of our intersection framework in these more complex scenarios could provide further insight into the neural code.

The statistical intersection information framework has been applied to investigate whether millisecond-scale spike timing of somatosensory cortical neurons provides information that is used for performing a whisker-based texture discrimination task (Fig. 6A-B), over and above that already carried by spike counts over time scales of tens of milliseconds (Zuo et al., 2015). The authors computed a spike-timing feature by projecting the single-trial spike train onto a timing template (constructed for optimal sensory discrimination) whose shape indicated the weight assigned to each spike depending on its timing (Fig. 6C). Computed spike counts corresponded to weighting the spikes with a flat template which assigns the same weight to spikes independent of their time. This provided timing and count features that had negligible correlation (the temporal distribution of spikes was largely independent from their total number). Both timing and count carried significant sensory (Fig. 6D), choice (Fig. 6E) and intersection (Fig. 6F) information, with timing carrying more information than count for all these types of information. The joint information was for all types of information larger than the one carried by either feature alone. These results indicate that in this task sensory information was complementarily multiplexed in spike counts and timing, and was also complementarily combined to perform the task. Of the two features, however timing carried both more sensory information and had a greater influence on the animal's choice. Thus, the statistical intersection framework helped form a very precise hypothesis that multiplexing spike timing and spike count information is the key neural code used to solve the task. A further application of the interventional intersection framework, not yet applied to this experiment, would strongly prove or disprove this multiplexing hypothesis for texture coding.

This example illustrates that the statistical analysis of information intersection may be critical to correctly interpret the results of an interventional experiment and to refine its design. In this case, profound texture-dependent spike timing differences were found even across nearby neurons (Zuo et al., 2015). The cellular-level and millisecond-scale temporal resolution of this information coding revealed by the statistical analysis strongly constrains the interventional experimental design, as it indicates that finely spatially patterned and temporally precise intervention must be used to test whether spike timing is part of the neural code. Also, this example shows how statistical intersection results are essential to interpret successes and failures of interventions. For example, in the presence of such profound neuron-to-neuron differences in spike timing responses to textures, a causal effect of spike timing on behavior would not have been detected using a wide-field optogenetic intervention that activated all neurons simultaneously (see also section "Considerations of interventional experimental design"). Statistical analysis would be essential to reveal that this failure would not have been because spike timing was not part of the neural code used to perform the task,

but because the optogenetically induced activity did not preserve the natural texture-dependent timing differences across neurons.

A study (O'Connor et al., 2013) that implemented an approach close in spirit to the intersection information framework both at the statistical-analysis and interventional level is a recent investigation of the role of spike timing and spike rate coding in whisker-based detection task of object location (Fig. 7A). The authors found, based on statistical measures, that both timing and rate carried both stimulus and choice information (Fig. 7B). The authors then probed the role of timing and rate by replacing the somatosensory object with optogenetic manipulation of layer 4 somatosensory neurons (Fig. 7C,D). The authors found that using optogenetics to induce neural activity with information in rate caused the animal to report the sensation of a “virtual pole” (Fig. 7C), whereas adding to this optogenetic manipulation information in spike timing relative to whisking did not elicit additional behavioral performance in virtual sensation (Fig. 7D). When interpreted within our framework, these results suggest that spike rates, but not timing, carry intersection information. An additional application of the statistical intersection framework to these neurophysiological recordings – not performed in that study – would allow a more precise evaluation of the impact of timing and rate codes on task performance (see previous section), and could provide an important independent confirmation of this hypothesis based on naturally evoked neural activity only.

Considerations of interventional experimental design

Interventional approaches may involve use of one or more experimental techniques such as optogenetic (Lerner et al., 2016) and chemogenetic (Sternson and Roth, 2014) manipulations, intra-parenchymal electrical stimulation (Tehovnik et al., 2006), transcranial direct current stimulation and transcranial magnetic stimulation (Woods et al., 2016), to name a few. Given its unique combination of high cell-type specificity and temporal resolution, below we focus mostly on optogenetics.

There are at least two dimensions over which experimental design may be varied. One is how intervention is coupled with sensory stimuli; the other is how intervention is performed. In the following we consider how these possible experimental variations along these dimensions relate to the intersection framework.

Virtual sensation interventional experiments vs experiments overriding or biasing natural sensory signals

Our framework assumes that we test sensory encoding and information readout using a perceptual discrimination task. An important experimental design question is how to incorporate interventional approaches. Our focus is on understanding the codes that arise from natural sensory cues, and so we mainly consider cases in which interventional trials are interleaved with non-interventional ones.

One practical question for experimental design is whether on intervention trials the sensory stimulus should also be presented, or if the intervention manipulation should be applied in isolation. One possibility is a “virtual sensation” experiment (Fig. 8A), in which patterns of neural activity are imposed by intervention in the absence of the sensory stimulus and the

animal is asked to report the perception of one of the two sensory stimuli. A classic example is the work of Romo and colleagues (Romo et al., 1998; Romo and Salinas, 2003) demonstrating that cortical microstimulation can entirely substitute for tactile stimulation in a frequency discrimination task. Another example of virtual sensation is the induction of an illusory sensation of pole touching during whisking using optogenetic stimulation of cortical primary somatosensory neurons, as discussed above (see Fig. 7 and (O'Connor et al., 2013)). The virtual sensation paradigm is very appealing because it can demonstrate the sufficiency of the considered neural code for creating sensation and for its direct relevance for the development of neural prosthetics.

