

AUTHOR QUERY FORM

Dear Author,

During the preparation of your manuscript for publication, the questions listedbelow have arisen. Please attend to these matters and return this form with yourproof.

Many thanks for your assistance.

Query	References Query	Remarks
1	AU: Please check highlighted text carefully; there seem to have been problems with file translation here.	

# 29 Noise, Correlations, and Information in Neural Population Codes

ELAD SCHNEIDMAN

**ABSTRACT** The brain relies on the spiking patterns of large groups of neurons to represent information. Uncovering the design principles of this “neural code” is fundamental to our understanding of information processing and computation in the brain. This task is made difficult by the fact that neural responses to their stimuli are highly selective and noisy, and because population activity patterns are often correlated and the number of potential population patterns is exponential in the size of the population. We present here the fundamental questions regarding the design and structure of probabilistic and correlated neural population codes and their implication for encoding and decoding. We show that correlations are neither good nor bad, but that the combination of stimulus-induced correlations and joint fluctuations of the cells balance the capacity of the code and its error-correction properties. We then show that the typically weak signal and noise correlations observed at the level of cell pairs can add up to give strong signal and noise correlations in large populations. Finally, we suggest that population codes may be designed to adapt to stimulus statistics in a way that would give collective coding that balances the information content and the code’s readability or learnability.

Information is represented in the brain by the action potentials of large groups of neurons. These basic symbols form temporal sequences of multi-neuron spiking and silences that are the universal language that our brain uses. Extending the language metaphor, to decipher this code we need to understand how this alphabet is used to form the vocabulary, what the grammar is, and, ultimately, the content of this neural code.

Much of our understanding of how information is “encoded” in the brain relies on studies of single neurons. As most of what we care about is the result of the joint activity of many neurons working together, the nature of the code of populations is central to neuroscience and our understanding of the brain. Experimentally, it has been hard to record the joint activity of many neurons over long periods of time. Since this task has become more feasible in recent years using multi-electrode arrays or imaging, we will most likely have in the not-too-distant future recordings of very large populations and even recordings from full large neural circuits. It is not clear, however, how we should interpret the code of such a large number of neurons.

Representation of information in the brain is necessarily noisy because of the stochastic nature of neurons and synapses. Overcoming this noise requires some form of error

correction, but our understanding of how the brain might do this is, at best, limited. While different brain areas show distinct spiking patterns and are coding stimuli of different modalities, universal features of the code include population spiking patterns that are sparse, noisy, correlated, and adaptive (Rieke, Warland, de Ruyter van Steveninck, & Bialek, 1997). We present here some of the fundamental design principles of the code of populations of neurons and, in particular, how noise and correlations are shaping it.

We suggest that neural noise means that we must rely on probability distributions in describing and analyzing the code, and that the brain must do so as well. We argue that the combinatorial number of potential activity patterns and correlations implies that understanding the code of large populations must rely on identifying simplifying principles of the design of neural codes. We present an information theory-based framework for assessing the role and impact of correlations and noise between cells and how it shapes the nature of the code. We review the dominant theoretical idea of redundancy reduction as a design principle of neural population codes and its limitations, and suggest instead that neural population codes rely on correlated collective coding that balances the amount of conveyed information with the readability or “learnability” of the code.

## *The neural code*

The universality of spikes as the basic symbols of neuronal activity has made it common to talk about the “neural code” (Rieke et al., 1997). But although the “letters” of the code are universal in the brain (i.e., spikes are the key carrier of information), the organization of these letters into words and their meaning differs between brain areas, between conspecifics, and presumably between species. Moreover, the mapping between stimuli and neural response is not stable even for the same cells over time due to adaptation and learning.

At the level of single cells, the study of the neural code has focused on what kind of information is carried by neurons of different types and in different brain areas, and how this information is encoded. In particular, a central question has been whether the main carrier of information is the

time-dependent spiking rates of neurons (rate coding), and how much is carried by the precise timing of the spikes (temporal coding) or the relative timing between spikes within the spiking patterns (Gerstner, Kreiter, Markram, & Herz, 1997; Perkel & Bullock, 1968; Rieke et al., 1997). At the level of populations, coding has often been explored in terms of averaging over cells by combining their activity linearly, but also in terms of the magnitude and structure of correlations between cells (Ahissar et al., 1992; Georgopoulos, Schwartz, & Kettner, 1986; Gray, König, Engel, & Singer, 1989; Hatsopoulos, Ojakangas, Paninski, & Donoghue, 1998; Martignon et al., 2000; Meister, Lagnado, & Baylor, 1995; Panzeri, Schultz, Treves, & Rolls, 1999; Pillow et al., 2008; Rieke et al., 1997; Schneidman, Berry, Segev, & Bialek, 2006; Stopfer, Bhagavan, Smith, & Laurent, 1997; Vaadia et al., 1995).

It seems almost obvious to expect that the structure of the code of retinal neurons would be different than that of the olfactory bulb, for example, due to differences in the stimulus structure, the time scales of stimulus changes, and the nature of spatiotemporal correlations in the stimulus. Similarly, the spiking patterns of neurons in the hippocampus and those in the visual cortex would probably have different organizational principles, due to the nature of information that is represented, time scales, neural noise, communication with other areas, and so on. Thus, it might be more appropriate to replace the notion of “the neural code” with a set of neural codes that are used in the brain. Moreover, even within the same neural circuit, the structure of the code may change when the stimulus statistics change. Despite these individual differences, many aspects of the code are universal, and must be related. Sensory systems are a natural place to explore the nature of neural codes since they are close to the stimulus, and we therefore have a sense of their primary role as conveying information about the outside world or of what is being encoded.

### *Neurons respond to stimuli with sparse and noisy spiking patterns*

Almost everywhere in the brain, neurons are silent most of the time (e.g., Shoham, O'Connor, & Segev, 2006), and respond with spikes only to a small fraction of the stimuli they are presented with. Because the mapping from stimuli to spikes at the level of a single neuron is so selective, it has been common (following Hartline, 1938) to characterize the properties of neurons in terms of the stimulus feature that a cell responds to, or “encodes.” This classic notion of the “receptive field”—which describes the part of stimulus space or the properties of the stimuli that a neuron responds to—is intuitive and simple to describe, but is often a little too simplistic. The inferred mapping between stimuli and responses of a neuron depends critically on the set of the stimuli used

to assess this mapping. Analysis of neural responses to simple artificial stimuli often suggests an easy way to describe the feature selectivity of a neuron. Yet the response of the same neuron to other kinds of stimuli, and in particular rich naturalistic stimuli, may differ considerably than the picture one might form based on the response to the artificial stimuli (David, Vinje, & Gallant, 2004). In particular, even stereotyped simple functions that one might infer from “standard” stimuli may have a qualitatively different feature selectivity when stimuli that are considered irrelevant are used (e.g., Geffen, de Vries, & Meister, 2007).

The common interpretation of feature selectivity of neurons is missing another aspect of neuronal encoding. The responses of neurons, even to their “preferred” stimuli, are not deterministic or reliable. Somewhat surprisingly, this “noise” in neural function and response is significant in almost every part of the nervous system, from synaptic unreliability to neural spiking, network activity, and behavior (Calvin & Stevens, 1968; De Ruyter van Steveninck, Lewen, Strong, Koberle, & Bialek, 1997; Harris & Wolpert, 1998; Körding & Wolpert, 2004; Reich, Victor, Knight, Ozaki, & Kaplan, 1997; Stevens & Zador, 1998; Warzecha & Egelhaaf, 1999). In particular, the variability of neural responses to repeated presentations of the same stimulus is apparent at different aspects of the spike trains—the number of spikes in a given window, the temporal structure of the spikes, and time delay with respect to the stimulus (figure 29.1).

Neural variability may arise from a variety of sources (Faisal, Selen, & Wolpert, 2008). One such source is sensor noise, such as bending of hair cells (Denk & Webb, 1992), fixational eye movements (Martinez-Conde, Macknik, & Hubel, 2004), and so on. These are usually associated with the mechanistic biophysical properties and the limits of the biological “hardware.” Another source of variability is synaptic fluctuations, which result in unreliable response on the postsynaptic side to incoming presynaptic spikes (Destexhe, Rudolph, Fellous, & Sejnowski, 2001; Seung, 2003; Deweese & Zador, 2004; Stevens & Wang, 1994). Yet another source is the stochastic nature of ion channel opening and closing (DeFelice, 1981; White, Rubinstein, & Kay, 2000), which means that even when presented with the exact same intracellular input, we can expect variability of the resulting spike train (Mainen & Sejnowski, 1995; Schneidman, Freedman, & Segev, 1998). In addition to inherent randomness or “thermal” noise, we might get variance in the response to the same exact inputs without true “randomness”. Variability may result from chaotic behavior at the level of single cells, or from the dynamics of neuronal networks (see Timme, Wolf, & Geisel, 2002; van Vreeswijk & Sompolinsky, 1996). The combination of all these sources of randomness and chaos shapes the probabilistic mapping from stimuli to responses, giving a

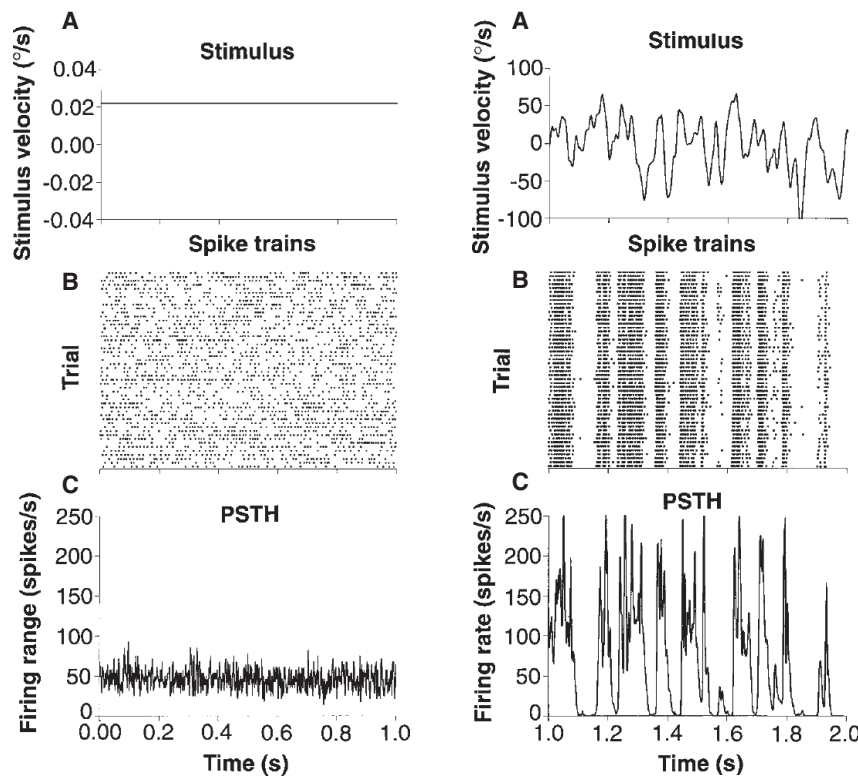


FIGURE 29.1 Reliability and reproducibility of in vivo spiking patterns in response to different stimuli. A fly (*Calliphora vicina*) viewed a pattern of random bars that was moved across the visual field either with constant or dynamic velocity. The spiking activity of H1, a motion-sensitive neuron in the fly's visual system, was recorded from an immobilized fly over multiple repetitions of the same visual stimulus. Left: Spike-pattern reliability for constant stimuli. (A) A random bar pattern was presented across the visual field at constant speed (0.022 degrees/sec) and in the preferred direction of the H1 neuron in the fly. (B) A set of 50 response traces to the stimulus in panel A, each lasting 1 sec, taken 20 sec apart. The occurrence of each spike is shown as a dot. The traces were taken from a segment of the experiment where transient responses have decayed. (C) The peristimulus time histogram (PSTH; bin width 3 ms, 96 presentations), which describes the rate at which spikes were generated in response to the stimulus shown in panel A. The

