Population-wide distributions of neural activity during perceptual decision-making

Adrien Wohrer [a,1], Mark D. Humphries [a,b,1], Christian K. Machens [a,c,*]

[a] Group for Neural Theory, INSERM U960, École Normale Supérieure (Département d'Études Cognitives), 29 rue d'Ulm, 75005 Paris, France
[b] Faculty of Life Sciences, University of Manchester, Manchester M60 1QD, UK
[c] Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown, Av. Brasília s/n, 1400-038 Lisboa, Portugal

ABSTRACT

Cortical activity involves large populations of neurons, even when it is limited to functionally coherent areas. Electrophysiological recordings, on the other hand, involve comparatively small neural ensembles, even when modern-day techniques are used. Here we review results which have started to fill the gap between these two scales of inquiry, by shedding light on the statistical distributions of activity in large populations of cells. We put our main focus on data recorded in awake animals that perform simple decision-making tasks and consider statistical distributions of activity throughout cortex, across sensory, associative, and motor areas. We transversally review the complexity of these distributions, from distributions of firing rates and metrics of spike-train structure, through distributions of tuning to stimuli or actions and of choice signals, and finally the dynamical evolution of neural population activity and the distributions of (pairwise) neural interactions. This approach reveals shared patterns of statistical organization across cortex, including: (i) long-tailed distributions of activity, where quasi-silence seems to be the rule for a majority of neurons; that are barely distinguishable between spontaneous and active states; (ii) distributions of tuning parameters for sensory (and motor) variables, which show an extensive extrapolation and fragmentation of their representations in the periphery; and (iii) population-wide dynamics that reveal rotations of internal representations over time, whose traces can be found both in stimulus-driven and internally generated activity. We discuss how these insights are leading us away from the notion of discrete classes of cells, and are acting as powerful constraints on theories and models of cortical organization and population coding.

© 2012 Elsevier Ltd. All rights reserved.

Contents

1. Introduction

Over the course of the last century, studies based on lesions, electrophysiology, fMRI, and other methods have quite successfully mapped out where different types of information are represented in the brain (Kandel et al., 2000). During the same period, our understanding of the single neuron and its role in the brain has increased substantially (Koch, 1999). While both research directions have created a wealth of knowledge about the organization of the brain, a large gap remains between them. At the center of this gap lie neural networks—thousands or millions of interconnected neurons, responding in myriad ways to whatever task an organism is engaged in. The number of degrees of freedom in these neural populations explodes, a phenomenon known as the "curse of dimensionality". This dimensionality explosion creates a tremendous challenge to unravel how information in populations of neurons is processed and represented. One challenge is to understand the structure and plasticity of these networks, one to link this structure and plasticity to the generated activity, and one challenge is to describe and interpret this activity. This latter problem will be the focus of our review.

What defines a neural population and how should we represent its activity? A neural population is a collection of single cells in a given region or area of the brain. Accordingly, the population activity is just the collection of the respective single cell activities. A common view is that if a large class of cells is activated by the same type of information, e.g., a feature of a visual stimulus, then any differences in their responses are noise that must be averaged over. At the other extreme, the details of every single neuron matter, and the activation of each neuron has to be considered separately. In this review, we will attempt to strike a balance between these extremes, and center on the statistical approach to characterizing population activity. The statistical approach aims to quantify the probability with which a set of features is represented in a population of neurons.

Such a quantitative, probabilistic description of the population activity will be useful on three fronts. First, it shows how information is embedded in population activity and may thereby expose widespread organizational principles or statistical patterns that are shared across brain areas. Second, such an understanding may help to study the computations carried out in a given circuit. Third, a statistical description of population activity imparts

important constraints to network models that seek to explain how these activities are generated through the interaction of neurons. To obtain these constraints, a description of the population activity for a particular experimental condition requires more detail than simply "N/500 cells were responsive (T-test, p<0.05)". Instead, we need to access the overall distribution of measured features in the population, and we need to provide a specific probabilistic model of the activity.

Here, we will review what we have learnt about population activities across the cortex, from sensory to executive to motor areas. We will restrict ourselves to the representation of simple stimuli or simple actions and put a strong emphasis on data that have been collected from animals performing well-controlled behavioral tasks. We will mostly leave aside parts of the literature that are concerned with complex stimuli (such as natural stimuli), complex movement sequences, etc. Focusing on one dimensionality explosion, namely the one we face when observing thousands of neurons, already provides ample perplexity. We will therefore disregard that second dimensionality explosion neuroscience is struggling with, namely the innumerable sensory stimuli or motor actions that may shape and modify the population activity.

The review is largely organized by model complexity. First, in Section 2, we discuss and specify the statistical approach to describing population activity, emphasizing its usefulness and its limitations. In Section 3, we review what is known about distributions of firing rates. In Section 4, we look beyond firing rates and review the distributions of second- and higher-order statistics of spike trains. In Section 5, we review how neural populations are tuned to stimuli or actions and look into distributions of tuning curves and other encoding features. In Section 6, we take time into consideration and review what is known about the dynamics of neural population activity. Finally, in Section 7, we discuss more recent attempts to yield full characterizations of the population activity, that include (pairwise) neural interactions and correlations.

2. The statistical approach to population activity

2.1. The statistical characterization of cell-to-cell variability

For concreteness, imagine an animal that is engaged in a task in which we monitor a set of sensory stimuli s and a set of movements or actions a (Fig. 1). In such a setting, we can often record the simultaneous activity of many neurons in a given patch of cortex. In analyzing the activity of these neurons in relation to the chosen set of stimuli or actions, we are likely to encounter strong cell-to-cell variability. Even if two cells are responsive to the same set of stimuli, their precise tuning is likely to differ. A third cell may show
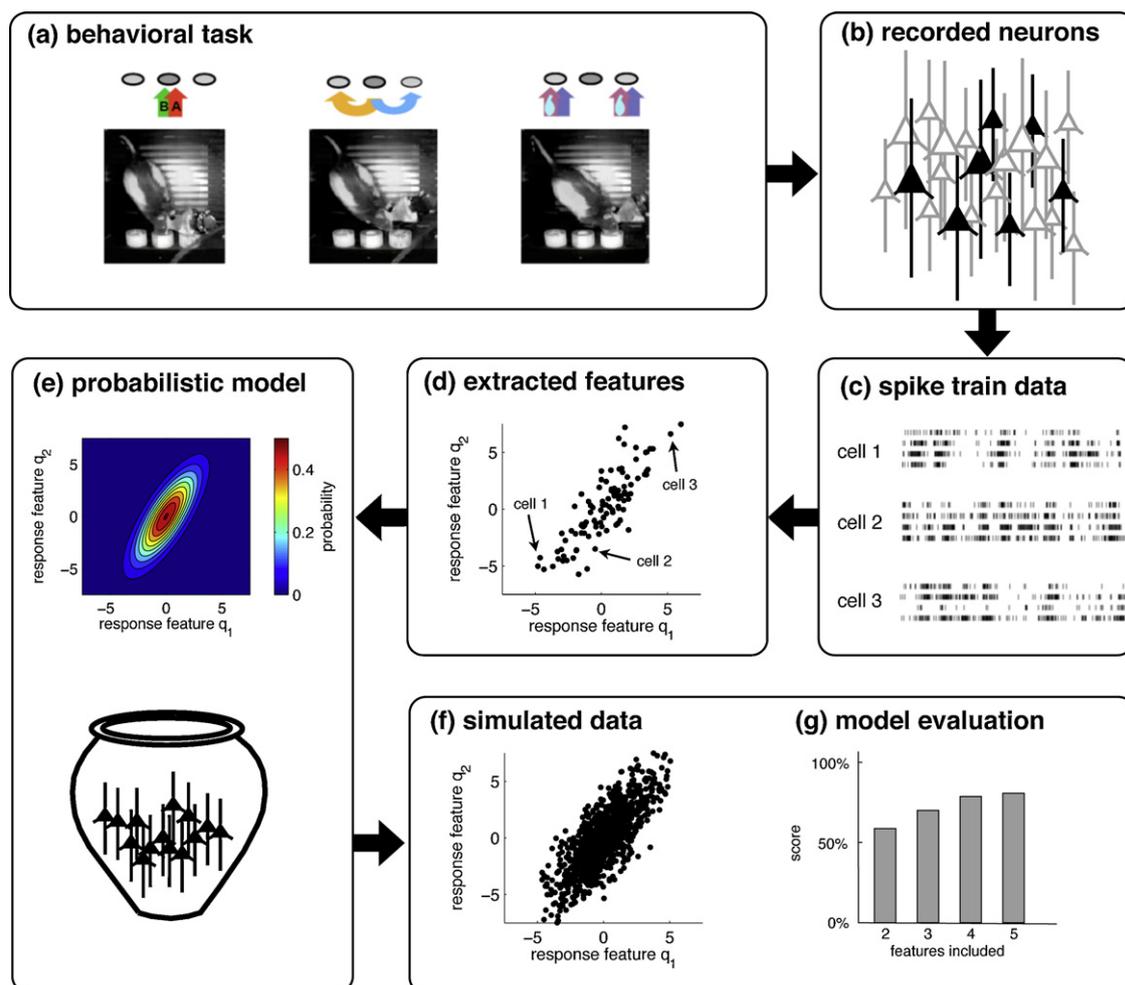
Fig. 1. The statistical approach to population responses. (a) We assume that data are recorded from awake behaving animals engaged in simple tasks (figure adapted from Feierstein et al., 2006, with permission from Elsevier). (b) The spike trains of a small subset of neurons in one area are recorded. (c) Recordings are sorted by trials over identical conditions. (d) For each neuron, we extract certain features from the spike trains, e.g., the stimulus-dependence of the firing rate or of the change in firing rate. (e) The distribution of these features across the recorded population leads to a probabilistic model of the population response. Such a model is usually generative, so that we can simulate data by randomly drawing neurons from the model, here depicted as an urn. (f) The simulated data should ideally be similar to the original data. (g) This similarity is then quantified to evaluate the probabilistic model.

yet another tuning to the same stimuli and so on. How are we to characterize these response differences?

In the statistical approach to population activity, we neglect the identities of the neurons and only characterize the statistics or probabilities of their activity or some feature thereof. Assume, for instance, that we are measuring spontaneous activity in a region of interest. Focusing on a time window of length $T = 10$ s, how likely are we to find cells that fire exactly $r = 24$ spikes? How likely to find cells that fire $r = 25$ spikes? To answer that question, we can simply construct a distribution $p(r)$ that measures the probability of observing the response r in the given scenario. The distribution will typically depend on one or more parameters, such as the mean and standard deviation in the case of a Gaussian, or just the mean in the case of an exponential distribution.

The situation becomes slightly more complicated if we consider that neural activities change with the presentation of a stimulus, the execution of an action, or simply the passage of time. Imagine we were to measure the responses of cells to two different stimuli—such as two visual gratings with different orientations. In this case, we can summarize a neuron's response in a two-dimensional vector $\mathbf{r} = (r_1, r_2)$, where r_1 denotes the response (such as a firing rate) to the first orientation, and r_2 denotes the response to the second orientation. In turn, we can ask what is the probability of observing a particular response $\mathbf{r}$ in the area we are recording from? Just as above, this amounts to characterizing the (now bi-variate) distribution $p(\mathbf{r}) = p(r_1, r_2)$.

As the number of stimuli or actions increases, as we start considering temporal changes in activity, or the precise timing of spikes, the distributions become distributions over many variables, and the naive construction of distributions $p(\mathbf{r})$ through histograms and the like becomes unfeasible. One way to address this problem is to construct quantitative models of the cells' responses. For instance, we could describe the tuning of each cell to the set of stimuli with a few parameters, such as the width of tuning, etc. In turn, we obtain a set of tuning parameters $\mathbf{q} = (q_1, q_2, \ldots)$ for each neuron. In a second step, we can then seek to quantify the distribution $p(\mathbf{q}) = p(q_1, q_2, \ldots)$ of these parameters over the population.

A problem of this two-step approach is that small modeling errors on the single-cell level may add up to large modeling errors on the population level. In other words, even though we may think that we are extracting the right features to describe the responses of single neurons, we may still miss important features of the population response. This problem can be avoided by directly modeling the population activity. For instance, we could use dimensionality reduction techniques such as principal component analysis to extract the relevant features from the data. In either case, we write $p(\mathbf{q})$ and use this distribution as a proper stochastic description of the population activity.

The statistical approach to characterizing population activity is illustrated in Fig. 1. We record responses from a subset of the actual population, Fig. 1b and c, extract the features of the cell's responses, plot these features for the whole population, Fig. 1d, and then fit a probabilistic model to the population data, Fig. 1e. In turn, if we were to draw random samples from this probabilistic model, Fig. 1f, these random samples should look similar—in a well-specified, statistical sense—to the recorded data. Importantly, in the last step we need to evaluate the probabilistic model, ideally on a new set of data. Such evaluations essentially compare how "likely" the new data are with reference to the model, when compared to the (known) likelihood of simulated data (Fig. 1g).

Why would such an approach, in which we neglect the identity of neurons, make any sense? Several reasons: first, if there are thousands or millions of neurons to consider, we can no longer tend to each individual neuron. Other methods are needed to study what is happening and the statistical approach is one of those

methods, presumably one of the simplest. Second, as we will see, the statistical approach can highlight commonalities in the properties of population responses across cortex, which in turn yields important insights into the cortical architecture, as well as constraints for network models of neural systems. Third, to the extent that there are different classes of cells, i.e., different types of neurons with clearly distinguishable response patterns, these will be identified by a statistical approach, which may allow us to link response patterns to anatomical features.

2.2. The sampling problem: recording methods and their biases

Recording a full neural population is virtually impossible with current technologies. We must therefore rely on inferring the properties of a population from a small sample of recorded neurons. This approach will generally work if the sampling is representative, so that the statistical properties of the few neurons we record are similar to the statistical properties of the population as a whole. Consequently, statistical approaches require a precise definition or boundary for the population, and a recording technique that samples neurons randomly and uniformly.

Historically, most studies used extracellular recordings with metal microelectrodes and relied on neural activity for localizing neurons and determining their responsiveness to stimuli or tasks. Consequently, electrophysiological studies are usually biased, in that the most active neurons are self-selective for participating in population analyses, and the rest of the neurons in a region of interest are omitted (see e.g., Olshausen and Field, 2005). This poses a serious problem since there is good reason to believe that most neurons in mammalian cortex are almost silent: from analyses of the expected numbers of neurons within the recording radius of a microelectrode, it has long been suspected that most neurons in cortex are not firing, otherwise an order of magnitude more spikes should be visible (see for a discussion of this "dark neuron" problem Shoham et al., 2006). However, these more silent neurons are largely ignored with classical electrophysiological methods.

The problem of biased sampling is illustrated in Fig. 2. Traditionally, population activity has been characterized by estimating a mean and standard error of the firing rate from pooled single neuron samples. However, any activity-dependent bias in the sampling of cells causes the sample mean to converge higher than the true population mean with an increasing number of samples. Consequently, the mean firing rates are too high and give little clue to the preponderance of low-firing neurons.

One way of avoiding this recording bias is to isolate neurons independently of their responsiveness to stimuli or tasks. This approach is common to tetrode recordings or multi-cell arrays, for instance. Nonetheless, in both cases neurons are still isolated by activity, even if that activity is not directly task-relevant, and less active neurons risk being overlooked. More recent recording methods such as patch-clamping and calcium-imaging circumvent the problem of activity bias completely. As we will see in Section 3.1, patch-clamp recordings provide a different picture of population activity, with many more silent neurons included, thereby confirming the suspicions mentioned above.

In principle, of course, any recording technique is likely to introduce a bias. Patch-clamping may bias the sampling toward larger neurons or reduce their firing rates, for instance, and calcium-imaging may bias the sampling toward neurons that favor expression or uptake of fluorescent markers. However, these biases are likely to be less important than the activity-dependent bias of traditional electrophysiology. Yet even without any recording bias, neurons that become active only under very specific circumstances are notoriously hard to characterize. In the visual system, for instance, neurons with smaller receptive fields are less likely to be
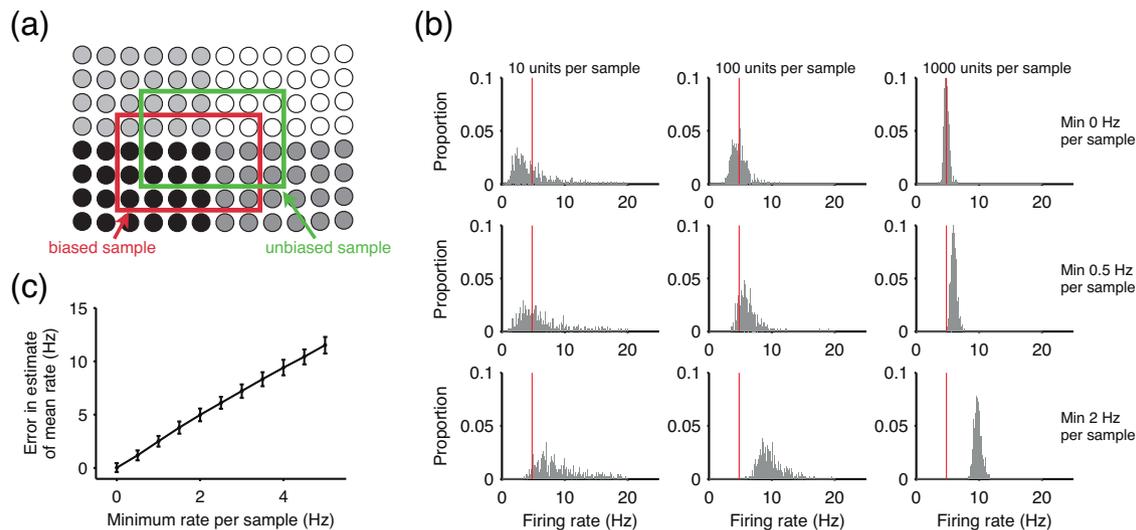
Fig. 2. The effect of recording biases on the estimation of distributions. (a) Statistical approaches rely on unbiased sampling. Within a population, neurons usually have different activity levels, here shown as different gray levels. Most extracellular recording techniques are more likely to find active cells, thus providing a biased sample (red box) as opposed to the desired unbiased sample of the population (green box). (b) A recording bias in single-cell recordings causes systematic errors in estimating population mean firing rates. Using a log-normal distribution of firing rates across neurons (see Section 3.1) to describe a simulated population, we took 500 samples of N individual firing rates; we repeated this for different levels of bias in the sample, by only taking samples above a certain firing rate, mimicking the bias in single-cell recordings. We plot the histograms of every sample's mean rate, for a range of sample sizes and biases; the red line gives the mean value of the population distribution. The top row, with no bias, shows the central limit theorem in action: increasing the sample size causes the estimates of the mean firing rate to converge around the true mean value. However any bias in the sampling causes the convergence to occur around an incorrect mean firing rate for the population. (c) Increasing bias increases error in the estimate. We plot the mean ± s.d. of the error in estimating the true firing rate mean, taken over every sample for $N = 1000$ units. The error increases linearly, as expected, and is approximately twice the size of the recording bias.

discovered or characterized than neurons with larger receptive fields, simply because they are less responsive (Olshausen and Field, 2005). These problems should be kept in mind when characterizing populations of neurons.

2.3. The sampling problem: simultaneous versus pooled recordings

In most studies, only a few neurons are recorded at the same time, and estimates of the population properties are obtained via pooling data across recording sessions. This approach is based on two crucial assumptions. The first assumption is that the population activity does not change over time. Since large shifts in population responses have only been observed in the context of learning, this assumption seems reasonable for a given task. The second assumption is that the activity of single neurons does not change from trial to trial. This, however, is usually not true as neurons exhibit strong trial-to-trial variability, thus leading to a potential confound as soon as we start characterizing populations by pooling over sequential experiments.

From a statistical point of view, we can remedy this problem by characterizing a neural response through a firing rate that is averaged over trials, and through higher-order statistics of the spike trains such as CVs, Fano factors, etc. However, even then the estimate of a neuron's firing rate will come with error bars. Since the errors differ between neurons, they will lead to a distribution of responses even if there were no cell-to-cell variability at all. Accordingly, whenever we compute neural responses by averaging over trials, we introduce artificial variability in the responses, an artifact that needs to be taken into account and corrected for in order to distinguish generic cell-to-cell variability from ordinary trial-to-trial variability.[2] While the effect is likely to be small if responses are averaged over a large amount of trials, it does become relevant if only a few trials are used for averaging (e.g., less

than ten) and may then lead to an artificial widening of the population distributions. That said, we note that this effect has largely been ignored in the literature, and we simply caution the reader that some of the distributions reported in our review may be slightly narrower than stated.

Our review focuses on characterizing cell-to-cell variability, and we will write $p(r)$ to refer to the probability of observing a certain response r within a population of cells. Once in a while, we also need to specify the above-mentioned trial-to-trial variability. In these cases, we will use the subscript t and write $p_t(r)$ to denote the probability of observing the response r of a given neuron within a set of (identical) trials.

3. Population distributions of firing rates

Perhaps the simplest characterization of a neuron population is the distribution of its activity, $p(r)$, when r is some measure of firing rate. Surprisingly, such distributions have only been characterized in a few cases. Standard practice in electrophysiological studies has been to report recorded firing rates across all neurons as the arithmetic mean ± standard error, rather than fit a probabilistic model $p(r)$. Seemingly few studies have directly reported the actual distribution of firing rates across all recorded neurons (whether single-neuron or multi-neuron recording). As we will see below, where they have been reported, the mean firing rate tells us little about the population, and firing rate distributions are clearly not Gaussian. Moreover, we will see that whether characterizing spontaneous activity, stimulus-evoked activity, or activity related to the execution of actions, the distributions are remarkably similar and robust.

3.1. Population activity has a (very) long-tailed distribution, across all of cortex

Studies explicitly reporting and examining distributions of firing rates in cortical populations are mostly recent

---

[2] In statistical literature, this correction procedure is known as *density deconvolution*.
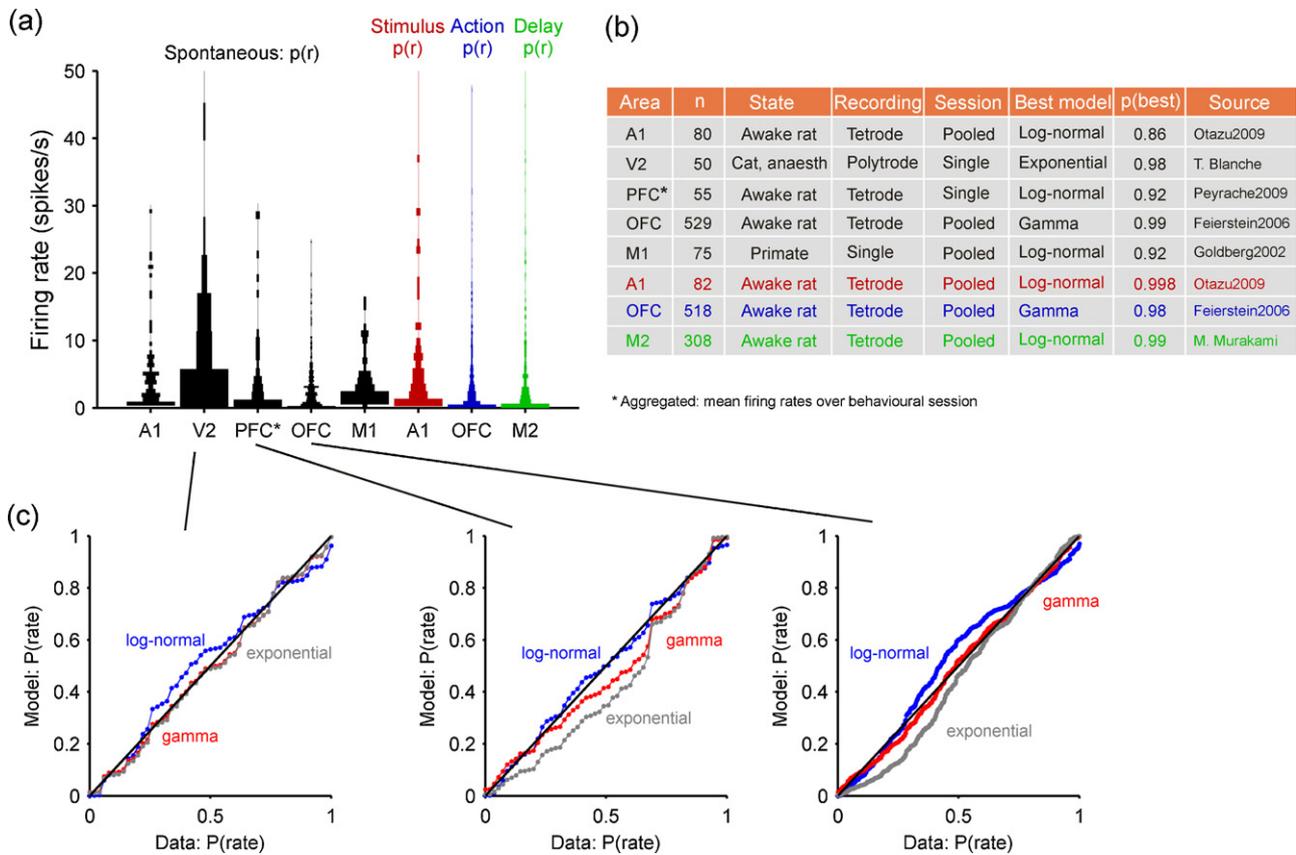
Fig. 3. Distributions of firing rates across neural populations in cortex. (a) Violin plot (Allen et al., 2012) of firing rate distributions across cortical regions, in both spontaneous and task-aligned conditions. Width of bar is proportional to the fraction of neurons in that bin. Optimal histogram widths separately derived for each data-set (Wasserman, 2004). (b) Summary of data-sources in panel (a) and best-fit model to each data-set. Exponential, gamma, and log-normal distributions were fitted using maximum likelihood; best-fitting model was chosen using the Bayesian Information Criterion; p(best): the posterior probability of that being the best model within those tested (Wagenmakers and Farrell, 2004; Wasserman, 2004). The Session column indicates whether the data were from a single recording session or pooled across sessions and/or animals. (c) Probability-probability plots showing the correspondence between the cumulative probability distributions of the firing rate data and of the fitted models. A perfect fit would lie on the diagonal. Left: the V2 data was the only data-set best-fit by the exponential distribution; Middle: the prefrontal cortex data was best-fit by a log-normal distribution; Right: the orbitofrontal cortex data-set was best-fit by a gamma distribution. These examples show not only the best fitting model, but also that it was a good fit to the data. Data sources: A1, from Otazu et al. (2009), data supplied by Gonzalo Otazu; V2: data recorded by Tim Blanche (see Blanche et al., 2005), available from crcns.org; PFC: from Peyrache et al. (2009), data supplied by Adrien Peyrache; OFC, from Feierstein et al. (2006), data supplied by Claudia Feierstein; M1, extracted from Fig. 6A of Goldberg et al. (2002); M2: data supplied by Masayoshi Murakami and Zachary Mainen, unpublished.

(see DeCharms and Zador, 2000 for an earlier discussion). They have consistently described distributions dominated by neurons with low firing rates—low in comparison to classical single electrode recording (Kerr et al., 2005)—with a small set of neurons scattered over orders of magnitude higher rates, forming a long tail.[3] Shafi et al. (2007) re-analyzed a data-set of extracellular single neuron recordings from awake monkeys, and reported long-tailed distributions of firing rates in both parietal and prefrontal cortices. Using cell-attached recording in awake mouse, O'Connor et al. (2010) reported that barrel cortex had a very long-tailed distribution of firing rates when averaged over trial epochs. In perhaps the only study to explicitly fit a probabilistic model p(r), Hromadka et al. (2008) used cell-attached recording in awake rat A1 to show that log-normal, rather than exponential, functions best-fit the distributions of firing rates across all task stages.

Bringing together the results of these studies with further multi-neuron recording data made available to us, we first survey the distributions of firing rates r and the probabilistic models p(r) we can use to describe those distributions, covering regions from

all across cortex. We show the ubiquity of the long-tailed distribution, and also that the most common model points to a non-negligible peak at very low firing rates.

3.1.1. Spontaneous and overall activity

To get a first statistical picture of the population, we start by building a simple distribution of the population activity p(r), where we ignore the animal's behavior. This covers a wide range of situations we label for convenience as "spontaneous": recordings under anesthesia, during sleep, and quiet wakefulness, as well as "baseline" epochs of recordings before a stimulus presentation or a movement. It also covers what we term here "aggregate" recordings, analyzed across presentations of stimuli or productions of actions. We group these as both are neural activity analyses without alignment to a component of a task (stimulus, action, or delay).

Fig. 3a shows how neural populations' distributions of spontaneous and aggregate firing rates are consistently right-skewed and long-tailed in examples from across the whole cortex: from primary sensory cortices (A1, V2), through "higher" prefrontal regions (OFC, PFC), to primary motor cortices (M1). These examples also show the consistency across species, recording technique and vigilance state—under anesthetic, awake and resting, or awake and behaving. In every case, the distributions are dominated by neurons firing less than 1 spike/s, with few neurons firing at rates an order of magnitude higher.

---

[3] In statistical literature, "long tails" often refer to precise mathematical definitions which essentially characterize a slower asymptotic drop-off than the exponential distribution. Since such notions are not applicable to finite, empirical distributions, we use the word "long-tailed" in a more informal meaning, to describe distributions which clearly drop off more slowly than a Gaussian.

Deciding which statistical models $p(r)$ best describe the firing rate distributions, whether these are different across cortical areas, and what that means, are open problems. For firing rate data available to us we fitted exponential, log-normal, and gamma probability distributions as examples of general potential models.[4] As firing rate data can only take zero or positive values, one might consider the exponential distribution to be the null model for such distributions: deviations from this distribution are then informative about the relative spread of firing rates into the lower-rate peak or higher-rate tail of the distribution.

We found that all three distributions occur in the cortical activity data-sets (Fig. 3b). Like Hromadka et al. (2008), we found log-normal models best-fit the firing rate distributions in A1, as well as PFC and M1, suggesting that there is a non-negligible peak of very low firing rates in each of these regions; both other models under-estimated $p(r)$ over the middle range of the distributions (Fig. 3c). The distribution of firing rates in V2 was best-fit by an exponential model, suggesting just a long tail (Fig. 3c). The distribution of firing rates in OFC was best-fit by a gamma model: the exponential model under-estimated the mid-range of the distribution and the log-normal model over-estimated the tail of distribution. Are these distinct distributions between cortical regions meaningful? This is unknown: they could represent differences between vigilance states (awake vs anesthetized), of recording type, of the sample size and construction (pooled vs single recording), or of the spatial sampling (see results on layer-specificity below). Nonetheless, they may also reflect different characteristics of the cortical micro-circuits in different areas.

As we noted in Section 2.2, such long-tailed distributions are surprising from the perspective of classical single neuron recording studies, but not surprising from the perspective that most neurons in cortex must be almost silent, given the likely recording radius of electrode technology (Shoham et al., 2006). Consequently, we expect most cortical neurons to be at the sub-1 Hz end of the firing rate distribution, and the few neurons that are more active ensure long-tailed distributions of firing rate everywhere in cortex. Such long-tailed distributions, dominated by ultra-low rate neurons, are consistent with estimates that mean rate over all cortex must be very low due to metabolic demands (Laughlin and Sejnowski, 2003; Lennie, 2003). What is perhaps genuinely surprising is the ubiquity of the rejection of an exponential model: low firing rates dominate, but not arbitrarily low.

3.1.2. Evoked activity and activity during task-engagement

Given the little firing in the spontaneous state, one may wonder what happens when the respective piece of cortex becomes engaged, e.g., through sensory stimulation, or through a particular behavioral task that an animal is carrying out. Ignoring the precise tuning of cells (see Section 5) or their temporal dynamics (see Section 6), this is simply asking what we know about the distribution of firing rates, $p(r)$, when a single stimulus is presented, or when a particular action taken. From current evidence, it seems all such distributions are also long-tailed—the majority of neurons have a low firing rate, and only a small subset display strong activity aligned to the task component.

Hromadka et al. (2008) pooled single cell-attached recordings to show that the A1 cortex in awake rats had long-tailed firing rate distributions aligned to stimuli presentations (Fig. 4a); similar to their recordings of spontaneous activity, they reported good fits by log-normal distributions. Using similar recording methods, O'Connor et al. (2010) reported long-tailed distributions of firing rates in

mouse barrel cortex (S1) aligned to stimulus contact by the whiskers (Fig. 4b, right). Shafi et al. (2007) pooled single-cell extracellular recordings from monkeys performing a delayed-match-to-sample task to show a long-tailed distribution of firing rates during the delay period.