Another possibility is to impose patterns in the presence of a sensory stimulus. This approach tests whether the imposed pattern can “override” or “bias” (Fig. 8B) the signal from the sensory stimulus. A classic example of this approach can be found in the work of Newsome and colleagues (Salzman et al., 1990) showing that MT microstimulation in a visual motion discrimination task can bias the animal’s perception toward the motion direction preferred by the neurons that were activated by microstimulation. A more recent example can be found in the study (also described above) examining the codes for sweet and bitter/salt taste sensation (Peng et al., 2015), where the authors showed that optogenetic activation of the sweet cortical field triggered fictive sweet sensation even in the presence of a salt stimulus. From the point of view of the formalism presented here, successfully overriding the signal from an opposite external stimulus is an appealing proof that the considered neural code provides information that is so crucial to the task that can even win over other contrasting sources of information, such as those that may come from different or parallel pathways conveying information from the sensory periphery that contradicts the one injected through intervention of neural activity.

Considerations on how to perform intervention on neural activity, and the advantages of patterned optogenetics

Imposing a pattern can be done in two conceptually distinct ways. In one, the experimenter mainly tries to “bias” (Guo et al., 2014; Li et al., 2015; O'Connor et al., 2013) the neural activity (Fig. 8B). This consists of shifting the endogenous activity in a certain direction (for example, lowering the firing rates of a neural population by imposing a slight hyperpolarization or by exciting a set of inhibitory neurons). This can be done, for example, using wide-field, single-photon optogenetic stimulation of a network of a number of opsin-expressing neurons (this is illustrated in Fig. 8C, note that the number of neurons in that sketch is limited to 7 for presentation purposes only). A problem with this approach is that it does not completely remove correlations of the patterns evoked by intervention with other brain variables that are present in the endogenous component of the activity (because the evoked activity adds to the endogenous one). This means that this intervention may not entirely break the correlations among features or between features and non-observed endogenous brain activity that the causal manipulations aim to remove. This is a concern particularly when investigating if intersection information is complementarily carried by more than one feature, as an interventional bias may affect all features in a correlated way. For example, a general hyperpolarization of the population may both lower the spike rate and delay the latency of neural activity. Given the highly synchronous generation of photocurrents in opsin expressing cells, wide-field optogenetics may even induce artificial correlations (Fig. 8C).

The second interventional approach is to try to impose, or “write down” (Peron and Svoboda, 2011) a target neural activity pattern on a neural population (Fig. 8D). This approach is, in principle, ideally suited to test hypotheses about the neural code, because it explicitly aims to overwrite endogenous activity, and so break down all sources of correlation. To crack a neural code, though, it needs to achieve high spatial and temporal precision. Recent optical developments (Bovetti and Fellin, 2015; Emiliani et al., 2015; Grosenick et al., 2015), termed patterned illumination, can deliver light to precise spatial locations (Fig. 8D, see also Supplemental Information). When combined with light-sensitive optogenetic actuators, patterned illumination can perturb electrical activity with near cellular resolution (Baker et al., 2016; Carrillo-Reid et al., 2016; Packer et al., 2015; Papagiakoumou et al., 2010; Rickgauer et al., 2014).

Taking full advantage of the intersection framework will depend crucially on further development of improved optogenetics methods to ‘write’ neural activity patterns. Current technologies target simultaneously few dozen cells with a temporal resolution of few milliseconds (Emiliani et al., 2015). Major areas of future developments include scaling up of the number of stimulated neurons while maintaining single-cell resolution, improving temporal resolution, performing large-scale 3D stimulation, and precisely quantifying tissue photodamage during intervention. In addition, it will, ultimately, be important to implement these technologies with a closed-loop system (Grosenick et al., 2015), so that intervention can be tied to behavior. This will be useful, among other things, to predict and discount residual effects of endogenous activity (Ahmadian et al., 2011). In fact, both the number of responsive neurons and their functional responses to the sensory stimulus and to the intervention may vary as a function of behavioral variables such as arousal, attention, or locomotion that are reflected in brain states and ongoing neural activity (see also next section “Potential confounds”). Coupling functional imaging with optogenetic intervention allows tracking these changes and adapting patterned photostimulation to brain dynamics. Moreover, because patterned illumination requires knowing where the cells to stimulate are, and what pattern to stimulate them with, it will be necessary to combine imaging with patterned photostimulation. Finally, taking full advantage of the intersection approach will require multimodal recording techniques. While electrophysiological recordings have millisecond time resolution, they currently lack the ability to determine accurately where the recorded cells are. Ideally, the best approach is to perform statistical analysis using both electrophysiology and functional imaging in the same area; that way, both high temporal and high spatial resolution could be achieved.