responses to constant stimuli gave highly irregular and noisy spiking patterns. Right: Spike-pattern reliability for dynamic stimuli. (A) The fly viewed the same spatial pattern as in the left panel but now moving with a time-dependent velocity, part of which is shown. The motion approximates a random walk with diffusion constant  $D \sim 14 \text{ degrees}^2/\text{sec}$ . For illustration, the waveform shown is low-pass filtered. In the experiment, a 10-sec waveform is presented 900 times, every 20 sec. (B) A set of 50 response traces to the stimulus waveform shown in panel A. (C) Averaged rate (PSTH) for the same segment. The rate is strongly modulated, but its time-average is very close to that in the left panel. Variability and reproducibility were much higher than for the constant stimulus, but still showed diverse, stimulus-dependent reliability. (Adapted from de Ruyter van Steveninck et al., 1997.) Such results have parallels in many experiments on sensory neurons.

surprisingly high level of variability in many cases (see more on this below).

Thus, rather than describing neural encoding in terms of their high selectivity (which is the basis of the notion of “grandmother cells”), the mapping from stimuli to responses must be described in terms of a probability distribution over the responses to each stimulus. This implies a computational cost in terms of learning the code and processing it, which then requires accurate sampling of the encoding distribution. Therefore, understanding the design and content of the code relies on the characterization of the variability of the temporal spiking patterns of a single cell and the spatiotemporal patterns of populations.

### *Noise and unreliability of the neural code limit its capacity and shape the structure of the code and processing*

The noise or unreliability of neural responses has a crucial role in determining how information is carried by neurons and how it may be interpreted or read by neurons “down the road.” If neurons were noise-free, each stimulus would result in a single particular neural “code word.” The capacity of the neural vocabulary could then be used efficiently, with every potential spiking pattern carrying distinct information. But, since the mapping from stimuli to responses is probabilistic, then for a given presentation of  $s$ , the response,  $r$ , is described by the probability distribution over the

possible responses  $p(r|s)$ . The readout of information from the neural response then depends on which particular  $r$  was “chosen.” The noise also means that different stimuli may result in the exact same neural response. Therefore, there is also a probabilistic mapping from responses to stimuli, and so observing a response  $r$ , our knowledge of the stimulus that has caused it is given by  $p(s|r)$ . This probabilistic relation sets the limits of the accuracy of decoding or reconstructing the stimulus, and in particular it means that it would generally be impossible to get perfect reconstruction of a stimulus. Furthermore, the noise also determines the similarity of neural responses in terms of their meaning and the similarity of stimuli: what makes stimuli similar for the brain is defined by the overlap of the distributions  $p(r|s)$  for different  $s$ , and the semantic similarity of responses is then given by the overlap of  $p(s|r)$  for different  $r$ . These can be considerably different from our intuition of similarity of the stimulus (Curto & Itskov, 2008; Ganmor, Segev, & Schneidman, 2011a; Kiani, Esteky, Mirpour, & Tanaka, 2007; Tkacik, Granot-Atedgi, Segev, & Schneidman, 2013).

To assess what do the neural responses convey about the stimulus and the impact of noise, we need to quantify the probabilistic nature of the relations between the stimuli  $\{s\}$  and responses  $\{r\}$ . In particular, we seek a measure that captures the relations in both directions—both the encoding distribution  $p(r|s)$ , and the inverse distribution  $p(s|r)$ , which can be viewed as the basis for decoding a neural response; these two distributions are related through Bayes’s rule  $p(s|r) = p(r|s)p(s)/p(r)$ . Such a measure of the relations between  $s$  and  $r$  should also avoid arbitrary assumptions about what features of the spiking patterns matter.

Information theory gives a first-principled way to do this using a minimal set of assumptions (Cover & Thomas, 2012). It does not assume any particular kind of dependency, such as linearity of relations between stimuli and responses, nor does it make assumptions about which stimuli are similar to one another (or the similarity between responses). Instead, it measures a general form of the relation between  $s$  and  $r$ . Importantly, the mutual information does not tell us which features of the stimulus are encoded or how—but it puts a bound on what can be inferred about the stimulus from the neural response, or vice versa (Rieke et al., 1997).

The fundamental measure of this framework is the entropy of a distribution  $p(x)$ , which gives a mathematically unique way to quantify our uncertainty about  $x$ , or how much we do not know about it (Cover & Thomas, 2012; the basis of the logarithm sets the units for the entropy; we use basis 2, which then gives the entropy in “bits”). For a stimulus set  $\{s\}$  where each particular  $s$  appears with probability  $p(s)$ , the entropy of  $H[p(s)]$  measures how much there is to know about the set of stimuli we might encounter. After observing the neural response, our knowledge about the stimulus has changed and is now given by the posterior distribution  $p(s|r)$ ,

which we can quantify by  $H[p(s|r)]$ . The difference between the two is the mutual information between stimulus and response,

$$I(s; r) = H[p(s)] - \langle H[p(s|r)] \rangle_r,$$

where  $\langle \rangle_r$  denotes averaging over  $p(r)$ . The mutual information between stimuli and responses is also equal to  $H[p(r)] - \langle H[p(r|s)] \rangle_s$ . This implies that  $I(s; r)$  can also be interpreted as the difference between the capacity or richness of the neural codebook, quantified by  $H[p(r)]$  and the average noisiness of the responses to a particular stimulus, given by the average over stimuli of the stimulus-dependent conditional entropy  $\langle H[p(r|s)] \rangle_s$ , or “noise entropy” (Rieke et al., 1997; Strong, Koberle, de Ruyter van Steveninck, & Bialek, 1998). (We note here that estimating these entropy terms and mutual information between stimuli and responses suffers from sampling issues and bias, and usually requires large amounts of data as well as bias correction and validation techniques (Strong et al., 1998; Treves & Panzeri, 1995)).

One can assess exactly how much of the neural response is “wasted” because of the noise by the ratio of the information that the neural response carries about the stimulus to the total response entropy, namely  $I(s; r)/H[p(r)]$ . If this ratio is close to 1, then almost all the capacity is used to convey information. Estimating this fraction accurately for temporal spiking patterns is experimentally and computationally demanding, but in the few cases where it has been quantified, almost 50% of the coding capacity of single cells was lost due to noise (see, e.g., Borst & Haag, 2001; Strong et al., 1998).

The noisiness of neural machinery seems wasteful in terms of energy and computational cost, raising the question: Why wasn’t noise selected against by evolution? In particular, despite the prevalence of noise at the level of receptors, synapses, and single neurons, there are examples of neural hardware that is extremely accurate, such as the low spontaneous rate of conformational change of rhodopsin molecules in photoreceptors (Rieke & Baylor, 1998). Moreover, some neural circuits are so finely tuned we can find neurons that respond with temporal accuracy of tens of microseconds (Simmons, Ferragamo, Moss, Stevenson, & Altes, 1990). At the behavioral level, we have examples of high performance despite the noise, which sometimes approaches physical limits, like photon counting (Bialek, 2012), sound localization (Simmons et al., 1990), or primate ability to follow a target at an accuracy limited by its sensory noise (Osborne, Lisberger, & Bialek, 2005).

Performing so well means that some neural circuits are very successful in overcoming their inherent variability and the noisy inputs they receive. One possibility is that these circuits have much more reliable components than other circuits, which in turn must have a cost in terms of design



and maintenance. Alternatively, using many noisy cells to encode the same stimulus would allow for some form of averaging or noise reduction. In most cases, it is not immediately unclear how the brain might use population codes to that effect (Dewese & Zador, 2004), in particular since neural noise is often nonadditive and depends on the specific stimulus presented.

### *How are population codes built?*

The brain relies on arrays of receptors and processing neurons to represent information about the outside world in almost all sensory systems. This population-based representation is inevitable given the richness of the natural sensory world that the brain is presented with, and the binary and noisy encoding capacity of single neurons. But what kind of population codes should we then expect to find in the brain? If neurons overlap in terms of the part of stimulus space they respond to, then population coding implies higher fidelity of coding, by having multiple cells covering the same part of the sensory space, and the potential for noise correction. Alternatively, every neuron might cover an almost unique part of stimulus space (with no or very little overlap with other cells), which would imply detailed coverage of the sensory space, although less potential for noise correction.

From a simple mathematical viewpoint, the potential coding capacity of populations is huge as a result of the combinatorial number of patterns they could use to encode information: If each neuron splits the stimulus space into “preferred” and “nonpreferred” parts and responds by spiking or silence when a stimulus appears in the corresponding parts of stimulus space, then with  $N$  such neurons, the space of potential responses has  $2^N$  distinct patterns. Thus, a handful of such neurons would be able to encode any stimulus ever encountered using a unique pattern, since even just 100 cells would have  $2^{100} \sim 10^{30}$  different activity patterns. However, such combinatorial codes which would efficiently utilize all potential code words would be hard to interpret and decode, because reading the meaning of a pattern would require detailed calculations to infer the contribution of each cell and the relations between them or an exponentially large lookup table. Moreover, because every population activity pattern could carry different information, a missing spike or noisy neuron could result in a very different message being represented, which would make the code even harder to interpret.