We confirmed the generality of these prior results by fitting distribution models to firing-rate data from each of the stimulus presentation, action, and delay parts of different tasks. For stimulus-aligned rates, we used recordings of rat A1 after presentation of a tone stimulus in a two-choice decision-making task (data from Otazu et al., 2009). For action-aligned rates, we used recordings of rat orbitofrontal cortex neurons during the movement epoch (move left or right) of a two-choice decision-making task (data from Feierstein et al., 2006). For delay-aligned rates, we used recordings from rat area M2 taken during a reinforcement task with delay, with firing rates calculated from a waiting period in contact with a nose-port before the rat exited to collect reward (unpublished data supplied by Masayoshi Murakami and Zachary Mainen). Fig. 3a shows that all three firing rate data-sets had long-tailed distributions. When we fitted the same exponential, gamma, and log-normal distribution models, we found best fits by the log-normal model for stimulus (A1) and delay (M2) components, and the gamma model for the action component (OFC) (Fig. 3b). Again, silence reigns, with the majority of neurons firing less than 1 spike per second.

Both the analyses of Hromadka et al. (2008) (Fig. 4a) and our analysis of data from Otazu et al. (2009) found clear best-fits of a log-normal distribution to cortical firing rates in rat A1. This consistency is reassuring as the two data-sets were collected respectively using single cell (patch-clamp) and multiple cell (tetrode) recording technologies, and in different behavioral states, passive receipt of stimuli against active engagement in a decision-making task. For rat A1 at least we thus have converging evidence for log-normal firing rate distributions.

3.1.3. Layer-specific differences in spontaneous activity

A confound in many studies of cortical activity in vivo is that we do not know the layer(s) from which the neurons are being recorded. This confound is important because different layers have different levels of activity. Barth and Poulet (2012) have provided a comprehensive review of differences in average firing rate between layers—we note a few key papers here. In a pioneering use of calcium-imaging in vivo, Kerr et al. (2005) reported that layer 2/3 in both A1 and S1 of the rat had very sparse firing rates under anesthesia, with a maximum of 0.1 Hz over 10 min of recording across 212 neurons (their Fig. 6G). Such low rates were confirmed by simultaneous cell-attached or whole-cell recordings from a sample of those neurons; later studies confirmed that layer 2/3 neurons have substantially lower rates than layer 4 and 5 neurons for both spontaneous and stimulus-evoked firing (de Kock et al., 2007; Sakata and Harris, 2009). Similarly, Beloozerova et al. (2003) reported significantly more activity in layer 5 neurons than layer 6 neurons of rabbit primary motor cortex during both rest and locomotion. It seems that the difference in average rates extends to higher moments too: distributions of $p(r)$ are different between different layers. Ringach et al. (2002) constructed a comprehensive map of the layer-by-layer differences in spontaneous firing rates of monkey V1 (Fig. 4c), showing that, while all layers were dominated by low firing rates of less than 1 Hz, layers 4–6 had markedly more neurons with higher firing rates than layers 2/3. A direct comparison by Hromadka et al. (2008) found that the firing rate distributions in rat A1 were consistently long-tailed in each layer (2/3, 4, 5 and 6), and, though not explicitly tested, that the distributions were not the same between layers. O'Connor et al. (2010) reported that barrel cortex had very long-tailed distributions in each layer (2/3, 4, 5, and 6); moreover,

---

[4] Models were fitted using maximum likelihood; to determine the best-fitting model, we computed a Bayesian Information Criterion for each model, and from that determined its posterior probability of having generated the data (Wagenmakers and Farrell, 2004).
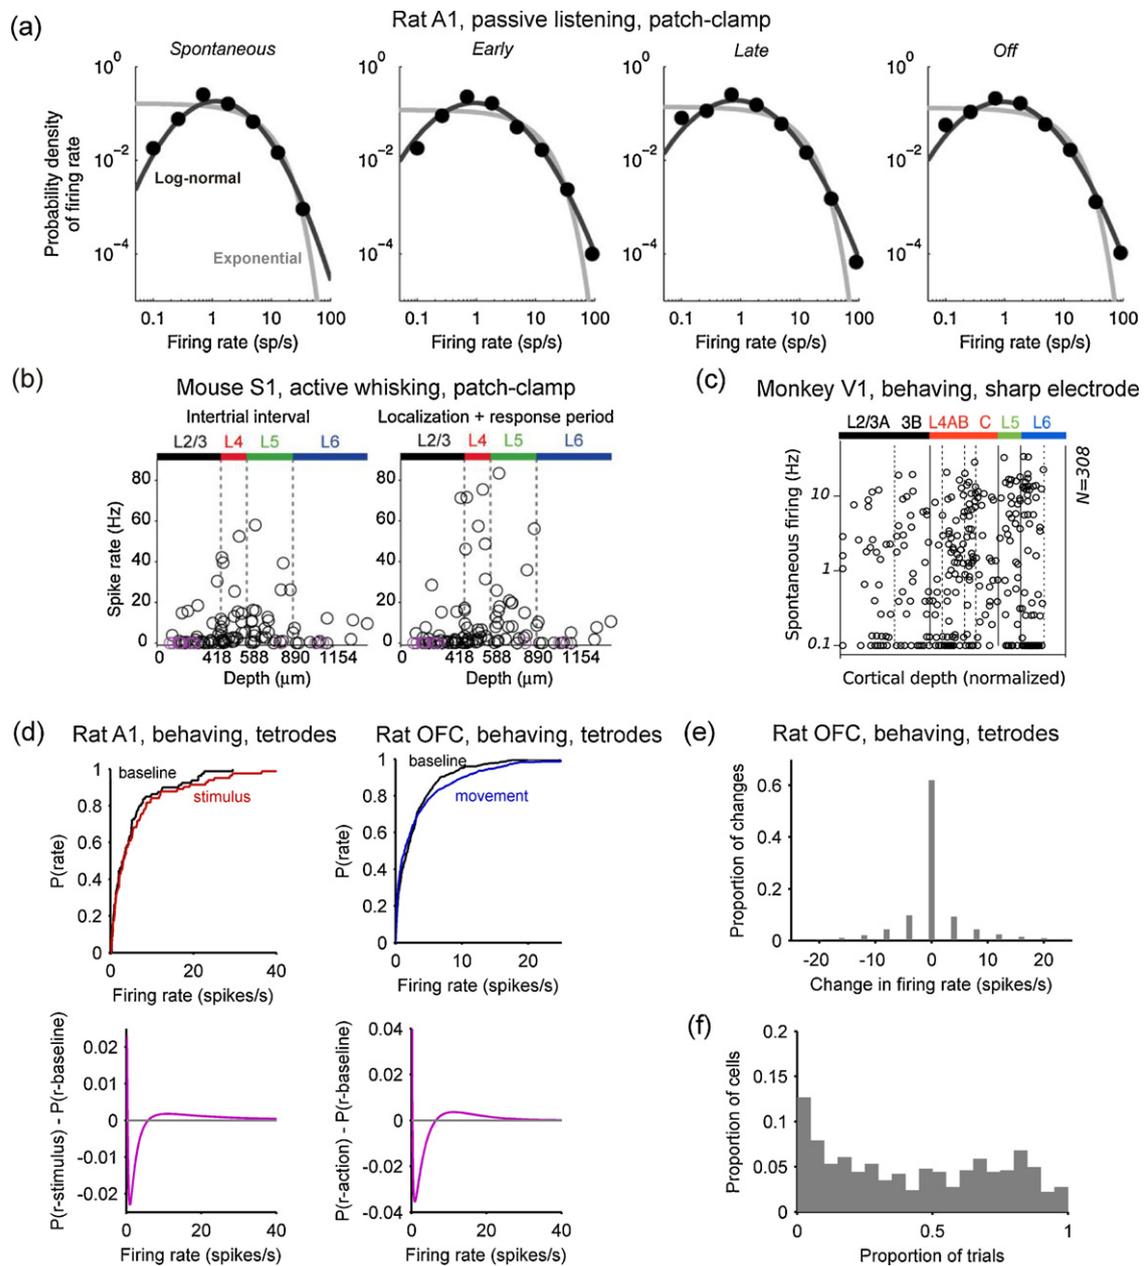
Fig. 4. Evoked changes from baseline firing-rate distributions are small. (a) Distribution of spike counts across neurons in the auditory cortex (A1) of awake rats responding to a pure tone pip. Activity distributions are shown before ('spontaneous'), during ('early', 'late') and after ('offset') the tone pip. Best model fits are provided, using respectively an exponential (gray) and a log-normal (black) distribution. The log-normal provides a markedly better fit. Loose patch clamp recordings, reproduced from Hromadka et al. (2008), with permission from the authors. (b) Distribution of spike counts in the layers of barrel cortex (S1) of awake mice involved in a tactile detection task. The mean level of activity is plotted for each neuron against its cortical depth, both during ('localization+response') and between ('intertrials') periods of whisker contact with the stimulus. Loose patch clamp recordings, reproduced from O'Connor et al. (2010), with permission from Elsevier. (c) Distribution of firing rates for spontaneous (baseline) activity in each layer of area V1 of monkeys observing a uniform-luminance screen; note firing rates are plotted on a log-scale (single extracellular micro-electrode recordings; replotted from Ringach et al., 2002). (d) Changes in mean firing rate distributions for two data-sets from Fig. 3: stimulus-evoked changes in A1 (tetrodes; data from Otazu et al., 2009), and movement-evoked changes in OFC (tetrodes; data from Feierstein et al., 2006). Top: the empirical cumulative probability distributions for the baseline ('spontaneous') and task-aligned firing rates. Bottom: difference in probability distribution functions for the best-fitting models to spontaneous activity, $p(r | baseline)$, stimulus-evoked activity, $p(r | stimulus)$, or action-related activity, $p(r | action)$ (the models are given in Fig. 3b), showing the smooth extent of the increase in higher firing rates in both A1 and OFC. (e) From rat OFC data (Feierstein et al., 2006), the distribution of rate changes between baseline and movement for every sampled firing rate (every neuron in every trial): no change occurred between baseline and movement in 62.1% of samples. (f) From rat OFC data (Feierstein et al., 2006), the distribution of the proportion of trials on which each neuron showed a difference in rate between baseline and movement; median proportion was 0.42.

the distributions were significantly different between layers, with layers 2/3 and 6 showing lower rates than layers 4 and 5 (the raw mean firing rate data are shown in Fig. 4b). Unfortunately, this is the current extent of layer-specific delineation of distributions; layer-specific distributions of the more complex models considered later in this review—of spike-train structure, feature encoding, temporal dynamics—have rarely been attempted.

Notable exceptions include the study of Ringach et al. (2002), which characterized the layer distributions of various metrics capturing V1 neurons' responses to moving gratings. Together, the work of Ringach et al. (2002), Hromadka et al. (2008), O'Connor et al. (2010) and the review of Barth and Poulet (2012) provide the first insights on layer-specific distributions, and indicate that far more information has yet to be discovered.

3.2. Spontaneous and evoked firing rate distributions are highly similar

If both spontaneous and task-aligned activity consists of long-tailed firing rate distributions, do these distributions differ? Answering this is the first application of taking a statistical model perspective of cortical activity: we can compare the probabilistic models for the different conditions to directly detect the effect of task events.

Stimulus representation in sensory neurons is probably one of the domains of electrophysiology where most literature is available, stemming from Adrian's first measurements of action potentials in nerve fibers responding to sensory stimulation (Adrian, 1928). Because many sensory neurons have always been found which respond vigorously to the presence of an adequate stimulus, it has generally been thought that population activity in visual cortices changes dramatically when a stimulus is presented (Barth and Poulet, 2012). Similar reasoning could be applied for the execution of an action, or the maintenance of working-memory during a delay period. We review recent results which have started to challenge this view thanks to new experimental techniques, revealing that population-wide statistics of task-driven and spontaneous activity may be much closer than previously expected.

3.2.1. Only small differences between distributions of spontaneous and evoked activity

Visual inspection of Fig. 4a and b shows that distributions of mean firing rates in rat A1 (Hromadka et al., 2008) and mouse S1 (O'Connor et al., 2010) barely changed between baseline and evoked firing. Similar results were obtained in PFC of monkeys performing a delayed match-to-sample task (Shafi et al., 2007), and the OFC of rats performing an olfactory discrimination task where distributions remained the same during epochs in which rats smelled an odor, moved, or received a reward (Feierstein et al., 2006). We could quantitatively reaffirm these results using the action-related activity in rat OFC (Feierstein et al., 2006) or stimulus-evoked activity in rat A1 (Otazu et al., 2009). Fig. 4d (top) visually confirms that the task-aligned distributions were not significantly different from their corresponding baseline distributions (two-sample Kolmogorov–Smirnov test, $p = 0.117$ for OFC; $p = 0.827$ for A1). Thus, there seems little change in the population firing rate distribution between spontaneous and evoked activity. Fig. 4d (bottom) illustrates the extent of the changes in both A1 and OFC by plotting the deviation between the models $p(r)$ fitted to the baseline and evoked firing rate distributions; from this perspective, we see a small and smooth increase in the probability of neurons appearing in the tail of the distributions of evoked firing rates.

3.2.2. The single neuron's perspective

These small changes in distribution obscure substantial changes at the single-neuron level between spontaneous (baseline) and evoked activity. O'Connor et al. (2010) reported that in mouse S1 34% neurons showed a significant change in the number of spikes evoked by whisker contact; the mean evoked change was between 2 and 4 spikes. In rat auditory cortex, Hromadka et al. (2008) always found around 5% of cells with strong responses to the stimulus, an increase of at least 20 spikes/s, independent of its exact nature (tone pips, white noise, natural stimuli). Still in rat auditory cortex, Bartho et al. (2009) presented a similar picture during the sustained response to tones, with most neurons returning to low activity levels while a few neurons kept on encoding stimulus value through their strong firing rates. Single neuron changes are thus both large (number of neurons showing a response over the whole experiment or individual neuron firing rate changes) and small (number of neurons showing a large response, or number of neurons simultaneously showing a response per trial). However, since firing rate changes of individual neurons can be both positive and negative, the overall effect of these changes on the firing rate distributions is rather small.

This large/small picture is also present in the action-evoked activity of rat OFC (Feierstein et al., 2006). The majority of the neurons in the data-set (349/515) had significant differences[5] (Mann–Whitney U-test, $p < 0.05$) between their rates in baseline and rates in movement epochs. However, on a trial-wise basis, these changes could be both negative and positive, and the dominant change from baseline to movement epochs was 0 spikes/s (Fig. 4e). Moreover, there was a wide-ranging distribution of the proportion of trials on which each neuron changed rate between baseline and movement epochs (Fig. 4f), with a median proportion of 42% of trials. Thus while the majority of neurons' rates were "significantly" altered by movement, the majority of neurons changed their rates on less than half the trials.

3.2.3. Whence the dramatic macro-scale changes in activity?

The above evidence argues that population distributions during spontaneous and stimulus-evoked responses of activity appear highly similar. How is this overall similarity of distributions compatible with the fluctuations in population-wide activity traditionally observed with imaging methods such as fMRI and EEG? Although the issue remains to be further explored, it seems that there is no fundamental contradiction between the two. Indeed, because overall firing rates in the population are so low, the few cells with strong response may actually be sufficient to significantly increase the overall number of spikes in the population in response to a stimulus. For example, Hromadka et al. (2008) report both that 16% of neurons contributed 50% of spikes and an overall increase of mean firing rate of almost 50% between pre-stimulus and early stimulus responses. Similarly, in mouse barrel cortex (S1) responding to a simple tactile discrimination task, O'Connor et al. (2010) report that around 10% of the cells in a single column account for 50% of all spikes fired. Their data also suggests that around 15% of the cells increased their firing rate by more than 10 spikes/s after the onset of active whisking (their Fig. S3, panel A1). Thus, the overall sparseness of responsive neurons may be compatible with significant increases of spiking activity, detectable using macro-scale recording methods. Since the distributions are long-tailed, this overall increase in average firing rate is compatible with only small changes in the distribution itself.

3.3. Firing rate distributions constrain theories of neural function

Why is it interesting that firing rate distributions are long-tailed? And why is it interesting that, statistically, the spontaneous and evoked activity distributions are so similar? We review here the implications of these first simple probabilistic models of cortical activity.

3.3.1. The theory of cell "classes"

From the perspective of probability distributions for firing rates, the classical physiological approach to defining distinct "cell classes" by firing rate is not appropriate. Traditional analyses make binary distinctions into two classes of neurons that fire at a significantly high or low rate, or that significantly fire to stimulus s or that significantly fire before action a and so on. But continuous long-tailed distributions show that there is nothing special about neurons with "significantly" higher activity; that amounts to

---

[5] 333/515 after adjustment for multiple comparisons using Benjamini–Hochberg test, with false discovery rate of 0.05.

drawing a threshold line somewhere on the distribution of firing rates. Moreover, these results show there are no grounds for demarcating responsive vs non-responsive neurons after stimulus, during delay, or before action. As Fig. 4 shows, when we fit a model to the distribution of rates, we can no longer distinguish such classes—the distribution shape changes smoothly. We can do statistical tests to distinguish responsive from non-responsive neurons, as we illustrated above for the OFC data, but from the distribution model perspective this is the wrong approach; to echo (Gelman and Stern, 2006; Nieuwenhuis et al., 2011): the difference between the significant and non-significant response is not itself significant.

3.3.2. Participation in computation

Knowledge of firing rate distributions provides constraints on and opens up new questions for theories of neural function. Any theories of neural function based on Gaussian distributions of population activity to encode stimuli or compute functions may require re-examining. Rather, the experimental results reviewed above support the view of *sparse* neural representations, where a relatively small number of active neurons mediate a population's function at any time.

Sparse representations constitute one of the first and most influential theoretical proposals concerning neural populations, whose roots can be traced back to ideas of redundancy reduction (Barlow, 1961, 1972) and of robust memory storage in neural networks (Marr, 1969; Kanerva, 1988). Sparse codes can be viewed as striking a balance between the two extreme views of population representation with 'grandmother cells' (one specific cell for each concept or object) at one extreme and fully distributed neural representations at the other. A strong conceptual quality of sparse-coding theories is to provide a functional explanation for the explosion in the number of neurons observed between subcortical structures and cortex: since a sparse code relies largely on the *identities* of active neurons rather than on their exact level of activity, the population size must be large enough to allow a rich and robust combinatorial repertoire of activities. In early sensory systems particularly, sparse coding is one of the most influential theories of population representation, further reviewed in Section 5.2.1.

While experimental observations of long-tailed population activity favorably echo the predictions of sparse coding theories, they also offer a promising glimpse on new theoretical perspectives. For example, the recurring observation of low but non-zero peaks in population distributions of activity (Figs. 3 and 4a), leading to better fits by log-normal rather than exponential models, is at odds with classical predictions of sparse coding models. And indeed, if most cortical neurons are barely active, then this raises the question of whether they meaningfully participate: either only the most active neurons are actually participating in the ongoing encoding or computation, and the sparse firing of occasional spikes in the remaining neurons is irrelevant (Shadlen and Newsome, 1998; Mazurek and Shadlen, 2002), or all active neurons participate. The latter view is supported by considerations of energy expenditure—the total number of spikes fired by the sparsely firing neurons is on a par with the total number of spikes fired by the few highly active neurons, thereby creating large energetic costs for cortex (Attwell and Laughlin, 2001). It is further supported by the observation that the dynamics of the local cortical circuit are exquisitely sensitive to the addition or deletion of a single spike (Izhikevich and Edelman, 2008; London et al., 2010).

3.3.3. Underlying cortical network

The long-tailed distributions also place strong constraints on the underlying cortical network and on the dynamical properties of cortical neurons. Earlier theoretical work predicted that simple balanced networks of excitatory and inhibitory neurons would have long-tailed distributions of firing rates (van Vreeswijk and Sompolinsky, 1996). A more complete theory was recently developed by Roxin et al. (2011), who showed that a long-tailed firing rate distribution in the population implies that each neuron's transfer function in vivo is an expansive non-linearity, such as a rising exponential or power-law function (Finn et al., 2007). They demonstrated that this result held across changes to the synaptic weight distribution, and held even if inhibitory interneurons were more active than the principal neurons.

This theory assumes a Gaussian distribution of input currents to each neuron, as generated in a uniformly random coupled network of neurons. An open question is whether the transfer function alone can generate long-tailed firing distributions when the Gaussian assumption is violated, whether by more realistic network topologies that promote localized synchrony (Galan et al., 2008; Haeusler et al., 2009), or by the non-Gaussian statistics of external inputs to the network, as has been demonstrated in rat A1 (DeWeese and Zador, 2006).

Side-stepping the assumption of Gaussian inputs, Koulakov et al. (2009) used a simple linear integration and linear transfer function model of an excitatory, recurrent cortical network to study how such long-tailed distributions of spontaneous activity might arise from aspects of network connectivity. They showed that a log-normal distribution of synaptic weights across the network is insufficient to generate log-normal firing rate distributions. Rather, a stable log-normal distribution of spontaneous firing rates arises if, in addition, synaptic inputs to each neuron are clustered by their mean weight, such that each neuron tends to receive either weak synaptic inputs or strong synaptic inputs. Open questions here are whether such a long-tailed distribution of rates would remain with any departure from linearity in the transfer function, and the impact of inhibitory neurons on the distribution.

3.3.4. Spontaneous activity as "browsing" through possible dynamical states

Several studies in the last decade have begun to investigate the functional and theoretical links between spontaneous and evoked activity. Joint studies of optical imaging and single-cell activities in anesthetized cat visual cortex have revealed a preserved link between a particular neuron's spiking and the overall activity in its embedding network, both during spontaneous and stimulus-evoked phases (Arieli et al., 1996; Tsodyks et al., 1999). In rat visual cortex, population-wide patterns of neural activation during a repeated stimulus presentation have been observed to replay during subsequent spontaneous activity (Han et al., 2008). In ferret visual cortex, Berkes et al. (2011) have found a quantitative similarity between the trial-to-trial probabilities of activity in the absence of visual stimulation and of activity averaged over natural visual stimulations. Furthermore, this similarity increased progressively during the early phases of visual development, pointing to sensory learning as a crucial shaping factor of spontaneous activity. Similarities have also been found in the temporal patterns of population activity between spontaneous and stimulus-evoked responses, both in anesthetized and awake states (MacLean et al., 2005; Luczak et al., 2009).

All these studies suggest that spontaneous activity may be akin to an internal (and experimentally uncontrolled) "browsing" through potential patterns of sensory-evoked responses. Consistent with this view, Luczak et al. (2009) have found that neural population activity in rat auditory cortex during sensory stimulation lives in a subset of possible activities observed during sensory spontaneous activity. From a more theoretical perspective, these results raise the exciting possibility that the organization of

spontaneous activity reflects an internal connectivity, learned by interaction with the environment, that would act as a probabilistic *prior* helping to interpret incoming stimuli (Ringach, 2009; Berkes et al., 2011). These theories have mostly operated on the level of few cells and considered the trial-to-trial variability of neural responses. What such theories predict in terms of cell-to-cell variability, and the respective population distributions of firing rates that we have discussed here, has yet to be determined.

4. Population distributions of spike-train statistics

In this section, we will review the population-wide statistics of spike-trains' structure across cortex. This is simply anything we know about distributions $p(r)$, when r is some quantification of each neuron's spike pattern beyond its rate. As the study of firing rate distributions alone is in a nascent stage, only a few studies have directly addressed this problem. So here we take a brief, broad survey of measures of spike-train structure, for which a population picture would provide us with useful additional constraints on neural theories and models.

4.1. Regularity of spike trains

The relative regularity of spikes emitted by a neuron is central to many theories of neural coding (Rieke et al., 1997; Gabbiani and Koch, 1998). Basic measures of neuron output irregularity such as the coefficient of variation ($CV = \sigma/\mu$) (e.g., Softky and Koch, 1993) and the Fano factor ($F = \sigma^2/\mu$) (e.g., Kara et al., 2000) are widely used. Both allow relative comparison between the output irregularity and a Poisson process, for which both CV and F equal 1 for different summary measures of the spike-train: the CV for inter-spike intervals, and the Fano factor for spike counts in time windows of equal length. However, these measures alone are difficult to compare over time or between neurons, as they are global measures that are sensitive to fluctuations in the neuron's firing rate (Softky and Koch, 1993; Holt et al., 1996; Ponce-Alvarez et al., 2010). That is, the standard deviation σ of the inter-spike intervals (for CV) or spike counts (for F) is computed using the global mean interval or count μ, but if that mean value is not locally stationary then irregularity tends to be over-estimated. Consequently, CV and F are not ideal for characterizing a population distribution $p(r)$ of spike-train regularity.

More advanced measures attempt to capture a rate-invariant irregularity, allowing meaningful comparisons between different neurons. One approach has been to compute a *rate-corrected CV* by computing separate CV values for inter-spike intervals grouped by sections of spike-train with the same firing rate (Softky and Koch, 1993; Maimon and Assad, 2009). Another has been to define alternative metrics that measure local changes in regularity, such as local variation (Shinomoto et al., 2009) and CV_2 (Holt et al., 1996). The CV_2 metric has been most widely used for detailed studies of spike-train regularity in cortex (Holt et al., 1996; Compte et al., 2003; Ponce-Alvarez et al., 2010; Hamaguchi et al., 2011). The comparative study of Ponce-Alvarez et al. (2010) suggested that of all tested metrics for regularity CV_2 was least affected by rate variation. We use it below, so note here its definition: if I_i and I_{i+1} are adjacent inter-spike intervals, then for that pair $CV_2(i) = 2|I_i - I_{i+1}|/(I_i + I_{i+1})$, and CV_2 is the average over all pairs; similar to other regularity metrics, $CV_2 = 1$ for a Poisson process.

4.1.1. Spike-train regularity changes across cortex and across tasks

Both the local-variation metric (Shinomoto et al., 2009) and the rate-corrected CV (Maimon and Assad, 2009) have been used to study the differences in spike-train regularity between cortical areas. Shinomoto et al. (2009) showed that, after accounting for rate differences, the average spike train changed from random

(Poisson-like) to bursty and then to regular patterning when moving from sensory to prefrontal and then to motor cortex in a data-set of multiple primate studies. Similarly, Maimon and Assad (2009) found that neurons in "sensory" parietal cortex (area MT) had irregular spike-trains, whereas neurons in "motor" parietal regions (LIP and Area 5) had more regular spike-trains. They also noted that the distributions of rate-corrected CV changed in shape and left-shifted from MT to LIP and then to Area 5 (their Fig. 4A). Thus in both these wide-ranging surveys of spike-train regularity across cortical areas, spike-trains were more regular in motor areas of cortex than in other regions, and both speculated that this was linked to the necessity for rate-coding in motor areas (Maimon and Assad, 2009; Shinomoto et al., 2009).

In a seminal paper, Compte et al. (2003) used measures of spike train structure including CV_2 to analyze single-unit recordings from the PFC of monkeys performing an oculomotor-delayed-response task. They reported that the spike trains of most neurons (around 60%) could be classified as Poisson-like, but that the rest of the neurons were more irregular (or bursty) than a Poisson model. Their most striking finding was that the population distribution of CV and CV_2 changed between the fixation and delay periods of the task (Fig. 5a). Thus, despite the overall increase in mean firing rate during the delay period,[6] the neurons' outputs were significantly more irregular than during the fixation period.

Hamaguchi et al. (2011) used CV_2 to analyze spike-train regularity of a database of single-unit recordings from the motor cortex of monkeys performing two delayed arm movement tasks. The population distribution of mean CV_2 was predominantly below $CV_2 = 1$ and left-skewed, suggesting marked regularity of single neuron firing in motor cortex (Fig. 5b). Moreover, when broken down by task stage, each neuron's CV_2 values were highly correlated between task stages ($r = 0.99$) and between the two arm movement tasks ($r = 0.86$), suggesting that each neuron had a fixed level of firing regularity. Also using single-unit recordings from motor cortex of a monkey performing a joystick task, Ponce-Alvarez et al. (2010) reported that CV_2 distributions did not differ between the two delay stages or arm movement stage of the task (though we note their reported CV_2 distributions were approximately symmetrical around 0.9, thus not consistent with the distributions in Hamaguchi et al., 2011).

Together, the analyses of Compte et al. (2003) and Hamaguchi et al. (2011) are consistent with the Shinomoto et al. (2009) and Maimon and Assad (2009) reports of greater regularity in motor regions of cortex than in prefrontal cortex. Moreover, the analyses of Compte et al. (2003) and Hamaguchi et al. (2011) together suggest that an additional difference in spike-train structure between motor cortex and PFC is that a neuron's irregularity is stable in motor cortex, but not in PFC.

The analyses of Compte et al. (2003), Ponce-Alvarez et al. (2010), and Hamaguchi et al. (2011) all pooled single-unit recordings using extracellular microelectrodes, and were all from primate studies. They were thus restricted to relatively high firing neurons in a single species. To further generalize these findings, we analyzed simultaneous recordings in rat PFC and in cat V2 for their distributions of CV_2. Fig. 5c shows that spike-train irregularity in rat PFC during a maze task had a symmetrical distribution around $CV_2 = 1$, and was thus consistent with the predominance of Poisson-like neurons reported by Compte et al. (2003) for primate PFC. Fig. 5d shows that spike-train regularity in cat V2 under anesthesia had a strongly left-skewed distribution, with a peak

---

[6] Note that this study used extracellular micro-electrode recordings and thus addressed only task-responsive neurons: the mean rate during fixation was 13 Hz, and during delay was 18 Hz; thus this increase was relatively small and likely confined to the long-tail of the population firing rate distribution, as illustrated in Fig. 4c.
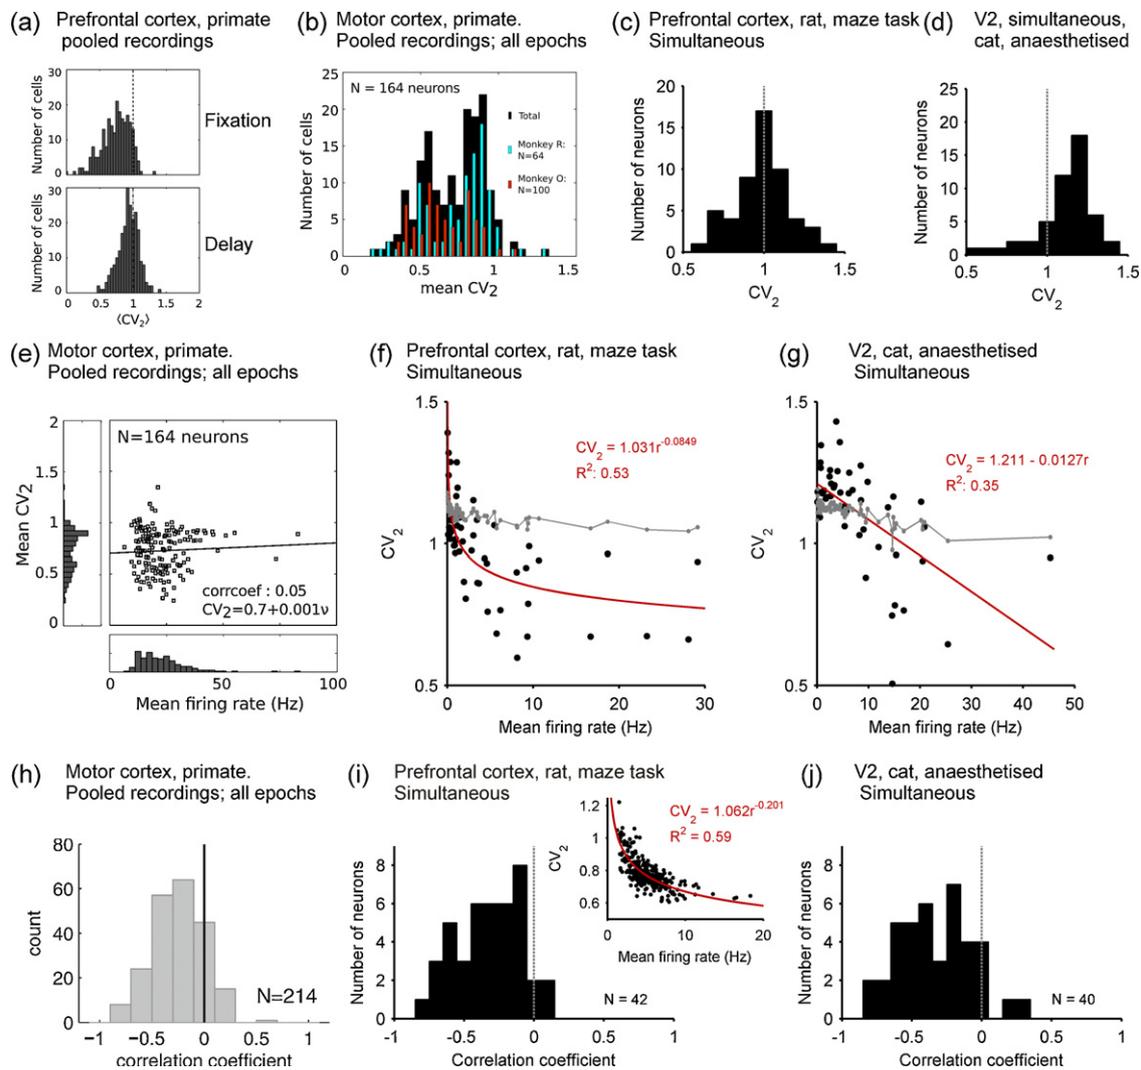
Fig. 5. Distributions of spike train regularity in cortex. (a) Distributions of rate-invariant irregularity measure CV_2 in primate PFC during the fixation (top) and delay (bottom) period of an oculomotor task; neurons pooled over multiple single-unit recordings. Replotted from Compte et al. (2003). (b) Distribution of CV_2 in primate motor cortex, averaged over all task stages; neurons pooled over multiple single-unit recordings. Replotted from Hamaguchi et al. (2011). (c) Distribution of CV_2 in a single tetrode recording from awake rat PFC. Data from study of Peyrache et al. (2009). (d) Distribution of CV_2 in a single polytrode recording from anesthetized cat V2. Data recorded by Tim Blanche (see Blanche et al., 2005), available from crcns.org. (e) Correlation between mean firing rate and CV_2 in primate motor cortex, calculated over all task stages. Note the lack of data-points below 10 Hz. Replotted from Hamaguchi et al. (2011). (f) Correlation between mean firing rate and CV_2 in a single recording from awake rat PFC; data: black symbols; best-fit model: red line. The gray symbols and lines give the predicted relationship if each spike-train was a rate-varying Poisson process: each point is the mean predicted CV_2 for a spike-train of that mean rate; 95% confidence intervals are too small to see on this scale. Data from study of Peyrache et al. (2009). (g) Correlation between mean firing rate and CV_2 in a single polytrode recording from anesthetized cat V2. Gray lines and symbols as in panel e. Data recorded by Tim Blanche (see Blanche et al., 2005), available from crcns.org. (h) Distribution of coefficients r for the correlation between rate and CV_2 for each neuron in a data-set of single-unit extracellular recordings from primate motor cortex during a joystick task. Reproduced from Ponce-Alvarez et al. (2010), with permission from Elsevier. (i) Distribution of coefficients r for the correlation between rate and CV_2 for a population of neurons simultaneously recorded in awake rat PFC; the inset shows an example neuron that had a power-law relationship between rate and CV_2. Data from study of Peyrache et al. (2009). (j) Distribution of coefficients r for the correlation between rate and CV_2 for a population of neurons simultaneously recorded in anesthetized cat V2. Data recorded by Tim Blanche (see Blanche et al., 2005), available from crcns.org.

around $CV_2 = 1.15$. Thus, under anesthesia many V2 neurons are more irregular than a Poisson process—while cortex was showing up/down state oscillations under this anesthesia (Humphries, 2011), the rate-invariance of the CV_2 measure suggests that the supra-Poisson spike-trains were independent of this oscillation. Nonetheless, these analyses of simultaneous recordings from rat PFC and cat V2 data confirm that, even when using recording techniques with less activity bias, much of cortical single neuron activity is still surprisingly well-described by a Poisson process, as both have peak CV_2 close to 1.