It is important to note that the framework of interventional intersectional information requires knowing precisely which values of the neural response features r_i are elicited by intervention in each trial. This in turn requires measurement of the neural response (r_i) on individual intervention trials. When this is not possible, confounds may arise. For example, in the absence of such measures it would be problematic to rule out a residual correlation between the interventionally-elicited neural response and other uncontrolled endogenous brain activity variables that would invalidate the rigor of the causal conclusions, or the elicited activity may be so un-natural (e.g. too synchronized with respect to natural activity patterns) that they may affect in an un-natural way downstream neural processes. When it is not possible to measure in each trial the elicited neural response, the study should be however accompanied by a rigorous quantification in separate trials or experiments of the precision of manipulated response under various conditions, that is adequate to allow extrapolation to individual trials during the behavioral task and that also characterizes the difference between manipulated and non-manipulated activity.

Potential confounds: when the framework may fail

The result of the intersection framework (and of any experimental approach combining neural recordings and interventional techniques) are potentially confounded by many limitations and factors which must be considered carefully to avoid reaching to the wrong conclusions. We have already discussed some of those confounds; in the following we discuss additional ones.

A key requirement for the intersection framework to succeed is that the animal uses the identified stimulus to make choices. This requirement can fail in two important ways. First, in some behaviors, there may be other sensory stimuli that co-vary with the stimuli of interest. In this case, it would be difficult to know which stimulus feature is being used by the animal to drive behavior, a problem exacerbated by the possibility that the stimulus features used by the animal might vary from trial to trial. Second, the animal might have fluctuations in attention, motivation, or arousal, or use non-stimulus features such as reward history, to drive choices. These factors may be present in the two alternatives forced-choice tasks that we discussed in this article, but are likely to be stronger in other task designs, such as go/no go tasks, where our framework could be in principle applied. In all these cases, factors other than the stimulus feature of interest would be involved in driving the animal's choice; that would compromise the proposed framework, since it assumes that the stimulus of interest drives behavior. Such factors can be conceptually formalized by assuming that both the sensory coding and the decision mechanisms may vary across trials, and/or that non-recorded or non-manipulated neurons, may vary across intervention and non-intervention trials (such as r_2 depicted in Fig. 4E,F when only intervening on feature r_1).

For these reasons, it is important to evaluate whether the variables describing behavior and the non-observed and non-manipulated endogenous variables are in a comparable state during intervention and non-intervention trials. In the presence of variations, a simple strategy could be to down-sample intervention and sensory-evoked trials so that only compatible brain or behavioral states are analyzed. A better solution, however, is to consider tasks in which it is known, based on high behavioral performance and good psychometric curves, that the stimulus feature of interest drives the animal's choice with high reliability. Similarly, the stimulus should be designed so that co-varying stimulus features are avoided. This will probably be easier with simple stimulus sets than with natural stimuli.

Variation of behavior and brain state variables across the experiment, on the other hand, offer an important opportunity to evaluate whether such variables have a "permissive" role on task performance. For example, in the virtual-pole sensation experiment of (O'Connor et al., 2013), the fact that virtual pole perception worked only when the animal whisked suggests a permissive role of whisker movements for active sensation. A strategy that could take advantage of these variations in state and behavior could be to include (using e.g. simple modeling techniques such as generalized Linear Models (Park et al., 2014)) behavioral factors such as slow variation across blocks of trials of motivation or reward history or brain states explicitly into the experimenter's sensory coding and decision boundary models. This could potentially lead to explaining the dynamic role of these factors in sensory coding.

Another significant confound can arise for interventional approaches when investigating partly parallel pathways. For example, suppose a behavior is generated by two brain areas that operate in a partly parallel or complementary way, as for example in an "OR" function

(Li et al., 2016). In such a case, when inactivating only one area with intervention, one may find little causal effect on behavior. However, interpreting this result as evidence that the inactivated region does not causally contribute much to behavior could be misleading. One way to alleviate these confounds would be to compute both statistical and interventional information intersection. One may use these measures to disambiguate the case in which the two areas contribute complementarily to behavior and so offer complementary intersection information from the case when the two areas operate entirely redundantly and so intersection information from both areas equals intersection information from one area alone. However, accurate interpretation of these measures would require knowledge of the functional anatomy, which, for example, inform the experimenter about the presence or absence of parallel and potentially redundant pathways. Moreover, completeness of activity monitoring and perturbation of the regions involved is also paramount, as this would be, for instance, useful to rule out that a failure to affect behavior by inactivating a region is due to incomplete control of all relevant neurons. We thus anticipate that as approaches move toward understanding larger and larger populations of neurons (Keller and Ahrens, 2015; Sofroniew et al., 2016) and the interconnections between neurons (Lichtman and Denk, 2011), these joint statistical and interventional approaches will become easier to interpret.

Determining the instructive vs permissive role of neural codes and neural circuit: from circuit dissection to circuit information flow

The intersection information framework (both statistical and interventional) has direct application for the dissection of neural circuits underlying behavior. Much work in systems neuroscience has used neurophysiology to identify neural correlates, and, due to recent optogenetics approaches, a wave of new studies has sought to identify which brain regions, cell types, axonal projection pathways, and circuits are required for accurate performance of behavioral tasks (Guo et al., 2014; O'Connor et al., 2013; Peng et al., 2015). It is essential to emphasize that simply measuring the effect of an intervention on choice without regard to stimulus coding precludes understanding a neural circuit's role in task performance. Here we propose that the use of intersection information is crucial to determine whether a neural circuit (or cell type or projection pathway) carries information that is instructive (Otchy et al., 2015) for task performance (contributes essential information for the task performance that is not provided elsewhere) or if the circuit is permissive for task performance (is required for, or modulates, the behavior but does not provide essential information).