On the other extreme, to overcome noise, population codes should rely on some form of redundancy in terms of the information conveyed about their stimuli. One simple possibility is duplication of neural responses, or high overlap by noisy samples of the same stimuli by neurons with similar response properties. This form of population coding could

use a very simple decoding mechanisms, such as population vector (Georgopoulos et al., 1986) and linear decoding (Bialek, Rieke, de Ruyter van Steveninck, & Warland, 1991; Dan, Atick, & Reid, 1996; Warland, Reinagel, & Meister, 1997), where the estimated or reconstructed version of the stimulus is given by  $s^{est}(t) = \sum_i \mathbf{k}_i * \mathbf{r}_i$ , where  $s^{est}(t)$  is the value

of the estimated stimulus at time  $t$ ,  $\mathbf{r}_i$  is a vector denoting the activity of neuron  $i$  up until time  $t$ , and  $\mathbf{k}_i$  for the each of the neurons are spatio-temporal filters, and  $*$  denotes the convolution between the filters and the neurons activity patterns. The filters  $\{\mathbf{k}_i\}$  are fit such that over all cells they minimize the reconstruction error between  $s$  and  $s^{est}$  over time.

However, the design of population codes may need to answer different and sometimes conflicting goals and constraints in terms of energy, error correction, cost of computation, or stimulus statistics. Thus, there is no universal optimal code. Different coding schemes might be efficient or even optimal for different combinations of stimuli, single cell properties, noise level, and connectivity between cells. In other words, to enable effective error correction and decoding of novel neural activity patterns, population codes must balance the richness of combinatorial codes with some form of overlap or dependency between the cells. This balance and therefore, the nature of the code is set by the relations between neurons.

### *Correlations between neurons*

The relations between the spiking patterns of different cells depend on the similarity of their respective stimuli and response properties (receptive fields), the synaptic connections between cells, and their joint circuitry or cells, from which they receive inputs. These different sources of dependencies between cells may result in plain similarity of their activity, linear correlation between their spiking patterns, or in more complicated statistical dependencies. For obvious experimental and practical reasons, the dependency between cells is most readily assessed in terms of the relations between pairs of cells. The discussion below presents the questions and measures of correlated neural codes at the level of pairs, but can be immediately extended to more cells.

The most obvious form of dependency between cells—and the one that the brain itself observes and can assess—is reflected in the activity of one cell given that of another cell, regardless of the details of the stimulus. This is often measured at the level of individual cell pairs by comparing the joint activity pattern of both cells to what would be expected from what we know about the activity of each cell, or the prediction of independent activity of the cells (figure 29.2A). Perhaps most common is the estimation of the Pearson correlation between the firing rates or the number of spikes of the cells over some time window  $\Delta t$  given by,

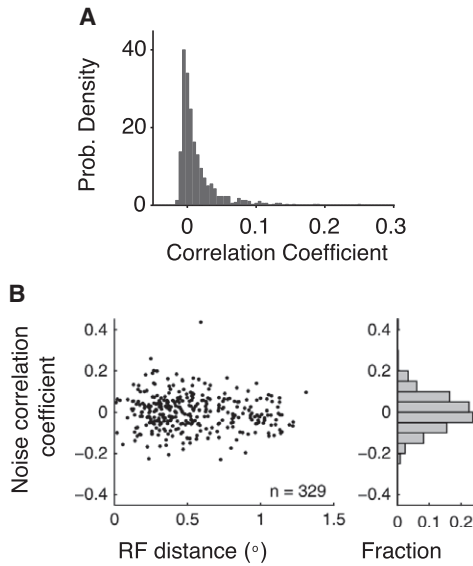


FIGURE 29.2 Distributions of pairwise signal correlations and pairwise noise correlations in two different neural systems. (A) A histogram of correlation coefficients for all pairs of 40 ganglion cells from a patch of an intact salamander retina, recorded using a multi-electrode array, which was presented with a long natural movie. These “signal correlations” are typically weak, even for cells with overlapping receptive fields. (Adapted from Schneidman et al., 2006.) (B) Example of typically weak pairwise noise correlations in the cortex. Left: Spike-count noise correlations of 329 pairs of neurons in the primary visual cortex (area V1) of awake monkeys were estimated using simultaneous recording by an array of chronically implanted tetrodes. Monkeys were presented with natural images that were each shown for 200 ms each, and the noise correlations for each stimulus were estimated based on the first 500 ms of each response. The average noise correlations (over stimuli) are shown as a function of the receptive field distance (mapped separately) for pairs of V1 cells. Right: Distribution of the average noise correlations over the population. (Adapted from Ecker et al., 2010.)

$\frac{(r_1 - \langle r_1 \rangle)(r_2 - \langle r_2 \rangle)}{\sigma_1 \sigma_2}$ , where  $r_1$  and  $r_2$  denote the responses of each of the cells,  $\langle \rangle$  denotes average over conditions, and  $\sigma_i$  denotes the standard deviation of  $r_i$ . However, the dependency between pairs of cells may be nonlinear, and moreover, the relation between cells may not necessarily manifest itself in terms of firing rates. Thus, in general, the dependency between cells should be measured in terms of the probability of seeing joint activity patterns of the two cells  $p(r_1, r_2)$  to the independent prediction  $p(r_1)p(r_2)$ , where  $r_1$  and  $r_2$  can be any form of activity (from spiking rates to more compound spiking patterns in time). Similarly, one can estimate the temporal dependency between cells in terms of  $p(r_1(t)r_2(t - \delta t))$ , where  $\delta t$  denotes the time difference between cells. This form of correlation between cells is commonly termed “signal correlation”, since the probability distributions over activity patterns are the result of summing over all stimuli and conditions that the cells observe together. This measure of correlation between the cells reflects what

knowing about cell 1 tells about the activity of cell 2, averaged over all stimuli, which is exactly what the brain can assess (since usually it has no knowledge of the specific stimulus that was presented). Various measures of similarity between  $p(r_1, r_2)$  and  $p(r_1)p(r_2)$  can be used to generalize over linear measures of correlations, but given the mathematical benefits of information theory, we can quantify the dependency between the two cells in terms of the mutual information between the cell spiking patterns, which is given by

$$I(r_1; r_2) = \sum_{r_1, r_2} p(r_1, r_2) \log_2 \frac{p(r_1, r_2)}{p(r_1)p(r_2)}.$$

We note that this way of expressing mutual information between the two cells is mathematically equivalent to the two other ways to define information between two variables presented above.

Another measure of dependency between cells is the relation between their activity patterns with respect to a specific stimulus (or stimuli). For a pair of cells and a particular stimulus,  $s$ , we can compare the distribution of responses of the two neurons to that particular stimulus, given by  $p(r_1, r_2 | s)$ , to the predicted distribution of joint responses if the cells were responding to the stimulus independently of one another, namely  $p(r_1 | s)p(r_2 | s)$ . This comparison between the empirical distribution of joint responses to the one that assumes that the cells are conditionally independent in responding to the stimulus measures how correlated are the fluctuations (or the “noise”) of the responses of the cells. These correlations are often estimated by presenting the same exact stimulus repeatedly, then shuffling the order of responses of one cell with respect to those of another. It has been common to average this measure of joint fluctuations in encoding over a range of different stimuli. This, then, measures the average nature of joint fluctuations in the responses, or the noisy nature of the joint responses, and is commonly termed “noise correlation”. It is noteworthy that this name is somewhat confusing, since these correlations depend on specific stimuli and then are averaged over stimuli. Thus, these noise correlations are generally not stimulus-independent, and have a strong impact on the way stimulus information is encoded and can be read. In particular, noise correlations for one stimulus may be very different than those of another stimulus. Moreover, while it is relatively easy to assess noise correlations for the case where no external stimulus is presented (ongoing activity), this tells us very little about noise correlations for different stimuli. In terms of quantifying these correlations, we can again use information theory to quantify the information that one cell tells us about another, given a particular stimulus value,  $I(r_1; r_2 | s) = \sum_{r_1, r_2} p(r_1, r_2 | s) \log_2 p(r_1, r_2 | s) / (p(r_1 | s)p(r_2 | s))$ , to measure the noise correlations for particular  $s$ . The average noise correlations would be then quantified by  $\langle I(r_1; r_2 | s) \rangle_s$ .

The two forms of dependencies between neurons we presented are coupled (the distinctions and relations between them have old roots in statistics; see, e.g., Good, 1953). At a fundamental level, the noise correlations are at the heart of the nature of the encoding of a stimulus by the population. Yet the brain has no direct access to these dependencies in general, and can only assess how cells work together over all stimuli (stimulus correlations). In other words, since our estimate of the signal correlations result from the joint probability  $p(r_1, r_2)$ , and this is based on the sum over stimuli of  $p(r_1, r_2 | s)$ , then signal correlations are in a sense a marginalized reflection of the nature of the joint response and joint noise of the cells to specific stimuli or noise correlations (figure 29.2B).

How do these correlations shape the neural population code? In terms of encoding, it is clear that any form of correlation or dependency between the cells must diminish the range of possible patterns the neurons may use to convey information. Yet it is this very dependency that would allow a “listener” to the code to correct errors. Thus, signal and noise correlations set a balance of coding properties. The dependencies between cells restrict the joint vocabulary, but the nature and magnitude of the stimulus-dependent encoding noise is the key for error correction.

A population coding scheme in which the cells encode their stimuli independently of one another (i.e., without noise correlations), would have several experimental and practical benefits, which would make the neural code simple to describe and to study. First, there would be no need to use simultaneous recording of multiple neurons to study the code, and we could combine cells from different recording sessions (although to know that, one would need to record the joint activity to verify that they are indeed independent). Second, the hard problem of sampling of the population code would be reduced to sampling the individual noise of each cell, which is exponentially easier. Third, analyzing the structure of the code would then reduce to understanding how a collection of neurons are encoding their stimuli, which again would be exponentially simpler than exploring joint coding schemes. However, the lack of noise correlation would also mean that the variance of joint responses of the population is larger than in the case of any correlated response of the cells to the stimulus.

While it is common to think of error correction in terms of averaging over multiple independent identical encoders, this is not the only way and in many cases not the optimal way to overcome noise. Critically, the dependencies between cells are a *double edge sword*. They reduce the richness or capacity of the population code, making it less diverse than it would be if the cells were encoding independently. But, this coupling between cells is also what allows for attenuation of noise or error correction by the cells themselves (Deweese & Zador, 2004). Thus, signal correlations or noise

correlations are not universally “good” or “bad” (see figure 29.3B for simple toy examples of the potential effect of noise and signal correlations on the code; Averbeck, Latham, & Pouget, 2006). Instead, the effect of correlations depends on the nature of single cell properties, the balance of the contributions of correlations to reducing the population variability, and richness of the population code (Abbott & Dayan, 1999; Amari & Nakahara, 2006; Averbeck & Lee, 2004; Averbeck et al., 2006; Latham & Nirenberg, 2005; Quiroga & Panzeri, 2009; Romo, Hernández, Zainos, & Salinas, 2003; Schneidman, Bialek, & Berry, 2003a).