4.1.2. Correlations of regularity and rate in cortical populations

Once we have separately characterized distributions of firing rate, $p(r_1)$, and distributions of spike-train structure, $p(r_2)$, for a population, we may want to take the further step in model complexity of characterizing their co-variation. We illustrate this here for rate and regularity, as rate-invariant measures of spike-train regularity have made it possible to meaningfully consider whether there is a relationship between a neuron's rate and the regularity of the spikes it emits. The proper characterization of this covariation amounts to estimating a joint distribution of rate and regularity statistics, which we will denote $p(r_1, r_2)$.

Both Compte et al. (2003) and Hamaguchi et al. (2011) reported that in, respectively, primate PFC and motor cortex there was no relationship between rate and CV_2 within the data-sets they analyzed. This is illustrated in Fig. 5e, taken from Hamaguchi et al. (2011), which clearly shows the absence of any correlation. However, as noted above, the analyses of Compte et al. (2003) and

Hamaguchi et al. (2011) pooled single-unit recordings using extracellular microelectrodes, with rates estimated over short portions of a trial, and were thus limited to neurons with relatively high firing rates; neither data-set contained neurons with rates less than 1 Hz. As we saw in Section 3.1, such low firing rates should dominate the population distribution. We thus also analyzed example simultaneous recordings with low-rate dominated distributions to check their joint population distribution of rate and regularity.

We found a clear negative correlation between mean firing rate and CV_2 in single recordings from both rat PFC (Fig. 5f) and cat V2 (Fig. 5g). While CV_2 is notionally rate-independent, to our knowledge its ability to cope with orders-of-magnitude firing rate fluctuations in ultra-low firing neurons (<1 Hz) has not been tested; thus, plausibly, this correlation at such low frequencies could have been an artifact of such rate fluctuations. To rule this out, we found the expected CV_2 for a null model in which each spike train was modeled as an inhomogenous Poisson process with varying rate derived from the data.[7] We found that this model predicted a relationship between mean firing rate and CV_2 (gray symbols and lines in panels f and g) that departed strongly from the data. Thus, the correlation between mean firing rate and CV_2 was not an artifact of rate fluctuations.

We also found the correlation differed between the rat PFC and cat V2: CV_2 was a power-law function of firing rate in awake rat PFC, but a linear function of firing rate in anesthetized cat V2.[8] The power-law fit in rat PFC shows that rate-regularity correlation occurs at very low firing rates and that there is little correlation above $\sim$5 Hz, where the power-law function flattens out. Thus in this range there appears no correlation between rate and regularity, consistent with this results of Compte et al. (2003) and Hamaguchi et al. (2011). The linear fit over all rates in cat V2 is potentially a result of the unique spike-train structure that occurs under slow-wave inducing anesthesia. Nonetheless, less-biased recordings dominated by low-firing neurons do reveal a different perspective on the joint distributions of rate and regularity.

4.1.3. The single neuron's perspective: joint distributions of regularity and rate modulation

An orthogonal perspective is to ask whether, over longer periods of time, each neuron has a correlation between its own rate and regularity, and thus determine the population distribution of rate-regularity modulation. (It is possible that the population show no correlation between the mean rate and mean CV_2 of its neurons, yet each neuron has rate-dependent modulation of its own regularity). Ponce-Alvarez et al. (2010) computed firing rate and CV_2 in windows of 200 ms for a data-set of 214 extracellular recordings in primate motor cortex, and found a large sub-set of neurons with a significantly negative correlation between their own rate and CV_2 (118/214 had $r < 0$ and $p < 0.01$). The distribution of correlation coefficients is plotted in Fig. 5h, and shows a left-skewed distribution with a peak close to zero. Thus neurons in primate motor cortex had a broad distribution of negative rate-regularity modulations.

Similar to the above problems with the population-level correlation between rate and regularity, the data-set in Ponce-Alvarez et al. (2010) was from extracellular recordings and thus biased toward high firing rates. To generalize their results, we applied a similar analysis to single recordings of population activity from awake rat PFC and anesthetized cat V2. For each neuron we calculated firing rate and CV_2 in sliding windows of 20 s advanced in steps of 5 s.[9] Similar to Ponce-Alvarez et al. (2010), we ensured sufficient data for a reliable correlation by omitting from each neuron any window with less than 20 inter-spike intervals, and omitted any neuron that had less than 40 windows as a result. Nonetheless, despite the dominance of low firing rates, we found a sub-set of neurons in both data-sets had strong negative correlations between their own spike rate and regularity: 26/42 in awake rat PFC and 20/40 in anesthetized cat V2 had $r < 0$ and $p < 0.01$. Fig. 5i and j shows that both PFC and V2 had a broad distribution of negative correlation coefficients, and thus seemingly a continuum of the strength of regularity modulations.

Moreover, we found that, beyond the simple linear correlation coefficient, most neurons with strong modulation of regularity actually have a non-linear correlation between rate and CV_2. The majority of strongly modulated neurons in both rat PFC and cat V2 had power-law relationships, as illustrated in Fig. 5i. Thus, similar to the population distribution, for individual neurons modulation of regularity happens at low firing rates.

4.2. Beyond spike-train regularity: self-correlation and precision

Measures of spike-train regularity are perhaps the most widely applied measures of spike-train structure. A wide-range of other measures are possible, but for which we lack knowledge of population distributions. We briefly review below two further aspects of spike-train structure that may yield insights into population coding. This is not exhaustive: further examples of measures amenable to population distributions include burst structure (e.g., Gourevitch and Eggermont, 2007), spiking coherence with an underlying oscillation (e.g., Benchenane et al., 2010), and of peak(s) in power spectra (e.g., Bair et al., 1994).

4.2.1. Cortical neurons have varying short time-scale self-correlation

Neurons can have detectable correlations between their inter-spike intervals (Perkel et al., 1967). Simulations of simple neuron models have shown these correlations may imply improved transmission of information about time-varying stimuli (Chacron et al., 2004). Such *serial* correlations can be computed up to any order n for the correlation between intervals I_i and intervals I_{i+n}, but are typically computed only for adjacent intervals ($n = 1$) because of the length of stationary recording required to compute sufficient higher-order interval correlations. Serial correlations between intervals are logically distinct from the properties of bursts and oscillations: a bursting or oscillating neuron is likely, but not guaranteed, to have serial correlations in its spike-train; a neuron with serial correlations in its spike-train does not imply that neuron is bursting or oscillating.

While serial correlations have been reported in a wide range of neural systems (Farkhooi et al., 2009), few studies have measured this property in mammalian cortex (Nawrot et al., 2007; Engel et al., 2008). Nonetheless, consistent with reports from other neural systems, this limited evidence suggests that some neurons in both rat S1 and entorhinal cortex have significant, negative serial correlations between adjacent intervals: thus long inter-spike intervals tend to be followed by short intervals, and vice-versa. As neurons with and without negative serial correlations may have different optimal inputs for transferring information (Lindner et al., 2005), understanding the prevalence and distribu-

---

[7] For each spike train we estimated its conditional intensity function $\lambda(t)$ to generate the inhomogenous Poisson process. We computed a rate histogram for the spike train using its mean inter-spike interval as the bin size, then fitted a cubic spline to estimate a smooth intensity function $\lambda(t)$. Using the estimated $\lambda(t)$, 100 realizations of the inhomogenous Poisson process were then created to compute the mean and 95% confidence intervals of the predicted CV_2 for each spike train.

[8] We fitted linear, single exponential (2 and 3 parameter), double exponential, power-law, truncated power-law and stretched exponential models to the data; the best model was chosen by lowest BIC score.

[9] We used a window and step-size two orders of magnitude larger than that used by Ponce-Alvarez et al., 2010), to account for the order-of-magnitude lower firing rates in both PFC and V2 data-sets (Fig. 3a).

tion of serial correlations within and between cortical populations could shed light on their information coding capacity.

4.2.2. Cortical neurons have variable firing precision

Separate from their intrinsic irregularity, a population of neurons may respond to repeated external stimuli or internal state transitions with different precision in their trial-to-trial repetitions of spike times. There is recent evidence that such changes in spike-time precision are similar for both internally and externally driven state changes. Luczak et al. (2007) reported that neurons in layer 5 of rat somatosensory cortex had a repeating sequential structure of each neuron's mean spike time at the onset of spontaneous UP-states; this ordering of spike-timing disappeared during the rest of the subsequent UP-state. The sequence of neuron firing during spontaneous activity was recapitulated following sensory input, in both anesthetized and awake animals (Luczak et al., 2009). Peyrache et al. (2010) showed the same sequential reinstatement of activity at the onset of spontaneous UP-states in neurons of medial prefrontal cortex in sleeping rats; they further confirmed that deriving these sequences from mean spike-time at the onset of the UP-states meant that the sequence of neuron firing was independent from the overall firing rate of each neuron.

As an approximation to such state transitions, we may also simply study the responses of neurons to the onset of noisy current inputs. The degree of input-driven precision is known to depend on both the magnitude of the noise in the current (Hunter et al., 1998), the resonant firing frequency of the post-synaptic neuron (Hunter et al., 1998; Fellous et al., 2001) and the time-scale of correlation in the noise (Galan et al., 2008). Such boundaries on reliability suggest there should be different distributions of spike-time reliability between different areas and layers, depending on the exact nature of the local microcircuit, and of the inputs from other cortical and sub-cortical regions. Consistent with this, data from cat V1 suggest that layer 4 neurons can produce a reliably timed first spike to a repeated stimulus, and that this reliability is subsequently reduced in both layers 2/3 and 5 (Cardin et al., 2010). While there have been efforts to quantify the reliability of individual neuron spike-timing using correlations between evoked spike patterns (Fellous et al., 2004; Humphries, 2011), to our knowledge we lack any *in vitro* or *in vivo* characterization of the population distribution of spike-timing precision.

4.3. Implications of distributions of spike-train metrics

In general, distributions of metrics for spike-train structure across cortical populations may place strong constraints on either or both of single neuron and network properties. We finish this survey by considering several implications for theories and models of cortex that might arise from knowing these distributions.

4.3.1. Cell "classes" are inconsistent with unimodal distributions

From the perspective of probability distributions for neural populations, the classical physiological approach to defining distinct "cell classes" based on spike-train structure is problematic. It is typical in electrophysiological studies to make sense of the data by demarcating the firing patterns of neurons into categories such as "regular", "irregular", "bursting" and so on. But by characterizing distributions of the spike-train measures used to define these categories (e.g., Fig. 5a–d), we can see that this approach amounts to drawing an arbitrary threshold line somewhere on the distribution—above that line is irregular, below is regular, and so on. Similarly, for joint distributions of rate and regularity, we can reasonably fit unimodal probabilistic models relating the two; and for single neurons there is a broad distribution of correlations between rate and regularity. Attempts

to define classes of neurons by spike-train structure may be more successful when using multiple defining metrics, as in Compte et al. (2003), or when more than one peak is clear in the distribution of a single metric, as suggested by the CV_2 data in Fig. 5b.

4.3.2. The irregularity of cortical firing and the balance of excitation and inhibition

The irregularity of cortical spike trains and its implication for single neurons or neural networks has been the focus of numerous studies. *In vitro* studies of cortical neurons have shown that they can reproduce spike-trains with millisecond-scale precision to repeated injections of the same noisy current, but not to injections of constant current (Mainen and Sejnowski, 1995; Fellous et al., 2001; Galan et al., 2008). Such data suggest that, under synaptic barrage *in vivo*, cortical neurons are capable of reproducing specific patterns of spike times. There is some evidence from a range of cortical areas for spike-timing precision in response to repeated stimuli (DeWeese et al., 2003; Fellous et al., 2004). Consequently, the irregularity of cortical neuron firing likely reflects inherently irregular ("noisy") inputs that in turn generate the irregular output.

Such constraints often act as a major spur for a rich vein of interacting computational and experimental work. For example, the report by Softky and Koch (1993) that cortical cells have consistently irregular spike-trains, inconsistent with a single neuron integrating random inhibitory and excitatory synaptic input, inspired a wealth of computational and experimental work to explain the source of the irregularity (Stevens and Zador, 1998). Computational models were able to account for the consistent irregularity through both single cell (Troyer and Miller, 1997; Gutkin and Ermentrout, 1998) and network properties (Tsodyks and Sejnowski, 1995; van Vreeswijk and Sompolinsky, 1996); one outcome of this work was the proposal of balanced excitation and inhibition in the inputs to cortical pyramidal cells (van Vreeswijk and Sompolinsky, 1996; Shadlen and Newsome, 1998; Salinas and Sejnowski, 2000). The theoretical studies in turn drove experimental work which confirmed that this balance exists (Shu et al., 2003; Haider et al., 2006; Renart et al., 2010). These series of studies provide probably the best explanation for the irregularity of firing in cortex and have helped to elucidate the "operating point" of cortical neurons *in vivo*.

The findings of Compte et al. (2003) that firing irregularity was enhanced in primate PFC during a delay period has spurred further theoretical work, as this result refuted the then prevalent models of working memory in cortical circuits. The impact of the Compte et al. (2003) results has been excellently reviewed by Barbieri (2008), whom we briefly summarize. The extant models of persistent activity were based on simple attractor networks: randomly connected, recurrent excitatory circuits that explained how transient stimuli could cause a sustained jump in the firing rates of a select sub-population of neurons, and thus act as a short-term memory for that stimulus. However, such networks predict that spike-train irregularity should dramatically decrease, not increase, during the persistent firing. By introducing either inhibitory feedback that scaled with the excitation or short-term depression at the excitatory synapses, a new wave of attractor models could account for the maintenance or increase in average irregularity during the persistent activity (Barbieri, 2008). The next step will be to find the constraints placed on such networks by attempting to capture the distributions of irregularity (Fig. 5b).

4.3.3. Implications of distributions of regularity for cortical microcircuits

All this previous work focused on the approximately Poisson output of single neurons. To our knowledge, no attempt has been made to capture the population distribution of irregularity in a

model of the cortex. Hamaguchi et al. (2011) have made a promising initial foray. From a data-set of single-unit recordings in primate motor cortex, they determined for each neuron the values of, and correlation between, rate and CV_2 over a set of short time windows in one task condition. They used a single LIF model neuron to replicate this data for each neuron, by searching over three parameters that accounted for the neuron's embedding in a sparse, random network: the balance of inhibitory and excitatory synaptic strength g, the maximal excitatory synaptic strength J, and the number of excitatory synapses C_e. The values of these parameters would thus predict the type of local network into which each neuron in the data-set was embedded. They found that most neurons in the data were well-fit by a wide range of values for these parameters, suggesting a broad class of cortical circuits that could account for the observed rate and CV_2 values. They did, however, deduce some general predictions from the range of best-fits: most neurons were inhibition-dominated ($g > 4$) with moderate numbers of excitatory synapses ($C_e > 100$), combined with high synaptic strengths J for irregular firing or low synaptic strengths J for regular firing. Consequently, their models predict that motor cortex neurons are mostly embedded in an inhibition-dominated network, whereas the most regular neurons have few and/or weak input connections. Again these data were restricted to high (>5 Hz) firing rates. The more challenging questions remain of how to account for strong correlations between rate and regularity at low firing rates (Fig. 5h–j), and what networks can actually reproduce the population distribution of CV_2 (Fig. 5a–d).

All of the above work is also focused on a single area of cortex. Differences in spike-train regularity between cortical areas imply either that the neurons in each area have differences in their dynamical properties that change the integration of their inputs (e.g., relative expression of receptors, or specific ion channels in dendrites) or that the local circuit into which they are embedded changes between areas to favor irregular or regular output, perhaps moving from a balanced to an unbalanced input regime (Softky and Koch, 1993; Stevens and Zador, 1998)—as also suggested by the models of Hamaguchi et al. (2011). Consequently, characterizing population distributions of spike-structure across cortex will do much to further our understanding of the cortical architecture.

5. Population distributions of feature encoding

The previous sections have described the overall distribution of neural activities (firing rates and higher order spiking statistics) in various brain areas and under different experimental conditions, including spontaneous activity, stimulus presentation and action execution. One key aspect of these distributions $p(r)$ seems to be their overall invariance to the exact feature—stimulus or motion—encoded (Section 3.2). How are we then to quantify *feature encoding* by neural populations, i.e., reallocations of activity in the population which allow the system to encode a particular stimulus, or action, differentially from other possible stimuli or actions?

5.1. Information is distributed heterogeneously across neurons

From a statistical perspective, characterizing such distributions amounts to a tremendously difficult problem, as it combines dimensionality explosions along two distinct axes: (i) the huge number of neurons involved, and (ii) the huge number of possible stimuli (or actions or tasks) that the system might be confronted with. In this review, we circumvent the second dimensionality explosion by focusing on low-dimensional stimulus spaces. Nonetheless, we warn the reader that not many studies (at the scale of sensory neuroscience literature) have quantified

distributions over stimulus encoding features. Where they have been quantified, they likely suffer from some of the biases described in Section 2.2 (Olshausen and Field, 2005). Apart from reviewing a few seminal studies, we also emphasize how important it is to obtain quantitative estimates for these distributions, both for our understanding of what cortex does, and for testing theories of feature encoding.

5.1.1. Population distributions of receptive field sizes

We will illustrate the main issues with a classical example. Consider one of the oldest characterizations of neural encoding: *receptive fields*, simply defined as the domain of physical space (retinal space for vision, frequency domain for audition, skin location for somatosensation) to which a sensory neuron is responsive. Population distributions of receptive field sizes have long been studied in early sensory cortical areas such as V1 (e.g., Hubel and Wiesel, 1962; Schiller et al., 1976; Dow et al., 1981), A1 (e.g., Schreiner et al., 2000; Moshitch et al., 2006) or S1 (e.g., Sur et al., 1980; DiCarlo et al., 1998). A common characteristic of these distributions is that they are quite broad (Fig. 6). In Fig. 6a (Dow et al., 1981), the sizes of V1 receptive fields are plotted against their retinal eccentricity (top curve, left ordinate axis), and typically display a large spread of sizes at each given eccentricity. In Fig. 6b (Schreiner et al., 2000), the (normalized) receptive bandwidths of A1 neurons are plotted for two intensity levels (10 dB and 40 dB). As in V1, the auditory receptive field sizes (or bandwidths) are heterogeneously distributed across the population, with over ten-fold variations in their (normalized) sizes irrespective of the stimulus' intensity; though increased intensity leads to a global broadening of the receptive field sizes. In Fig. 6c (DiCarlo et al., 1998), similar histograms are built for receptive field sizes in S1 (plus their excitatory and inhibitory subfields), again revealing the same diversity. Visual inspection of Fig. 6b and c suggests that the distributions are somewhat close to Gaussians in logarithmic space, and therefore potentially long-tailed.

Surprisingly, the observed variability in receptive field sizes is not a trivial consequence of mapping properties at the level of receptors. Visual and tactile spaces are known to follow an inhomogeneous mapping with a high density of representation in key areas (fovea for vision, hand and face for somatosensation) and lower density in other areas (visual periphery, limbs and trunk for somatosensation). Similarly, from the cochlea on, auditory space is naturally endowed with a logarithmic structure over frequencies. However, these factors cannot explain the heterogeneity in Fig. 6. In panel (a), cortical magnification—the ratio of interneuronal distances on the cortex and in sensory space—shows much less variation at each given eccentricity (bottom curve, right ordinate axis), so it cannot account for the variation of receptive field sizes. In panel (b), all quantities are treated in log-scale, and each neuron's effective bandwidth (i.e., receptive field size) is defined with respect to its preferred characteristic frequency, yet variability remains the same. Finally in panel (c), all receptive fields come from distal fingertips and have therefore rather constant magnifications (Sur et al., 1980). Hence, no matter the sensory modality, over ten-fold variations can be observed in the receptive field sizes of the cortical neurons covering any given point in stimulus space.

The diversity of receptive field sizes is therefore larger than expected from peripheral processing. In fact, receptive field sizes in the retina are rather homogeneously distributed at each retinal eccentricity (Dacey and Petersen, 1992). To a lesser extent, similar principles may hold for the cochlea. Especially at lower frequencies, the bandwidths may be more narrowly distributed (Carney et al., 1999, reproduced in Fig. 9e), although a strict, quantitative comparison between peripheral and cortical auditory tuning has yet to be carried out. Similarly, we are not aware of quantitative
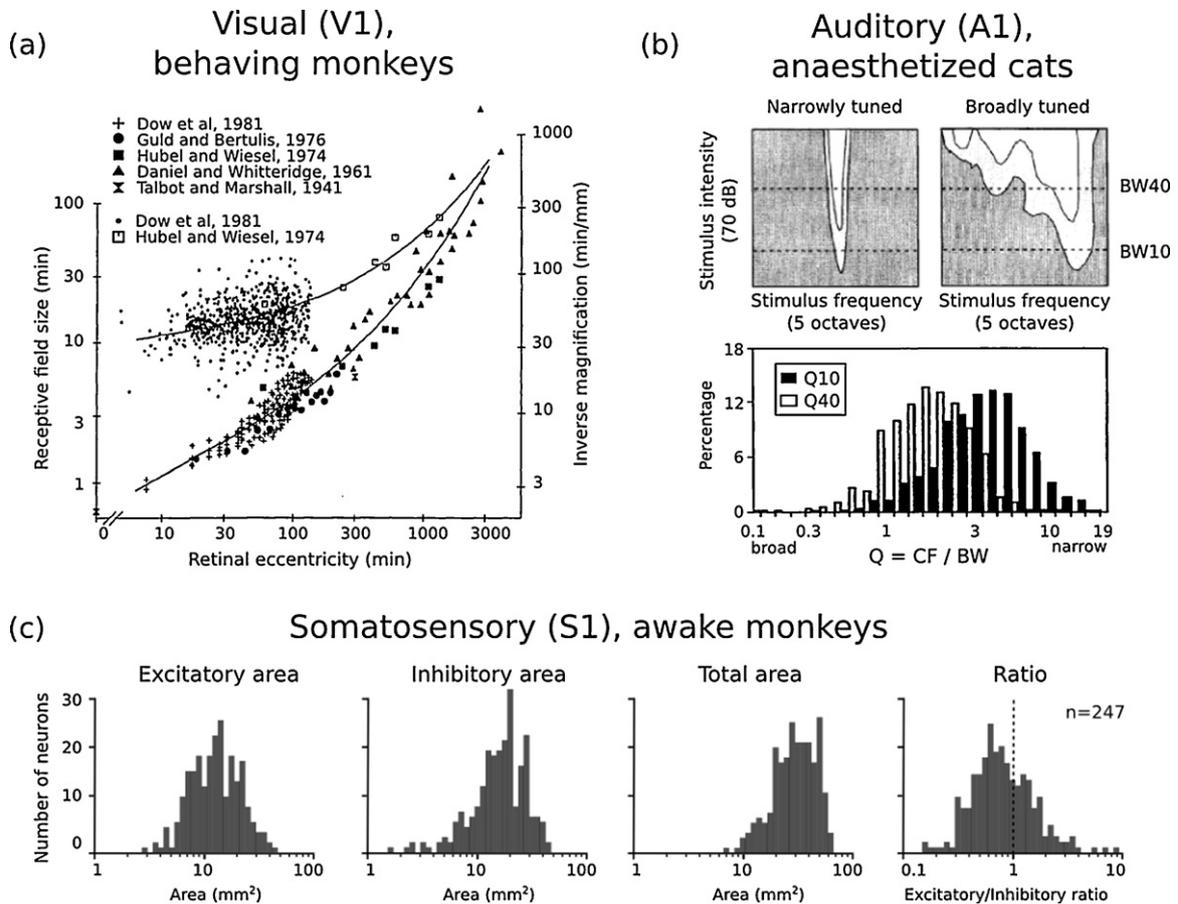
Fig. 6. Cortical receptive fields sizes are broadly distributed. (a) Distribution of receptive field sizes (top curve, left ordinates) and inverse cortical magnification (bottom curve, right ordinates) in monkey V1 neurons as a function of retinal eccentricity; reproduced from Dow et al. (1981), with permission from Springer. (b) Top panels: sample responses of A1 neurons to pure tones, as a function of tone frequency, and intensity (measured in dB above the neuron's threshold of response). Bottom panel: distribution of relative response bandwidths (best frequency divided by bandwidth) in cat A1 neurons, at respectively 10 dB and 40 dB above threshold; reproduced from Schreiner et al. (2000), with permission from Annual Reviews, Inc. (c) Distribution of receptive field sizes in monkey S1 neurons from the distal digit pads. From left to right, population histograms are shown for the excitatory subfields, inhibitory subfields, full receptive fields, and for the exc/inh ratio of areas; replotted from DiCarlo et al. (1998).

distributions of receptive field sizes for mechanoreceptors. A visual inspection of typical sizes in Talbot et al. (1968) reveals some level of heterogeneity, as can be expected from their embedding in the irregular skin surface, but the mean and standard deviations of receptive field sizes reported suggest less heterogeneity than in S1 (Talbot et al., 1968). It thus seems that in each case, stimulus information is actively transformed from a relatively homogeneous tiling of stimulus space at the periphery to a relatively heterogeneous tiling in sensory cortex.

5.1.2. The tiling of feature spaces

Naturally, receptive field sizes are just one particularly simple example out of the many features which can be encoded in cortical neurons. Their tiling of the physical stimulus space, however, provides a natural intuition for the characterization of stimulus encoding in neuron populations. Fig. 7 illustrates the basic ideas by focusing on a two-dimensional stimulus space, i.e., a space in which each stimulus is defined by just two numbers, $\mathbf{s} = (s_1, s_2)$. This example maps directly on the receptive field example if we
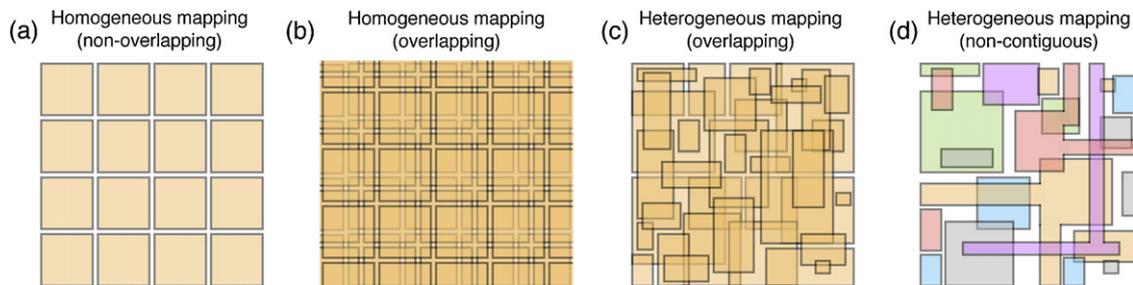
Fig. 7. Schematics for tiling of stimulus spaces by a neural population. Each tile corresponds to the part of stimulus space to which a given neuron is responsive. (a and b) Representations of peripheral systems are often quite homogeneous. (c) Cortical representations appear usually more heterogeneous, resulting in broader population distributions. (d) In higher cortical areas, individual neurons may develop fragmented representations of the original stimulus space. Here, different neurons are depicted with different colors. For visual clarity, only a few neurons are shown.

assume that the stimulus is the position of a point in the physical space, such as an illuminated dot in (two-dimensional) retinal coordinates. Each neuron covers a part of this stimulus space, illustrated by the tiles in Fig. 7, and responds whenever the stimulus falls into that part of the space. This way, the whole stimulus space is covered by the population of neurons, and this covering (or tiling) can be arranged in many different ways. In Fig. 7a, only one neuron will fire for a given stimulus, while all other neurons remain silent. In Fig. 7b and c, several neurons cover each location in stimulus space, and hence several neurons will fire for a given stimulus.

Of course, stimulus space is not two-dimensional. In a complex visual image, for instance, every single pixel can have a different light intensity, so that every pixel functions as a stimulus dimension. Nonetheless, the intuitive picture developed above generalizes to this high-dimensional space. (It also extends to non-sensory information, such as the encoding of a specific movement sequence.) The essential question of information encoding in populations is always how the neurons tile the respective space. Is the tiling homogeneous (Fig. 7a and b) or heterogeneous (Fig. 7c)? Is it complete (Fig. 7a) or over-complete (Fig. 7b and c)? Do individual neurons cover simple and contiguous (Fig. 7a–c), or complex and non-contiguous (Fig. 7d) regions of stimulus space?

In each case, the contribution of a neuron to the tiling of the respective space can be summarized by a few parameters, $\mathbf{q} = (q_1, q_2, \ldots)$, that measure, for instance, the position of the neuron's tile, its extent, shape, and so on. The probability of finding a neuron with the respective properties, $p(\mathbf{q})$, can then be characterized by constructing a statistical model of the tiling. In the receptive field example discussed above, we studied the sizes of the different tiles, thereby providing a first example of how information is distributed in populations of neurons. Clearly, the examples in Fig. 6 show that tiling is both overcomplete and heterogeneous as shown in Fig. 7c.