Fig. 4A-D schematizes four cases of different neural circuit architectures in which two neural features, r_1 and r_2 (for example, the activity of neurons in two different brain areas, cell types, or projection pathways), may inform choice. In all four cases, feature r_2 contributes essential information to choice and task performance (r_2 is thus instructive), but the role of r_1 varies. An interventional intersection framework would correctly identify these four circuit architectures. The case of “parallel” information flow (Fig. 4A), in which r_1 and r_2 both provide complementary instructive information, could be revealed by finding that the intersection information provided by r_1 and r_2 jointly is larger than that provided by r_1 alone (that is, r_1 and r_2 provide complementary stimulus information to choice). The case of “serial” information flow (Fig. 4C), in which r_1 provides instructive information to r_2 and r_2 informs choice, could be discovered by finding that the intersection information provided by r_1 and r_2 jointly equals that provided by r_2 alone. The case in which r_1 provides permissive – but not instructive – information (Fig. 4D), could be identified by finding that r_1 carries

interventional choice information but not intersection information. Finally, the case when r_l is not used for choice (Fig. 4B) corresponds to the absence of both choice and intersection information in r_l .

It is important to note that the framework we discussed here is general and can in principle be applied not only to determine how sensory information carried by different codes is used to produce behavior, but it can also be used to study how stimulus information flows across neural populations. For example, the same reasoning expressed above applies to considering a group of some brain regions r_1, \dots, r_n (whose activity we can record and manipulate, not necessarily at the same time) and a downstream area c (whose activity we assume we can record at the same time when we manipulate or record r_1, \dots, r_n). In this case the meaning of intersection information would be that of the amount of information about stimulus s carried by population r_1, \dots, r_n that is transmitted downstream to area c . In essence, we have replaced choice by activity in area c . We could therefore identify the neural response features that influence activity in downstream regions, leading to hypotheses about the mechanisms of information flow in neural circuits.

Ideally, these statistical and causal measures of information flow should be integrated with information about anatomy, response timing, and information dynamics. For example, in the presence of a partly feedforward or hierarchical architecture, anatomy could be used to identify the earliest areas where sensory, choice and intersection information are developed (Koulakov et al., 2005), and thus better track the computations leading to task performance. Similarly, the timing of stimulus and choice information in neural activity could be used to infer whether, for example, choice signals reflect a neuron's causal effect on behavioral choice or rather a top-down signal (Nienborg and Cumming, 2009).

Concluding remarks

We presented a new framework to crack the neural code underlying sensory perception. The framework emphasizes neural response features that both carry sensory information and lead to appropriate actions, with the emphasis on “appropriate”. These are the neural response features with a large intersection between sensory information and readout. Based on this framework, we provided a first attempt to formalize statistical ways to identify these features from recordings of neural activity, and to design interventional experiments that can causally test the degree of intersection information. This approach can resolve open debates about the nature of the neural code. Moreover, the ideas we proposed in this framework can guide researchers in the design of experiments, in the design of new statistical tools, and in the development of the new technology, that will lead us to crack the neural code.

Acknowledgements

We thank the referees for their insightful comments, and members of our laboratories for precious feedback. We are particularly indebted to M.E. Diamond, H. Safaai, D. Chicharro, C. Kayser and C. Runyan for valuable collaborations on these topics and/or comments on the manuscript. SP and EP acknowledge the warm hospitality of Harvard Medical School while writing this paper. This work was supported by the Fondation Bertarelli, the European

Research Council (ERC, NEURO-PATTERNS), the Flag-Era JTC Human Brain Project (SLOW-DYN), a Burroughs-Wellcome Fund Career Award at the Scientific Interface, the Searle Scholars Program, the New York Stem Cell Foundation, and NIH grants R01-MH107620, R01-NS089521 and U01NS090576. C.D.H. is a New York Stem Cell Foundation Robertson Neuroscience Investigator. Funding for PEL was provided by the Gatsby Charitable Foundation.

Figure captions

Figure 1: Intersection information helps combining statistics, neural recordings, behavior and intervention to crack the neural code for sensory perception.

A) Schematic showing two crucial stages in the information processing chain for sensory perception: sensory coding and information readout. In this example an animal must discriminate between two stimuli of different color ($s=1$, green and $s=2$, blue) and make an appropriate choice ($c=1$, pink and $c=2$, red). Sensory coding expresses how different stimuli are encoded by different neural activity patterns. Information readout is the process by which information is extracted from single-trial neural population activity to inform behavioral choice. The intersection between sensory coding and information readout is defined as the features of neuronal activity that carry sensory information that is read out to inform a behavioral choice. Note that, as explained in the main text, a neural feature may show both sensory information and choice information but have no intersection information; this is visualized here by plotting the intersection information domain in the space of neural features as smaller than the overlap between the sensory coding and information readout domains. **B)** Only information at the intersection between sensory coding and readout contributes to task performance. Neural population response features that belong to this intersection can be identified by statistical analysis of neural recordings during behavior. Interventional (e.g. optogenetics) manipulations of neural activity informed by statistical analysis of sensory information coding can then be used to causally probe the contribution of neural features to task performance at this intersection.