### *From correlations to redundancy, independence, and synergy in neural population codes*

Intuitively, the design of information coding by a neural population can be described in terms of the information that is over-represented by several cells and can be read from each of these cells individually (redundant coding), information that is carried uniquely by some cells but not by others (independent coding), and information that is carried by the joint activity of cells in a way that can only be read from their joint activity (synergistic coding; Amari & Nakahara, 2006; Averbeck & Lee, 2004; Latham & Nirenberg, 2005; Schneidman et al., 2003a). It is not immediately clear, though, how to translate the correlation structure among cells, and between the cells and stimuli, into these notions of redundant, synergistic, and independent coding. First, the relation between cells and the stimuli they respond to can be quantified using various correlation measures. For example, while it is straightforward (and common) to assess the Pearson correlation between stimuli and responses, it is not clear how to use such linear measures for the correlation between the stimulus and the combined responses of two cells (or more). Second, and more importantly, it is not clear how these measures relate to one another and to these coding notions.

Information theory gives a principled and self-consistent way to combine these notions of dependency between stimuli and cells, without making arbitrary assumptions—since the mutual information between variables measures their dependency in a way that is not limited to linear correlations, or any specific metric on similarity. We can then compare the information that two cells convey together about the stimulus  $I(s; r_1, r_2)$  and the information that each of the cells carries on its own,  $I(s; r_1)$  and  $I(s; r_2)$ . The synergy or redundancy of the code can be then expressed as the difference between the two, namely

$$SR(r_1, r_2; s) = I(s; r_1, r_2) - (I(s; r_1) + I(s; r_2)).$$

Then, the cells are *redundant* when  $SR < 0$ , since in this case they must be carrying redundant information. The cells are *synergistic* if  $SR > 0$ , since together they convey more



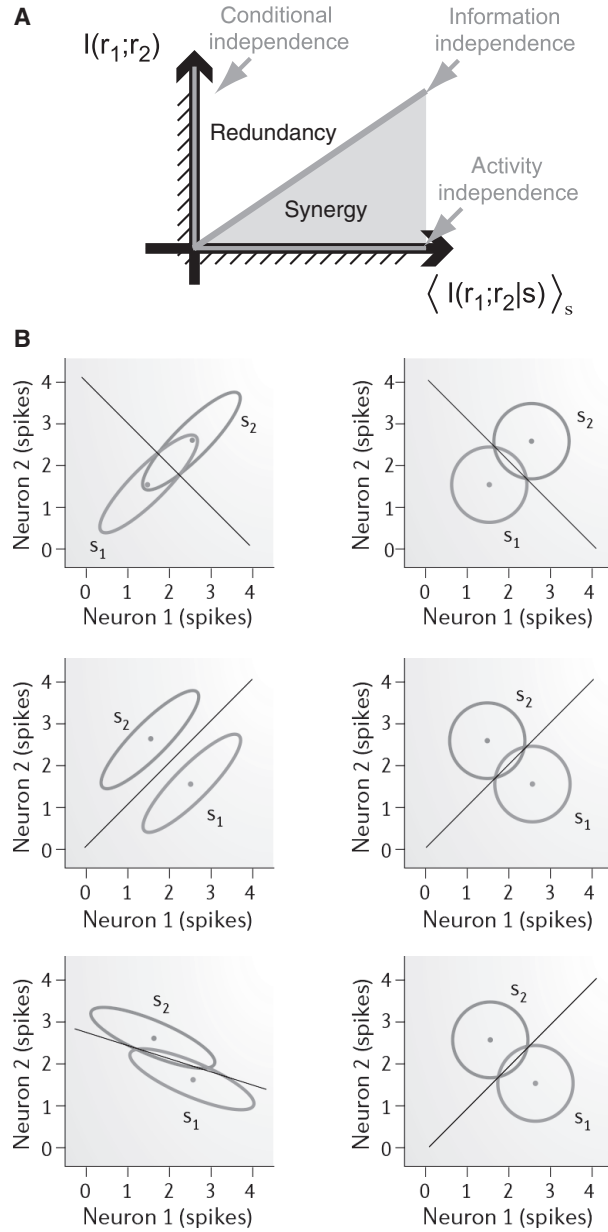


FIGURE 29.3 Graphical presentations of the complex effects of combinations of signal and noise correlations on the information content of population codes. (A) Graphical presentations of synergy as a combination of other measures of non-independence. As discussed in the text, we can represent the synergy or redundancy of a pair of cells as a point in a plane with the axes  $\langle I(r_1; r_2 | s) \rangle_s$  and  $I(r_1; r_2)$ . Because both of these measures are nonnegative, only the top right quadrangle is allowed. Neuronal pairs whose activity is independent, i.e. they have no signal correlations, would be represented by point along the abscissa with  $I(r_1; r_2) = 0$ . Neuronal pairs whose response is independent given the stimulus, i.e. they have no noise correlations, would result in points along the ordinate, with  $I(r_1; r_2 | s) = 0$ . The case where the neurons carry independent information corresponds to the diagonal that separates the synergistic values from the redundant ones. (Adapted from Schneidman et al., 2003a.) (B) A toy model example of the effects of correlation on information encoding. In all three cases (rows), the response distributions for two neurons that respond to two different stimuli are

information than taken separately. If  $SR = 0$ , then the cells carry independent information. This framework also allows us to make direct link to the relations between the activity of the cells and the stimulus  $I(r_1, r_2; s)$ , the signal correlations between cells  $I(r_1; r_2)$ , and the average noise correlations over stimuli  $I(r_1; r_2 | s)$ . The synergy/redundancy (SR) measure defined above is also equal to the difference between the correlations between neurons measured over all stimuli measured as the information between cells, and the average noise correlations, measured as the average over stimuli of the stimulus-dependent information between cells (figure 29.3A)

$$SR(r_1, r_2; s) = I(s; r_1, r_2) - (I(s; r_1) + I(s; r_2)) \\ = I(r_1; r_2) - \langle I(r_1; r_2 | s) \rangle_s$$

In relative terms we can quantify the synergy or redundancy in terms of the ratio  $SR(r_1, r_2; s) / \min(I(s; r_1), I(s; r_2))$ , which measures what fraction of the information that is carried by

shown. The panels on the left show the joint responses of the cells, whereas those on the right show the conditionally independent version of the same responses, also known as “shuffled response,” which can be achieved from the joint responses by shuffling the order of responses of each cell to different presentations of the same stimulus; i.e., left panels show the result of  $p(r_1, r_2 | s)$  and those on the right show the result of  $p(r_1 | s)p(r_2 | s)$ . Each ellipse (which appears as a circle in the uncorrelated plots) indicates the 95% confidence interval for the responses. Each diagonal line shows the optimal decision boundary—that is, responses falling above the line are classified as stimulus 2 and responses below the line are classified as stimulus 1. The  $x$ -axis is the response of neuron 1, the  $y$ -axis the response of neuron 2. (Note that in this particular simple case, signal correlations correspond to the relative positions of the mean responses and noise correlations control the shape and orientation of the ellipses, and thus whether correlations result in more or less encoded information). Top row: In this case the responses were chosen such that the signal and noise correlations are both positive, and this leads to larger overlap between the ellipses for the correlated rather than for the uncorrelated responses. A larger fraction of the ellipses lie on the “wrong” side of the decision boundary for the correlated responses than for the independent responses. In this case  $I(r_1, r_2; s) < I(r_1; s) + I(r_2; s)$ , and so the cells carry more redundant information compared to the conditional independent case. Middle row: Here the signal correlations are negative and noise correlations are positive, and there is less overlap in the correlated than uncorrelated responses. A smaller fraction of the ellipses lie on the wrong side of the decision boundary for the correlated responses, and in this case  $I(r_1, r_2; s) > I(r_1; s) + I(r_2; s)$ . This also means that the cells in this case carry synergistic information. Bottom row: The same fraction of the ellipses lies on the wrong side of the decision boundary for both the correlated and independent responses, so in this case  $I(r_1, r_2; s) = I(r_1; s) + I(r_2; s)$ . Thus, the presence of correlations does not guarantee an effect on the amount of information encoded—it is the detailed correlation structure of the signal and noise correlations that shape the content and organization of the code. (Adapted from Averbeck et al., 2006.)

the cells is redundant, and is often denoted as the fractional redundancy.

The different mutual information terms we presented and the relations between them do not tell us which stimuli are associated with which responses. Nor do they tell us what is the redundant or synergistic information that is encoded by the cells. Other kinds of measures must be used to give a breakdown of the encoded information and relations between cells or stimuli, which would have to rely on additional assumptions, such as metric approaches to the stimuli or decoding errors (Brunel & Nadal, 1998; Kang & Sompolinsky, 2001; Shamir & Sompolinsky, 2004), or measures of information carried by specific patterns of activity (Brenner, Strong, Koberle, Bialek, & de Ruyter van Steveninck, 2000; Deweese & Meister, 1999; Pola, Thiele, Hoffmann, & Panzeri, 2003; Schneidman et al., 2011). The power and benefit of using information-theory based measures here is exactly their assumption- and metric-free nature, which gives bounds on the relations between cells and stimuli.

Somewhat surprisingly, the correlations between pairs of cells are typically weak—both in terms of the signal correlations, and the average of noise correlations over stimuli—in many different sensory systems (Bair, Zohary, & Newsome, 2001; Ecker et al., 2010; Nirenberg, Carcieri, Jacobs, & Latham, 2001). This has often been interpreted to suggest that we could either neglect noise correlations, or signal ones. Moreover, estimates of pairwise redundancy between cells have also given relatively low values of redundancy, typically on the order of a few percent (Chechik et al., 2006; Gawne & Richmond, 1993; Narayanan, 2005; Puchalla, Schneidman, Harris, & Berry, 2005; Reich, Mechler, & Victor, 2001; see an example in figure 29.4). Yet even these typically weak correlations can shape the reliability of neural coding and computation (Cafaro & Rieke, 2010; Lee, Port, Kruse, & Georgopoulos, 1998).