There are, however, two complications that we need to take into account. One complication is that a neuron's response to a stimulus is not binary. Rather, its response should be quantified through the average firing rate, or, if we want to be more general, through the probabilities of its spikes. In either case, the response becomes analogue, and instead of a region or tile in stimulus space, a neuron's response is better visualized as a cloud that is thick wherever the averaged response is large, and thin wherever the averaged response is small. In turn, we speak of the *tuning* of the neuron to the different stimuli, and call its response a tuning curve whenever the stimulus space is small. Such a tuning curve is typically characterized by a small set of parameters that measures the width of the tuning, the amplitude of the response, etc.

A second, quite fundamental problem appears when we apply the overall approach to higher sensory areas. Within the confines of the receptive field example, for instance, the tiling becomes more and more complex as we move up the sensory hierarchy. Specifically, individual neurons start responding to different, non-contiguous regions in stimulus space. Indeed, the receptive fields of neurons are known to have multiple peaks in A1 (Sutter and Schreiner, 1991) or S2 (Fitzgerald et al., 2006). Hence, the simple tiling pictures breaks down and the tiles of individual neurons acquire quite complex shapes (Fig. 7d), making it extremely difficult to characterize the stimulus-response function of individual neurons. The reason for this complexity is, of course, that the respective cortical systems are extracting complex features from the sensory stimuli.

A solution to this problem, at least in principle, is to change the definition of stimulus space. As we change the stimulus space (e.g., from sound pressure level to the amplitude of different frequency bands, from the light intensities of pixels to the angle and size of a given object), the respective tiling changes. In other words, how information is distributed within a population of neurons depends on how we choose to represent that information. This problem makes it far more difficult to observe commonalities in information representation across different cortical populations, as any observed differences (or commonalities) interact with the chosen definition of the stimulus space. That said, we believe that there are commonalities to be found, and that it will be important to quantify them. An exhaustive review of such a complex topic is impossible, and we will just illustrate a few examples.

5.1.3. Population distributions of tuning curves

A particularly common example is given by stimuli or actions that have the topology of an angle (orientation, direction of movement, direction of reach, etc.). These have been extensively studied because of their convenience: they live on a one-dimensional "ring", so suffer from no problems linked to defining boundaries. Experimentally, tuning curves to such stimuli are generally simple, displaying a single bump of tuning around the neuron's preferred angle (e.g., Fig. 8a and c); this makes them suitable for easy population characterizations. In V1, Busse et al. (2009) have recently derived the population distribution of tuning in neurons responding to gratings of different orientations and contrasts (Fig. 8a and b). They found the neurons' width of tuning and semi-saturation contrast to be independently distributed across the population (panel b), and suggested that this independence promotes a global contrast-invariance in the population response (their Fig. 1D). Similarly, in primary motor cortex encoding of hand movement direction (Georgopoulos et al., 1982), Amirikian et al. (2000) reported a broad distribution of tuning curve widths (Fig. 8c and d) with values markedly narrower than the 90° used in traditional studies of population decoding in hand movement tasks (Georgopoulos et al., 1982). (We note, though, that 39% of the recorded neurons were yet more heterogeneous, with bimodal or asymmetric tuning curves that are not included in Fig. 8d.) Because of their simplicity, ring topologies are also favored by theoretical works on neural populations, such as attractor networks (Compte et al., 2000) or probabilistic population codes (see Section 5.2.2). Ultimately, population distributions of tuning like those presented here will set benchmarks to test these theories.

In some cases, population distributions of tuning parameters can help to infer the nature of information encoding in a population. For example, neurons in somatosensory cortex S1 are known to be tuned in two ways to the frequency of vibrating stimuli (Salinas et al., 2000). Some neurons transmit the frequency value through the periodicity of their own spike trains ("periodic neuron" in Fig. 8e), while others transmit this value by firing more spikes as the frequency increases ("rate neuron" in Fig. 8e). When plotted as a distribution, both types of tuning were found to be present in the population (Fig. 8f), with all possible strengths and combinations. Mechanoreceptor afferents typically display strong periodic tuning and weak rate tuning (as in the top left corner of Fig. 8f); however in S1, modulations in rate, not periodicity, have been shown to best covary with behavior on a trial-to-trial basis (Salinas et al., 2000). Hence, the appearance of aperiodic, rate-tuned neurons in cortex (bottom right corner in Fig. 8f) is likely due to an active reshaping of periodicity information into firing rate information. Fig. 8f suggests that this reshaping is a population phenomenon, and sets quantitative constraints on future models of how it is implemented in the brain.

As a second example, we consider the activity of neurons in various sensory areas during two-alternative behavioral tasks. Many studies have shown a population-wide correlation between neurons' sensitivity to a stimulus and their *choice probability* measure (Fig. 8h). Choice probability (CP) quantifies the trial-to-trial covariations between the activity of a sensory neuron during
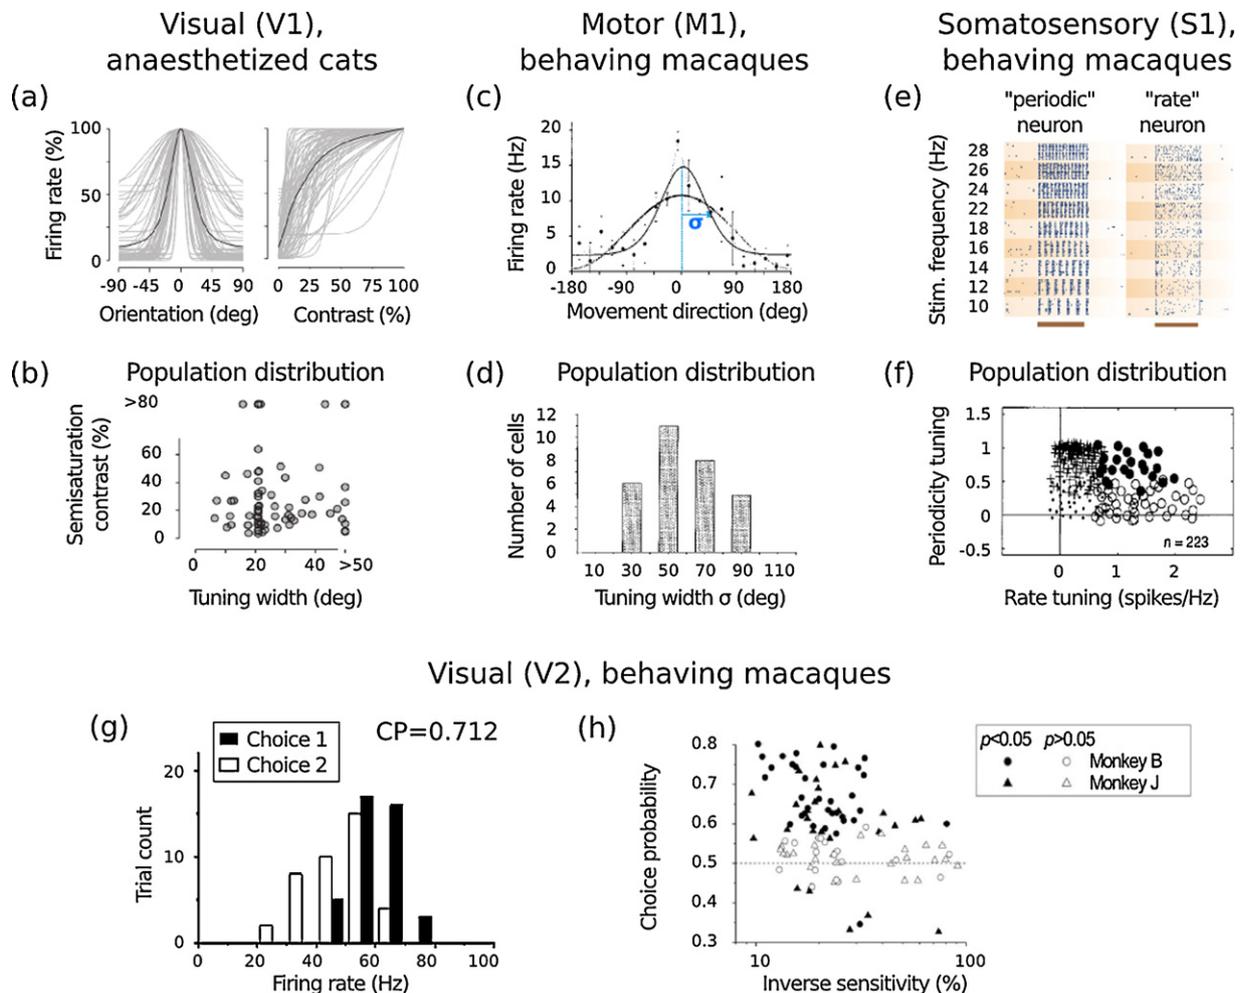
Fig. 8. Population-wide distributions of tuning parameters provide qualitative hints about the nature of feature encoding in different areas. (a) Population response of V1 neurons to oriented gratings (anesthetized cats, multi-unit electrode array). Left: Single neuron tuning curves to grating orientation. Right: Single neuron tuning curves to grating contrast. (b) Population distribution of two tuning features, namely sharpness of orientation tuning and contrast sensitivity, from the data in (a). Reproduced from Busse et al. (2009), with permission from Elsevier. (c) Example of a "sharp" tuning curve for arm movement in primate M1 cortex. Symbols: the neuron's mean firing rate as a function of the angular distance from the neuron's preferred direction (at 0 °). Black line: best-fit of a von Mises function; gray line: approximate fit of the classic cosine tuning curve for comparison. (d) Distribution of tuning curve width in primate M1 cortex, taken as the half-width of the von Mises function (blue line in panel c). A half-width of $\sigma = 90$ ° is the cosine tuning curve assumed in Georgopoulos et al. (1982). Reproduced from Amirikian et al. (2000), with permission from Elsevier. (e) Population response of S1 neurons to vibratory stimulation on the fingertip (awake macaques, multi-electrode recordings). S1 neurons can display tuned responses to the vibration frequency, either through conservation of the stimulus periodicity ("periodic" neuron), or through their overall amount of firing ("rate" neuron). (f) Joint distribution of these two forms of tuning across recorded S1 neurons. Rate tuning is not correlated with periodic tuning, supporting a plausible specific role of the former in tactile perception. Reproduced from Hernández et al. (2000), Copyright (2000) National Academy of Sciences, U.S.A. (g) Trial-to-trial distributions of activity for an MT cell conditioned on the animal's subsequent decision, when the stimulus is ambiguous (0% binocular disparity). Choice probability quantifies the amount of separation between the two distributions. (h) Choice probability and neuronal sensitivity are correlated at the population level, indicating that the animal's ultimate decision correlates more with neurons tuned to the stimulus. Macaque MT during a binocular discrimination task, reproduced from Uka and DeAngelis (2004), with permission from Elsevier.

the presentation of a behaviorally relevant, yet noisy and uninformative stimulus, and the animal's subsequent decision (Britten et al., 1996; Dodd et al., 2001; Romo et al., 2002; Uka and DeAngelis, 2004; Nienborg and Cumming, 2006; de Lafuente and Romo, 2006; Price and Born, 2010; Hernández et al., 2010). Such covariations are intriguing, as they imply that sensory neurons are somewhat tuned to the animal's future decision, even when a stimulus does not provide any information about the required choice. Britten et al. (1996) postulated a purely "bottom-up" explanation, whereby trial-to-trial fluctuations in sensory neurons' activities result in trial-to-trial fluctuations of the animal's *percept*, which can thus influence the subsequent action. A key argument for this hypothesis was precisely the type of population-wide correlation displayed in Fig. 8h, which shows that the most stimulus-sensitive (and thus, task-relevant) neurons are more predictive of the animal's future choice. Later, a similar population-wide correlation argument was used to demonstrate the effect of

task learning on choice probability values (Uka and DeAngelis, 2004). Further supporting the bottom-up perspective, choice probability signals seemingly increase going up the perceptual/decision pathway from primary sensory areas to prefrontal cortex (de Lafuente and Romo, 2006; Fig. 10b). However, while all these findings are necessary predictions of the bottom-up interpretation, they are not a sufficient proof. And indeed, by investigating the dynamics of CPs, Nienborg and Cumming (2009) showed a top-down contribution to the CP signals in primate V2, whereby the animal's forming decision appears to feedback on sensory areas and influence their firing (see Section 6.1.3).

In the examples above, the putative functionality of an area is revealed by studying the organization of tuning at the level of the population, rather than the individual neuron. These results therefore illustrate the benefits of analyzing population distributions of encoding features. As a caveat, we notice that in the above examples, as in the vast majority of the literature of sensory

systems, the plotted distributions provide only qualitative insights. In particular, unlike the recent experiments reviewed in Section 3, the distributions come from populations that are sampled with a strong recording bias for task-related activity. Hopefully, future studies will provide us with unbiased and precisely quantified tuning distributions.

5.2. Theories of feature encoding in neural populations

A number of theories of neural population coding are now available, which likely provide the future keys to quantitative interpretations of population-wide distributions of activity and tuning. In this section, we briefly review two important theoretical frameworks describing feature encoding in neural populations. *Sparse coding* (Section 5.2.1) is a theoretical framework which explicitly accounts for the heterogeneity of tuning across neurons, and associates it with a clear functional explanation. *Probabilistic population representations* (Section 5.2.2) is the term we choose to describe frameworks that assume sensory populations encode not precise stimulus values, but probabilistic distributions over possible stimulus values. Although these probabilistic frameworks have not yet provided quantitative predictions of population distributions to our knowledge, they have the potential to do so in the near future.

5.2.1. Sparse neural coding

In Section 3.3.2 we saw that the long tails of firing rate distributions evoke the influential theoretical concept of *sparse* neural representations. For sensory systems, the theory of *sparse coding* (Barlow, 1972; Field, 1994) proposes that each stimulus triggers only a small and very specific subset of neurons. The neurons in turn constitute the basic "dictionary" elements that describe a stimulus in a non-redundant fashion. Different stimuli, even if relatively similar, are encoded by different sets of active neurons, which correspond to different (albeit similar) dictionary elements. Cortical areas could generate such sparse representations because their total number of neurons often exceeds the dimensionality of the original stimuli. Thanks to their sparseness, cortical patterns of activity could then be more amenable to long-range transmission of information, to downstream processing, and even to learning (Olshausen and Field, 2004; Ganguli and Sompolinsky, 2012). The theory of sparse coding successfully accounts for several properties of sensory systems, such as (i) the huge number of cortical neurons compared to subcortical structures, (ii) the qualitative features of receptive fields in early sensory cortical areas (Olshausen and Field, 1996; Lewicki, 2002), and (iii) nonlinear suppressive interactions between sensory neurons, which are assumed to promote the sparse representation of stimuli (Olshausen and Field, 1997, 2004; Asari et al., 2006; Smith and Lewicki, 2006; Rozell et al., 2008).

As a result of these successes, the past decade has seen several experimental efforts to test the predictions of sparse coding. A first, straightforward test is to assess the levels of sparseness in sensory areas. In the neuroscience literature, "sparseness" may either refer to a sparse distribution of active neurons across a neural population (*population sparseness*) or to a sparse distribution of stimuli to which a particular neuron is responsive (*lifetime sparseness*). These two indicators measure different properties, and are not necessarily correlated (Willmore and Tolhurst, 2001). However, the theory of sparse coding requires both forms of sparseness.[10] In the insect olfactory system, a sparsification of neural responses to odors was observed between primary olfactory

relays (PN cells in the antennal lobe) and higher level brain areas (KC cells in the mushroom body) (Laurent, 2002), both at the lifetime and population level. In mammalian visual cortex, the observed population and lifetime sparseness are much higher than predicted by standard LN (linear-nonlinear) models of sensory processing (Weliky et al., 2003), and are strongest in the case of complex ("natural") stimuli covering all of the cells' non-classical receptive fields (Vinje and Gallant, 2002; Haider et al., 2010; Wolfe et al., 2010). In rat auditory cortex responding to tone pips, Bartho et al. (2009) computed population-wide distributions of lifetime sparseness and found a proportion of neurons whose stimulus sparseness increased between the early and late phase of response, possibly revealing the dynamics of an active sparsification mechanism.

Importantly, the theory of sparse coding reproduces several aspects of the population-wide distributions of tuning parameters. Olshausen and Field (1997) first showed how the principle of sparse coding, applied to encoding natural images, predicts the structure and distribution of V1 receptive fields. Ringach (2002) measured the distribution of receptive fields in cat V1 (Fig. 9a), revealing several commonalities but also some discrepancies with the predictions of Olshausen and Field (1997). In particular, bump-shaped receptive fields with little spatial opponency (Fig. 9b, blue dots close to origin) occurred far more often in V1 than in the prediction (Fig. 9b, red dots). Recently, Rehn and Sommer (2007) have shown how such bump-shaped receptive fields can be predicted under the assumption of "hard" sparse coding, with total silence imposed in all but the most active units (Fig. 9c). In the auditory domain, Smith and Lewicki (2006) demonstrated that a sparse coding principle, applied to encoding natural sounds, can reproduce the filter characteristics of auditory nerve fibers, both at the single cell (Fig. 9d) and population level (Fig. 9e). Importantly, the predicted distribution matched experimental data only if trained on a full set of natural sounds including environmental sounds and animal vocalizations.

New experimental challenges to sparse coding theory may arise from studies using unbiased methods for recording neural populations (Section 2.2). Concerning population sparseness, it is currently unclear to what extent sparse coding models can account for the largely log-normal distributions of sensory-evoked activity (Section 3.1). Particularly intriguing in this regard are the respective peaks of low but non-zero activity (Fig. 4a) which contradict current sparse coding models, since these predict peaks at zero. Furthermore, we currently do not know whether sparse coding models can predict the global invariance of the activity distributions, irrespective of the exact stimulus value (Section 3.2). Concerning lifetime sparseness, we would like to know whether sparse coding models can quantitatively predict the distribution of lifetime sparseness across neurons.[11] Experimentally answering this question is quite challenging since, the more specific a neuron (for example, with a very small receptive field), the less likely it is to be detected on the basis of its response to an arbitrary test stimulus (Olshausen and Field, 2005). Unbiased recording methods such as patch-clamping thus offer a novel opportunity to detect these very specific neurons.[12]

---

[10] High population sparseness by definition, and high lifetime sparseness because each neuron is silenced by better-tuned neurons in all stimulus conditions but those extremely close to the dictionary element represented by that neuron.

[11] Essentially, the levels of lifetime sparseness predicted by a sparse model depend on the number of neurons participating in the representation: the more neurons are present, the less often each one is likely to be active. Furthermore, sparse models of visual processing naturally produce a heterogeneous distribution of lifetime sparseness, where units associated with low spatial frequencies are often active (low lifetime sparseness) while units associated with higher spatial frequencies have higher lifetime sparseness (Bruno Olshausen, personal communication).

[12] In V1, spiny stellate cells from layer IV could be a good candidate for such future investigations, as their large number and low activity rates suggest a sparse-coding strategy (Bruno Olshausen, personal communication).
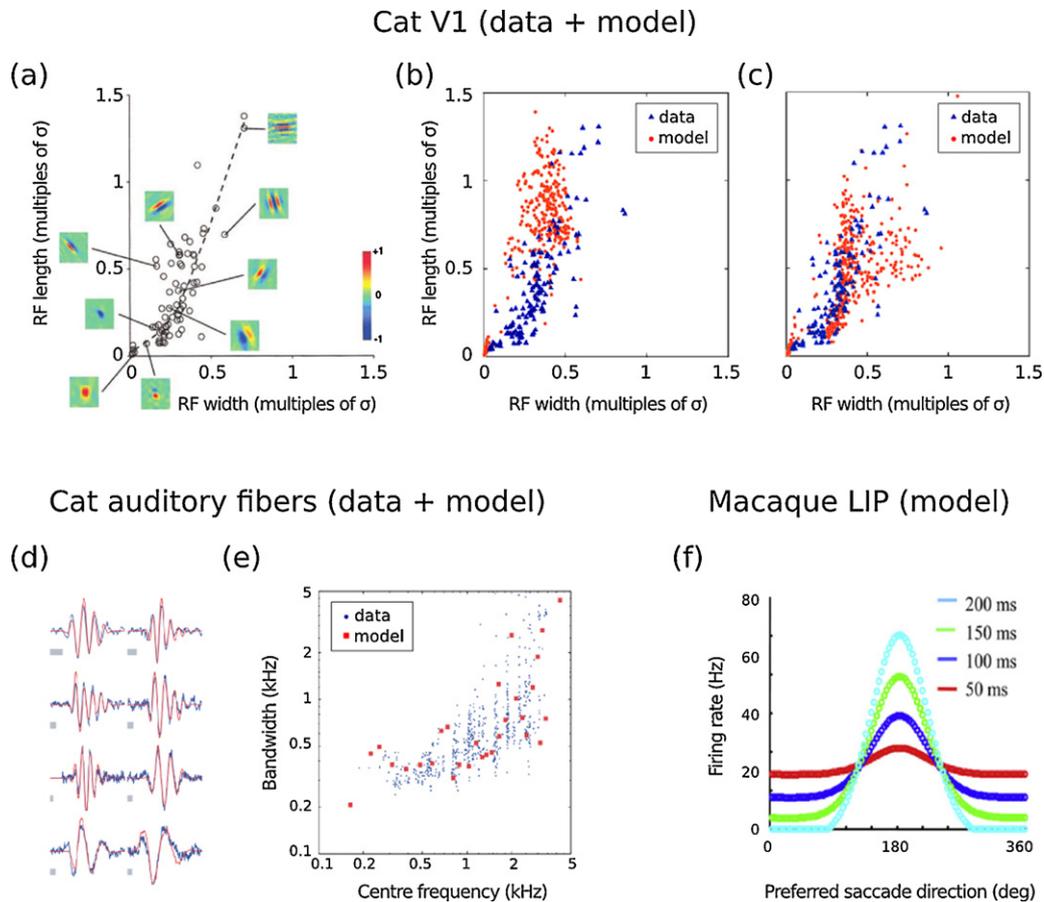
Fig. 9. Theories of feature encoding yield predictions for population-wide distributions of activity or tuning. (a) Receptive fields of V1 cells are fitted by Gabor functions, and plotted in the "shape domain" consisting of each field's length and width divided by its characteristic spatial frequency. Receptive fields with a single bump fall near the origin, while receptive fields with several oscillations fall away from the origin. Reproduced from Ringach (2002), with permission from the American Physiological Society. (b) Distribution, in the same shape domain, of receptive fields predicted by the sparse coding model of Olshausen and Field (1997) (red dots); blue dots are the experimental data from Ringach (2002). (c) Same as (b), but red dots are receptive fields predicted by the "hard" sparse coding model of Rehn and Sommer (2007). Panels b and c reproduced from Rehn and Sommer (2007), with permission from Springer. (d) Transfer filters measured from cat auditory nerve fibers (blue) have similar shapes as the transfer filters predicted by a sparse analysis of natural auditory signals (red). Gray bars represent 5 ms. (e) Population distribution of center frequency and bandwidth for cat auditory nerve fibers (blue dots) and for sparse kernels learned from natural auditory signals (red dots). Panels d and e reproduced from Smith and Lewicki (2006), with permission from Nature Publishing Group. (f) Evolution of population activity in a probabilistic population code model of sensory integration in area LIP. Population sparseness is predicted to increase as a result of the inferred distribution over stimuli becoming sharper. Reproduced from Beck et al. (2008), with permission from Elsevier.

We note that sparse coding is mostly a theory of sensory systems; one may rightfully ask whether similar concepts also apply to descriptions of motor systems. Much less data is available regarding the sparseness of motor representation. Firing rate distributions in motor-related cortices also appear to be well fit by log-normal models (Fig. 3), speaking in favor of high population sparseness. Furthermore, motor-related cortices seemingly display the exquisite sensitivity required for only a few neurons to reliably drive motor commands. Brecht et al. (2004) showed that intracellular nano-stimulation, invoking action potentials in a single layer 5 or layer 6 rat M1 whisker neuron, caused detectable changes in the movement of specific whiskers. Consistent with this, multi-electrode recordings have established that the activity of only a few tens of neurons are necessary to accurately decode movement in rat (Laubach et al., 2000) and primate (Carmena et al., 2005; Rickert et al., 2009). However, sparsity in decoding does not imply sparsity of representation (the focus of this review); indeed, motor-related activity in cortex appears to contain a degree of redundancy (Narayanan et al., 2005). For example, Carmena et al. (2005) found that even after removing the top 50 neurons contributing to the decoding of hand position, 60% of the total ensemble contribution could still be achieved from the residual population. Whether such high levels of information in marginally active neurons agree with the predictions of sparse coding still has to be tested.

5.2.2. Probabilistic population representations

Behavioral studies have long suggested that perception is akin to a form of probabilistic inference. When faced with ambiguous stimuli (e.g., a noisy image), humans and animals alike are capable of combining sensory evidence with *a priori* statistical knowledge about the world (e.g., "what am I likely to see in my current situation?") in order to take optimal decisions. The ability to perform probabilistic inference implies that the brain encodes not only the value of a stimulus, but also information about its confidence in the specific value. In the extreme case, sensory systems may represent a full (posterior) probability distribution over all possible stimulus values. And indeed, correlates of confidence and probabilistic inference have been found at the single neuron level (Yang and Shadlen, 2007; Morgan et al., 2008; Kiani and Shadlen, 2009).

Two closely related frameworks have been proposed which explain how probabilistic information can be encoded in neural populations. *Probabilistic population codes* (Pouget et al., 2000, 2003; Ma et al., 2006) assume a simple correspondence between the instantaneous activity profile of a neural population and the

probability distribution $p(f)$ over a set of (stimulus) features f. The correspondence itself can vary, depending on the system being modeled,[13] but the underlying intuition is always the same: Each neuron's instantaneous activity, through inversion of its tuning curve, can be interpreted as a noisy "vote" for a given feature value f. When all neural votes are concatenated, the resulting histogram can either be peaked around a single value (corresponding to high level of confidence), or much more widespread (corresponding to low level of confidence). *Sampling-based codes* (Fiser et al., 2010; Berkes et al., 2011) assume that probability distributions are represented implicitly—through samples drawn from them—rather than explicitly. In particular, the framework suggests that each successive snapshot of the instantaneous population activity provides a new sample from the probability distribution. Such sampling-based codes provide an appealing functional explanation of trial-to-trial variability, and allow a straightforward implementation of probabilistic learning and inference (Fiser et al., 2010).

These frameworks make strong conceptual claims regarding the embedding of information in sensory populations, which should ultimately impact their predicted distributions of activity, tuning, and temporal dynamics across populations. These predictions, though, are yet to be made. Here, we simply delimit their reach and possible forms. A first important remark is that, since these theories deal with the neural representation of uncertainty, they can only be assessed in situations where stimulus uncertainty is routinely manipulated. In this spirit, Beck et al. (2008) have shown how probabilistic population codes can account for the integrative properties of individual LIP neurons during a random dot motion discrimination task. While the study focused on single neuron activities, it made several qualitative predictions regarding population sparseness: in the context of sensory integration tasks, the distribution of neural activities should concentrate around fewer active neurons as time passes, reflecting the sharpening of the encoded distribution $p(f)$ thanks to the integrated sensory evidence (Fig. 9f). Similar arguments can be made for sampling-based codes, so that both frameworks predict that population sparseness increases with increased confidence in stimulus value. There is currently no evidence for or against such a prediction. On the level of single neurons, both frameworks predict that a neuron's trial-averaged activity increases as the feature represented by the neuron becomes more likely, thus resulting in population-wide dynamics of ramping up or down. Such ramping behaviors have indeed been observed at the population level, as reviewed in Section 6.

Strictly speaking, neither probabilistic population codes nor the sampling-based hypothesis are self-contained theories of neural encoding, in the sense that they do not specify the way a stimulus s—even with high confidence—is embedded in a population. As a result, to yield predictions regarding distributions of tuning and lifetime sparseness, they must be coupled with a specific representational hypothesis. So far, these stimulus representations have mostly been restricted to simple, homogeneous settings (typically the ring model, see Ma et al., 2006; Beck et al., 2008). However, Lochmann et al. (2012) have recently proposed a framework that merges aspects of probabilistic population representation with sparse coding. In their theory, the neural population still represents a probability distribution, yet individual neurons compete via lateral inhibition to represent objects, which increases the overall sparseness of the representation. Such a framework has the potential to predict quantitative distributions of activity and tuning features, and how they vary with uncertainty.

---

[13] Most often, $p(f)$ is obtained as a linear projection from the population activity profile.

6. Population distributions of time-varying neural activities

In the previous sections, we have considered simple measures of neural responses such as average firing rates (or spike counts), CVs, Fano Factors, etc. These measures are highly reduced descriptions of neural activity. In particular, they do not characterize the *time course* of neural responses. In this section, we present recent approaches that have sought to describe the distribution of such time-resolved neural activities in a given population of neurons, both for stimulus-free and stimulus-driven conditions.

6.1. The temporal structure of population responses to a stimulus

The simplest measure of a neuron's temporal course of response is its trial-averaged firing rate. In the case of a sensory response, the stimulus onset acts as the initial temporal alignment ($t = 0$), and the subsequent trial-averaged firing rate corresponds to the peri-stimulus time histogram (PSTH) of the neuron. For a single neuron, we denote by $r(t, s)$ this trial-averaged, time-varying, and stimulus-dependent firing rate. In the perspective of this review, we then seek to describe how the PSTHs of different neurons are distributed across the population. We note that such a stimulus-dependent time course consists of many different variables, one for each time step and stimulus condition, so that characterizing a distribution over these PSTHs is a formidable task. We will, however, start with something much simpler: the distribution of response latencies across neurons.

6.1.1. Population distributions of response latencies

The response latency of a neuron is the time between the onset of a stimulus and the response of the neuron. The organization of neural latencies at the population level is of high interest to systems neuroscience. It can help determine a global hierarchy of processing through respective area latencies (Schmolesky et al., 1998; de Lafuente and Romo, 2006), and also potentially reveal neural codes based on spike timing (Abeles, 1991; VanRullen and Thorpe, 2001; Gollisch and Meister, 2008). Ideally, we would want to know how these response latencies are distributed across a population for a given, fixed stimulus. There are, however, two problems that need to be taken into account. First, since most neurons are active even before a stimulus is shown, response latencies have to be measured with respect to a baseline. However, as shown in Section 3.2, the overall invariance of population distributions of activity implies that for every neuron whose activity increases in response to a given stimulus, another neuron's activity decreases. Furthermore, some of the neurons may be unresponsive. Since response latencies are traditionally defined with respect to a measurable increase in activity, they are ill-defined for part of the population. Second, the response latency of any particular neuron depends on the stimulus presented: for some stimuli, a neuron may respond faster than for other stimuli.

In the literature, these problems have often been addressed by reporting "minimum latencies". These are defined as the smallest measurable latency across a set of stimuli. For instance, Heil (1997) tested neurons in the primary auditory cortex of anesthetized cats with a range of sounds with different onset envelopes. The author found a distribution of minimum response latencies that ranged from 6 to 37 ms, with most of the neurons falling within 9–15 ms. In the primary visual cortex of awake behaving monkeys, Maunsell and Gibson (1992) found minimum latencies ranging from about 20–70 ms using high contrast gratings, with most neurons falling in the range 35–55 ms. In what is probably the most comprehensive study of visual response latencies, Schmolesky et al. (1998) measured distributions of minimum latencies across several areas of the visual cortex in anesthetized monkeys, including LGN, V1,
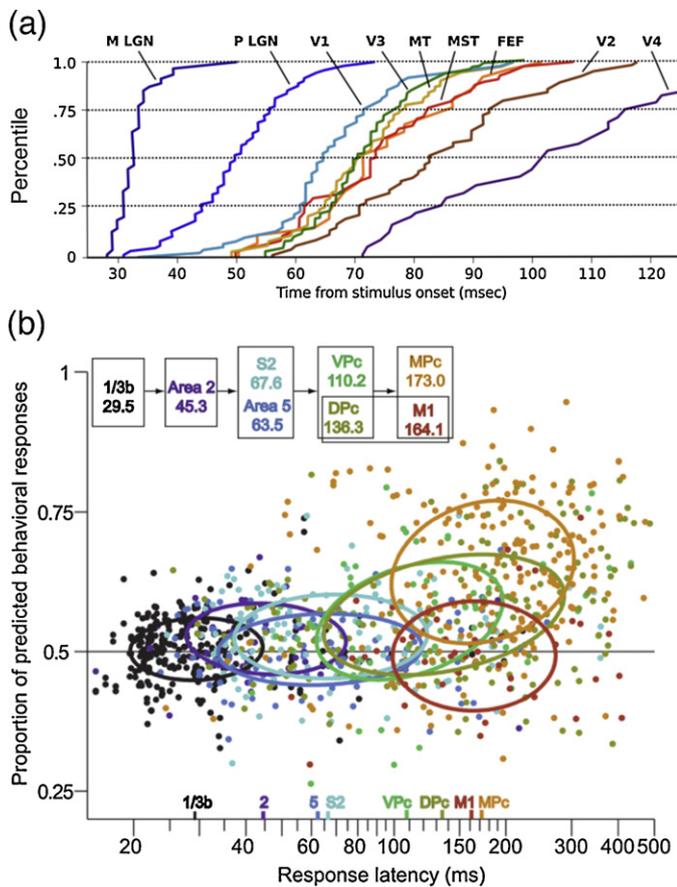
Fig. 10. Distribution of response latencies across cortical areas. (a) Cumulative distributions of (minimum) response latencies in several areas of the visual system of anesthetized monkeys. Replotted from Schmolesky et al. (1998). (b) Latencies and choice probabilities in different areas involved in a tactile discrimination task. Latencies and choice probabilities increase along the putative processing pathway of information. Reproduced from de Lafuente and Romo (2006), Copyright (2006) National Academy of Sciences, U.S.A.