Figure 2. Schematic of possible pairs of neural population features involved in sensory perception

A) Features r_1 and r_2 are the pooled firing rates of two neuronal populations (yellow and cyan) that encode two different visual stimuli ($s=1$, green; and $s=2$, blue). Values of single-trial responses of each population response feature can be represented as dots in the two-dimensional plot of spike count variables in the r_1, r_2 space (rightmost panel in A). **B)** Features r_1 and r_2 are low-dimensional projections of large-population activity (computed for example with PCA as weighted sum of the rates of the neurons). **C)** Features r_1 and r_2 are spike timing and spike count of a neuron. **D)** Features r_1 and r_2 are the temporal regularity of the spike train of a neuron and spike count.

Figure 3. Impact of response features on sensory coding, readout, and intersection information

In the left panels A_1, B_1, C_1 we illustrate stimulus and choice dependences of two hypothetical neural features, r_1 and r_2 , with scatterplots of simulated neural responses to two stimuli, $s=1$

or $s=2$. The dots are color-coded: green if $s=1$ and blue if $s=2$. Dashed black and red lines represent the sensory and decision boundaries, respectively. The region below the sensory boundary corresponds to responses that are decoded correctly from features r_1, r_2 if the green stimulus is shown; the region above the sensory boundary corresponds to responses that are decoded correctly if the blue stimulus is shown. Filled circles correspond to correct behavioral choices; open circles to wrong choices. Panels A_2, B_2, C_2 plot only the trials that contribute to the calculation of intersection information. Those are the behaviorally correct trials (filled circles) in the two regions of the r_1, r_2 plane regions in which the decoded stimulus \hat{s} and the behavioral choice are both correct. Each region is color coded with the color of the stimulus color that contributes to it. White regions indicate portion of the r_1, r_2 plane that cannot contribute to the intersection because for these responses either the decoded stimulus or choice is incorrect. The larger the colored areas and the number of dots included in panels A_2, B_2, C_2 the larger the intersection information. Panels A_3, B_3, C_3 plot a possible neural circuit diagram that could lead to the considered result. In these panels s indicates the sensory stimulus, r_i indicate the neural features and c the readout neural system, and arrows indicate directed information transfer **A1-3**) No intersection information (the sensory and decision boundary are orthogonal). **B1-3**) Intermediate intersection information (the sensory and decision boundary are partly aligned). **C1-3**) Large intersection (the sensory and decision boundary are fully aligned).

Figure 4. Causal manipulations to study the permissive and instructive roles in coding and information flow

A-D) Interventional approaches can be used to disambiguate among different conditions: **A**) the neural feature r_1 and r_2 carry significant information about the stimulus, s , and provide essential stimulus information to the decision readout, c . **B**) r_1 does not send information to c , but only receives a copy of the information via r_2 , which does send stimulus information to c . **C**) r_1 provides instructive information about s to r_2 and r_2 informs c instructively; **D**) r_1 influences c but does not directly carry information about s . **E-F**) Interventional approaches can be used to reveal cases in which r_1 informs c but does not send stimulus information that contributes to task performance (black arrow in E) from cases in which r_1 sends stimulus information used for decisions (colored arrow in F).

Fig 5: Schematic of an experimental design to probe intersection information with intervention

Three examples of neural responses (quantified by features r_1, r_2) to two stimuli, with conventions as in Fig. 3. We assume that some patterns of neural activity are evoked by interventional manipulation in some other trials. The “lightning bolts” indicate activity patterns in r_1, r_2 space evoked by intervention: they are color-coded with the choice that they elicited (as determined by the decision boundary – the dashed red line). Choice $c=1$ is color-coded as pink, and $c=2$ as dark red. The choices evoked by the intervention can be used to determine, in a causal manner, the position of the decision boundary (as the line separating different choices). The correspondence between the stimulus that would be decoded from the neural responses to the intervention-induced choice can be used to compute interventional intersection information. **A**) A case with no interventional intersection information (the sensory and decision boundary are orthogonal). **B**) A case with intermediate intersection (the sensory and decision boundary are partly aligned). **C**) A case with large intersection (the sensory and decision boundary are fully aligned).

Figure 6. Examples of statistical intersection measures in a texture discrimination task

This figure shows how spike timing and spike count in primary somatosensory cortex encode textures of objects, and how this information contributes to whisker-based texture discrimination. **A-B)** Schematic of the texture discrimination task. **A)** On each trial, the rat perched on the edge of the platform and extended to touch the texture with its whiskers. **B)** Once the animal identified the texture, it turned to the left or the right drinking spout, where it collected the water reward. **C)** Schematic of the computation of spike count and spike timing signals in single trials. **D-F)** The mean \pm SEM (over $n=459$ units recorded in rat primary somatosensory cortex) of texture information (D), choice information (E), and fraction of intersection information fII (F). Modified with permission from (Zuo et al., 2015).