The results for pairs of cells seem to be consistent with one of the more influential theoretical ideas in neuroscience, known as “redundancy reduction”. Attneave (1954) and Barlow (1961, 2001) suggested that neural populations “aim” to encode information efficiently, and since redundancy is prevalent in natural sensory scenes, neurons would decorrelate their responses with respect to one another, to optimally use their joint coding capacity. This idea of a design principle for neural populations was the basis for exploration of possible structures of the neural code (van Hateren, 1992) and used to predict the receptive field properties of single cells in the visual system (Atick & Redlich, 1990, 1992; Dan et al., 1996), and at the population level, corresponded to the notion of independent component analysis and coding (Bell & Sejnowski, 1995, 1997; Hyvärinen & Hoyer, 2001).

However, removing redundancy and achieving decorrelation would be problematic for two reasons—as Barlow

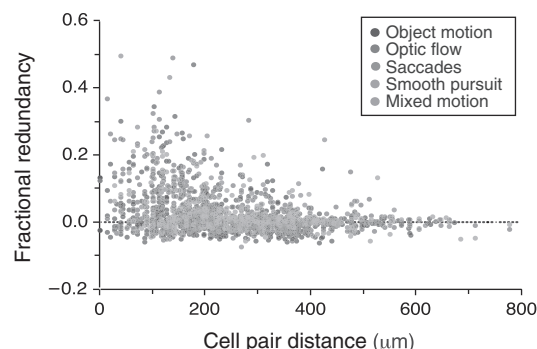


FIGURE 29.4 Pairwise redundancy in the retina, responding to natural scenes. Four salamander retina patches were presented with natural movie clips, and the ganglion cell responses were recorded extracellularly using a multi-electrode array. The fractional redundancy (see main text) for 1,838 cell pairs is shown (each dot represents one pair of cells), as a function of the distance between the centers of the cells’ receptive-fields. The type of motion present in each movie clip is shown by the dot: object motion, saccades, optic flow, smooth pursuit, and combinations of motion. Most pairs show weak redundancy or information independence. (Adapted from Puchalla et al., 2005.)

himself argued later (Barlow, 2001). First, redundancy is crucial for overcoming noise, and so reducing the redundancy, and making the cells independent encoders, would make the code highly susceptible to noise. Second, learning from examples must rely on identifying structure in patterns themselves. An efficient code in terms of its information rate would have no such redundancy, making the patterns seem almost random in structure. This would make it a hard code to learn from examples, since there would be no apparent dependencies from which one could infer the rules or structure (which is what one uses when learning a new language, for example).

What, then, is the nature of correlation and redundancy in large populations of cells? Naively, if every cell has 5% correlation with 100 of its neighbors, then every piece of information is repeated on average 5 times (Puchalla et al., 2005). Thus, even before we consider high-order dependencies, the network implications of the correlation and redundancy at the level of pairs depend on the number of pairs and the relations between the pairs.

### *Signal and noise correlations and the population codes of large groups of neurons*

Despite the typically weak correlation and redundancy at the level of pairs, the correlation among large groups of neurons can be surprisingly strong. Even small groups of ~10 neurons in the retina already show strongly correlated activity patterns that deviate significantly from what might be expected from their weak pairwise correlations (Schneidman

et al., 2006; Schnitzer & Meister, 2003; Shlens et al., 2006). These deviations are clear at the level of specific patterns, showing orders of magnitude mismatch between the rates of appearance of patterns that one observes experimentally and what would be expected from independent cells. Over all patterns, the difference between the joint distribution of activity patterns of  $N$  cells,  $p(r_1, r_2, \dots, r_N)$ , and the independent model  $p_{\text{independent}}(r_1, r_2, \dots, r_N) = p(r_1)p(r_2)\dots p(r_N)$ , quantifies the strength of network correlations, implying that the population vocabulary is far less diverse than it would have been if cells were independent (Schneidman et al., 2006).

To understand the nature of the code of large neural populations, it is imperative to use tools that would enable us to dissect the nature of the different orders of dependency and correlations between cells and their coding role (Amari, 2001; Martignon et al., 2000; Nakahara & Amari, 2002; Schneidman, Still, Berry, & Bialek, 2003b). Somewhat surprisingly, strong network correlations could arise from the collective effect of the quadratic number of weak pairwise dependencies. This can be seen by building the minimal model that relies on the pairwise correlations between cells: since the entropy of a distribution measures how random or structured it is, the distribution  $p(r_1, r_2, \dots, r_N)$  with maximal entropy (Jaynes, 1957) that is consistent with the observed firing rates  $r_i$  and the pairwise correlations between cells (which is equivalent to giving the average of all products  $r_i r_j$ ) is the minimal or most parsimonious model that relies only on pairwise relations between cells (Schneidman et al., 2003). If we discretize neural responses into small time bins, where in each bin a neuron either spikes  $r_i = 1$  or is silent  $r_i = 0$ , then this minimal pairwise model is given by a distribution of the form

$$p^{(2)}(r_1, r_2, \dots, r_N) = \frac{1}{Z} \exp \left( \sum_i \alpha_i r_i + \sum_{i < j} \beta_{ij} r_i r_j \right).$$

Here,  $\alpha_i$  (for each neuron) and  $\beta_{ij}$  (for each pair) are Lagrange multipliers that are found to obey the required constraints, and  $Z$  is a normalization factor, or the *partition function*. This solution, which is equivalent to the Ising model from physics (Landau & Lifshitz, 1996), is mathematically unique and can be found numerically, giving the most parsimonious, or minimal pairwise model of the population codebook.

In the vertebrate retina, cultured cortical neurons, cortical slices, and in vivo cortical recordings, it was found that such pairwise maximum-entropy models give an extremely accurate description of the population activity vocabulary for both spatial and temporal population activity patterns (see figure 29.5A and Amari, Nakahara, Wu, & Sakai, 2003; Ganmor, Segev, & Schneidman, 2011b; Marre, El Boustani, Frégnac, & Destexhe, 2009; Ohiorhenuan et al., 2010; Schneidman et al., 2006; Shlens et al., 2006, 2009; Tang et al., 2008).

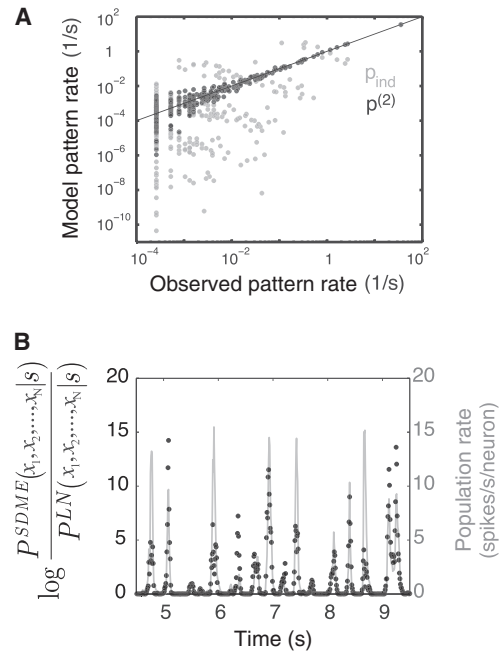


FIGURE 29.5 Strong signal and noise correlations at the level of large populations of neurons. (A) An example of strong signal correlations in the vertebrate retina and the accuracy of pairwise maximum-entropy models. The response of 10 simultaneously recorded ganglion cells in an intact patch of salamander retina, in response to a natural movie, was recorded using a multi-electrode array. Neural responses were discretized into 20 ms bins, and the activity of the group was represented in each time bin as a binary word of length 10, with  $x_i = "1"$  if a neuron was active, and  $"0"$  if it were silent. The rate of occurrence of each pattern predicted if all cells are independent is plotted against the measured rate is shown in gray. Each dot stands for one of the  $2^{10}$  possible binary activity patterns for the 10 cells. Black line shows equality. For the same group of cells, the rate of occurrence of each firing pattern predicted from the pairwise maximum-entropy model  $p^{(2)}$  that takes into account all pairwise correlations is plotted against the measured rate (black dots). The rates of commonly occurring patterns are predicted with better than 10% accuracy, and scatter between predictions and observations is confined largely to rare events for which the measurement of rates is itself uncertain. (Adapted from Schneidman et al., 2006.) (B) The stimulus-dependent maximum-entropy (SDME) model (see main text) was fit to describe the response of 100 retinal ganglion cells to full-field flicker Gaussian stimulus. The pairwise SDME model predicts population activity patterns for  $N = 100$  neurons better than a model that assumes that the cells are independent encoders, given by a set of conditionally independent LN models for each of the cells. The log-likelihood ratio of the population firing patterns under the pairwise SDME model and under the conditionally-independent model is shown as a function of time (black dots, scale on left) for an example stimulus repeat (which was not used to train the model). For reference, the average firing rate of the population is shown in grey (scale on right). Thus, the noise correlations in the population are very strong at particular times along the stimulus. (Adapted from Granot-Atedgi et al., 2013.)

Going beyond the overall vocabulary of the code or signal correlation, the immediate question is what is the role of correlations in encoding specific stimuli, or what is the nature of population noise correlations? Similar to the case of the overall vocabulary, also for the case of the nature of the responses of large populations of neurons to the same stimulus are often poorly described by a model that assumes no noise correlations (Granot-Atedgi, Tkacik, Segev, & Schneidman, 2013; Pillow et al., 2008; Schneidman et al., 2006). Explicitly, this means that  $p(r_1, r_2, \dots, r_N | s)$  can be significantly different from  $p(r_1 | s)p(r_2 | s) \dots p(r_N | s)$ . This difference between these distributions reflects the impact of noise correlation at the level of the network for a particular stimulus (Granot-Atedgi et al. 2013), which for a population of  $N$  cells we can quantify by  $I(r_1, r_2, \dots, r_N | s) = H[p(r_1 | s)p(r_2 | s) \dots p(r_N | s)] - H[p(r_1, r_2, \dots, r_N | s)]$ . Importantly, while  $I(r_1, r_2, \dots, r_N | s)$  can be very high for particular  $s$ , it is often the case that averaging over stimuli, gives a relatively low value of  $\langle I(r_1, r_2, \dots, r_N | s) \rangle$ , since for many other stimuli the noise correlations at the level of the group can be low (Granot-Atedgi et al. 2013). As it has been customary in the literature to report the average noise correlations over stimuli, and these reports often suggested weak noise correlations (see e.g., Ecker et al. 2010), one should note that these averages may indeed mask strong noise correlations for particular stimuli (see e.g. Granot-Atedgi et al. 2013).