V3, MT, MST, V2, and V4, see Fig. 10a. They find widely distributed minimum latencies in all areas. In the primary visual cortex, the minimum latencies range from 34 to 97 ms, the bulk of which falls within 55–75 ms. In awake macaques during a somatosensory detection task, de Lafuente and Romo (2006) measured the distribution of response latencies for several areas, from sensory to motor systems (S1, S2, VPc, DPc, M1, MPs). The latency distributions were widespread for each area, but overall the mean latency was found to increase from sensory to prefrontal areas, and to correlate with increased encoding of the animal's decision (choice probability, see Section 5.1.3)—as opposed to pure sensory encoding in the earlier-responding areas (Fig. 10b). These results clearly point out that decision formation is a distributed process, both across areas and in time.

As mentioned above, minimum latencies are determined by pooling over a large set of stimuli. Consequently, they are suited to measure the speed of information transmission, but they do not depict an accurate picture of the population response to a given stimulus. One way of obtaining such a picture is to first construct a model for the responses of single neurons, and then use this model as a basis for describing the population statistics. For sensory systems, the most prominent model is the linear-nonlinear model, which has classically been estimated using reverse correlation analysis. The linear part of these models is sometimes equated with the receptive field, and we have already discussed the static case in Section 5. When equipped with temporal filter

components, the linear-nonlinear model is called the spatio-temporal or spectro-temporal receptive field (STRF). While many studies have measured and characterized receptive fields, few studies have explicitly provided a distribution of the receptive field parameters (for exceptions see Valois et al., 1982; Parker and Hawken, 1988; DeAngelis et al., 1993). In a seminal study, DeAngelis et al. (1993) investigated cells in the primary visual cortex of cats and estimated both the temporal and spatial profiles of neurons' receptive fields. They found latencies to peak responses distributed in a range of 30–150 ms, with an average of 68 ms. Although the authors do not specify the onset response of the population for a given stimulus, the respective population statistics could be reverse engineered from the distribution of STRF parameters provided.

6.1.2. The population vector and other linear models

Even though the study by DeAngelis et al. (1993) provides a complete, model-based picture of the population response, there are two caveats. First, as noted in Section 2.2, the sampling of neurons in this (and most other) studies is likely to be biased, therefore providing a potentially inaccurate picture of the population dynamics. This problem has to be kept in mind; however, since it has not been fully addressed in the literature, we will ignore it here. Second, the distribution of the temporal dynamics is based on a two-step approach in which a model is first fit to the neural response, and then the population statistics are evaluated. A major disadvantage of such a two-step approach is that small systematic deviations of the single neuron model from the actual response can potentially add up to major deviations on the population level. In the case of receptive field models, this problem is exacerbated by their poor performance in predicting a neuron's response to arbitrary stimuli (Linden et al., 2003; David et al., 2004; Machens et al., 2004; Olshausen and Field, 2005). To circumvent this problem and get a better hold of the population dynamics, one should ideally model the population response directly. This problem has been addressed in the literature, mostly by reducing the number of stimuli considered while expanding the number of neurons, thereby modeling the complete (recorded) population in a single shot.

One basic idea for such a population-based approach is to model the time-varying firing rate of each neuron as a linear combination of a few "response modes" $A_i(t, s)$ with $i = 1 \ldots k$ where k is a small number, typically less than ten. These response modes have to be chosen such that they capture the main trends in the population, both in terms of the temporal dynamics of the response, and in terms of the stimulus-dependencies. The firing rate of the jth neuron can then be written as

$$r_j(t, s) = w_{j1}A_1(t, s) + w_{j2}A_2(t, s) + \cdots + w_{jk}A_k(t, s) \quad (1)$$

where the parameters w_{ji}, sometimes called the "factor loadings", determine the weights of the different modes i for neuron j. Accordingly, each neuron is characterized by its own set of parameter values which can be summarized in vector form as $\mathbf{w}_j = (w_{j1}, \ldots, w_{jk})$. The population response can then be represented by the response modes and their distribution within the population, as measured by $p(\mathbf{w})$. In terms of linear algebra, the hypothesis that the k-dimensional response-mode description is "accurate enough" means that it captures all the temporal and stimulus-induced variations of the firing rates in all recorded neurons.

The response modes are usually built as linear combination of the original neural activities, $\mathbf{r}(t, s) = (r_1(t, s), r_2(t, s), \ldots, r_N(t, s))$. In linear algebra formulation, we thus seek a $k \times N$ projection matrix $\mathbf{U}$ such that the vector of response modes, $\mathbf{A}(t, s) = (A_1(t, s), \ldots, A_k(t, s))$, is computed as $\mathbf{A}(t, s) = \mathbf{Ur}(t, s)$. In the classical "population vector" approach, due to Georgopoulos et al. (1986), the matrix $\mathbf{U}$ is empirically fixed based on each cell's individual tuning properties,

and on the type of information sought by the researcher. For example, the original work of Georgopoulos et al. (1986) in rhesus monkey motor cortex aimed at reconstructing or decoding an arm trajectory in a plane (so $k = 2$) from the motor neuron population activities. In this case, the reconstruction of the two spatial components x and y of the movement was given by a linear weighted sum of the neural activities. In our notation, these linear weights correspond to the two rows of the matrix $\mathbf{U}$, whereas the reconstructed (x, y)-position of the movement over time corresponds to the evolution of the two first response modes so that $x(t, s) = A_1(t, s)$ and $y(t, s) = A_2(t, s)$, where t is time and s the movement target (or stimulus). In the population vector approach, the factor loadings w_{ji} have usually not been computed, although they could be determined in principle, for instance, by using linear regression.

Since this approach makes empirical assumptions about the nature of neural encoding, it may lose considerable information compared to the real population code, similar to the two-step approaches explained above. As a result, other approaches have rather focused on *unsupervised* descriptions, where matrix $\mathbf{U}$ is imposed by the data itself. One such method is Principal Component Analysis (PCA), in which case the matrix $\mathbf{U}$ is constructed of the first k eigenvectors of the $N \times N$ temporal covariance matrix of $\mathbf{r}(t, s)$. Each reconstructed component $A_j(t)$ is then the jth *principal component* of the activity. Taken together, the first principal components generally account for a majority of the variance in the population activity and thereby reveal an intrinsically low-dimensional structure in the population activity (see Chapin and Nicolelis, 1999 for an early application of PCA in this sense; Optican and Richmond, 1987 and McClurkin et al., 1991 for earlier, but single-neuron based applications of PCA). In PCA, the factor loadings are simply given by $\mathbf{W} = \mathbf{U}^T$ where $\mathbf{W}$ is a matrix with rows $\mathbf{w}_j$, and $\mathbf{U}^T$ denotes the matrix transpose of $\mathbf{U}$.

6.1.3. Temporal dynamics of trial-averaged, stimulus-evoked firing rates

Not many studies have investigated population dynamics or their statistics in neural systems. We now review how the above linear frameworks have been applied to study population dynamics in a variety of sensory systems responding to a stimulus (Fig. 11). Some of the first applications come from insect systems. While outside the range of topics considered in this review, we find it nonetheless useful to briefly consider these results. Fig. 11a shows the population response of principal neurons in the insect antennal lobe, responding to presentations of two different odors (Mazor and Laurent, 2005). Here, the three coordinates correspond to the first three response modes, A_1, A_2, and A_3, or equivalently, to a projection of the neural population response onto the first three principal component axes. The population responses display strong phasic excursions at odor onset and offset, whereas the sustained response to odors ("fixed point" FP in the figure) is generally closer to baseline activity, with less stimulus differentiation (Mazor and Laurent, 2005; Stopfer et al., 2003). Trajectories for different odors diverge strongly and right from stimulus onset, indicating that stimulus information is the dominant factor shaping the response of antennal lobe principal neurons. Haddad et al. (2010) have similarly analyzed population responses in various insect and mammalian olfactory systems, and found that the first PCA components have preserved characteristics and interpretation: the first principal component represents the overall strength of odor-induced activity in the population and covaries with the animal's approach or withdrawal from the odor, while the second component covaries with odor toxicity.

In contrast, activity in mammalian sensory cortices seems to be dominated by a purely temporal component, independent of stimulus identity. Fig. 11b (Bartho et al., 2009) represents the two first principal components of population responses in primary auditory cortex of anesthetized rats. The three lines (or "trajectories") correspond to the presentation of three tone pips of different frequencies. Trajectories appear very similar for different stimuli, dominated instead by a stereotyped phasic excursion corresponding to stimulus onset. A similar picture arises in the primate somatosensory cortex of awake monkeys responding to tactile vibratory stimuli (S2 cortex, Fig. 11d–i, unpublished analysis of previously published data; Hernández et al., 2010). The first principal component is again a stereotypical transient response (Fig. 11f) accounting for an overwhelming amount of the population's total variance (Fig. 11d, principal eigenvalues).

In both these experiments, the precise stimulus value is read out in a different linear subspace than the temporal transient at onset, and accounts for a much smaller amount of total variance. In the macaque S2 experiment, stimulus information appears clearly in the second principal component (and also in the third, not shown), taking the form of a stationary shift in activity depending on the exact stimulus value (Fig. 11g). In the rat A1 experiment, the authors must explicitly search for another linear projection matrix $\mathbf{U}$ emphasizing separation of different stimuli (Fig. 11c, Linear Discriminant Analysis), because the amount of population activity associated with precise stimulus information is so small that it scarcely appears in traditional PCA. To complete this population picture, the authors find that population activity undergoes a strong rotation in neural activity space between the "transient" and "sustained" phases of response. Moreover, the angular separation (of the respective population activities) between these two phases is much larger than the angular separation between two stimuli during the same phase (Bartho et al., 2009). What could then be the functional signification of the strong initial temporal component? One possibility is that it simply represents stimulus onset. Another possibility is that it reflects coarse stimulus information (e.g., stimulation of the finger, or stimulation within a certain frequency band). Indeed, in macaque IT neurons responding to objects and faces, Matsumoto et al. (2005) report PCA components which progressively differentiate object identities, from coarse identity in the early phase of response (object vs monkey face vs human face) to finer identity (face expression, object shape, etc.) in the latter phase.

To summarize, these results hint at a multiplexing of neuronal responses in mammalian cortex, between strong initial dynamics directly after stimulus onset, and precise stimulus differentiation involving a smaller number of fired spikes afterwards. The relative impact of these two components appears to depend on the exact place of the area in cortical processing, with early sensory areas (A1, V1, S1, etc.) potentially displaying a stronger initial temporal transient and less differentiation than higher-level areas which are involved in more abstract features (MT, LIP, IT, etc.).

Most studies of population dynamics focus exclusively on characterizing the response modes $A_i(t, s)$, and ignore the factor loadings or weighting coefficients $\mathbf{w}_j$ and their distribution across the population. These coefficients are likely to provide important additional information. This is exemplified in a study by Chapin and Nicolelis (1999), for instance, in which simultaneous recordings in the thalamus are used to analyze population responses to whisker stimulation in awake rats. They found distributed factor loadings that could be related to the topology of the whisker system: neurons coding for different whiskers often had quite different factor loadings; however, due to the small (but simultaneously recorded) ensembles considered, the statistics of the factor loadings could not be evaluated. In macaque IT neurons responding to objects and faces, Matsumoto et al. (2005) find factor loadings which seem to follow a long-tailed distribution, but no actual statistics (such as kurtosis) are provided. It remains therefore an open question how the distribution of factor loadings compares to the long-tailed distributions of neural activities
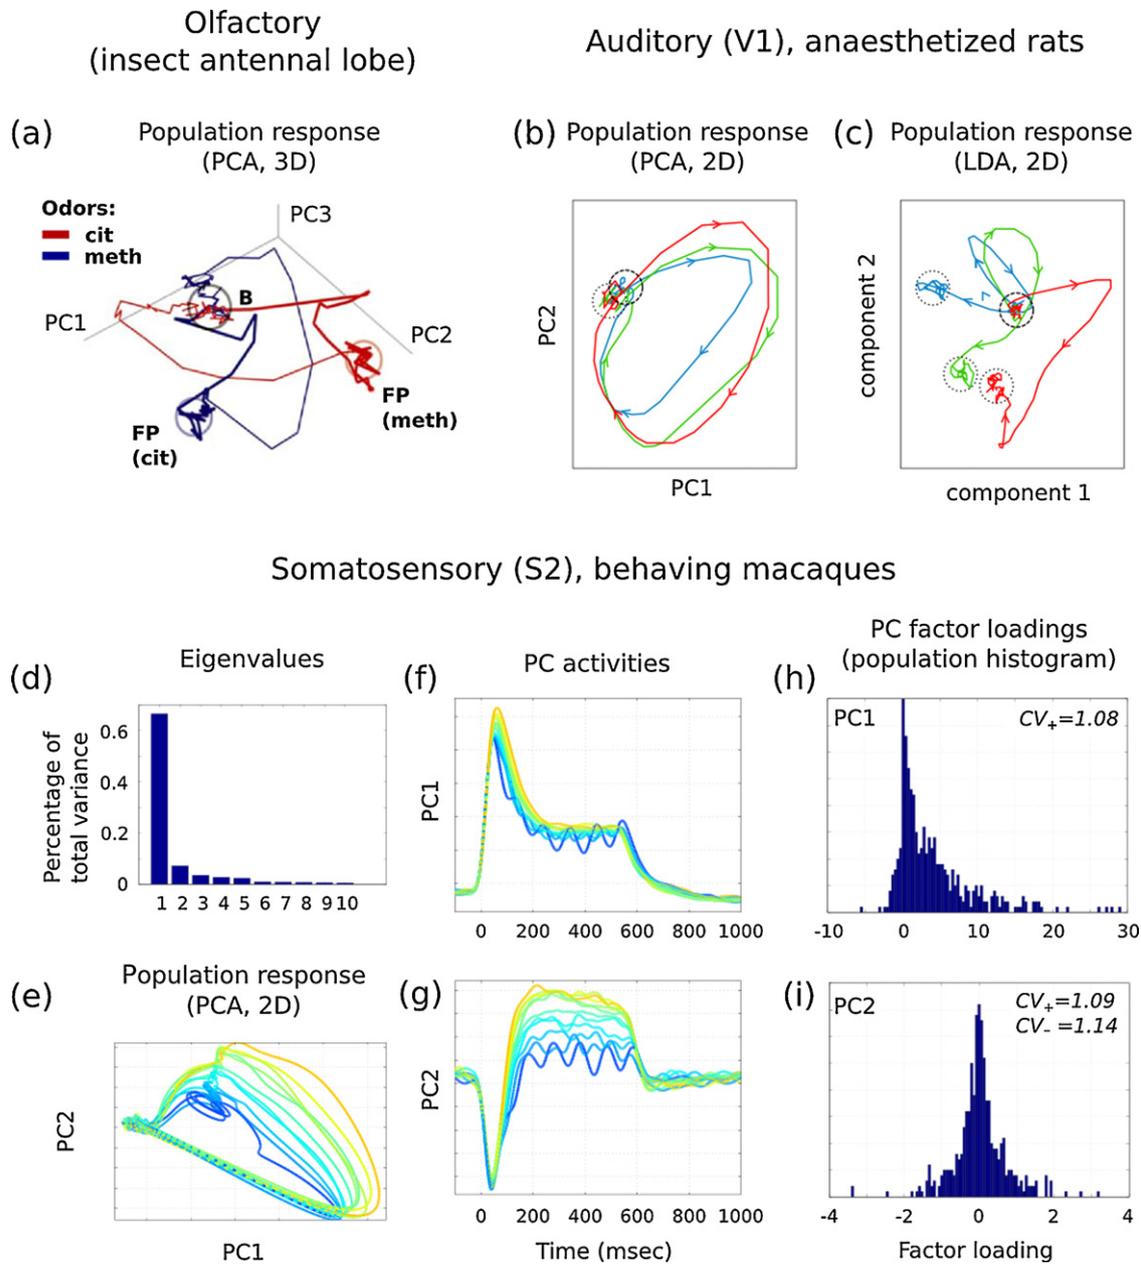
Fig. 11. The temporal structure of the population response to sensory stimuli. (a) Principal Component Analysis (PCA) of insect antennal lobe principal neurons responding to two odors. Reproduced from Mazor and Laurent (2005), with permission from Elsevier. B: baseline, bold traces: stimulus onset, FP: fixed point (persistent activity), thin traces: stimulus offset. (b) and (c) Population dynamics of rat auditory cortex responding to tone pips is visualized thanks to PCA (panel b) and Linear Discriminant Analysis (panel c, segregates responses to different stimuli). Replotted from Bartho et al. (2009). (d)–(i) Population dynamics (determined using PCA) of macaque S2 neurons responding to a vibratory frequency on the fingertip. Colors from blue to orange indicate increasing frequency values. (d): First 10 eigenvalues of signal covariance matrix. (e): Population activity in the 2 first Principal Components (PC). (f),(g): Temporal activity in the first two PCs. (h),(i): Population histograms of neural factor loadings on the two first PCs. CV+ and CV− are the respective coefficients of variation for the positive and negative parts of the distributions. Reanalysis of published data (Hernández et al., 2010).

described in Section 3. Indeed, the strong phasic component revealed with PCA in the examples above may appear incompatible with the notion of a sparse population response to the stimulus.

To explore this question a bit more, we computed the factor loadings, $\mathbf{w}_j$, for individual neurons in the macaque S2 data (Fig. 11h and i). We find that most neurons contribute positively to the first principal component, experiencing a rise in activity at stimulus onset (whereas negative components indicate a dip in activity). In contrast, neural contributions to the second, stimulus-dominated component are balanced, reflecting the long-known existence of positively and negatively tuned neurons in this area

(Romo et al., 2002). Finally, we find a common feature to these two distributions: they are long-tailed, as revealed for example by their large (one-sided) CVs. So here again, it seems that relatively few neurons take on the largest contributions to the population's temporal dynamics, consistent with the statistics of activity observed during spontaneous activity. All in all, although several studies have used PCA or other dimensional reduction methods to study the dynamics of sensory populations, little emphasis has been put on interpreting the distribution of factor loadings, and in particular how it compares to the distribution arising from unstructured (shuffled) data.
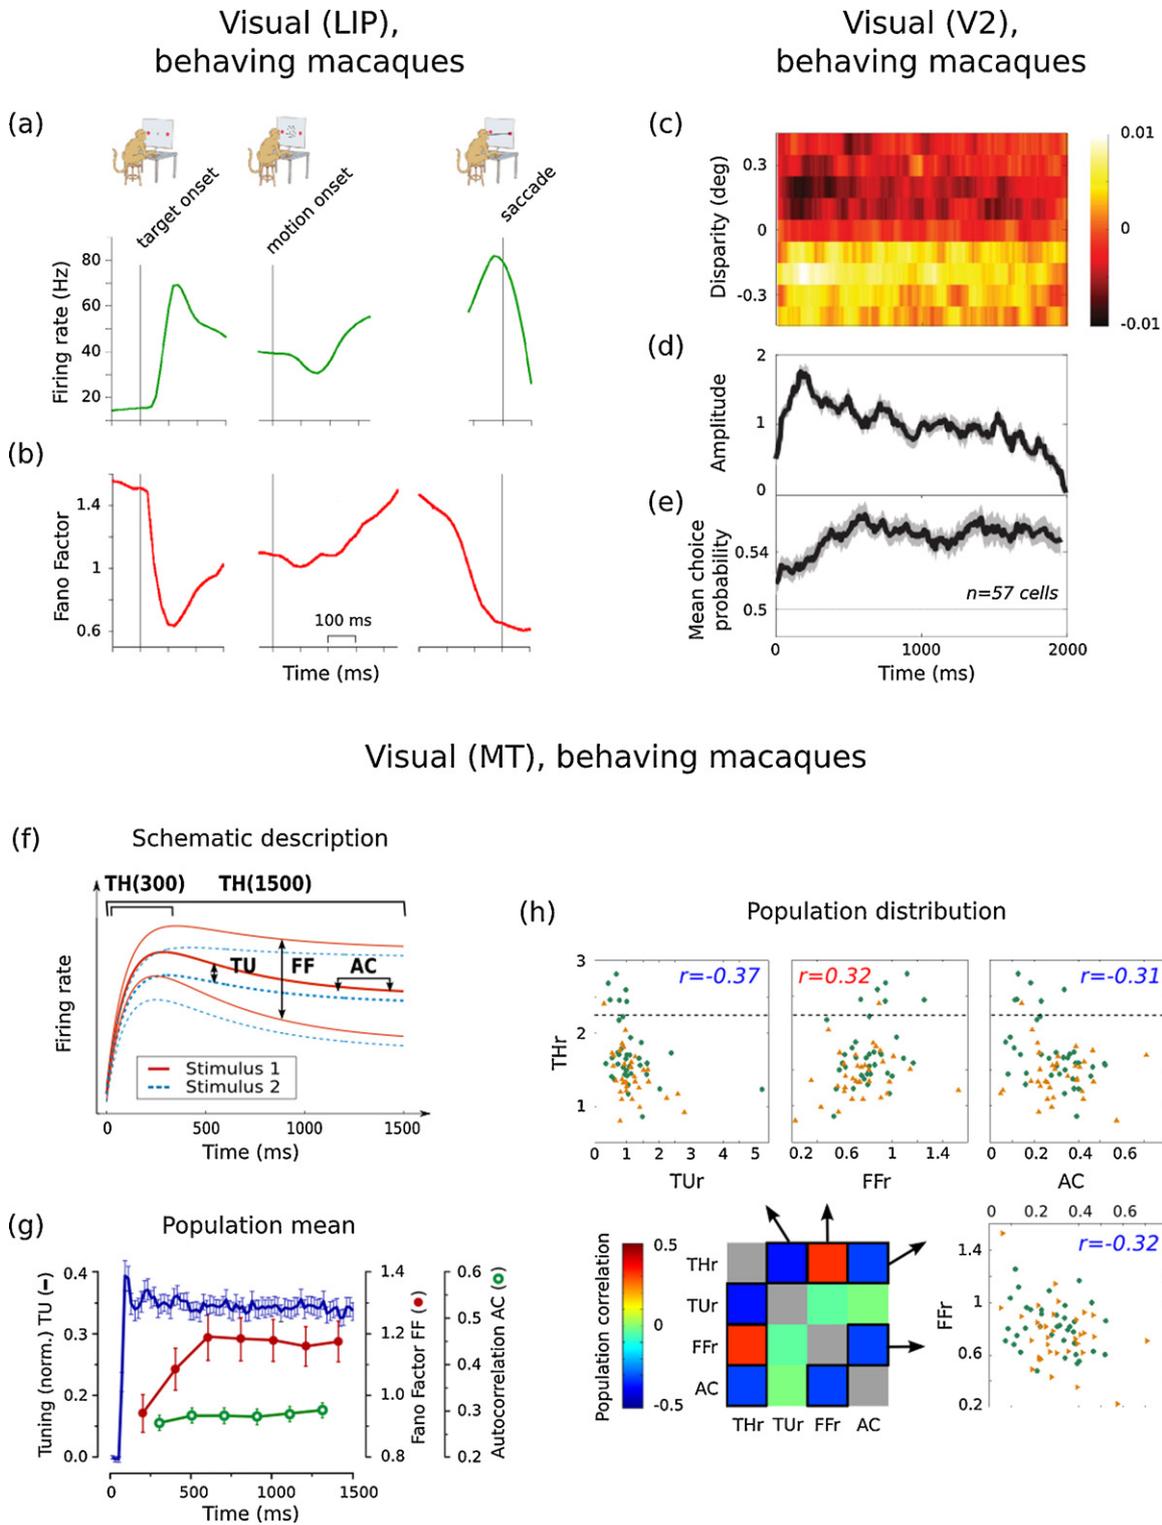
Fig. 12. Temporal dynamics of trial-to-trial variance and stimulus discriminability. (a) and (b) Mean firing rate and Fano Factor amongst LIP neurons with a given preferred direction, responding to a random dot motion stimulus. The Fano Factor is not constant as predicted by Poisson-like statistics (panel b). Instead, it reveals variance quenching at stimulus onset and offset, and a rise of variance during the course of sensory stimulation. Reproduced from Churchland et al. (2011), with permission from Elsevier. (c)–(e) Causal influence of stimulus and choice probabilities in area V2, in macaques during a binocular disparity discrimination task. The psychophysical kernel for the animal (panel c, overall amplitude in panel d), which reflects the causal influence of stimulus disparity across time on the animal's ultimate decision, has a different temporal evolution than choice probability signals (panel e, average CP in the V2 population). This is inconsistent with a purely bottom-up interpretation of choice probability signals. Reproduced from Nienborg and Cumming (2009), with permission from Nature Publishing Group. (f)–(h) Temporal evolution of neurometric thresholds and other spiking statistics in macaque MT neurons during a binocular depth discrimination task. (f): Overall scheme of the analysis. A neuron responds to two possible stimuli with different binocular depths (stimulus 1: plain red curves, mean response plus trial-by-trial standard deviation, stimulus 2: dashed blue curves). From its spike counts on a trial-by-trial basis, a neurometric threshold (TH) can be derived measuring the neuron's sensitivity to the stimulus, based either on the neuron's spike count over the whole stimulation period (1500 ms), or over the sole initial response (first 300 ms). Candidate statistics influencing the value of the threshold are: neural tuning (TU) to the stimulus, variance-to-mean ratio (Fano Factor, FF), and noise autocorrelation (AC) for the neuron on a trial-by-trial basis. (g): Population mean across neurons for the three statistics. The mean amount of tuning is constant across the response, as well as the mean noise autocorrelation (with a value around 0.3). The Fano Factor increases in the second half of the stimulation,

6.1.4. Temporal dynamics of trial-to-trial variance and stimulus discriminability

In the previous section we have only considered population studies based on individual neurons' trial-averaged firing rates. However, the response of an individual neuron can also be characterized by higher-order statistics, such as its trial-to-trial variance: How much does the neuron's activity in a fixed window of time vary from trial to trial? Neglecting the aspect of time, we already addressed this question in Section 4.2.2. Here we consider the time-resolved version of the Fano Factor (FF, variance-to-mean ratio) and ask how the Fano Factor (or other measures of variance) evolve in time, when computed at different instants of the sensory stimulation. This question has recently met renewed interest, with the findings that, in awake-behaving animals, (i) reduction in trial-to-trial variance at stimulus onset is a hallmark of (sensory) cortices (Harris et al., 2011; Churchland et al., 2010), and (ii) the exact temporal evolution of variance can give hints about the nature of the underlying neural computations (Churchland et al., 2011).

The first finding—quenching of variance at stimulus onset—bears a natural interpretation given the relationships between spontaneous and sensory activity described above (Section 3.2). From the point of view of the trial-to-trial variance, switching from spontaneous activity (and its trial distribution $p_t(r)$) to stimulus-driven activity (with trial-distribution $p_t(r|s)$) is akin to a probabilistic conditioning operation, which leads to a decrease in variance. Such "variance quenching" can be observed in Fig. 12a, as a marked decrease in Fano Factor around stimulus onset. Other studies have extended this study of variance quenching to populations of simultaneously recorded neurons (Luczak et al., 2009; Yu et al., 2009; Churchland et al., 2010).

The second finding—interpreting the temporal course of variance as a trace of the underlying neural computations—has been the subject of a recent article by Churchland et al. (2011), whose main point is illustrated in Fig. 12a and b. In area LIP of macaques performing a random dot discrimination task, trial-to-trial variance grows continuously during the course of sensory stimulation (Fano Factor, panel b). This behavior is at odds with Poisson-like spiking processes, for which the Fano Factor is constant (and equal to 1 in the Poisson case). Instead, such a temporal growth of FF is expected in the case of internally generated dynamics which vary on a trial-to-trial basis. Two main candidate mechanisms are: (i) integrative mechanisms of sensory signals, with random-walk-like dynamics, and (ii) attractor dynamics resulting from a top-down biasing of sensory signals on a trial-by-trial basis. Model simulations by Churchland et al. (2011) suggest that either mechanism could be at play in the present data.

Fig. 12a and b only depicts the mean behavior across all neurons tuned to a particular movement direction, and we found that population descriptions of the "growing variance" effect are rare in the literature. Uka and DeAngelis (2003) have led such an analysis in area MT of macaques performing a binocular depth discrimination task, summarized in Fig. 12f–h. They considered the temporal evolution of three spike statistics for each recorded neuron (Fig. 12f): (i) its *tuning* to binocular depth (TU, difference in trial-averaged firing rate between the two binocular conditions), (ii) its *Fano Factor* (FF), and (iii) its *noise auto-correlation* (AC), assessed from the trial-to-trial covariance between spike counts of disjoint, adjacent 200 ms time windows. Fig. 12g depicts the temporal evolution of the population mean for the three indicators. Note the rise in Fano Factor during stimulation, reminiscent of Fig. 12b.

A fourth statistic was also computed for each neuron: its neurometric threshold (TH), measuring how precisely different stimuli can be discriminated on a trial-to-trial basis by looking at the spike count of the neuron—high TH corresponding to poor sensitivity. If the neuron behaves as a stationary Poisson process, TH is expected to scale as $T^{-1/2}$, where T is the length of the temporal window used to count spikes. A striking finding of the authors was that for most neurons this scaling does not occur: when computing thresholds using either the initial 300 ms period, or the full 1500 ms stimulation period, one finds that the full period adds surprisingly little additional information, resulting in a ratio TH(300)/TH(1500) often much lower than its predicted value of $\sqrt{5}$ for Poisson statistics (dashed lines in Fig. 12h). The greater importance of early neural responses for stimulus discrimination has been found in other systems as well (Luna et al., 2005; Price and Born, 2010; Stanford et al., 2010).

All these results hint at neural population statistics which vary qualitatively between the *early* and *late* phases of stimulus response. To assess these effects, Uka and DeAngelis (2003) summarized each neuron's behavior by four scalar indicators: (i) its *tuning ratio* (TUr) between early and late period, (ii) its *Fano Factor ratio* (FFr) between early and late period, (iii) its average level of noise autocorrelation, AC, across the stimulation, and (iv) its *threshold ratio* (THr) between early and late periods. Fig. 12h shows the population distribution for these four indicators. The neurons' relative loss in sensitivity during late stimulation (THr $< \sqrt{5}$) correlates, at the population level, with a decrease in tuning (TUr >1), an increase in trial-to-trial variance (FFr <1), and a strong level of noise autocorrelation implying non-Poisson statistics (AC >0). Furthermore, a significant correlation is found between indicators AC and FFr: neurons with strong autocorrelation levels (i.e., large AC) tend to experience more drift on a trial-by-trial basis, leading to higher variance toward the end of the stimulation (i.e., small FFr). Again, this is the type of behavior expected if internally generated dynamics develop in the population during the late period. Finally, we note that the two variance indicators (FFr and AC) are not correlated with the tuning indicator (TUr), showing that the putative internally generated dynamics are not *directly* related to each neurons' individual sensory content.