Figure 7. Examples of statistical and interventional intersection measures with sensory and illusory touches

This figure shows results of the statistical and interventional test of the role of cortical spike timing and spike count in the neural code for whisker-based object location. The test involved closed-loop optogenetic stimulation causing illusory perception of object location. **A)** Schematic of the task: four trial types during a closed-loop optogenetic stimulation behavior session depending on pole location and optogenetic stimulation (cyan lightning bolts). A “virtual pole” (magenta) was located within the whisking range (gray area). Mice reported object location by licking or not licking. **B)** Decoding object location and behavioral choice from electrophysiologically recorded spikes in layer 4 of somatosensory cortex. Each dot corresponds to the decoding performance (fraction correct) of one neuron. **C)** Optogenetically-imposed spike rates evoked virtual pole sensation. Left: Optogenetic stimulation (blue circles) coupled to whisker movement (gray, whisking angle θ) during object location discrimination. Asterisk, answer lick. Middle: Responses in the four trial types across one behavioral session. Green, yes responses; gold, no responses. Right: Optogenetic stimulation in NO trials (red), but not in YES trials (blue), in barrel cortex increases the fraction of yes responses. Lightning bolt and “no stim” labels indicate the presence and absence of optogenetic stimulation, respectively. Error bars, s.e.m. Each line represents an individual animal. **D)** Adding timing information in the optogenetically evoked activity did not improve virtual pole perception. Top: delayed optogenetic stimulation was triggered by whisker crossing with variable delays, Δt . Middle: whisker movements with whisker crossing (red circles) and corresponding optogenetic stimuli (cyan circles) for $\Delta t=50$ ms. Bottom: fooling index (fraction of trials reporting sensing of a virtual pole) as function of Δt . Modified with permission from (O'Connor et al., 2013).

Figure 8. Experimental configurations for interventional optogenetic approaches.

A) In a virtual sensation experiment, the animal behavior is tested applying the optogenetic intervention in the absence of the external sensory stimulus. **B)** Alternatively, optogenetic intervention can be paired with sensory stimulation with the aim to overriding or biasing neural activity evoked by the sensory stimulus. **C)** In the wide-field configuration for optogenetic manipulation, light is delivered with no spatial specificity within the illuminated area, resulting in the activation (red cells) of most opsin-positive neurons. Stimulation in this regime may lead to over-synchronous neural responses (right panel). The orange lightning bolts in the right panel indicate the time at which successive stimuli are applied. The neurons displayed in panels C-D are meant to represent a population of N neurons expressing the opsins and their number is here limited to 7 for presentation purposes only. **D)** Patterned illumination permits the delivery of photons precisely in space. When multiple and diverse light patterns are consecutively delivered (orange lightning bolts), optical activation of neural networks with complex spatial and temporal patterns becomes possible (right panel).

References

- Ahmadian, Y., Packer, A.M., Yuste, R., and Paninski, L. (2011). Designing optimal stimuli to control neuronal spike timing. *J Neurophysiol* 106, 1038-1053.
- Averbeck, B.B., Latham, P.E., and Pouget, A. (2006). Neural correlations, population coding and computation. *Nat Rev Neurosci* 7, 358-366.
- Baden, T., Berens, P., Franke, K., Roson, M.R., Bethge, M., and Euler, T. (2016). The functional diversity of retinal ganglion cells in the mouse. *Nature* 529, 345-349.
- Baker, C.A., Elyada, Y.M., Parra, A., and Bolton, M.M. (2016). Cellular resolution circuit mapping with temporal-focused excitation of soma-targeted channelrhodopsin. *eLife* 5.
- Bovetti, S., and Fellin, T. (2015). Optical dissection of brain circuits with patterned illumination through the phase modulation of light. *J Neurosci Meth* 241, 66-77.
- Boyden, E.S., Zhang, F., Bamberg, E., Nagel, G., and Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. *Nature neuroscience* 8, 1263-1268.
- Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13, 87-100.
- Buonomano, D.V., and Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nat Rev Neurosci* 10, 113-125.
- Carrillo-Reid, L., Yang, W., Bando, Y., Peterka, D.S., and Yuste, R. (2016). Imprinting and recalling cortical ensembles. *Science* 353, 691-694.
- Chen, J.L., Carta, S., Soldado-Magraner, J., Schneider, B.L., and Helmchen, F. (2013). Behaviour-dependent recruitment of long-range projection neurons in somatosensory cortex. *Nature* 499, 336-+.
- Chen, X.K., Gabitto, M., Peng, Y.Q., Ryba, N.J.P., and Zuker, C.S. (2011). A Gustotopic Map of Taste Qualities in the Mammalian Brain. *Science* 333, 1262-1266.
- Cunningham, J., and Yu, B.M. (2014). Dimensionality reduction for large-scale neural recordings. *Nature neuroscience* 15, 1752-1758.
- Deisseroth, K., and Schnitzer, M.J. (2013). Engineering approaches to illuminating brain structure and dynamics. *Neuron* 80, 568-577.