Two families of models have been used to show how taking the relations between cells into account gives a significant improvement in describing the correlated nature of population encoding of specific stimuli. The generalized linear model (GLM; Pillow et al., 2008) describes the response of each cell in a population to a stimulus  $s$  by its instantaneous firing rate, given by

$$\lambda_i(t) = \exp \left( \mathbf{k}_i * \mathbf{s} + \mathbf{h}_i * \mathbf{r}_i + \sum_{j \neq i} \mathbf{l}_{ji} * \mathbf{r}_j + \mu_i \right) \text{ where } * \text{ denotes}$$

the convolution between two vectors,  $\mathbf{r}_i$  is a vector (over time bins) denoting the spike train history of cell  $i$  at time  $t$ , and  $\{\mathbf{r}_j\}$  are the spike-train histories of the other cells at time  $t$ ,  $\mathbf{k}_i$  is the stimulus filter of neuron  $i$ ,  $\mathbf{h}_i$  is the history filter for neuron  $i$ ,  $\mathbf{l}_{ji}$  are the coupling filters between neuron  $j$  to neuron  $i$ , and  $\mu_i$  is the baseline log-firing rate of cell  $i$  that were found to maximize the likelihood of a spiking pattern of each of the cells. GLMs have been shown to be more accurate in capturing noise correlations among small populations of neurons (Pillow et al., 2008; Truccolo, Hochberg, & Donoghue, 2009). Maximum entropy models that are also stimulus-dependent can give the minimal models that are consistent with stimulus-dependent features of the neuronal response. A simple version of such stimulus-dependent maximum entropy models (Tkacik et al. 2010) is the minimal model that has the empirical stimulus-dependent firing rate of the cells  $r_i(s)$  (where the average is over multiple

presentations of the same stimulus) and the correlation between cells  $r_{ij}$  over all stimuli. This model is given by a distribution of the form

$$p^{(2)}(r_1, r_2, \dots, r_N | s) = \frac{1}{Z(s)} \exp \left( \sum_i \alpha_i(s) r_i + \sum_{i < j} \beta_{ij} r_i r_j \right)$$

where the individual stimulus parameters for each of the cells,  $\alpha_i(s)$  and the stimulus-independent pairwise interaction terms  $\beta_{ij}$  are set to match the measured firing rates  $r_i(s)$  and the pairwise correlations between cells;  $Z(s)$  is a normalization factor or partition function for each stimulus  $s$ , given by

$$Z(s) = \sum_{\{r_i\}} \exp \left( \sum_i \alpha_i(s) r_i + \sum_{i < j} \beta_{ij} r_i r_j \right).$$

These models (Granot-Atedgi et al. 2013) show orders of magnitude improvement in capturing the noise correlations for different stimuli for large populations of cells (figure 29.5B). Thus, also in the case of noise correlations, the typically weak noise correlations at the level of pairs add up to strong noise correlations at the level of the network for particular stimuli. Importantly, if we average the measure of noise correlation at population level over stimuli

These modeling approaches are not only beneficial but critical for the analysis and understanding the nature of the code of large neural populations. Again, since the number of potential patterns that  $N$  cells can transmit is exponential in  $N$ , any hope of direct sampling the joint activity of even just 100 neurons, which have  $2^{100}$  potential patterns, is doomed to fail. In fact, we often find if we take the population activity patterns of large group of cells observed in a long experiment, then model that we fit to one half of the data is a better predictor of the other half, than the original data itself. Hence, even describing the data would be more accurate using a model. Ultimately, understanding the code of large groups of neurons and in particular the relations between the correlations between cells and the stimulus, must therefore rely on models that would allow us to identify the design principles of the code, which the modeling frameworks like maximum entropy and GLMs aim to give us.

## Discussion

The emerging picture of the design of population codes is one in which large groups of noisy and highly selective neurons are strongly correlated. These strong network correlations rely on a large set of typically weak dependencies that add up and result in strongly correlated collective behavior at the level of the population.

The correlations between neurons are almost inevitable, given the overlap in receptive fields, correlations in the stimuli, and synaptic connections between neurons. An overlap in the receptive field of cells (which would be reflected in signal correlations) is imperative for the ability



to overcome noise, but it also means that some of the potential coding capacity of the neurons is lost due to their overlap. Weak noise correlations, or lack thereof, would suggest an obvious way to average over independent encoders of the same stimulus, but this would prevent the cells from correcting or reducing the noise through collective coding. The role and value of correlations are set by the balance of correlations in the stimulus and neural noise, while connections between the cells shape the nature of population codes. In particular, the combination of noise and signal correlations may enable neural populations to operate at a point that may give a code that has high information capacity by matching it to the statistics of stimuli and noise (figure 29.6; Tkacik, Prentice, Balasubramanian, & Schneidman, 2010). The redundancy in the code may have an additional role in creating repeating structure in the code, which is important for learning from examples (Barlow, 2001). Interestingly, recent results on the nature of correlations in large neural populations under naturalistic stimuli suggest that strong correlations result from a sparse network of high-order interactions, making the population code simple to learn (Barlow, 2001; Ganmor et al., 2011a).

Why does noise play such a significant role in shaping the neural code? One possible reason is that evolution has not optimized the neural code as well as other features of our brains. However, with  $10^{11}$  neurons and  $10^{15}$  synapses in the human brain, it is clear that if this were the key aspect of optimality, the evolutionary pressure for more reliable neural circuits would be immense. The brain's performance despite its noisy hardware suggests that neural noise may be a "feature" rather than a "bug." Noisy or probabilistic-based coding and computation may enable the brain to overcome local minima in learning (Seung, 2003), and underlie innate exploration of new solutions. It may even allow the computational benefits of probabilistic algorithms, which can be better than any known deterministic ones.

The nature of noise and correlations has been suggested to imply different design principles of information processing and computation that one brain area may perform on the population-activity patterns of noisy correlated neurons it receives from another. We have presented here the notion of strongly correlated population codes that rely on typically weak pairwise relations, which give noise-tolerant and possibly learnable codes. Other, nondistinct features that have been reported include codes from which important behavioral information can be read fast and reliably, even if not perfectly (Gollisch & Meister, 2008; Haddad et al., 2010; van Rullen & Thorpe, 2001); predictive coding of stimulus features by neural populations (Hosoya, Baccus, & Meister, 2005; Palmer, Marre, & Berry, 2013; Srinivasan, Laughlin, & Dubs, 1982); sparseness and overcompleteness of the code (Olshausen & Field, 1997; Simoncelli & Olshausen, 2001); time scale of representation (Burak & Fiete, 2012); and

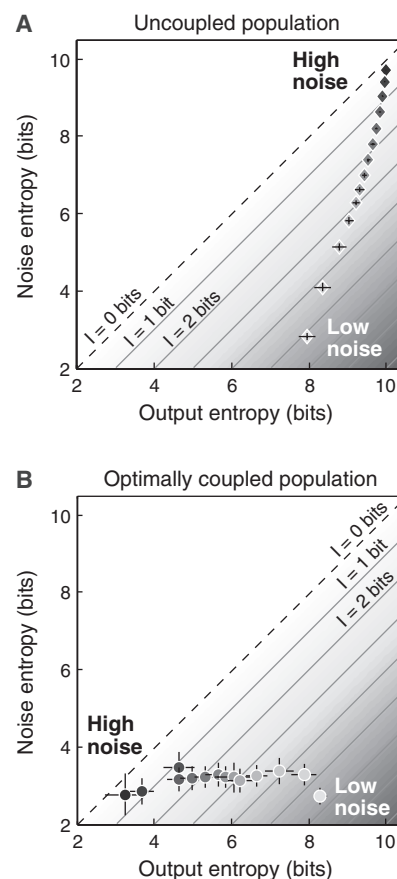


FIGURE 29.6 The effect of noise and correlation on the nature of optimal information encoding by model networks. (A and B) A toy SDME model with 10 neurons was studied under Gaussian correlated stimuli for the case of no coupling between the cells, aside from signal correlations (in panel A), and with the optimal couplings that would maximize the encoded information (in panel B). The output entropy (total population entropy) and the population noise entropy are shown for uncoupled and optimal coupled networks parametrically as neural reliability (or  $1/\text{noise}$ ) changes from high (bright symbols) to low (dark symbols). The information that the population transmits  $I$ , is the difference between output and noise entropies and is shown by the gray scale gradient. Networks with low reliability transmit less information and thus lie close to the diagonal, while networks that achieve high information rates lie close to the lower right corner. The optimal network uses its couplings to maintain a consistently low network noise entropy despite a 10-fold variation in neural reliability. Error bars are computed over 30 replicate optimizations for each reliability ( $1/\text{noise}$ ) level. Thus, these two models show distinct information coding capability, based on optimal coupling between cells, but also distinct profiles of coding vocabulary and coding noise. (Adapted from Tkacik et al., 2010.)

energy efficiency (Laughlin, 2001; Laughlin & Sejnowski, 2003; Levy & Baxter, 1996). It is not unlikely that the principles would be different in different parts of the brain. In particular, the effect of noise in neural circuits goes beyond the capacity and accuracy of representing sensory information that we presented here. It also imposes limits on the



nature of persistent activity in networks and propagation in the brain (London, Roth, Beeren, Häusser, & Latham, 2010), and is likely to have a crucial role in shaping the computations that underlie decision making based on population activity (Beck et al., 2008; Churchland et al., 2011; Lee, 2008) and the representation of probabilistic information in the brain (Beck et al., 2008; Berkes et al., 2011; Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; Deneve, Latham, & Pouget, 1999; Lee et al., 1998; Pouget, Dayan, & Zemel, 2000).

The ability to record large populations of neurons over long time periods would enable us in the not-so-distant future to address the design and content of very large population codes and their dynamics. Recent experimental studies suggest that network correlation structure may change during development or learning (Berkes, Orban, Lengyel, & Fiser, 2011; Tkacik et al., 2010) to match stimulus statistics. Theoretical analysis and experimental evidence suggest that very different correlation structures would be optimal for different stimuli and noise conditions (Chechik et al., 2006; Pitkow & Meister, 2012; Tkacik et al., 2010), and that the ongoing activity of neural populations may be reflecting the result of network optimization to match stimulus statistics (Berens, Ecker, Gerwinn, Tolia, & Bethge, 2011; Berkes et al., 2011; Ganguli & Simoncelli, 2010; Tkacik et al., 2010). Similar frameworks to the one we presented here would allow us to ask how correlations change during adaptation and learning, as well as decipher the design of neural population codes and the network computations that rely on them.

**ACKNOWLEDGMENTS** This work was supported by grants from the European Research Council grant (No. 311238 Neuro-Popcode), the Israeli Science Foundation, the SFARI foundation, and the estate of Toby Bieber.