The internally generated dynamics shown in these examples are likely to reflect computations, e.g., about an upcoming decision or action. In sensory areas, such computations can be measured by choice probabilities (see Section 5.1.3; Figs. 8g and 10a). The dynamics of these decision-related signals in turn allow us to learn something about the information flow in cortex. An intriguing result came from a study of Nienborg and Cumming (2009) in primate V2 cortex, which revealed an intrinsically top-down contribution to the neurons' choice probability signals (Fig. 12c–e). Studying a disparity discrimination task, the authors found a dissociation between (i) the causal influence of a stimulus on the animal's decision and (ii) the measurement of choice probability. In the bottom-up interpretation, choice probabilities merely represent the former causal influence, so both (i) and (ii) should display a similar temporal evolution. Instead, their results indicated that the causal influence of a stimulus on the animal's decision is larger in the early period of stimulation (Fig. 12d), while choice probabilities are larger in the late period of stimulation (Fig. 12e). These findings can only be explained if the forming decision itself—and not the mere stimulus values—progressively influences the firing of V2 neurons. Decision formation is thus a likely candidate to explain the growing trial-to-trial variance and loss of stimulus discriminability shown in the rest of Fig. 12.

indicating a global divergence of the population response on a trial-by-trial basis. (h) Correlation analysis for early/late ratios across the population. Different symbols indicate the two animals in the experiment. The dashed line is the predicted threshold ratio (THr) for homogenous Poisson processes. The whole analysis hints at a progressive loss of sensitivity to the stimulus in the neural population, as resulting from the influence of internally generated dynamics on a trial-by-trial basis (see main text). From Uka and DeAngelis (2003) (panel g replotted from article; panel h, unpublished recomputation from the data).

For the moment, these statistical accounts of the "growing variance" effect are purely descriptive. Ultimately, we need population assessments of the effect that are more quantitative and model-based, and again focus on tools of dimensionality reduction to get a better understanding of the internally generated dynamics. Likely, such analysis will have to operate on trial-to-trial recordings of correlated neural populations (Section 7). However, the examples of Fig. 12 show that even independent cell recordings, when combined in population distributions, can start to unveil the complex dynamics at play in the late phases of sensory stimulation in awake behaving animals.

6.2. The temporal structure of internally generated population responses

In contrast to stimulus-locked responses, neural activity that maintains past sensory information, plans future actions, or causes movements, is generated internally. We will here briefly review what we know about such internally produced population activity, considering both activity that precedes an event as well as activity that follows an event. In both cases, the trial-averaged and time-resolved neural response is often called a PETH (peri-event time histogram).

6.2.1. Temporal dynamics of anticipation

Probably the first to investigate the population dynamics of internally generated activity were Georgopoulos et al. (1989). They studied population activity in the motor cortex of a monkey deliberating and performing an arm movement. As described in Section 6.1.2, every cell in their recordings was characterized by its preferred direction, $\mathbf{u}_i$, in the physical space. In turn, the direction that the population of neurons is encoding at any particular time can be read out by the population vector, $\mathbf{A}(s, t) = \sum_i \mathbf{u}_i r_i(s, t)$, i.e., a weighted linear combination of the firing rates r_i of the individual neurons. Here, t is the time and s the desired direction of the arm movement, and the population vector may change with either variable. As before, we can also write this as a matrix-vector multiplication, $\mathbf{A}(s, t) = \mathbf{U}\mathbf{r}(s, t)$, where the columns of $\mathbf{U}$ are given by the preferred directions $\mathbf{u}_i$. The authors show that this population vector, $\mathbf{A}(s, t)$, initially points in the direction of a light signal that indicates where the monkey has to move its hand. In trials where the light signal indicates a movement in a perpendicular direction, the population vector still points into the original direction initially, but then rotates over the course of 260 ± 30 ms to point toward the perpendicular direction. Immediately after this rotation, the monkey moves its arm.

Interestingly, this finding of a rotating population vector is not specific to the motor cortex. Using an oculomotor delayed-response task, Takeda and Funahashi (2004) constructed an analog version of the reaching task, and showed that a similar rotation of the population vector takes place during a three-second delay period in the PFC of a monkey (Fig. 13a). In this task, as well as in the reaching task of Georgopoulos et al. (1989), the population vector provides a straightforward interpretation of the information represented in the population activity—it corresponds to a location in space. However, its gradual change over time requires that the firing rates of some neurons decrease, namely those tuned to the initial direction, while those of other neurons increase, namely those tuned to the final, perpendicular direction. Such increasing and decreasing firing rates, often called "ramping", have been found in many cortical areas whenever an animal actively anticipates an event in time (Durstewitz, 2004), suggesting that the observed population dynamics may generalize to other periods of waiting, deliberating,

or anticipating, beyond the confines of the specific arm reaching or oculomotor tasks.

Such temporal dynamics were probably first reported by Kojima and Goldman-Rakic (1982). A systematic study of the varying dynamics of firing rates in a population of prefrontal neurons was carried out in Brody et al. (2003). The authors focused on the delay period of a task in which monkeys were discriminating two frequencies of a vibrating stimulus applied to the fingertip, and separated by 3 s. Brody et al. (2003) report neurons whose firing rates either increase or decrease. They show that these increases (or decreases) can be ramp-like, with a smooth change of the firing rate over 3 s, or jump-like, meaning that neurons fire at a relatively constant level, and then suddenly shift their firing rates to a higher (or lower) level at some point in the delay period. Since then, several authors have confirmed these findings, both in the PFC (Mita et al., 2009) and OFC of monkeys (Ichihara-Takeda and Funahashi, 2007), in the dmPFC of rats (Narayanan and Laubach, 2009), and in the rat's medial agranular cortex (Matell et al., 2011). In all cases, cells' firing rates increase, decrease, and sometimes follow more complex dynamics while anticipating an event. Generally, the activity of individual neurons is extremely heterogenous (Brody et al., 2003).

In many of these tasks, there is no simple equivalence to the population vector. In the original interpretation put forward by Georgopoulos et al. (1989), neurons have a fixed and preferred tuning, and the rotation of the population vector indicates a change in the spatial information represented (or mused upon) by the monkey. An alternative interpretation is to assume that the represented information remains the same, while the neurons' preferred direction changes over time relative to the movement onset (Sergio et al., 2005; Hatsopoulos et al., 2007; Rickert et al., 2009). Such a change in a neuron's preferred tuning can be evaluated in other tasks as well. Machens et al. (2010) reinvestigated population activity in the prefrontal cortex, using the same delay-period data set as Brody et al. (2003). The authors used a version of PCA that seeks to capture the maximum amount of variance while also demixing the dependencies on time t and stimuli s in the task (Machens et al., 2010; Brendel et al., 2011). They find that six components (or response modes A_i) suffice to represent more than 95% of the variance in the population activity during the delay period of a task (see Fig. 13c). Using their method, the authors find three response modes that only carry information about the task timing (upper three panels in Fig. 13c) and three response modes that carry information about the stimulus frequency (lower three panels in Fig. 13c). The latter three response modes capture the temporal evolution of stimulus tuning in the population.

Similar to the population vector rotation, the authors find a rotation of the original stimulus tuning in anticipation of the upcoming second stimulus frequency. This rotation can be visualized by replotting components 5 and 6 from Fig. 13c, as done in Fig. 13b, left panel. The angular change from the initial tuning direction to the direction at the end of the delay period is plotted in the right panel, showing a similar, almost orthogonal rotation as in Takeda and Funahashi (2004), albeit counterclockwise and with different temporal dynamics. In their work, due to the constraints of the task, it must be the neurons' tuning that changes, rather than the information stored in the monkey brain. We note that the PCA-based analysis also illustrates that the rotation of the stimulus tuning does not capture the complete dynamics in the population. Rather, a large part of the overall activity is related to purely temporal components (Fig. 13c, first row) which do not exhibit any tuning to the stimuli.

The separation of timing information and tuning information into orthogonal subspaces, as shown in Fig. 13c, is similar to the separation of time and stimulus-information shown in Fig. 11b and
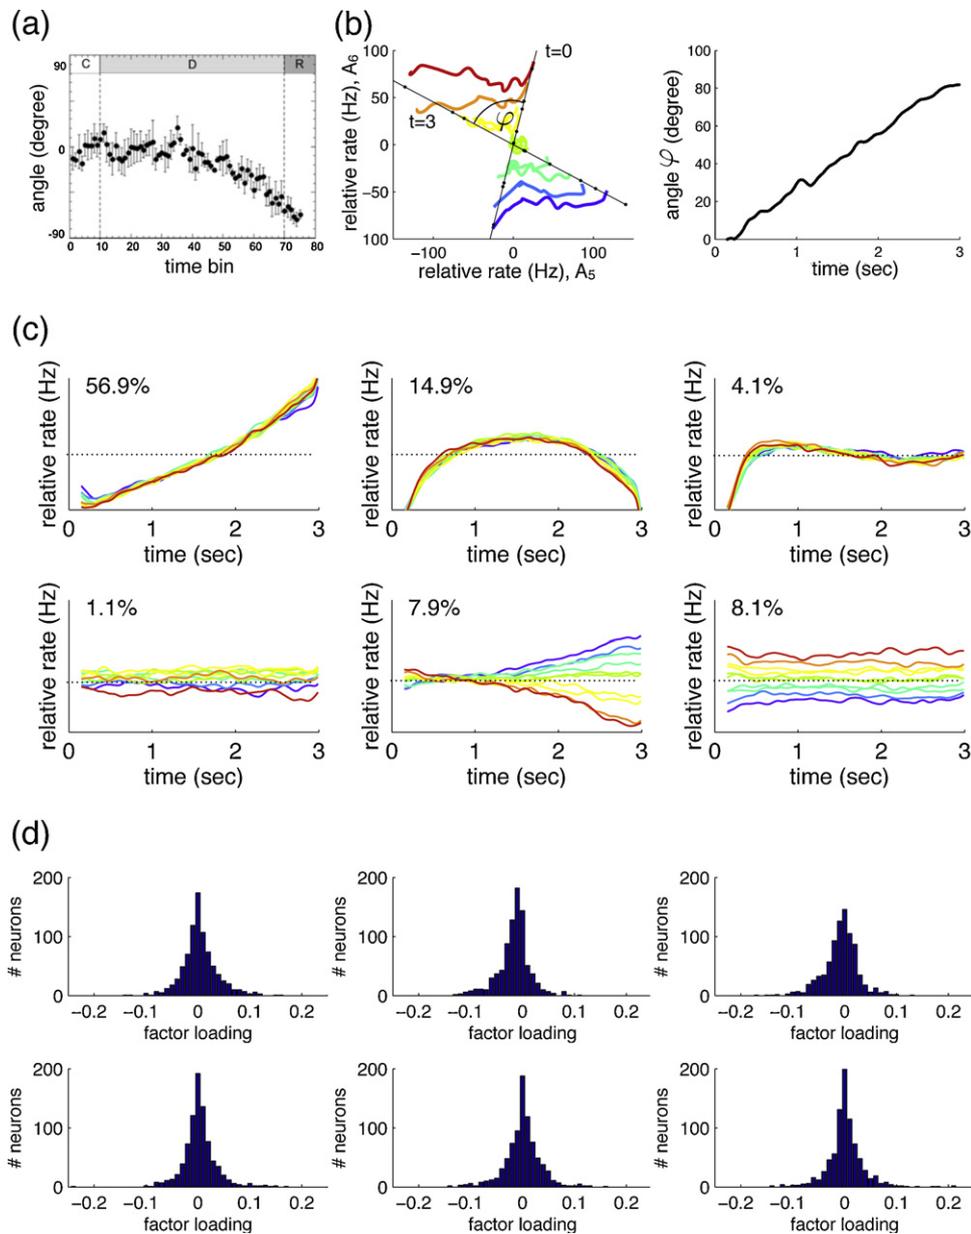
Fig. 13. Population statistics of delay activity. (a) Rotation of the population vector in monkey PFC during an oculomotor-delayed-response task. Shown is the angle of the population vector over time with respect to the population vector at time $t = 0$. Reproduced from Takeda and Funahashi (2004), with permission from Oxford University Press. (b) Rotation of neural tuning in monkey PFC during the delay period of a somatosensory discrimination task. The left panel shows the population activity during the delay period for different stimuli (colored). Shown is the activity of the recorded population ($N = 842$ neurons) projected onto the two response modes with the strongest stimulus tuning. The right panel shows the angle between the original direction of tuning at time $t = 0$, and subsequent tunings. (c) Demixed principal components of PFC population activity during a delay period, from a monkey performing a working memory task. The components show that temporal and stimulus-related activities can be separated (demixed) into orthogonal subspaces. Inset numbers are the percentages of total variance accounted for by each mode. Adapted from Machens et al. (2010). (d) Factor loadings for the principal components from above. The factor loadings show that there are no clear separate classes of cells. For instance, the first panel shows that some cells have large positive loadings (and hence firing rates that ramp up according to the first panel in (c)), some cells have large negative loadings (thereby ramping down), but most cells cluster around zero, so that their firing rates ramp neither up or down (adapted from Machens et al., 2010).

c. Indeed, a similar separation of timing, stimulus, and decision-related information for a whole task was recently reported by Brendel et al. (2011), suggesting that these types of orthogonal representations could be a general organization principle for higher-order areas in the brain.

6.2.2. The population statistics of temporal dynamics and the lacking evidence for cell classes

The description of firing rates in terms of decreasing versus increasing "ramps" implies that there may be different classes of cells, such as up-ramping versus down-ramping neurons. There is, however, no evidence for that if all of the recorded neurons are included in the analysis. For instance, Narayanan and Laubach (2009) analyze delay-period activity in the dmPFC of rats using PCA. They show that most of the firing rate variance in the delay period data is captured by a few components. The factor loadings of these components, however, follow a unimodal distribution. A similar result was obtained by Machens et al. (2010) in their analysis of delay-period activity in monkey PFC (described above), showing

sparse and unimodal distributions of factor loadings with exponential tails (see Fig. 13d).

Of course, the existence of cell classes could require more sophisticated analysis than the plotting of factor loadings from PCA. We note that to the extent that researchers have found distinct classes of cells, this has often been based on a priori exclusion of neurons that showed "insignificant" responses with respect to the feature under investigation. This exclusion of cells, however, is prone to turn a unimodal distribution into a bimodal distribution, thereby artificially introducing classes of cells. Systematic attempts to find classes of cells should be based on clustering analysis. This was done in a recent study by Jin et al. (2009), for instance, who identified 66 classes of cells by analyzing the responses of more than 2000 neurons in the dlPFC of monkeys performing a sequential saccades task, i.e., a task that stretched over more than one delay period. Apart from the large number of cell classes, we note that about one third of the neural responses could not be attributed to any of these classes. Accordingly, if there are different classes of cell responses, these are not likely to be easily interpretable.

6.2.3. Single trials and the anticipation of variable delays

Most of the work reviewed in the last section was concerned with delay activity recorded from animals performing tasks with a fixed delay period. What happens when the length of the delay period changes? Once the animal has adjusted to the new length, several authors have found that the dynamics of the activity rescales (Komura et al., 2001; Brody et al., 2003; Mita et al., 2009; Machens et al., 2010; Shinomoto et al., 2011). While the extent to which these activity changes constitute a real rescaling still has to be quantified, the evidence clearly suggests that the population of neurons reaches exactly the same state, i.e., combination of firing rates, at the end of the longer delay period that it reached at the end of the shorter period. These observations suggest that the system uses the delay period to move between two states in state space, one at the onset of the delay and one at the offset. We note that the end of the delay period is usually identical with the onset of a stimulus or a movement.

One can make similar observations if the delay period is variable as well. This can happen if the delay is not known to the animal beforehand, or if the animal itself has some influence on the length of the delay period. In these cases, it is much harder to compare activity across trials, and the study of population activity essentially has to rely on simultaneous recordings of cells. Nonetheless, a common observation in these studies concerns what happens with the population activity in the neighborhood of a measured or controlled event (such as a movement onset). In this case, firing rates converge to the same point independent of their time course beforehand (Churchland et al., 2010). This type of threshold crossing is important for proactive movements (Maimon and Assad, 2006; Churchland and Shenoy, 2007; Afshar et al., 2011), and suggests that part of the computations carried out during the delay period may be a mechanistic preparation of the system for the arrival of a new stimulus or the initiation of a motor act (Churchland et al., 2010).

We note that even in studies with a fixed delay period, averaging over trials to construct PETHs will obscure part of the population dynamics. Such averaging generally assumes that there are no systematic differences between trials, i.e., that the dynamics of the system unfolds equally every single time. This, however, is unlikely to be the case, since it would require very precise timing of changes in neural activity, far more precise than what is known from psychophysics. Evidence from the motor systems clearly points to significant trial-to-trial differences in the unfolding of anticipatory delay activity (Churchland and Shenoy, 2007; Churchland et al., 2010; Afshar et al., 2011). Even in the absence

of stimulation, we therefore run into the same phenomena that we studied in Section 6.1.4. The measured PETHs could, therefore, be generated by various combinations of heterogeneous single trials. Several authors have suggested, for instance, that neurons with ramping activity could simply be neurons that show stepwise increases in firing rates, yet these increases occur at different times of the delay period for different trials (Durstewitz, 2004; Mongillo et al., 2003). Indeed, Brody et al. (2003) showed that some neurons have sudden changes in firing rate. While a systematic investigation of ramping activity on the single-trial level is still to come, we note that the answer is likely to be both: there are neurons that rapidly change firing rates during the delay period at different time points, and there are neurons whose firing rates increase smoothly and almost linearly during the same period (Kesseli, Romo, Machens, personal communication).

7. Population distributions of correlations

So far we have limited our survey to descriptions of population activity where neurons are considered independently; in the "urn" metaphor of Fig. 1e, we randomly picked neurons one by one. This approach, although powerful, neglects functional (and anatomical) relationships which may exist between neurons as a result of their embedding in the same connected network. These inter-neuron relationships generally bear the name of neural *correlations*. Because of combinatorial explosion, correlations are most often studied on a pairwise basis; in the image of Fig. 1e, this is equivalent to drawing pairs of neurons from the urn, rather than individual neurons.

Neurons can be correlated along several dimensions. First, *anatomical* correlations can be defined, e.g., based on the spatial position of neurons in cortex. Second, correlations can be defined at a functional level, on the basis of individual neurons' trial-averaged firing rates. Such correlations take the generic name of *signal* correlations, from the traditional notion that signals are encoded only in trial-averaged firing rates. Third, correlations can be studied on a trial-to-trial basis. Trial-to-trial covariations of activity between neurons have long been observed and take the generic name of *noise* correlations; they arise because the neurons are involved in common pathways of information processing, e.g., through shared connectivity. Measuring such noise correlations is a technical challenge, as it requires the simultaneous recording of several neurons. In Section 7.1 we discuss the distribution of anatomical and signal correlations, and justify the interest in studying them. We then turn in Section 7.2 to the harder, yet more widely studied subject of noise correlations.

7.1. Anatomical and signal correlations

Since anatomical and signal correlations can be derived from individually recorded neurons, one may question whether they contain any useful information *per se*. For example, anatomical information can directly be translated into a physical mapping of neurons, which completely determines the distribution of inter-neuronal distances. Similarly, signal information (i.e., trial-averaged responses) can directly be translated into a distribution of tuning (Section 5) and/or dynamics (Section 6) which completely determines the structure of signal correlations. In primary sensory cortices, such anatomical and signal (tuning) information have been combined into a fair understanding of feature mapping (orientation columns in V1, tonotopy in A1, body maps in S1), without any need to resort to pairwise correlations.

However, precisely the simplicity of anatomical and signal correlations renders them useful in complex situations where full population distributions are difficult to construct. For example, signal correlation matrices are the typical objects on which

Principal Component Analysis is performed to extract the dominant dynamics in a population's trial-averaged response (Section 6.1.2). Moreover, signal correlations can prove useful in higher sensory areas, since here the objects encoded bear complex relations to the raw stimulus information, and the concept of a tuning curve becomes unwieldy. Fig. 14a shows a typical situation where the full distribution of population tuning is out of reach.

Here, Sato et al. (2009) recorded from macaque IT neurons responding to a variety of objects and faces. Each blue cross in Fig. 14a represents a *pair* of neurons and their signal correlation, computed as the linear correlation between the two neurons' tuning curves to the complex set of stimuli. Exact encoding properties are lost in this process, but not their repartition across neurons; the authors could thus reveal the presence of "activity
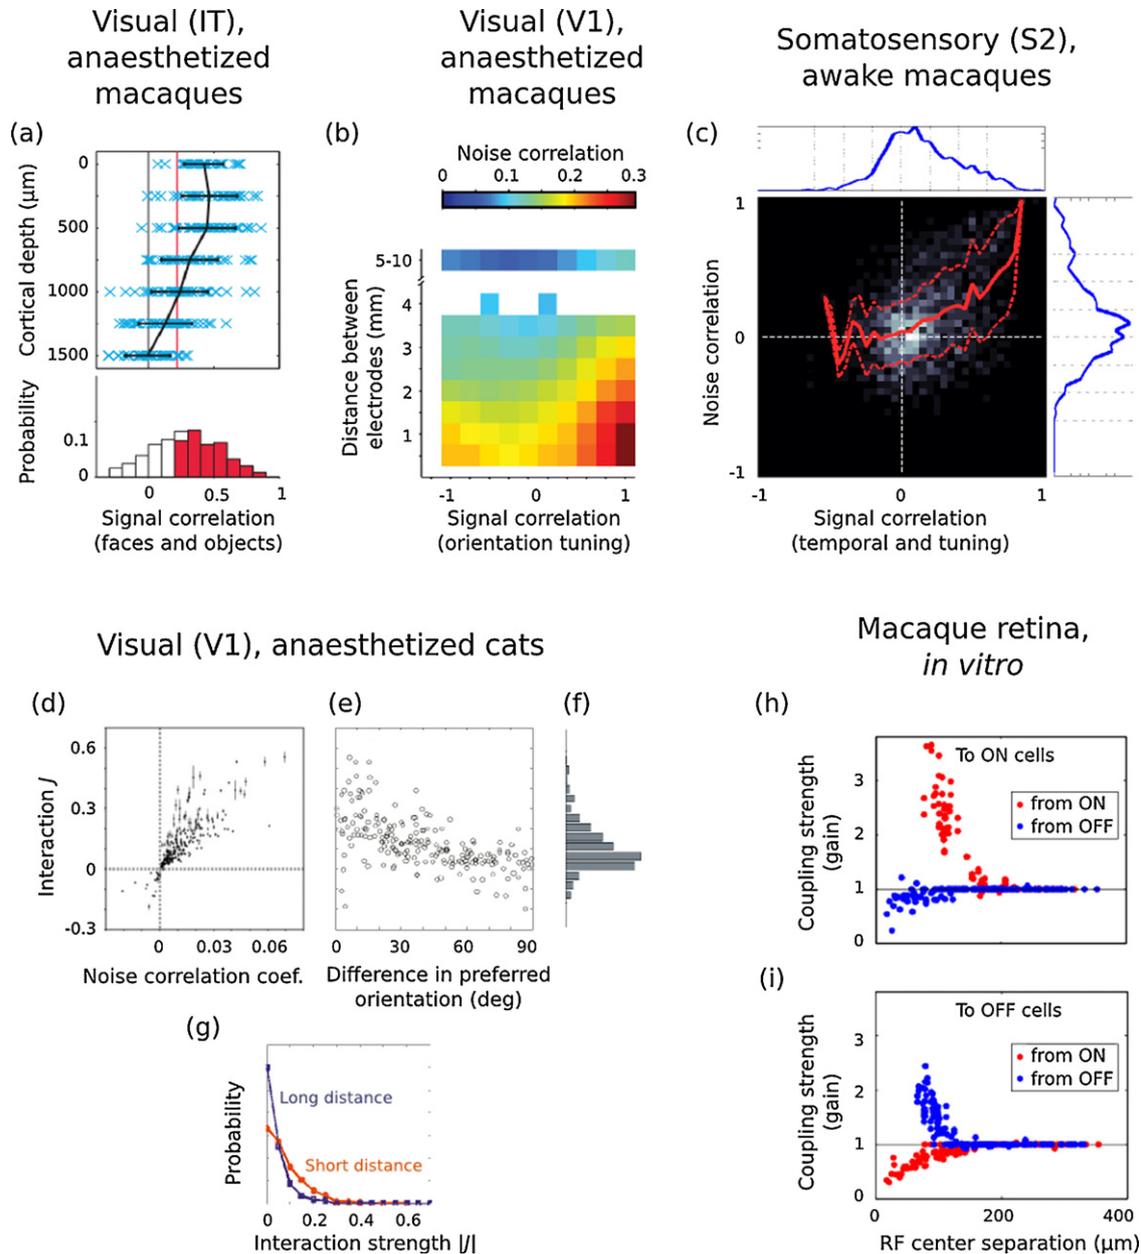
Fig. 14. Population distributions of correlations. (a) Distribution of signal correlations in IT neurons responding to various visual objects. Each blue cross depicts a pair of recording sites at the same cortical depth (ordinate), with the abscissa encoding signal correlation strength (similarity of tuning to a set of 60 objects) between multiunit activities at the two sites. The black curve and bars are mean and standard deviation of pairwise tuning similarity. The red line is the $p = 0.05$ significance threshold. The clear dependency between cortical depth and signal correlations is the mark of "activity spots", whose neurons have broadly similar tuning and which are located mostly in the upper layers of cortex. Reproduced from Sato et al. (2009), with permission from Oxford University Press. (b) Average pairwise noise correlation strength in macaque V1, plotted as a function of distance and tuning similarity (orientation) between the recorded neurons; reproduced from Smith and Kohn (2008), with permission from the Society for Neuroscience. (c) Joint population distribution of signal correlation and noise correlation coefficients between pairs of neurons in secondary somatosensory cortex of macaques during a tactile discrimination task. The signal correlation between two neurons is computed as the overall similarity of their trial-averaged firing rates across time and stimuli (as in Wohrer et al., 2010). Blue histograms on the sides are marginalized distributions. Red curves represent conditional mean and standard deviation of noise correlation value at a given level of signal correlation. The width of the standard deviation plays an active role in estimating the overall sensitivity to stimulus in the population. Replotted from Wohrer et al. (2010), a re-analysis of data from Hernández et al. (2010). (d)–(g) Statistics of coupling strengths in an Ising model fit to cat V1 neurons. Coupling strengths relate, but are not equivalent, to classical noise correlation coefficients (panel d). Coupling is higher for similarly tuned neurons (orientation, panel e) and for nearby neurons (panel g). Reproduced from Yu et al. (2008), with permission from Oxford University Press. (h) and (i) Statistics of coupling strengths in a Generalized Linear Model fit to retinal ganglion cell activity. On and Off cells are mostly interacting through mutual inhibition, whereas cells of the same polarity are positively coupled. Mean coupling strength decreases with interneuronal distance. Reproduced from Pillow et al. (2008), with permission from Nature Publishing Group.

spots" in the upper layers of IT, that is, localized regions with broadly similar tuning properties.

Even in primary sensory cortices, the well-known feature maps have strong limitations. First, they do not apply uniformly across species—for example there are no orientation columns in rodent V1. Second, recent calcium imaging studies have revealed that these mappings can display strong discontinuities (Ohki et al., 2005, direction maps in cat V1) or even totally break down at the single-cell resolution (Rothschild et al., 2010, tonotopy in mouse A1). In such settings too, signal correlations and inter-cellular distance become a useful way of statistically describing the embedding of tuning properties on the cortical surface (see e.g., Rothschild et al., 2010).

Finally, the study of anatomical and signal correlations acquires further importance when considering their relation to noise correlations. Since all three forms of correlations have the same pairwise structure, one may ask how they are linked, and how they constrain each other. In particular, signal correlations play an important role in studies of population coding: coupled with measures of noise correlation, they allow us to assess the decoding capacities of a neural population. These issues are reviewed in the next section.

7.2. Noise correlations

Noise correlations have attracted a lot of interest, because their precise form can completely change the overall picture of neural activity in the population. For instance, noise correlations can reveal clustered patterns of activation between neighboring neurons; they can show balanced influences between mutually inhibitory populations; and they can uncover synchronous waves of activation across large stretches of cortex. Furthermore, the amount of (Fisher) information about a stimulus which can be read out from a sensory population relies on the relationship between the signal and noise covariance matrices (explained below, see Abbott and Dayan, 1999; Averbeck et al., 2006), so knowledge of the latter is mandatory for any quantitative estimate. Also, noise correlations allow us to track population-wide modulations of activity resulting from attentional effects (Roelfsema et al., 2004; Cohen and Newsome, 2009; Mitchell et al., 2009; Cohen and Maunsell, 2011), thus explaining some of the uncontrolled variability inherent to single neuron recordings.

Unfortunately, noise correlations are also a difficult subject to study. First, there are technical limitations on the number of neurons which can be simultaneously recorded, although powerful techniques such as microelectrode arrays, tetrodes, or optical imaging have increased this number during recent decades. Second, it is essentially impossible to reliably estimate all the statistics which fully define a population's joint firing process.[14] It generally takes more trials to estimate the pairwise distribution of two neuron's firing than to estimate their respective mean firing rates, and again more trials to estimate the distributions for triplets of neurons, and so on. When time is added in the process this combinatorial explosion is yet larger, and even the pairwise structure of neural activity, assessed through the neurons' *Joint Peri-Stimulus Histograms* (Aertsen et al., 1989), becomes difficult to measure with the number of trials usually granted to experimenters.

As a result, although noise correlations are central to population descriptions of neural activity, we know little of their population-wide distributions. While it is straightforward to estimate noise correlations from simultaneous recordings, this is not the same as estimating their distribution. To clarify this point, imagine we record simultaneous activities from an ensemble of 20 neurons using tetrodes, and compute for each pair of neurons (i, j) the Pearson index of noise correlation, ρ_{ij}. What do these measures tell us about the noise correlations we should expect to find in the next session run, when moving the tetrodes to another site from the same neural area? In this section, we will distinguish a neuron *ensemble*, simultaneously recorded neurons during a session run, from a neuron *population*, which denotes a much larger set of neurons with anatomical and functional coherence. To understand the structure of noise correlations in the full population, one must assess their *statistical* organization from the set of recorded ensembles.

Because estimating noise correlations is challenging in itself, most experiments limit themselves to the simplest account of neural activity, spike counts. The object describing neural activity is then the N-dimensional vector $\mathbf{r} = (r_1, r_2, \ldots, r_N)$ collecting the spike count from each neuron in the ensemble. Because of noise correlations, the values taken by $r_1, r_2, \ldots$ on each trial are not independent, but undergo common variations. The most straightforward way to encompass these covariations is the linear noise correlation (Pearson) coefficient, expressed for each pair of neurons (i, j) in the ensemble as:

$$\rho_{ij} = \frac{\langle (r_i - \langle r_i \rangle)(r_j - \langle r_j \rangle) \rangle}{\sqrt{\langle (r_i - \langle r_i \rangle)^2 \rangle \langle (r_j - \langle r_j \rangle)^2 \rangle}},$$

where the symbol $\langle \cdot \rangle$ denotes averaging over trials. The Pearson correlation coefficient ρ_{ij} ranges between -1 and 1, with value 1 indicating complete covariation across trials, and value -1 indicating complete anti-covariation across trials. A value of zero indicates the absence of any (linear) dependency between the two neurons on a trial-to-trial basis. The concatenation of all these measures in a single matrix $\boldsymbol{\rho} = (\rho_{11}, \rho_{12}, \ldots)$ defines the so-called *noise correlation matrix* of the population.

7.2.1. Noise correlations relate to anatomical and functional structures of cortex

Such noise correlation coefficients are routinely estimated for pairs of neurons in present-day experiments. However, less work has concentrated on determining their statistical patterns across the population. Where this has been done, noise correlations have generally been measured in relation to two other indicators: distance between neurons, a measure of anatomical correlation, and similarity of tuning, a measure of signal correlation. Fig. 14b (Smith and Kohn, 2008) depicts an example of such analysis: the two-dimensional plot shows the average strength of noise correlations between pairs of recorded neurons, as a function of the distance between the two recording electrodes, and of similarity of orientation tuning. The resulting pattern is stereotypical and has been observed in several studies: noise correlations are statistically stronger (i) for nearby neurons and/or (ii) for neurons with similar tuning properties.