Doron, G., von Heimendahl, M., Schlattmann, P., Houweling, A.R., and Brecht, M. (2014). Spiking irregularity and frequency modulate the behavioral report of single-neuron stimulation. *Neuron* 81, 653-663.

Emiliani, V., Cohen, A.E., Deisseroth, K., and Hausser, M. (2015). All-Optical Interrogation of Neural Circuits. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 35, 13917-13926.

Engineer, C.T., Perez, C.A., Chen, Y.H., Carraway, R.S., Reed, A.C., Shetake, J.A., Jakkamsetti, V., Chang, K.Q., and Kilgard, M.P. (2008). Cortical activity patterns predict speech discrimination ability. *Nature neuroscience* 11, 603-608.

Grosenick, L., Marshel, J.H., and Deisseroth, K. (2015). Closed-loop and activity-guided optogenetic control. *Neuron* 86, 106-139.

Guo, Z.C.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G.P., and Svoboda, K. (2014). Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* 81, 179-194.

Haefner, R.M., Gerwinn, S., Macke, J.H., and Bethge, M. (2013). Inferring decoding strategies from choice probabilities in the presence of correlated variability. *Nature neuroscience* 16, 235-242.

Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* 484, 62-68.

Harvey, M.A., Saal, H.P., Dammann, J.F., 3rd, and Bensmaia, S.J. (2013). Multiplexing stimulus information through rate and temporal codes in primate somatosensory cortex. *PLoS biology* 11, e1001558.

Hernandez, A., Nacher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., Vazquez, Y., Camarillo, L., and Romo, R. (2010). Decoding a Perceptual Decision Process across Cortex. *Neuron* 66, 300-314.

Histed, M.H., and Maunsell, J.H. (2014). Cortical neural populations can guide behavior by integrating inputs linearly, independent of synchrony. *Proceedings of the National Academy of Sciences of the United States of America* 111, E178-187.

Houweling, A.R., and Brecht, M. (2008). Behavioural report of single neuron stimulation in somatosensory cortex. *Nature* 451, 65-68.

Huber, D., Petreanu, L., Ghitani, N., Ranade, S., Hromadka, T., Mainen, Z., and Svoboda, K. (2008). Sparse optical microstimulation in barrel cortex drives learned behaviour in freely moving mice. *Nature* 451, 61-64.

Jacobs, A.L., Fridman, G., Douglas, R.M., Alam, N.M., Latham, P.E., Prusky, G.T., and Nirenberg, S. (2009). Ruling out and ruling in neural codes. *Proceedings of the National Academy of Sciences of the United States of America* 106, 5936-5941.

Jia, X.X., Tanabe, S., and Kohn, A. (2013). Gamma and the Coordination of Spiking Activity in Early Visual Cortex. *Neuron* 77, 762-774.

Kayser, C., Montemurro, M.A., Logothetis, N.K., and Panzeri, S. (2009). Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron* 61, 597-608.

Keller, P.J., and Ahrens, M.B. (2015). Visualizing whole-brain activity and development at the single-cell level using light-sheet microscopy. *Neuron* 85, 462-483.

Koulakov, A.A., Rinberg, D.A., and Tsigankov, D.N. (2005). How to find decision makers in neural networks. *Biol Cybern* 93, 447-462.

Lerner, T.N., Ye, L., and Deisseroth, K. (2016). Communication in Neural Circuits: Tools, Opportunities, and Challenges. *Cell* 164, 1136-1150.

Li, N., Chen, T.W., Guo, Z.V., Gerfen, C.R., and Svoboda, K. (2015). A motor cortex circuit for motor planning and movement. *Nature* 519, 51-U88.

Li, N., Daie, K., Svoboda, K., and Druckmann, S. (2016). Robust neuronal dynamics in premotor cortex during motor planning. *Nature* 532, 459-464.

Lichtman, J.W., and Denk, W. (2011). The big and the small: challenges of imaging the brain's circuits. *Science* 334, 618-623.

Luczak, A., McNaughton, B.L., and Harris, K.D. (2015). Packet-based communication in the cortex. *Nat Rev Neurosci* 16, 745-755.

Luna, R., Hernandez, A., Brody, C.D., and Romo, R. (2005). Neural codes for perceptual discrimination in primary somatosensory cortex. *Nature neuroscience* 8, 1210-1219.

Meister, M.L., Hennig, J.A., and Huk, A.C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 33, 2254-2267.

Morcos, A.S., and Harvey, C.D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. *Nature neuroscience* 19, 1672-1681.

Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. *Nature* 341, 52-54.

Nienborg, H., and Cumming, B.G. (2009). Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* 459, 89-92.

Nikolic, D., Fries, P., and Singer, W. (2013). Gamma oscillations: precise temporal coordination without a metronome. *Trends Cogn Sci* 17, 54-55.

O'Connor, D.H., Hires, S.A., Guo, Z.V., Li, N., Yu, J., Sun, Q.Q., Huber, D., and Svoboda, K. (2013). Neural coding during active somatosensation revealed using illusory touch. *Nature neuroscience* 16, 958-965.

O'Keefe, J., and Recce, M.L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3, 317-330.