## REFERENCES

- ABBOTT, L. F., & DAYAN, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Comput*, 11, 91–101.
- AHISSAR, E., VAADIA, E., AHISSAR, M., BERGMAN, H., ARIELI, A., & ABELES, M. (1992). Dependence of cortical plasticity on correlated activity of single neurons and on behavioral context. *Science*, 257, 1412–1415.
- AMARI, S. I. (2001). Information geometry on hierarchy of probability distributions. *IEEE Trans Inform Theory*, 47, 1701–1711.
- AMARI, S.-I., & NAKAHARA, H. (2006). Correlation and independence in the neural code. *Neural Comput*, 18, 1259–1267.
- AMARI, S.-I., NAKAHARA, H., WU, S., & SAKAI, Y. (2003). Synchronous firing and higher-order interactions in neuron pool. *Neural Comput*, 15, 127–142.
- ATICK, J. J., & REDLICH, A. N. (1990). Towards a theory of early visual processing. *Neural Comput*, 2, 308–320.
- ATICK, J. J., & REDLICH, A. N. (1992). What does the retina know about natural scenes? *Neural Comput*, 4, 196–210.
- ATTNEAVE, F. (1954). Some informational aspects of visual perception. *Psychol Rev*, 61, 183–193.
- AVERBECK, B. B., LATHAM, P. E., & POUGET, A. (2006). Neural correlations, population coding and computation. *Nat Rev Neurosci*, 9, 3633–3644.
- AVERBECK, B. B., & LEE, D. (2004). Coding and transmission of information by neural ensembles. *Trends Neurosci*, 27, 225–230.
- BAIR, W., ZOHARY, E., & NEWSOME, W. T. (2001). Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J Neurosci*, 21, 1676–1697.
- BARLOW, H. (2001). Redundancy reduction revisited. *Network*, 12, 241–253.
- BARLOW, H. B. (1961). Possible principles underlying the transformation of sensory messages. In W. Rosenblith (Ed.), *Sensory communication* (pp. 217–234). Cambridge, MA: MIT Press.
- BECK, J. M., MA, W. J., KIANI, R., HANKS, T., CHURCHLAND, A. K., ROITMAN, J., ... POUGET, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron*, 60, 1142–1152.
- BELL, A. J., & SEJNOWSKI, T. J. (1995). An Information-maximization approach to blind separation and blind deconvolution. *Neural Comput*, 7, 1129–1159.
- BELL, A. J., & SEJNOWSKI, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vis Res*, 37(23), 3327–38.
- BERENS, P., ECKER, A. S., GERWINN, S., TOLIAS, A. S., & BETHGE, M. (2011). Reassessing optimal neural population codes with neurometric functions. *Proc Natl Acad Sci USA*, 108, 4423–4428.
- BERKES, P., ORBAN, G., LENGUEL, M., & FISER, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, 331, 83–87.
- BIALEK, W. (2012). *Biophysics*. Princeton, NJ: Princeton University Press.
- BIALEK, W., RIEKE, F., DE RUYTER VAN STEVENINCK, R. R., & WARLAND, D. (1991). Reading a neural code. *Science*, 252, 1854–1857.
- BORST, A., & HAAG, J. (2001). Effects of mean firing on neural information rate. *J Comput Neurosci*, 10, 213–221.
- BRENNER, N., STRONG, S. P., KOBERLE, R., BIALEK, W., & DE RUYTER VAN STEVENINCK, R. R. (2000). Synergy in a neural code. *Neural Comput*, 12, 1531–1552.
- BRUNEL, N., & NADAL, J.-P. (1998). Mutual information, Fisher information, and population coding. *Neural Comput*, 10, 1731–1757.
- BURAK, Y., & FIETE, I. R. (2012). Fundamental limits on persistent activity in networks of noisy neurons. *Proc Natl Acad Sci USA*, 109, 17645–17650.
- CAFARO, J., & RIEKE, F. (2010). Noise correlations improve response fidelity and stimulus encoding. *Nature*, 468(7326), 964–967.
- CALVIN, W. H., & STEVENS, C. F. (1968). Synaptic noise and other sources of randomness in motoneuron interspike intervals. *J Neurophysiol*, 31, 574–587.
- CHECHIK, G., ANDERSON, M. J., BAR-YOSEF, O., YOUNG, E. D., TISHBY, N., & NELKEN, I. (2006). Reduction of information redundancy in the ascending auditory pathway. *Neuron*, 51, 359–368.
- 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, 818–831.
- COVER, T. M., & THOMAS, J. A. (2012). *Elements of information theory*. New York, NY: Wiley.
- CURTO, C., & ITRSKOV, V. (2008). Cell groups reveal structure of stimulus space. *PLoS Comput Biol*, 4, e1000205.

- DAN, Y., ATICK, J. J., & REID, R. C. (1996). Efficient coding of natural scenes in the lateral geniculate nucleus: Experimental test of a computational theory. *J Neurosci*, 16, 3351–3362.
- DAVID, S. V., VINJE, W. E., & GALLANT, J. L. (2004). Natural stimulus statistics alter the receptive field structure of V1 neurons. *J Neurosci*, 24, 6991–7006.
- DAW, N. D., O'DOHERTY, J. P., DAYAN, P., SEYMOUR, B., & DOLAN, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, 441, 876–879.
- DE RUYTER VAN STEVENINCK, R. R., LEWEN, G. D., STRONG, S. P., KOBERLE, R., & BIALEK, W. (1997). Reproducibility and variability in neural spike trains. *Science*, 275, 1805–1808.
- DEFELICE, L. J. (1981). *Introduction to membrane noise*. New York, NY: Plenum.
- DENEVE, S., LATHAM, P. E., & POUGET, A. (1999). Reading population codes: A neural implementation of ideal observers. *Nat Neurosci*, 2, 740–745.
- DENK, W., & WEBB, W. W. (1992). Forward and reverse transduction at the limit of sensitivity studied by correlating electrical and mechanical fluctuations in frog saccular hair cells. *Hear Res*, 60, 89–102.
- DESTEXHE, A., RUDOLPH, M., FELLOUS, J. M., & SEJNOWSKI, T. J. (2001). Fluctuating synaptic conductances recreate in vivo-like activity in neocortical neurons. *Neuroscience*, 107, 13–24.
- DEWESE, M. R., & MEISTER, M. (1999). How to measure the information gained from one symbol. *Network*, 10, 325–340.
- DEWESE, M. R., & ZADOR, A. M. (2004). Shared and private variability in the auditory cortex. *J Neurophysiol*, 92, 1840–1855.
- 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, 584–587.
- FAISAL, A. A., SELEN, L. P. J., & WOLPERT, D. M. (2008). Noise in the nervous system. *Nat Rev Neurosci*, 9, 292–303.
- GANGULI, D., & SIMONCELLI, E. P. (2010). Implicit encoding of prior probabilities in optimal neural populations. *Adv Neural Inf Process Syst*, 23, 658–666.
- GANMOR, E., SEGEV, R., & SCHNEIDMAN, E. (2011a). Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proc Natl Acad Sci USA*, 108, 9679–9684.
- GANMOR, E., SEGEV, R., & SCHNEIDMAN, E. (2011b). The architecture of functional interaction networks in the retina. *J Neurosci*, 31, 3044–3054.
- GAWNE, T. J., & RICHMOND, B. J. (1993). How independent are the messages carried by adjacent inferior temporal cortical neurons? *J Neurosci*, 13, 2758–2771.
- GEFFEN, M. N., DE VRIES, S. E. J., & MEISTER, M. (2007). Retinal ganglion cells can rapidly change polarity from off to on. *PLoS Biol*, 5, e65.
- GEORGIOPOULOS, A. P., SCHWARTZ, A. B., & KETTNER, R. E. (1986). Neuronal population coding of movement direction. *Science*, 233, 1416–1419.
- GERSTNER, W., KREITER, A. K., MARKRAM, H., & HERZ, A. V. (1997). Neural codes: Firing rates and beyond. *Proc Natl Acad Sci USA*, 94, 12740–12741.
- GOLLISCH, T., & MEISTER, M. (2008). Rapid neural coding in the retina with relative spike latencies. *Science*, 319, 1108–1111.
- GOOD, I. J. (1953). The population frequencies of species and the estimation of population parameters. *Biometrika*, 40, 237–264.
- GRANOT-ATEDGI, E., TKACIK, G., SEGEV, R., & SCHNEIDMAN, E. (2013). Stimulus-dependent maximum entropy models of neural population codes. *PLoS Comput Biol*, 9, e1002922.
- GRAY, C. M., KÖNIG, P., ENGEL, A. K., & SINGER, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338, 334–337.
- HADDAD, R., WEISS, T., KHAN, R., NADLER, B., MANDAIRON, N., BENSATI, M., ... SOBEL, N. (2010). Global features of neural activity in the olfactory system form a parallel code that predicts olfactory behavior and perception. *J Neurosci*, 30, 9017–9026.
- HARRIS, C. M., & WOLPERT, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, 394, 780–784.
- HARTLINE, H. K. (1938). The response of single optic nerve fibers of the vertebrate eye to illumination of the retina. *Am J Physiol*, 121, 400–415.
- HATSPOPOULOS, N. G., OJAKANGAS, C. L., PANINSKI, L., & DONOGHUE, J. P. (1998). Information about movement direction obtained from synchronous activity of motor cortical neurons. *Proc Natl Acad Sci USA*, 95(26), 15706–15711.
- HOSOGA, T., BACCUS, S. A., & MEISTER, M. (2005). Dynamic predictive coding by the retina. *Nature*, 436, 71–77.
- HYVÄRINEN, A., & HOYER, P. O. (2001). A two-layer sparse coding model learns simple and complex cell receptive fields and topography from natural images. *Vision Res*, 41, 2413–2423.
- JAYNES, E. (1957). Information theory and statistical mechanics. *Phys Rev*, 106, 620–630.
- KANG, K., & SOMPOLINSKY, H. (2001). Mutual information of population codes and distance measures in probability space. *Phys Rev Lett*, 86, 4958–4961.
- KIANI, R., ESTEKY, H., MIRPOUR, K., & TANAKA, K. (2007). Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *J Neurophysiol*, 97, 4296–4309.
- KÖRDING, K. P., & WOLPERT, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427, 244–247.
- LANDAU, L. D., & LIFSHITZ, E. M. (1996). *Statistical physics*. Oxford, UK: Elsevier.
- LATHAM, P. E., & NIRENBERG, S. (2005). Synergy, redundancy, and independence in population codes, revisited. *J Neurosci*, 25, 5195–5206.
- LAUGHLIN, S. B. (2001). Energy as a constraint on the coding and processing of sensory information. *Curr Opin Neurobiol*, 11, 475–480.
- LAUGHLIN, S. B., & SEJNOWSKI, T. J. (2003). Communication in neuronal networks. *Science*, 301, 1870–1874.
- LEE, D. (2008). Game theory and neural basis of social decision making. *Nat Neurosci*, 11, 404–409.
- LEE, D., PORT, N. L., KRUSE, W., & GEORGIOPOULOS, A. P. (1998). Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci*, 18, 1161–1170.
- LEVY, W. B., & BAXTER, R. A. (1996). Energy efficient neural codes. *Neural Comput*, 8, 531–543.
- 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.
- MAINEN, Z., & SEJNOWSKI, T. (1995). Reliability of spike timing in neocortical neurons. *Science*, 268, 1503–1506.
- MARRE, O., EL BOUSTANI, S., FRÉGNAC, Y., & DESTEXHE, A. (2009). Prediction of spatiotemporal patterns of neural activity from pairwise correlations. *Phys Rev Lett*, 102, 138101.
- MARTIGNON, L., DECO, G., LASKEY, K., DIAMOND, M., FREIWALD, W., & VAADIA, E. (2000). Neural coding: Higher-order temporal patterns in the neurostatistics of cell assemblies. *Neural Comput*, 12, 2621–2653.