The dependency of noise correlation strength on interneuronal distance reflects the increased chance that nearby neurons share common inputs, leading to covariations in their activities on a trial-to-trial basis. This dependency is a robust result which has been observed in a variety of settings (Lee et al., 1998; Petersen et al., 2001; Constantinidis and Goldman-Rakic, 2002; Kerr et al., 2005; Yu et al., 2008; Ohiorhenuan et al., 2010; Rothschild et al., 2010), suggesting that cortical activity often involves the co-occurrent activities of local circuits of cells. Noise correlation strengths have also been related to other anatomical indicators. In two-photon imaging of rat barrel cortex responding to whisker deflection, Kerr

---

[14] The so-called *Janóssy measures*, in the nomenclature of point-processes (Daley and Vere-Jones, 2007). For example, the neurons' PSTHs define the first order Janóssy measure of the population process, whereas all pairwise JPSTHs define the second order Janóssy measure. Full description of a population of neurons would require knowledge of Janóssy measures up to order ∞.

et al. (2005) reported increased noise correlations between pairs of neurons closer to the barrel center for the stimulated whisker, likely resulting from increased levels of shared task-related inputs. In primary visual cortex of anesthetized macaques, Constantinidis and Goldman-Rakic (2002) reported different strengths of noise correlations between (putative) excitatory and inhibitory cells, with higher levels of noise correlation between inhibitory neurons.[15]

The dependency of noise correlation strength on tuning similarity observed in Fig. 14b is a first example of the numerous studies which have explored the statistical link between noise correlations and various measures of signal correlations. Although the precise definition of the term signal correlation varies between authors,[16] all versions aim at the same goal: characterizing neurons with similar responses to the task. Similarly, Fig. 14c presents the joint distributions of signal and noise correlations in macaque S2 cortex (data re-analyzed from Wohrer et al., 2010; Hernández et al., 2010). Fig. 14b and c both display the most commonly found pattern: an increased chance of finding positive noise correlation in pairs of neurons with positive signal correlation (Lee et al., 1998; Bair et al., 2001; Constantinidis et al., 2001; Constantinidis and Goldman-Rakic, 2002; Kohn and Smith, 2005; Rothschild et al., 2010). This observation again has an intuitive explanation: *shared tuning reveals shared input*. Similarity in tuning between two neurons is likely to occur from their involvement in a shared processing pathway, resulting from shared connectivity. In turn, this shared connectivity will often result in covariations of activity on a trial-by-trial basis. Ko et al. (2011) have recently provided explicit evidence for this intuition in mouse V1 cortex, where the probability of finding a direct synaptic connection between two neighboring neurons increases with the strength of their signal and/or noise correlation.

7.2.2. The statistical distribution of correlations constrains the neural code

Signal and noise correlation have a particular importance in the field of neural coding, as their joint organization defines the form, and overall amount, of information conveyed by the population in the task. For example, the (Fisher) information about a stimulus that is embedded in a sensory population depends on the respective strengths and orientations of the signal covariance matrix and the noise covariance matrix (Abbott and Dayan, 1999; Averbeck et al., 2006; Shamir and Sompolinsky, 2006).[17] The "signal/noise" distributions shown in Fig. 14c can be interpreted as an element-wise statistical link between these two matrices. Thus, these distributions are likely to encompass the core information needed for estimates of coding capacity in the population.

Following this intuition, Wohrer et al. (2010) proposed a specific method to generate noise correlation matrices on the basis of a signal correlation matrix, while respecting the element-wise statistical relationship between the two (Fig. 14c); using this extrapolation method, they could predict the coding efficiency of a full population of recorded neurons in macaque primary somatosensory cortex. In contrast to a common assumption (Shadlen et al., 1996; Sompolinsky et al., 2001), noise correlations were not detrimental to coding efficiency. The overall distribution of sensitivity in the population actually resembled that estimated when assuming neuronal independence, although the population

did display an important amount of noise correlations. As an important note, predicted results changed significantly whether one considered only the *mean* signal/noise relationship (plain red curve in Fig. 14c) or also incorporated its statistical *variations* across the population (dotted red curves in Fig. 14c). Indeed, neglecting these statistical variations underestimated the number of neural pairs with signal- and noise-correlations of opposite signs (i.e., the upper left and bottom right quadrants in Fig. 14c), whereas these pairs typically induce high synergetic encoding (Romo et al., 2003; Averbeck et al., 2006). This result exemplifies the need to strive for more precise statistical descriptions of the noise correlation structure, instead of simply computing population averages.

In summary, even limiting oneself to spike counts, understanding the structure of cortical activity on a trial-to-trial basis will require establishing better population-wide relationships between (i) cortical anatomy, (ii) tuning properties, and (iii) noise correlations. Much work is left on this road; even more since the object of study itself—pairwise spike count correlation—is a somewhat tricky measure[18] which can be subject to a number of experimental artifacts. First, de la Rocha et al. (2007) have shown how pairwise spike count correlations do not only reflect a degree of correlation in the neurons' *inputs*, but also systematically correlate with the two neurons' output firing rates. This phenomenon, due to the integrate-and-fire nature of neurons, was systematically observed both in models and *in vitro* recordings. Second, measures of spike count correlations are very sensitive to various experimental artifacts which may artificially increase their values (Brody, 1999), such as attentional states (Mitchell et al., 2009; Cohen and Maunsell, 2011), up and down states during anesthesia (Harris and Thiele, 2011), reaction times, electrode migrations, etc. Various lines of work have actually suggested that active neuronal mechanisms may promote either cortical decorrelation (Montani et al., 2007; Ecker et al., 2010; Renart et al., 2010) or, more generally, correlation structures that are transparent to downstream neural decoders—in the sense that a decoder assuming neural independence does not lose much information compared to a full decoder (Nirenberg et al., 2001; Latham and Nirenberg, 2005, but see also Graf et al., 2011). Several breakthroughs are to be expected in the near future on the subject of noise correlations. We believe that these breakthroughs will ultimately bear a statistical expression across full populations.

7.2.3. Beyond spike count correlations

There are more features of correlated activity than simply spike count correlations. Studies having addressed these issues can grossly be divided according to two main focuses: (i) to characterize assemblies of simultaneously active cells, or (ii) to characterize sequential structures of firing in a population. These are complex problems whose study now represents large and dynamic areas of research in neuroscience. Yet, to date, population-wide statistical descriptions of the phenomena are still rare.

Several recent lines of work have striven to better characterize synchronous co-activations in cell ensembles, in a variety of experimental settings (Schneidman et al., 2006; Shlens et al., 2006; Osborne et al., 2008; Yu et al., 2008, 2011; Peyrache et al., 2009; Humphries, 2011; Ohiorhenuan et al., 2010; Lopes-Dos-Santos et al., 2011). These studies generally consider the ongoing succession of ensemble activity patterns $(r_1, r_2, \ldots, r_N)$ computed in short successive time windows, and search to characterize the overall distribution of occurrences of specific patterns across long streaks of little constrained brain activity (e.g., spontaneous

---

[15] We note however that these higher correlation coefficients may also be a confound of higher firing rates in inhibitory cells—see de la Rocha et al. (2007), also detailed further down.

[16] For example, signal correlation can be computed across different stimuli and/or different instants in time.

[17] More precisely, Fisher information can generally be expressed as the largest eigenvalue of matrix SN^{-1}, where S and N are respectively the signal- and noise-covariance matrices pertaining to the system.

[18] For example it is highly nonlinear, since it involves division by the variance terms.

activity, or ongoing responses to natural stimuli). In particular, a number of studies have tried to characterize "instantaneous" distributions $p(r_1, r_2, \ldots, r_N)$ where all r_i are binary activations in very short time windows. The Ising model has recently been proposed as a natural fit for this distribution (Schneidman et al., 2006; Shlens et al., 2006; Osborne et al., 2008; Yu et al., 2008), since it is the model which accounts for pairwise correlations with the least model-related assumptions (maximum entropy distribution). It has further been suggested that the pairwise coupling strengths J_{ij} in the fitted Ising model may have a more direct link with underlying neural connectivity, than corresponding spike count correlations ρ_{ij} (Schneidman et al., 2006; Yu et al., 2008). Fig. 14d–g (Yu et al., 2008) shows population distributions for coupling strengths J_{ij} in several 10-neuron ensembles from V1 cortex. Values for J_{ij} covary with, but are not equivalent to, corresponding spike count correlations ρ_{ij} (panel d). One finds again the functional relationship to tuning similarity (signal correlation, panel e), and the anatomical relationship to interneuronal distance (panel g). Further work along this line has found discrepancies between neural activities and the Ising model, especially for nearby neurons (Ohiorhenuan et al., 2010). An alternative model based on a dichotomized Gaussian process has been proposed to yield better fits to observed pattern occurrences (Macke et al., 2009; Yu et al., 2011); we know yet of no population distributions for the resulting model parameters.

The temporal aspects of noise correlations have also met renewed interest, and specific attempts have been made to unveil causal dependencies between the spikes of different neurons in an ensemble. The traditional and most straightforward way of estimating these dependencies is to build a cross-correlation histogram for each pair of simultaneously recorded neurons. Oddly, although building such cross-correlograms is very common practice, we have not found any statistical descriptions of how the *temporal* features of these cross-correlograms (width of peak, time of peak) are distributed in a neural population. Recently, a particular family of point processes (Generalized Linear Models, Paninski, 2004; Pillow et al., 2008) has been proposed to account for spiking dependencies between correlated neurons, including temporally structured dependencies. Generalized Linear Models have the advantage of bearing a convenient mathematical expression allowing easy fits from large sets of data (unlike Ising models). We know of no work having studied the population statistics of the temporal structure embedded in Generalized Linear Model descriptions of neural ensembles. A population description of retinal interaction *strengths*, based on Generalized Linear Models, is pictured in Fig. 14h–i (Pillow et al., 2008). More advanced descriptions of the kind appear a promising way of describing spiking interactions at the population level.

8. Conclusion

The perplexing heterogeneity of cortical responses has long been known to electrophysiologists. We have here described straightforward statistical ways of characterizing this heterogeneity, be it through simple distributions of firing rates, through distributions of tuning parameters, or through the more complex distributions or statistics of the temporal dynamics of activity. A majority of these results took the form of purely empirical measures, where (i) experimental features (such as firing rates or tuning parameters) are purely descriptive statistics of activity, and (ii) the population distributions are scatter-plots of experimental data. This approach is simple and often leads to results directly interpretable by eye or in terms of simple statistics, e.g., a positive correlation between two indicators across the population.

Although many classical studies (e.g., Hubel and Wiesel, 1962) have taken great care in describing and characterizing this

heterogeneity, we were surprised in our survey of the literature how atypical such statistical descriptions have become. The most common approach nowadays seems to be to calculate means over classes of neurons. However, we have found little evidence for distinct, functional classes of neurons. The prevalent notion of cell classes seems to stem from two faulty concepts: (1) an a priori exclusion of cells based on the "insignificant responses". This exclusion tends to eliminate cells at the center of a distribution, thus turning a unimodal distribution into a bi- or multi-modal distribution; and (2) an a priori determination of cell classes by eye. While our eyes can help us in identifying features of data, any presumed cell classes should nonetheless by verified by statistical methods. Cell classes do exist in certain cases (see e.g., Cohen et al., 2012), but should not be presupposed. We note that even in V1 the distinction between simple and complex cells is less clear than thought, as shown in Ringach et al. (2002).

As statistical descriptions get finer, one may start considering a move from histogram and scatter plots to a quantitative characterization of neural response heterogeneity. For instance, is a distribution of firing rates exponential or log-normal? Such quantitative characterizations rely on statistical methods and the fitting of specific probabilistic models to a set of data. At the population level, we can assemble the response features of individual neurons, q, and then construct and fit a distribution $p(q)$ to the data, as exemplified in Fig. 1. As we have seen, such quantitative fits are still very rare.

Such model-based descriptions of neural activity bear a number of practical and conceptual advantages that could boost the study of population distributions. First, model-based analysis allows for better statistical assessment. It can help produce error bars for complex experimental measures, or quantify the goodness of a fit thanks to various techniques of model assessment and selection (Hastie et al., 2009). Second and more importantly, model-based analysis allows us to *generate new data*. If a neural population is totally described by statistical models for its population-wide distributions, including its trial-to-trial variability, then it becomes possible to pick random samples from these distributions, and thus consider neural ensembles from that population of any size, possibly much larger than the actual number of recorded neurons. Model-based descriptions could thus help to extrapolate from the recorded data to the whole population, allowing us to make quantitative predictions at the level of the full neural population. Third, the statistical quantification of population activity allows us to compare neural activity across areas, observing commonalities and differences in the population responses. Firing rate distributions, for instance, seem remarkably conserved across areas, from sensory to decision-making to motor areas. Only a surprisingly small subset of neurons seems to be active at any point in time. There are likely similar preserved features in the tuning of neurons across a population, some of which we have pointed out in Section 5. Even on the level of dynamics, there are interesting similarities that can be observed after the onset of sensory stimuli, during the anticipation of stimuli or movements, or during delay periods. We believe that the investigation and clarification of these commonalities will continue to yield important constraints for any theory of the cerebral cortex.

Last but not least, in reviewing such a large section of the literature we are likely to have missed some studies. To the extent that this is the case, it has been an oversight.

References

Abbott, L.F., Dayan, P., 1999. The effect of correlated variability on the accuracy of a population code. Neural Computation 11 (January (1)), 91–101., http://www.ncbi.nlm.nih.gov/pubmed/9950724.

Abeles, M., 1991. Corticonics: Neural Circuits of the Cerebral Cortex. Cambridge University Press, Cambridge, UK.

Adrian, E., 1928. The Basis of Sensation. WW Norton & Co, New York, USA.

Aertsen, a.M., Gerstein, G.L., Habib, M.K., Palm, G., 1989. Dynamics of neuronal firing correlation: modulation of "effective connectivity". Journal of Neurophysiology 61 (May (5)), 900–917., http://www.ncbi.nlm.nih.gov/pubmed/2723733.

Afshar, A., Santhanam, G., Yu, B.M., Ryu, S.I., Sahani, M., Shenoy, K.V., 2011. Single-trial neural correlates of arm movement preparation. Neuron 71 (August (3)), 555–564., http://www.ncbi.nlm.nih.gov/pubmed/21835350.

Allen, E.A., Erhardt, E.B., Calhoun, V.D., 2012. Data visualization in the neurosciences: overcoming the curse of dimensionality. Neuron 74, 603–608.

Amirikian, B., Georgopoulos, A.P., Georgopulos, A.P., 2000. Directional tuning profiles of motor cortical cells. Neuroscience Research 36 (January (1)), 73–79., http://www.ncbi.nlm.nih.gov/pubmed/10678534.

Arieli, A., Sterkin, A., Grinvald, A., Aertsen, A., 1996. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. Science 273, 1868–1871.

Asari, H., Pearlmutter, B.A., Zador, A.M., 2006. Sparse representations for the cocktail party problem. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 26 (July (28)), 7477–7490., http://www.ncbi.nlm.nih.gov/pubmed/16837596.

Attwell, D., Laughlin, S.B., 2001. An energy budget for signaling in the grey matter of the brain. Journal of cerebral blood flow and Metabolism 21, 1133–1145.

Averbeck, B.B., Latham, P.E., Pouget, A., 2006. Neural correlations, population coding and computation. Nature Reviews. Neuroscience 7 (May (5)), 358–366., http://www.ncbi.nlm.nih.gov/pubmed/16760916.

Bair, W., Koch, C., Newsome, W., Britten, K., 1994. Power spectrum analysis of bursting cells in area MT in the behaving monkey. Journal of Neuroscience 14, 2870–2892.

Bair, W., Zohary, E., Newsome, W.T., 2001. Correlated firing in macaque visual area MT: time scales and relationship to behavior. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 21 (March (5)), 1676–1697., http://www.ncbi.nlm.nih.gov/pubmed/11222658.

Barbieri, F., 2008. Can attractor network models account for the statistics of firing during persistent activity in prefrontal cortex? Frontiers in Neuroscience 2 (1), 114–122., http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2570072/.

Barlow, H., 1961. Possible principles underlying the transformation of sensory messages. Sensory Communication 217–234.

Barlow, H.B., 1972, January. Single units and sensation: a neuron doctrine for perceptual psychology? Perception 1, 371–394., http://www.ncbi.nlm.nih.gov/pubmed/19806956.

Barth, A.L., Poulet, J.F.A., 2012. Experimental evidence for sparse firing in the neocortex. Trends in Neurosciences 35, 345–355.

Bartho, P., Curto, C., Luczak, A., Marguet, S.L., Harris, K.D., 2009. Population coding of tone stimuli in auditory cortex: dynamic rate vector analysis. The European Journal of Neuroscience 30 (November(9)), 1767–1778., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2861167&tool=pmcentrez&rendertype=abstract.

Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E., Pouget, A., 2008. Probabilistic population codes for Bayesian decision making. Neuron 60 (December (6)), 1142–1152., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2742921&tool=pmcentrez&rendertype=abstract.

Beloozerova, I.N., Sirota, M.G., Swadlow, H.a., 2003. Activity of different classes of neurons of the motor cortex during locomotion. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 23 (February (3)), 1087–1097., http://www.ncbi.nlm.nih.gov/pubmed/12574439.

Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P.L., Gioanni, Y., Battaglia, F.P., Wiener, S.I., 2010. Coherent theta oscillations and reorganization of spike timing in the hippocampal-prefrontal network upon learning. Neuron 66 (June (6)), 921–936., http://www.ncbi.nlm.nih.gov/pubmed/20620877.

Berkes, P., Orbán, G., Lengyel, M., Fiser, J., 2011. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. Science 331 (January (6013)), 83–87., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3065813&tool=pmcentrez&rendertype=abstract.

Blanche, T.J., Spacek, M.A., Hetke, J.F., Swindale, N.V., 2005. Polytrodes: high-density silicon electrode arrays for large-scale multiunit recording. Journal of Neurophysiology 93 (May (5)), 2987–3000., http://www.ncbi.nlm.nih.gov/pubmed/15548620.

Brecht, M., Schneider, M., Sakmann, B., Margrie, T.W., 2004. Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. Nature 427 (February (6976)), 704–710., http://www.ncbi.nlm.nih.gov/pubmed/14973477.

Brendel, W., Romo, R., Machens, C.K., 2011. Demixed principal component analysis. In: Advances in Neural Information Processing Systems 24: Proceedings of the 2011 Conference (NIPS 2011). pp. 1–9.

Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., Movshon, A.J., 1996. A relationship between behavioral choice and the visual response of neurons in macaque MT. Visual Neuroscience 13, 87–100.

Brody, C., 1999. Correlations without synchrony. Neural Computation 11 (October (7)), 1537–1551., http://www.ncbi.nlm.nih.gov/pubmed/10490937.

Brody, C.D., Hernández, A., Zainos, A., Romo, R., 2003. Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. Cerebral Cortex 13 (November (11)), 1196–1207, http://dx.doi.org/10.1093/cercor/bhg100.

Busse, L., Wade, A.R., Carandini, M., 2009. Representation of concurrent stimuli by population activity in visual cortex. Neuron 64 (December (6)), 931–942., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2807406&tool=pmcentrez&rendertype=abstract.

Cardin, J.A., Kumbhani, R.D., Contreras, D., Palmer, L.A., 2010. Cellular mechanisms of temporal sensitivity in visual cortex neurons. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 30 (March (10)), 3652–3662., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2880457&tool=pmcentrez&rendertype=abstract.

Carmena, J.M., Lebedev, M.A., Henriquez, C.S., Nicolelis, M.A.L., 2005. Stable ensemble performance with single-neuron variability during reaching movements in primates. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 25 (November (46)), 10712–10716., http://www.ncbi.nlm.nih.gov/pubmed/16291944.

Carney, L.H., McDuffy, M.J., Shekhter, I., 1999. Frequency glides in the impulse responses of auditory-nerve fibers. The Journal of the Acoustical Society of America 105 (April (4)), 2384–2391., http://www.ncbi.nlm.nih.gov/pubmed/10212419.

Chacron, M., Lindner, B., Longtin, A., 2004. Noise shaping by interval correlations increases information transfer. Physical Review Letters 92 (February (8)), 1–4, http://dx.doi.org/10.1103/PhysRevLett.92.080601.

Chapin, J.K., Nicolelis, M.A., 1999. Principal component analysis of neuronal ensemble activity reveals multidimensional somatosensory representations. Journal of Neuroscience Methods 94 (December (1)), 121–140., http://www.ncbi.nlm.nih.gov/pubmed/10638820.

Churchland, A.K., Kiani, R., Chaudhuri, R., Wang, X.-J., Pouget, A., Shadlen, M.N., 2011. Variance as a signature of neural computations during decision making. Neuron 69 (February (4)), 818–831., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3066020&tool=pmcentrez&rendertype=abstract.

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I., Shenoy, K.V., 2010. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? Neuron 68 (November (3)), 387–400., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2991102&tool=pmcentrez&rendertype=abstract.

Churchland, M.M., Shenoy, K.V., 2007. Delay of movement caused by disruption of cortical preparatory activity. Journal of Neurophysiology 97 (January (1)), 348–359., http://www.ncbi.nlm.nih.gov/pubmed/17005608.

Churchland, M.M., Yu, B.M., Cunningham, J.P., Sugrue, L.P., Cohen, M.R., Corrado, G.S., Newsome, W.T., Clark, A.M., Hosseini, P., Scott, B.B., Bradley, D.C., Smith, M.A., Kohn, A., Movshon, J.A., Armstrong, K.M., Moore, T., Chang, S.W., Snyder, L.H., Lisberger, S.G., Priebe, N.J., Finn, I.M., Ferster, D., Ryu, S.I., Santhanam, G., Sahani, M., Shenoy, K.V., 2010. Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nature Neuroscience 13 (March (3)), 369–378., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2828350&tool=pmcentrez&rendertype=abstract.

Cohen, J., Haesler, S., Vong, L., Lowell, B., Uchida, N., 2012. Neuron-type-specific signals for reward and punishment in the ventral tegmental area. Nature 482, 85–88.

Cohen, M.R., Maunsell, J.H.R., 2011. Using neuronal populations to study the mechanisms underlying spatial and feature attention. Neuron 70 (June (6)), 1192–1204., http://www.ncbi.nlm.nih.gov/pubmed/21689604.

Cohen, M.R., Newsome, W.T., 2009. Estimates of the contribution of single neurons to perception depend on timescale and noise correlation. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 29 (May (20)), 6635–6648., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2705937&tool=pmcentrez&rendertype=abstract.

Compte, A., Brunel, N., Goldman-rakic, P.S., Wang, X.-J., 2000. Dynamics underlying spatial working memory in a cortical network model E. Cerebral Cortex 910–923.

Compte, A., Constantinidis, C., Tegner, J., Raghavachari, S., Chafee, M.V., Goldman-Rakic, P.S., Wang, X.-J., 2003. Temporally irregular mnemonic persistent activity in prefrontal neurons of monkeys during a delayed response task. Journal of

Neurophysiology 90 (November (5)), 3441–3454., http://www.ncbi.nlm.nih.gov/pubmed/12773500.

Constantinidis, C., Franowicz, M.N., Goldman-Rakic, P.S., 2001. Coding specificity in cortical microcircuits: a multiple-electrode analysis of primate prefrontal cortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 21 (May (10)), 3646–3655., http://www.ncbi.nlm.nih.gov/pubmed/11331394.

Constantinidis, C., Goldman-Rakic, P.S., 2002. Correlated discharges among putative pyramidal neurons and interneurons in the primate prefrontal cortex. Journal of Neurophysiology 88 (December (6)), 3487–3497., http://www.ncbi.nlm.nih.gov/pubmed/12466463.

Dacey, D.M., Petersen, M.R., 1992. Dendritic field size and morphology of midget and parasol ganglion cells of the human retina. Proceedings of the National Academy of Sciences of the United States of America 89 (October (20)), 9666–9670, http://dx.doi.org/10.1073/pnas.89.20.9666.

Daley, D., Vere-Jones, D., 2007. An Introduction to the Theory of Point Processes, vol. 1. Springer Verlag, New York, USA.

David, S.V., Vinje, W.E., Gallant, J.L., 2004. Natural stimulus statistics alter the receptive field structure of v1 neurons. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 24 (August (31)), 6991–7006., http://www.ncbi.nlm.nih.gov/pubmed/15295035.

de Kock, C.P.J., Bruno, R.M., Spors, H., Sakmann, B., 2007. Layer- and cell-type-specific suprathreshold stimulus representation in rat primary somatosensory cortex. Journal of Physiology 581, 139–154.

de la Rocha, J., Doiron, B., Shea-Brown, E., Josić, K., Reyes, A., 2007. Correlation between neural spike trains increases with firing rate. Nature 448 (August (7155)), 802–806., http://www.ncbi.nlm.nih.gov/pubmed/17700699.

de Lafuente, V., Romo, R., 2006. Neural correlate of subjective sensory experience gradually builds up across cortical areas. Proceedings of the National Academy of Sciences of the United States of America 103 (September (39)), 14266–14271., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1551903&tool=pmcentrez&rendertype=abstract.

DeAngelis, G., Ohzawa, I., Freeman, R., 1993. Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. I. General characteristics and postnatal development. Journal of Neurophysiology 69 (4), 1091–1117., http://jn.physiology.org/content/69/4/1091.short.

DeCharms, R.C., Zador, A.M., 2000. Neural representation and the cortical code. Annual Review of Neuroscience 23 (March (1)), 613–647.

DeWeese, M.R., Wehr, M., Zador, A.M., 2003. Binary spiking in auditory cortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 23 (August (21)), 7940–7949., http://www.ncbi.nlm.nih.gov/pubmed/12944525.

DeWeese, M.R., Zador, A.M., 2006. Non-Gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 26 (November (47)), 12206–12218., http://www.ncbi.nlm.nih.gov/pubmed/17122045.

DiCarlo, J.J., Johnson, K.O., Hsiao, S.S., 1998. Structure of receptive fields in area 3b of primary somatosensory cortex in the alert monkey. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 18 (April (7)), 2626–2645., http://www.ncbi.nlm.nih.gov/pubmed/9502821.

Dodd, J.V., Krug, K., Cumming, B.G., Parker, A.J., 2001. Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 21 (July (13)), 4809–4821., http://www.ncbi.nlm.nih.gov/pubmed/11425908.

Dow, B.M., Snyder, A.Z., Vautin, R.G., Bauer, R., 1981. Magnification factor and receptive field size in foveal striatal cortex of the monkey. Experimental Brain Research 44, 213–228.

Durstewitz, D., 2004. Neural representation of interval time. Neuroreport 15 (April (5)), 745–749., http://www.ncbi.nlm.nih.gov/pubmed/15073507.

Ecker, A.S., Berens, P., Keliris, G.A., Bethge, M., Logothetis, N.K., Tolias, A.S., 2010. Decorrelated neuronal firing in cortical microcircuits. Science 327 (January (5965)), 584–587., http://www.ncbi.nlm.nih.gov/pubmed/20110506.

Engel, T.A., Schimansky-Geier, L., Herz, A.V.M., Schreiber, S., Erchova, I., 2008. Subthreshold membrane-potential resonances shape spike-train patterns in the entorhinal cortex. Journal of Neurophysiology 100 (September (3)), 1576–1589., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2544463&tool=pmcentrez&rendertype=abstract.

Farkhooi, F., Strube-Bloss, M., Nawrot, M., 2009. Serial correlation in neural spike trains: experimental evidence, stochastic modeling, and single neuron variability. Physical Review E 79 (February (2)), 1–10, http://dx.doi.org/10.1103/PhysRevE.79.021905.

Feierstein, C.E., Quirk, M.C., Uchida, N., Sosulski, D.L., Mainen, Z.F., 2006. Representation of spatial goals in rat orbitofrontal cortex. Neuron 51 (August (4)), 495–507., http://www.ncbi.nlm.nih.gov/pubmed/16908414.

Fellous, J., Houweling, A., Modi, R., 2001. Frequency dependence of spike timing reliability in cortical pyramidal cells and interneurons. Journal of Neurophysiology 85, 1782–1787., http://jn.physiology.org/content/85/4/1782.short.

Fellous, J.-M., Tiesinga, P.H.E., Thomas, P.J., Sejnowski, T.J., 2004. Discovering spike patterns in neuronal responses. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 24 (March (12)), 2989–3001., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2928855&tool=pmcentrez&rendertype=abstract.

Field, D.J., 1994. What is the goal of sensory coding? Neural Computation 559–601.

Finn, I.M., Priebe, N.J., Ferster, D., 2007. The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. Neuron 54 (April (1)), 137–152., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1993919&tool=pmcentrez&rendertype=abstract.

Fiser, J., Berkes, P., Orbán, G., Lengyel, M., 2010. Statistically optimal perception and learning: from behavior to neural representations. Trends in Cognitive Sciences 14 (March (3)), 119–130., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2939867&tool=pmcentrez&rendertype=abstract.

Fitzgerald, P.J., Lane, J.W., Thakur, P.H., Hsiao, S.S., 2006. Receptive field (RF) properties of the macaque second somatosensory cortex: RF size. Shape, and Somatotopic Organization 26 (24), 6485–6495.

Gabbiani, F., Koch, C., 1998. Principles of spike train analysis. In: Koch, C., Segev, I. (Eds.), Methods of Neuronal Modelling: From Ions to Networks. MIT Press, Cambridge, MA, pp. 313–360.

Galan, R.F., Ermentrout, G.B., Urban, N.N., 2008. Optimal time scale for spike-time reliability: theory, simulations, and experiments. Journal of Neurophysiology 99, 277–283.

Ganguli, S., Sompolinsky, H., 2012. Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. Annual Review of Neuroscience 35, 485–508.

Gelman, A., Stern, H., 2006. The difference between "significant" and "not significant" is not itself statistically significant. The American Statistician 60, 328–331.

Georgopoulos, A., Lurito, J., Petrides, M., Schwartz, A., Massey, J., 1989. Mental rotation of the neuronal population vector. Science 243 (4888), 234., http://www.sciencemag.org/content/243/4888/234.short.

Georgopoulos, A., Schwartz, A., Kettner, R., 1986. Neuronal population coding of movement direction. Science 233 (September (4771)), 1416–1419, http://dx.doi.org/10.1126/science.3749885.

Georgopoulos, A.P., Kalaska, J.F., Caminiti, R., Massey, J.T., 1982. On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. Journal of Neuroscience 2, 1527–1537.

Goldberg, J.A., Boraud, T., Maraton, S., Haber, S.N., Vaadia, E., Bergman, H., 2002. Enhanced synchrony among primary motor cortex neurons in the 1-methyl-4-phenyl-1,b,c,6-tetrahydropyridine primate model of Parkinson's disease. Journal of Neuroscience 22, 4639–4653.

Gollisch, T., Meister, M., 2008. Rapid neural coding in the retina with relative spike latencies. Science 319 (5866), 1108–1111.

Gourevitch, B., Eggermont, J.J., 2007. A nonparametric approach for detection of bursts in spike trains. Journal of Neuroscience Methods 160, 349–358.

Graf, A.B.A., Kohn, A., Jazayeri, M., Movshon, J.A., 2011. Decoding the activity of neuronal populations in macaque primary visual cortex. Nature Neuroscience 14 (February (2)), 239–245., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3081541&tool=pmcentrez&rendertype=abstract.

Gutkin, B.S., Ermentrout, G.B., 1998. Dynamics of membrane excitability determine interspike interval variability: a link between spike generation mechanisms and cortical spike train statistics. Neural Computation 10, 1047–1065.

Haddad, R., Weiss, T., Khan, R., Nadler, B., Mandairon, N., Bensafi, M., Schneidman, E., Sobel, N., 2010. Global features of neural activity in the olfactory system form a parallel code that predicts olfactory behavior and perception. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 30 (July (27)), 9017–9026., http://www.ncbi.nlm.nih.gov/pubmed/20610736.

Haeusler, S., Schuch, K., Maass, W., 2009. Motif distribution, dynamical properties, and computational performance of two data-based cortical microcircuit templates. Journal of Physiology, Paris 103, 73–87.

Haider, B., Duque, A., Hasenstaub, A.R., McCormick, D.A., 2006. Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. Journal of Neuroscience 26, 4535–4545.

Haider, B., Krause, M., Duque, A., Yu, Y., Touryan, J., Mazer, A., Mccormick, D.A., 2010. Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. Neuron 65 (1), 107–121.

Hamaguchi, K., Riehle, A., Brunel, N., 2011. Estimating network parameters from combined dynamics of firing rate and irregularity of single neurons. Journal of Neurophysiology 105, 487–500.

Han, F., Caporale, N., Dan, Y., 2008. Reverberation of recent visual experience in spontaneous cortical waves. Neuron 60 (October (2)), 321–327., http://www.ncbi.nlm.nih.gov/pubmed/18957223.

Harris, K.D., Bartho, P., Chadderton, P., Curto, C., de la Rocha, J., Hollender, L., Itskov, V., Luczak, A., Marguet, S.L., Renart, A., Sakata, S., 2011. How do neurons work together? Lessons from auditory cortex. Hearing Research 271 (January (1–2)), 37–53., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2992581&tool=pmcentrez&rendertype=abstract.

Harris, K.D., Thiele, A., 2011. Cortical state and attention. Nature Reviews. Neuroscience 12 (September (9)), 509–523., http://www.ncbi.nlm.nih.gov/pubmed/21829219.

Hastie, T., Tibshirani, R., Friedman, J., 2009. The Elements of Statistical Learning. Springer, New York, USA.

Hatsopoulos, N.G., Xu, Q., Amit, Y., 2007. Encoding of movement fragments in the motor cortex. Journal of Neuroscience 27, 5105–5114.