Otchy, T.M., Wolff, S.B.E., Rhee, J.Y., Pehlevan, C., Kawai, R., Kempf, A., Gobes, S.M.H., and Olveczky, B.P. (2015). Acute off-target effects of neural circuit manipulations. *Nature* 528, 358-363.

Packer, A.M., Russell, L.E., Dagleish, H.W.P., and Hausser, M. (2015). Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo. *Nat Methods* 12, 140-U186.

Panzeri, S., Brunel, N., Logothetis, N.K., and Kayser, C. (2010). Sensory neural codes using multiplexed temporal scales. *Trends Neurosci* 33, 111-120.

Papagiakoumou, E., Anselmi, F., Begue, A., de Sars, V., Gluckstad, J., Isacoff, E.Y., and Emiliani, V. (2010). Scanless two-photon excitation of channelrhodopsin-2. *Nat Methods* 7, 848-U117.

Park, I.M., Meister, M.L., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature neuroscience* 17, 1395-1403.

Pearl, J. (2009). *Causality* (Cambridge: Cambridge University Press).

Peng, Y.Q., Gillis-Smith, S., Jin, H., Trankner, D., Ryba, N.J.P., and Zuker, C.S. (2015). Sweet and bitter taste in the brain of awake behaving animals. *Nature* 527, 512-515.

Peron, S., and Svoboda, K. (2011). From cudgel to scalpel: toward precise neural control with optogenetics. *Nat Methods* 8, 30-34.

Quiñ Quiroga, R., and Panzeri, S. (2009). Extracting information from neuronal populations: information theory and decoding approaches. *Nat Rev Neurosci* 10, 173-185.

Reich, D.S., Mechler, F., and Victor, J.D. (2001). Independent and redundant information in nearby cortical neurons. *Science* 294, 2566-2568.

- Rickgauer, J.P., Deisseroth, K., and Tank, D.W. (2014). Simultaneous cellular-resolution optical perturbation and imaging of place cell firing fields. *Nature neuroscience* 17, 1816-1824.
- Rigotti, M., Barak, O., Warden, M.R., Wang, X.J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature* 497, 585-590.
- Romo, R., Hernandez, A., Zainos, A., and Salinas, E. (1998). Somatosensory discrimination based on cortical microstimulation. *Nature* 392, 387-390.
- Romo, R., and Salinas, E. (2003). Flutter discrimination: neural codes, perception, memory and decision making. *Nat Rev Neurosci* 4, 203-218.
- Salzman, C.D., Britten, K.H., and Newsome, W.T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. *Nature* 346, 174-177.
- Shadlen, M.N., Britten, K.H., Newsome, W.T., and Movshon, J.A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 16, 1486-1510.
- Shamir, M. (2014). Emerging principles of population coding: in search for the neural code. *Current Opinion in Neurobiology* 25, 140-148.
- Shriki, O., Kohn, A., and Shamir, M. (2012). Fast coding of orientation in primary visual cortex. *PLoS computational biology* 8, e1002536.
- Shusterman, R., Smear, M.C., Koulakov, A.A., and Rinberg, D. (2011). Precise olfactory responses tile the sniff cycle. *Nature neuroscience* 14, 1039-U1136.
- Sofroniew, N.J., Flickinger, D., King, J., and Svoboda, K. (2016). A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *eLife* 5.
- Sternson, S.M., and Roth, B.L. (2014). Chemogenetic tools to interrogate brain functions. *Annual review of neuroscience* 37, 387-407.
- Tehovnik, E.J., Tolias, A.S., Sultan, F., Slocum, W.M., and Logothetis, N.K. (2006). Direct and indirect activation of cortical neurons by electrical microstimulation. *J Neurophysiol* 96, 512-521.
- Victor, J.D. (2000). How the brain uses time to represent and process visual information(1). *Brain Research* 886, 33-46.
- Victor, J.D., and Nirenberg, S. (2008). Indices for testing neural codes. *Neural computation* 20, 2895-2936.

Wilson, N.R., Runyan, C.A., Wang, F.L., and Sur, M. (2012). Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature* 488, 343-348.

Woods, A.J., Antal, A., Bikson, M., Boggio, P.S., Brunoni, A.R., Celnik, P., Cohen, L.G., Fregni, F., Herrmann, C.S., Kappenman, E.S., *et al.* (2016). A technical guide to tDCS, and related non-invasive brain stimulation tools. *Clinical neurophysiology : official journal of the International Federation of Clinical Neurophysiology* 127, 1031-1048.

Yang, Y., DeWeese, M.R., Otazu, G.H., and Zador, A.M. (2008). Millisecond-scale differences in neural activity in auditory cortex can drive decisions. *Nature neuroscience* 11, 1262-1263.

Yang, Y., and Zador, A.M. (2012). Differences in sensitivity to neural timing among cortical areas. *The Journal of neuroscience : the official journal of the Society for Neuroscience* 32, 15142-15147.

Zuo, Y.F., Safaai, H., Notaro, G., Mazzone, A., Panzeri, S., and Diamond, M.E. (2015). Complementary Contributions of Spike Timing and Spike Rate to Perceptual Decisions in Rat S1 and S2 Cortex. *Current Biology* 25, 357-363.