- MARTINEZ-CONDE, S., MACKNIK, S. L., & HUBEL, D. H. (2004). The role of fixational eye movements in visual perception. *Nat Rev Neurosci*, 5, 229–240.
- MEISTER, M., LAGNADO, L., & BAYLOR, D. A. (1995). Concerted signaling by retinal ganglion cells. *Science*, 270, 1207–1210.
- NAKAHARA, H., & AMARI, S.-I. (2002). Information-geometric measure for neural spikes. *Neural Comput*, 14, 2269–2316.
- NARAYANAN, N. S. (2005). Redundancy and synergy of neuronal ensembles in motor cortex. *J Neurosci*, 25, 4207–4216.
- NIRENBERG, S., CARCIERI, S. M., JACOBS, A. L., & LATHAM, P. E. (2001). Retinal ganglion cells act largely as independent encoders. *Nature*, 411, 698–701.
- OHORHENUAN, 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, 617–621.
- OLSHAUSEN, B. A., & FIELD, D. J. (1997). Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Res*, 37, 3311–3325.
- OSBORNE, L. C., LISBERGER, S. G., & BIALEK, W. (2005). A sensory source for motor variation. *Nature*, 437, 412–416.
- PALMER, S. E., MARRE, O., & BERRY II, M. J. (2013). Predictive information in a sensory population. arXiv: 1307.0225v1
- PANZERI, S., SCHULTZ, S. R., TREVES, A., & ROLLS, E. T. (1999). Correlations and the encoding of information in the nervous system. *Proc R Soc Lond B Biol Sci*, 266, 1001–1012.
- PERKEL, D. H., & BULLOCK, T. H. (1968). Neural coding. *Neurosciences Research Program Bulletin*, 6, 221–348.
- 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, 995–999.
- PITKOW, X., & MEISTER, M. (2012). Decorrelation and efficient coding by retinal ganglion cells. *Nat Neurosci*, 15, 628–635.
- POLA, G., THIELE, A., HOFFMANN, K. P., & PANZERI, S. (2003). An exact method to quantify the information transmitted by different mechanisms of correlational coding. *Network*, 14, 35–60.
- POUGET, A., DAYAN, P., & ZEMEL, R. (2000). Information processing with population codes. *Nat Rev Neurosci*, 1, 125–132.
- PUCHALLA, J. L., SCHNEIDMAN, E., HARRIS, R. A., & BERRY, M. J. (2005). Redundancy in the population code of the retina. *Neuron*, 46, 493–504.
- QUIROGA, R. Q., & PANZERI, S. (2009). Extracting information from neuronal populations: Information theory and decoding approaches. *Nat Rev Neurosci*, 10, 173–185.
- REICH, D. S., MECHLER, F., & VICTOR, J. D. (2001). Independent and redundant information in nearby cortical neurons. *Science*, 294, 2566–2568.
- REICH, D. S., VICTOR, J. D., KNIGHT, B. W., OZAKI, T., & KAPLAN, E. (1997). Response variability and timing precision of neuronal spike trains in vivo. *J Neurophysiol*, 77, 2836–2841.
- RIEHLE, A. (1997). Spike synchronization and rate modulation differentially involved in motor cortical function. *Science*, 278, 1950–1953.
- RIEKE, F., & BAYLOR, D. A. (1998). Single-photon detection by rod cells of the retina. *Rev Mod Phys*, 70, 1027.
- RIEKE, F., WARLAND, D. K., de RUYTER VAN STEVENINCK, R. R., & BIALEK, W. (1997). *Spikes: Exploring the neural code*. Cambridge, MA: MIT Press.
- ROMO, R., HERNÁNDEZ, A., ZAINOS, A., & SALINAS, E. (2003). Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. *Neuron*, 38, 649–657.
- SCHNEIDMAN, E., BERRY, M. J., SEGEV, R., & BIALEK, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*, 440, 1007–1012.
- SCHNEIDMAN, E., BIALEK, W., & BERRY, M. J. (2003a). Synergy, redundancy, and independence in population codes. *J Neurosci*, 23, 11539–11553.
- SCHNEIDMAN, E., FREEDMAN, B., & SEGEV, I. (1998). Ion channel stochasticity may be critical in determining the reliability and precision of spike timing. *Neural Comput*, 10, 1679–1703.
- SCHNEIDMAN, E., PUCHALLA, J. L., SEGEV, R., HARRIS, R. A., BIALEK, W., & BERRY, M. J. (2011). Synergy from silence in a combinatorial neural code. *J Neurosci*, 31, 15732–15741.
- SCHNEIDMAN, E., STILL, S., BERRY, M. J., & BIALEK, W. (2003b). Network information and connected correlations. *Phys Rev Lett*, 91, 238701.
- SCHNITZER, M. J., & MEISTER, M. (2003). Multineuronal firing patterns in the signal from eye to brain. *Neuron*, 37, 499–511.
- SEUNG, H. S. (2003). Learning in spiking neural networks by reinforcement of stochastic synaptic transmission. *Neuron*, 40, 1063–1073.
- SHAMIR, M., & SOMPOLINSKY, H. (2004). Nonlinear population codes. *Neural Comput*, 16, 1105–1136.
- SHLENS, J., FIELD, G. D., GAUTHIER, J. L., GRESCHNER, M., SHER, A., LITKE, A. M., & CHICHILNISKY, E. J. (2009). The structure of large-scale synchronized firing in primate retina. *J Neurosci*, 29, 5022–5031.
- SHLENS, J., FIELD, G. D., GAUTHIER, J. L., GRIVICH, M. I., PETRUSCA, D., SHER, A., ... CHICHILNISKY, E. J. (2006). The structure of multi-neuron firing patterns in primate retina. *J Neurosci*, 26, 8254–8266.
- SHOHAM, S., O'CONNOR, D. H., & SEGEV, R. (2006). How silent is the brain: Is there a “dark matter” problem in neuroscience? *J Comp Physiol A*, 192, 777–784.
- SIMMONS, J. A., FERRAGAMO, M., MOSS, C. F., STEVENSON, S. B., & ALTES, R. A. (1990). Discrimination of jittered sonar echoes by the echolocating bat, *Eptesicus fuscus*: The shape of target images in echolocation. *J Comp Physiol A*, 167, 589–616.
- SIMONCELLI, E. P., & OLSHAUSEN, B. A. (2001). Natural image statistics and neural representation. *Annu Rev Neurosci*, 24, 1193–1216.
- SRINIVASAN, M. V., LAUGHLIN, S. B., & DUBS, A. (1982). Predictive coding: A fresh view of inhibition in the retina. *Proc R Soc Lond B Biol Sci*, 216, 427–459.
- STEVENS, C. F., & WANG, Y. (1994). Changes in reliability of synaptic function as a mechanism for plasticity. *Nature*, 371, 704–707.
- STEVENS, C. F., & ZADOR, A. M. (1998). Input synchrony and the irregular firing of cortical neurons. *Nat Neurosci*, 1, 210–217.
- STOPFER, M., BHAGAVAN, S., SMITH, B. H., & LAURENT, G. (1997). Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature*, 390, 70–74.
- STRONG, S. P., KOBERLE, R., DE RUYTER VAN STEVENINCK, R. R., & BIALEK, W. (1998). Entropy and information in neural spike trains. *Phys Rev Lett*, 80, 197.
- TANG, A., JACKSON, D., HOBBS, J., CHEN, W., SMITH, J. L., PATEL, H., ... BEGGS, J. M. (2008). A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. *J Neurosci*, 28, 505–518.
- TIMME, M., WOLF, F., & GEISEL, T. (2002). Coexistence of regular and irregular dynamics in complex networks of pulse-coupled oscillators. *Phys Rev Lett*, 89, 258701.

- TKACIK, G., GRANOT-ATEDGI, E., SEGEV, R., & SCHNEIDMAN, E. (2013). Retinal metric: A stimulus distance measure derived from population neural responses. *Phys Rev Lett*, 110, 058104.
- TKACIK, G., PRENTICE, J. S., BALASUBRAMANIAN, V., & SCHNEIDMAN, E. (2010). Optimal population coding by noisy spiking neurons. *Proc Natl Acad Sci USA*, 107, 14419–14424.
- TREVES, A., & PANZERI, S. (1995). The upward bias in measures of information derived from limited data samples. *Neural Comput*, 7, 399–407.
- TRUCCOLO, W., HOCHBERG, L. R., & DONOGHUE, J. P. (2009). Collective dynamics in human and monkey sensorimotor cortex: Predicting single neuron spikes. *Nat Neurosci*, 13, 105–111.
- VADIA, E., HAALMAN, I., ABELES, M., BERGMAN, H., PRUT, Y., SLOVIN, H., & AERTSEN, A. (1995). Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*, 373, 515–518.
- VAN HATEREN, J. H. (1992). A theory of maximizing sensory information. *Biol Cybernetics*, 68, 23–29.
- VAN RULLEN, R., & THORPE, S. J. (2001). Rate coding versus temporal order coding: What the retinal ganglion cells tell the visual cortex. *Neural Comput*, 13, 1255–1283.
- VAN VREESWIJK, C. A., & SOMPOLINSKY, H. (1996). Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274, 1724–1726.
- WARLAND, D. K., REINAGEL, P., & MEISTER, M. (1997). Decoding visual information from a population of retinal ganglion cells. *J Neurophysiol*, 78, 2336–2350.
- WARZECHA, A. K., & EGELHAAF, M. (1999). Variability in spike trains during constant and dynamic stimulation. *Science*, 283, 1927–1930.
- WHITE, J. A., RUBINSTEIN, J. T., & KAY, A. R. (2000). Channel noise in neurons. *Trends Neurosci*, 23, 131–137.