Heil, P., 1997. Auditory cortical onset responses revisited. I. First-spike timing. Journal of Neurophysiology 77 (5), 2616–2641.

Hernández, A., Nácher, V., Luna, R., Zainos, A., Lemus, L., Alvarez, M., Vázquez, Y., Camarillo, L., Romo, R., 2010. Decoding a perceptual decision process across cortex. Neuron 66 (April (2)), 300–314., http://www.ncbi.nlm.nih.gov/pubmed/20435005.

Hernández, A., Zainos, A., Romo, R., 2000. Neuronal correlates of sensory discrimination in the somatosensory cortex. Proceedings of the National Academy of Sciences of the United States of America 97 (May (11)), 6191–6196., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=18580&tool=pmcentrez&rendertype=abstract.

Holt, G.R., Softky, W.R., Koch, C., Douglas, R.J., 1996. Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. Journal of Neurophysiology 75, 1806–1814.

Hromadka, T., DeWeese, M., Zador, A., 2008. Sparse representation of sounds in the unanesthetized auditory cortex. PLoS Biology 6 (1) http://dx.plos.org/10.1371/journal.pbio.0060016.

Hromadka, T., Deweese, M.R., Zador, A.M., 2008. Sparse representation of sounds in the unanesthetized auditory cortex. PLoS Biology 6, e16.

Hubel, D., Wiesel, T., 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. Journal of Physiology 160, 106–154.

Humphries, M.D., 2011. Spike-train communities: finding groups of similar spike trains. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 31 (February (6)), 2321–2336., http://www.ncbi.nlm.nih.gov/pubmed/21307268.

Hunter, J.D., Milton, J.G., Thomas, P.J., Cowan, J.D., 1998. Resonance effect for neural spike time reliability. Journal of Neurophysiology 80, 1427–1438.

Ichihara-Takeda, S., Funahashi, S., 2007. Activity of primate orbitofrontal and dorsolateral prefrontal neurons: task-related activity during an oculomotor delayed-response task. Experimental Brain Research 181, 409–425.

Izhikevich, E.M., Edelman, G.M., 2008. Large-scale model of mammalian thalamocortical systems. Proceedings of the National Academy of Sciences of the United States of America 105, 3593–3598.

Jin, D.Z., Fujii, N., Graybiel, A.M., 2009. Neural representation of time in cortico-basal ganglia circuits. Proceedings of the National Academy of Sciences of the United States of America 106 (November (45)), 19156–19161., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2776432&tool=pmcentrez&rendertype=abstract.

Kandel, E.R., Schwartz, J.H., Jessell, T.M., 2000. Principles of Neural Science. McGraw-Hill, New York, USA.

Kanerva, P., 1988. Sparse Distributed Memory. MIT Press, Cambridge, MA, USA.

Kara, P., Reinagel, P., Reid, R.C., 2000. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. Neuron 27 (September (3)), 635–646., http://www.ncbi.nlm.nih.gov/pubmed/11055444.

Kerr, J.N.D., Greenberg, D., Helmchen, F., 2005. Imaging input and output of neocortical networks in vivo. Proceedings of the National Academy of Sciences of the United States of America 102 (September (39)), 14063–14068., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1201343&tool=pmcentrez&rendertype=abstract.

Kiani, R., Shadlen, M.N., 2009. Representation of confidence associated with a decision by neurons in the parietal cortex. Science 324 (May (5928)), 759–764., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2738936&tool=pmcentrez&rendertype=abstract.

Ko, H., Hofer, S.B., Pichler, B., Buchanan, K.A., Sjöström, P.J., Mrsic-Flogel, T.D., 2011. Functional specificity of local synaptic connections in neocortical networks. Nature 473 (May (7345)), 87–91., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3089591&tool=pmcentrez&rendertype=abstract.

Koch, C., 1999. Biophysics of Computation. Oxford University Press, Oxford, UK.

Kohn, A., Smith, M.A., 2005. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 25 (April (14)), 3661–3673., http://www.ncbi.nlm.nih.gov/pubmed/15814797.

Kojima, S., Goldman-Rakic, P.S., 1982. Delay-related activity of prefrontal neurons in rhesus monkeys performing delayed response. Brain Research 248, 43–49., http://linkinghub.elsevier.com/retrieve/pii/0006899382911453.

Komura, Y., Tamura, R., Uwano, T., Nishijo, H., Kaga, K., Ono, T., 2001. Retrospective and prospective coding for predicted reward in the sensory thalamus. Nature 412 (August (6846)), 546–549., http://www.ncbi.nlm.nih.gov/pubmed/11484055.

Koulakov, A.A., Hromádka, T., Zador, A.M., 2009. Correlated connectivity and the distribution of firing rates in the neocortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 29 (March (12)), 3685–3694., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2784918&tool=pmcentrez&rendertype=abstract.

Latham, P.E., Nirenberg, S., 2005. Synergy, redundancy, and independence in population codes, revisited. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 25 (May (21)), 5195–5206., http://www.ncbi.nlm.nih.gov/pubmed/15917459.

Laubach, M., Wessberg, J., Nicolelis, M.A., 2000. Cortical ensemble activity increasingly predicts behaviour outcomes during learning of a motor task. Nature 405 (June (6786)), 567–571., http://www.ncbi.nlm.nih.gov/pubmed/10850715.

Laughlin, S.B., Sejnowski, T.J., 2003. Communication in neuronal networks. Science 301, 1870–1874.

Laurent, 2002. Olfactory network dynamics and the coding of multidimensional signals. Nature Reviews. Neuroscience 3 (November (11)), 884–895., http://www.ncbi.nlm.nih.gov/pubmed/12415296.

Lee, D., Port, N.L., Kruse, W., Georgopoulos, A.P., 1998. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 18 (February (3)), 1161–1170., http://www.ncbi.nlm.nih.gov/pubmed/9437036.

Lennie, P., 2003. The cost of cortical computation. Current Biology 13, 493–497.

Lewicki, M.S., 2002. Efficient coding of natural sounds. Nature Neuroscience 5 (April (4)), 356–363., http://www.ncbi.nlm.nih.gov/pubmed/11896400.

Linden, J.F., Liu, R.C., Sahani, M., Schreiner, C.E., Merzenich, M.M., 2003. Spectrotemporal structure of receptive fields in areas AI and AAF of mouse auditory cortex. Journal of Neurophysiology 90 (4), 2660–2675.

Lindner, B., Chacron, M.J., Longtin, A., 2005. Integrate-and-fire neurons with threshold noise: a tractable model of how interspike interval correlations affect neuronal signal transmission. Physical Review E 72, 021911.

Lochmann, T., Ernst, U.A., Denève, S., 2012. Perceptual inference predicts contextual modulations of sensory responses. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 32 (March (12)), 4179–4195., http://www.ncbi.nlm.nih.gov/pubmed/22442081.

London, M., Roth, A., Beeren, L., Häusser, M., Latham, P.E., 2010. Sensitivity to perturbations in vivo implies high noise and suggests rate coding in cortex. Nature 466, 123–127.

Lopes-Dos-Santos, V., Conde-Ocazionez, S., Nicolelis, M.A.L., Ribeiro, S.T., Tort, A.B.L., 2011. Neuronal assembly detection and cell membership specification by principal component analysis. PloS One 6 (January (6)), e20996., http://www.ncbi.nlm.nih.gov/pubmed/21698248.

Luczak, A., Barthó, P., Harris, K.D., 2009. Spontaneous events outline the realm of possible sensory responses in neocortical populations. Neuron 62 (May (3)), 413–425., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2696272&tool=pmcentrez&rendertype=abstract.

Luczak, A., Bartho, P., Marguet, S.L., Buzsáki, G., Harris, K.D., 2007. Sequential structure of neocortical spontaneous activity in vivo. Proceedings of the National Academy of Sciences of the United States of America 104, 347–352.

Luna, R., Hernández, A., Brody, C.D., Romo, R., 2005. Neural codes for perceptual discrimination in primary somatosensory cortex. Nature Neuroscience 8 (September (9)), 1210–1219., http://www.ncbi.nlm.nih.gov/pubmed/16056223.

Ma, W.J., Beck, J.M., Latham, P.E., Pouget, A., 2006. Bayesian inference with probabilistic population codes. Nature Neuroscience 9 (November (11)), 1432–1438 http://www.ncbi.nlm.nih.gov/pubmed/17057707 http://www.ncbi.nlm.nih.gov/pubmed/15254513.

Machens, C.K., Romo, R., Brody, C.D., 2010. Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 30 (January (1)), 350–360., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2947945&tool=pmcentrez&rendertype=abstract.

Machens, C.K., Wehr, M.S., Zador, A.M., 2004. Linearity of cortical receptive fields measured with natural sounds. Journal of Neuroscience 24, 1089–1100.

Macke, J.H., Berens, P., Ecker, A.S., Tolias, A.S., Bethge, M., 2009. Generating spike trains with specified correlation coefficients. Neural Computation 21 (February (2)), 397–423., http://www.ncbi.nlm.nih.gov/pubmed/19196233.

MacLean, J.N., Watson, B.O., Aaron, G.B., Yuste, R., 2005. Internal dynamics determine the cortical response to thalamic stimulation. Neuron 48 (December (5)), 811–823., http://www.ncbi.nlm.nih.gov/pubmed/16337918.

Maimon, G., Assad, J.A., 2006. A cognitive signal for the proactive timing of action in macaque LIP. Nature Neuroscience 9 (July (7)), 948–955., http://www.ncbi.nlm.nih.gov/pubmed/16751764.

Maimon, G., Assad, J.A., 2009. Beyond Poisson: increased spike-time regularity across primate parietal cortex. Neuron 62 (May (3)), 426–440., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2743683&tool=pmcentrez&rendertype=abstract.

Mainen, Z.F., Sejnowski, T.J., 1995. Reliability of spike timing in neocortical neurons. Science 268, 1503–1506.

Marr, D., 1969. A theory of cerebellar cortex. Journal of Physiology 202 (2), 437.

Matell, M.S., Shea-Brown, E., Gooch, C., Wilson, A.G., Rinzel, J., 2011. A heterogeneous population code for elapsed time in rat medial agranular cortex. Behavioral Neuroscience 125 (February (1)), 54–73., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3078766&tool=pmcentrez&rendertype=abstract.

Matsumoto, N., Okada, M., Sugase-Miyamoto, Y., Yamane, S., Kawano, K., 2005. Population dynamics of face-responsive neurons in the inferior temporal cortex. Cerebral Cortex 15 (August (8)), 1103–1112., http://www.ncbi.nlm.nih.gov/pubmed/15563724.

Maunsell, J., Gibson, J., 1992. Visual response latencies in striate cortex of the macaque monkey. Journal of Neurophysiology 68 (4), 1332–1344.

Mazor, O., Laurent, G., 2005. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. Neuron 48 (November (4)), 661–673., http://www.ncbi.nlm.nih.gov/pubmed/16301181.

Mazurek, M.E., Shadlen, M.N., 2002. Limits to the temporal fidelity of cortical spike rate signals. Nature Neuroscience 5, 463–471.

McClurkin, J.W., Optican, L.M., Richmond, B.J., Gawne, T.J., 1991. Concurrent processing and complexity of temporally encoded neuronal messages in visual perception. Science 253 (5020), 675–677.

Mita, A., Mushiake, H., Shima, K., Matsuzaka, Y., Tanji, J., 2009. Interval time coding by neurons in the presupplementary and supplementary motor areas. Brain Research 12 (4), 502–507.

Mitchell, J.F., Sundberg, K.A., Reynolds, J.H., 2009. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. Neuron 63 (September (6)), 879–888., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2765230&tool=pmcentrez&rendertype=abstract.

Mongillo, G., Amit, D., Brunel, N., 2003. Retrospective and prospective persistent activity induced by Hebbian learning in a recurrent cortical network. European Journal of Neuroscience 18 (7), 2011–2024, http://dx.doi.org/10.1046/j.1460-9568.2003.02908.x/full.

Montani, F., Kohn, A., Smith, M.A., Schultz, S.R., 2007. The role of correlations in direction and contrast coding in the primary visual cortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 27 (February (9)), 2338–2348., http://www.ncbi.nlm.nih.gov/pubmed/17329431.

Morgan, M.L., Deangelis, G.C., Angelaki, D.E., 2008. Multisensory integration in macaque visual cortex depends on cue reliability. Neuron 59 (August (4)),

662–673., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2601653&tool=pmcentrez&rendertype=abstract.

Moshitch, D., Las, L., Ulanovsky, N., Bar-Yosef, O., Nelken, I., 2006. Responses of neurons in primary auditory cortex (A1) to pure tones in the halothane-anesthetized cat. Journal of Neurophysiology 95 (June (6)), 3756–3769., http://www.ncbi.nlm.nih.gov/pubmed/16554513.

Narayanan, N., Laubach, M., 2009. Delay activity in rodent frontal cortex during a simple reaction time task. Journal of Neurophysiology 101, 2859–2871., http://jn.physiology.org/cgi/content/abstract/101/6/2859.

Narayanan, N.S., Kimchi, E.Y., Laubach, M., 2005. Redundancy and synergy of neuronal ensembles in motor cortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 25 (April (17)), 4207–4216., http://www.ncbi.nlm.nih.gov/pubmed/15858046.

Nawrot, M.P., Boucsein, C., Rodriguez-Molinad, V., Aertsen, A., Grun, S., Rotter, S., 2007. Serial interval statistics of spontaneous activity in cortical neurons in vivo and in vitro. Neurocomputing 70, 1717–1722.

Nienborg, H., Cumming, B.G., 2006. Macaque V2 neurons, but not V1 neurons, show choice-related activity. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 26 (September (37)), 9567–9578., http://www.ncbi.nlm.nih.gov/pubmed/16971541.

Nienborg, H., Cumming, B.G., 2009. Decision-related activity in sensory neurons reflects more than a neuron's causal effect. Nature 459 (May (7243)), 89–92., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2917918&tool=pmcentrez&rendertype=abstract.

Nieuwenhuis, S., Forstmann, B.U., Wagenmakers, E.-J., 2011. Erroneous analyses of interactions in neuroscience: a problem of significance. Nature Neuroscience 14, 1105–1107.

Nirenberg, S., Carcieri, S.M., Jacobs, A.L., Latham, P.E., 2001. Retinal ganglion cells act largely as independent encoders. Nature 411 (June (6838)), 698–701., http://www.ncbi.nlm.nih.gov/pubmed/11395773.

O'Connor, D.H., Peron, S.P., Huber, D., Svoboda, K., Connor, D.H.O., 2010. Neural activity in barrel cortex underlying vibrissa-based object localization in mice. Neuron 67 (September (6)), 1048–1061., http://www.ncbi.nlm.nih.gov/pubmed/20869600.

Ohiorhenuan, I.E., Mechler, F., Purpura, K.P., Schmid, A.M., Hu, Q., Victor, J.D., 2010. Sparse coding and high-order correlations in fine-scale cortical networks. Nature 466 (July (7306)), 617–621., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2912961&tool=pmcentrez&rendertype=abstract.

Ohki, K., Chung, S., Ch'ng, Y.H., Kara, P., Reid, R.C., 2005. Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. Nature 433 (February (7026)), 597–603., http://www.ncbi.nlm.nih.gov/pubmed/15660108.

Olshausen, B., Field, D., 1996. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature.

Olshausen, B., Field, D., 1997. Sparse coding with an overcomplete basis set: a strategy employed by V1? Vision Research 37 (23), 3311–3325., http://www.sciencedirect.com/science/article/pii/s0042698997001697.

Olshausen, B.A., Field, D.J., 2004. Sparse coding of sensory inputs. Current Opinion in Neurobiology 14 (August (4)), 481–487., http://www.ncbi.nlm.nih.gov/pubmed/15321069.

Olshausen, B.A., Field, D.J., 2005. What is the other 85% of V1 doing? In: van Hemmen, J.L., Sejnowski, T.J. (Eds.), 23 Problems in Systems Neuroscience. OUP, Oxford.

Optican, L.M., Richmond, B.J., 1987. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III. Information theoretic analysis. Journal of Neurophysiology 57 (January (1)), 162–178., http://www.ncbi.nlm.nih.gov/pubmed/3559670.

Osborne, L.C., Palmer, S.E., Lisberger, S.G., Bialek, W., 2008. The neural basis for combinatorial coding in a cortical population response. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 28 (December (50)), 13522–13531 http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2693376&tool=pmcentrez&rendertype=abstract.

Otazu, G.H., Tai, L.-H., Yang, Y., Zador, A.M., 2009. Engaging in an auditory task suppresses responses in auditory cortex. Nature Neuroscience 12 (May (5)), 646–654., http://www.ncbi.nlm.nih.gov/pubmed/19363491.

Paninski, L., 2004. Maximum likelihood estimation of cascade point-process neural encoding models. Network: Computation in Neural Systems 15 (November (4)), 243–262 http://www.informaworld.com/openurl?genre=article&doi=10.1088/0954-898X/15/4/002&magic=crossref||D404A21C5BB053405B1A640AFF-D44AE3.

Parker, A.J., Hawken, M.J., 1988. Two-dimensional spatial structure of receptive fields in monkey striate cortex. Journal of the Optical Society of America A 5, 598–605.

Perkel, D.H., Gerstein, G.L., Moore, G.P., 1967. Neuronal spike trains and stochastic point processes. I. The single spike train. Biophysical Journal 7, 391–418.

Petersen, R.S., Panzeri, S., Diamond, M.E., 2001. Population coding of stimulus location in rat somatosensory cortex. Neuron 32 (November (3)), 503–514., http://www.ncbi.nlm.nih.gov/pubmed/12459298.

Peyrache, A., Benchenane, K., Khamassi, M., Wiener, S.I., Battaglia, F.P., 2010. Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution. Journal of Computational Neuroscience 29 (August (1–2)), 309–325., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2940043&tool=pmcentrez&rendertype=abstract.

Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S.I., Battaglia, F.P., 2009. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. Nature Neuroscience 12 (July (7)), 919–926., http://www.ncbi.nlm.nih.gov/pubmed/19483687.

Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., Simoncelli, E.P., 2008. Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature 454 (August (7207)), 995–999., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2684455&tool=pmcentrez&rendertype=abstract.

Ponce-Alvarez, A., Kilavik, B.E., Riehle, A., 2010. Comparison of local measures of spike time irregularity and relating variability to firing rate in motor cortical neurons. Journal of Computational Neuroscience 29, 351–365.

Pouget, A., Dayan, P., Zemel, R., 2000. Information processing with population codes. Nature Reviews. Neuroscience 1 (November (2)), 125–132., http://www.ncbi.nlm.nih.gov/pubmed/11252775.

Pouget, A., Dayan, P., Zemel, R.S., 2003. Inference and computation with population codes. Annual Review of Neuroscience 26 (January), 381–410., http://www.ncbi.nlm.nih.gov/pubmed/12704222.

Price, N.S.C., Born, R.T., 2010. Timescales of sensory- and decision-related activity in the middle temporal and medial superior temporal areas. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 30 (October (42)), 14036–14045., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3059109&tool=pmcentrez&rendertype=abstract.

Rehn, M., Sommer, F.T., 2007. A network that uses few active neurones to code visual input predicts the diverse shapes of cortical receptive fields. Journal of Computational Neuroscience 22 (April (2)), 135–146., http://www.ncbi.nlm.nih.gov/pubmed/17053994.

Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., Harris, K.D., 2010. The asynchronous state in cortical circuits. Science 327 (January (5965)), 587–590., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2861483&tool=pmcentrez&rendertype=abstract.

Rickert, J., Riehle, A., Aertsen, A., Rotter, S., Nawrot, M.P., 2009. Dynamic encoding of movement direction in motor cortical neurons. Journal of Neuroscience 29, 13870–13882.

Rieke, F., Warland, D., de Ruyter van Stevninck, R., Bialek, W., 1997. Spikes: Exploring the Neural Code. MIT Press, Cambridge, MA.

Ringach, D.L., 2002. Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. Journal of Neurophysiology 88 (July (1)), 455–463., http://www.ncbi.nlm.nih.gov/pubmed/12091567.

Ringach, D.L., 2009. Spontaneous and driven cortical activity: implications for computation. Current Opinion in Neurobiology 19 (August (4)), 439–444 http://www.ncbi.nlm.nih.gov/pubmed/19647992 http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3319344&tool=pmcentrez&rendertype=abstract.

Ringach, D.L., Shapley, R.M., Hawken, M.J., 2002. Orientation selectivity in macaque V1: diversity and laminar dependence. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 22 (July (13)), 5639–5651., http://www.ncbi.nlm.nih.gov/pubmed/12097515.

Roelfsema, P.R., Lamme, V.A.F., Spekreijse, H., 2004. Synchrony and covariation of firing rates in the primary visual cortex during contour grouping. Nature Neuroscience 7 (September (9)), 982–991., http://www.ncbi.nlm.nih.gov/pubmed/15322549.

Romo, R., Hernández, A., Zainos, A., Lemus, L., Brody, C.D., 2002. Neuronal correlates of decision-making in secondary somatosensory cortex. Nature Neuroscience 5 (November (11)), 1217–1225., http://www.ncbi.nlm.nih.gov/pubmed/12368806.

Romo, R., Hernández, A., Zainos, A., Salinas, E., 2003. Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. Neuron 38 (May (4)), 649–657., http://www.ncbi.nlm.nih.gov/pubmed/12765615.

Rothschild, G., Nelken, I., Mizrahi, A., 2010. Functional organization and population dynamics in the mouse primary auditory cortex. Nature Neuroscience 13 (March (3)), 353–360., http://www.ncbi.nlm.nih.gov/pubmed/20118927.

Roxin, A., Brunel, N., Hansel, D., Mongillo, G., van Vreeswijk, C., 2011. On the distribution of firing rates in networks of cortical neurons. Journal of Neuroscience 31 (November (45)), 16217–16226, http://dx.doi.org/10.1523/jneurosci.1677-11.2011.

Rozell, C.J., Johnson, D.H., Baraniuk, R.G., Olshausen, B.A., 2008. Sparse coding via thresholding and local competition in neural circuits. Neural Computation 20 (October (10)), 2526–2563., http://www.ncbi.nlm.nih.gov/pubmed/18439138.

Sakata, S., Harris, K.D., 2009. Laminar structure of spontaneous and sensory-evoked population activity in auditory cortex. Neuron 64 (November (3)), 404–418., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2778614&tool=pmcentrez&rendertype=abstract.

Salinas, E., Hernandez, A., Zainos, A., Romo, R., 2000. Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 20 (July (14)), 5503–5515., http://www.ncbi.nlm.nih.gov/pubmed/10884334.

Salinas, E., Sejnowski, T.J., 2000. Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. Journal of Neuroscience 20, 6193–6209.

Sato, T., Uchida, G., Tanifuji, M., 2009. Cortical columnar organization is reconsidered in inferior temporal cortex. Cerebral Cortex 19 (August (8)), 1870–1888., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2705700&tool=pmcentrez&rendertype=abstract.

Schiller, P.H., Finlay, B.L., Volman, S.F., 1976. Quantitative studies of single-cell properties in monkey striate cortex. I. Spatiotemporal organization of receptive fields. Journal of Neurophysiology 39 (November (6)), 1288–1319., http://www.ncbi.nlm.nih.gov/pubmed/825621.

Schmolesky, M., Wang, Y., Hanes, D., Thompson, K., Leutgeb, S., Schall, J., Leventhal, A., 1998. Signal timing across the macaque visual system. Journal of Neurophysiology 79 (6), 3272–3278.

Schneidman, E., Berry, M.J., Segev, R., Bialek, W., 2006. Weak pairwise correlations imply strongly correlated network states in a neural population. Nature 440 (April (7087)), 1007–1012., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=1785327&tool=pmcentrez&rendertype=abstract.

Schreiner, C.E., Read, H.L., Sutter, M.L., 2000. Modular organization of frequency integration in primary auditory cortex. Annual Review of Neuroscience 501–529.

Sergio, L.E., Hamel-Paquet, C., Kalaska, J.F., 2005. Motor cortex neural correlates of output kinematics and kinetics during isometric-force and arm-reaching tasks. Journal of Neurophysiology 94, 2353–2378.

Shadlen, M.N., Britten, K.H., Newsome, W.T., Movshon, A.J., 1996. Behavioral analysis of the relationship responses to visual motion neuronal. Journal of Neuroscience 76 (4), 1486–1510.

Shadlen, M.N., Newsome, W.T., 1998. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. Journal of Neuroscience 18, 3870–3896.

Shafi, M., Zhou, Y., Quintana, J., Chow, C., Fuster, J., Bodner, M., 2007. Variability in neuronal activity in primate cortex during working memory tasks. Neuroscience 146 (May (3)), 1082–1108., http://www.ncbi.nlm.nih.gov/pubmed/17418956.

Shamir, M., Sompolinsky, H., 2006. Implications of neuronal diversity on population coding. Neural Computation 18 (August (8)), 1951–1986., http://www.ncbi.nlm.nih.gov/pubmed/16771659.

Shinomoto, S., Kim, H., Shimokawa, T., Matsuno, N., Funahashi, S., Shima, K., Fujita, I., Tamura, H., Doi, T., Kawano, K., Inaba, N., Fukushima, K., Kurkin, S., Kurata, K., Taira, M., Tsutsui, K.-I., Komatsu, H., Ogawa, T., Koida, K., Tanji, J., Toyama, K., 2009. Relating neuronal firing patterns to functional differentiation of cerebral cortex. PLoS Computational Biology 5 (July (7)), e1000433., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2701610&tool=pmcentrez&rendertype=abstract.

Shinomoto, S., Omi, T., Mita, A., Mushiake, H., Shima, K., Matsuzaka, Y., Tanji, J., 2011. Deciphering elapsed time and predicting action timing from neuronal population signals. Frontiers in Computational Neuroscience 5 (June), 29., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3122070&tool=pmcentrez&rendertype=abstract.

Shlens, J., Field, G.D., Gauthier, J.L., Grivich, M.I., Petrusca, D., Sher, A., Litke, A.M., Chichilnisky, E.J., 2006. The structure of multi-neuron firing patterns in primate retina. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 26 (August (32)), 8254–8266., http://www.ncbi.nlm.nih.gov/pubmed/16899720.

Shoham, S., O'Connor, D.H., Segev, R., 2006. How silent is the brain: is there a "dark matter" problem in neuroscience. Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology 192, 777–784.

Shu, Y., Hasenstaub, A., McCormick, D.A., 2003. Turning on and off recurrent balanced cortical activity. Nature 423, 288–293.

Smith, E.C., Lewicki, M.S., 2006. Efficient auditory coding. Nature 439 (February (7079)), 978–982., http://www.ncbi.nlm.nih.gov/pubmed/16495999.

Smith, M.A., Kohn, A., 2008. Spatial and temporal scales of neuronal correlation in primary visual cortex. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 28 (November (48)), 12591–12603., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2656500&tool=pmcentrez&rendertype=abstract.

Softky, W.R., Koch, C., 1993. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. Journal of Neuroscience 13, 334–350.

Sompolinsky, H., Yoon, H., Kang, K., Shamir, M., 2001. Population coding in neuronal systems with correlated noise. Physical Review E 64 (October (5)), 1–11, http://dx.doi.org/10.1103/PhysRevE.64.051904.

Stanford, T.R., Shankar, S., Massoglia, D.P., Costello, M.G., Salinas, E., 2010. Perceptual decision making in less than 30 milliseconds. Nature Neuroscience 13 (March (3)), 379–385., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2834559&tool=pmcentrez&rendertype=abstract.

Stevens, C.F., Zador, A.M., 1998. Input synchrony and the irregular firing of cortical neurons. Nature Neuroscience 1, 210–217.

Stopfer, M., Jayaraman, V., Laurent, G., 2003. Intensity versus identity coding in an olfactory system. Neuron 39 (September (6)), 991–1004., http://www.ncbi.nlm.nih.gov/pubmed/12971898.

Sur, M., Merzenich, M.M., Kaas, J.H., 1980. Magnification, receptive-field area, and "hypercolumn" size in areas 3b and 1 of somatosensory cortex in owl monkeys.

Journal of Neurophysiology 44 (August (2)), 295–311., http://www.ncbi.nlm.nih.gov/pubmed/7411189.

Sutter, M.L., Schreiner, C.E., 1991. Physiology and topography of neurons with multipeaked tuning curves in cat primary auditory cortex. Journal of Neurophysiology 65 (May (5)), 1207–1226., http://www.ncbi.nlm.nih.gov/pubmed/1869913.

Takeda, K., Funahashi, S., 2004. Population vector analysis of primate prefrontal activity during spatial working memory. Cerebral Cortex 14 (December (12)), 1328–1339., http://www.ncbi.nlm.nih.gov/pubmed/15166104.

Talbot, W., Darian-Smith, I., Kornhuber, H., Mountcastle, V., 1968. The sense of flutter-vibration: comparison of the human capacity with response patterns of mechanoreceptive afferents from the monkey hand. Journal of Neurophysiology 31 (2).

Troyer, T.W., Miller, K.D., 1997. Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. Neural Computation 9, 971–983.

Tsodyks, M., Kenet, T., Grinvald, A., Arieli, A., 1999. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. Science 286 (5446), 1943., http://www.sciencemag.org/content/286/5446/1943.short.

Tsodyks, M.V., Sejnowski, T., 1995. Rapid state switching in balanced cortical network models. Network 6, 111–124.

Uka, T., DeAngelis, G.C., 2003. Contribution of middle temporal area to coarse depth discrimination: comparison of neuronal and psychophysical sensitivity. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 23 (April (8)), 3515–3530., http://www.ncbi.nlm.nih.gov/pubmed/12716961.

Uka, T., DeAngelis, G.C., 2004. Contribution of area MT to stereoscopic depth perception: choice-related response modulations reflect task strategy. Neuron 42 (April (2)), 297–310., http://www.ncbi.nlm.nih.gov/pubmed/15091344.

Valois, R.L.D., Albrecht, D.G., Thorell, L.G., 1982. Spatial frequency selectivity of cells in macaque visual cortex. Vision Research 22, 545–559.

van Vreeswijk, C., Sompolinsky, H., 1996. Chaos in neuronal networks with balanced excitatory and inhibitory activity. Science 274, 1724–1726.

VanRullen, R., Thorpe, S., 2001. The time course of visual processing: from early perception to decision-making. Journal of Cognitive Neuroscience 13 (4), 454–461.

Vinje, W.E., Gallant, J.L., 2002. Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1. The Journal of Neuroscience: The Official Journal of the Society for Neuroscience 22 (April (7)), 2904–2915., http://www.ncbi.nlm.nih.gov/pubmed/11923455.

Wagenmakers, E.-J., Farrell, S., 2004. AIC model selection using Akaike weights. Psychonomic Bulletin & Review 11, 192–196.

Wasserman, L., 2004. All of Statistics: A Concise Course in Statistical Inference. Springer, New York.

Weliky, M., Fiser, J., Hunt, R.H., Wagner, D.N., 2003. Coding of natural scenes in primary visual cortex. Neuron 37 (February (4)), 703–718., http://www.ncbi.nlm.nih.gov/pubmed/12597866.

Willmore, B., Tolhurst, D.J., 2001. Characterizing the sparseness of neural codes. Network (Bristol, England) 12 (August (3)), 255–270., http://www.ncbi.nlm.nih.gov/pubmed/11563529.

Wohrer, A., Romo, R., Machens, C., 2010. Linear readout from a neural population with partial correlation data. In: Advances in Neural Information Processing Systems (NIPS). No. 1. pp. 2–10.

Wolfe, J., Houweling, A.R., Brecht, M., 2010. Sparse and powerful cortical spikes. Current Opinion in Neurobiology 20 (3), 306–312, http://dx.doi.org/10.1016/j.conb.2010.03.006.

Yang, T., Shadlen, M.N., 2007. Probabilistic reasoning by neurons. Nature 447 (June (7148)), 1075–1080., http://www.ncbi.nlm.nih.gov/pubmed/17546027.

Yu, B.M., Cunningham, J.P., Santhanam, G., Ryu, S.I., Shenoy, K.V., Sahani, M., 2009. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. Journal of Neurophysiology 102 (July (1)), 614–635., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2712272&tool=pmcentrez&rendertype=abstract.

Yu, S., Huang, D., Singer, W., Nikolic, D., 2008. A small world of neuronal synchrony. Cerebral Cortex 18 (December (12)), 2891–2901., http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2583154&tool=pmcentrez&rendertype=abstract.

Yu, S., Yang, H., Nakahara, H., Plenz, D., Santos, G.S., Nikolic, D., 2011. Higher-order interactions characterized in cortical activity. Journal of Neuroscience 31 (48), 17514–17526.