

Published in final edited form as:

Curr Opin Neurobiol. 2007 August ; 17(4): 397–400. doi:10.1016/j.conb.2007.07.002.

Analyzing the activity of large populations of neurons: how tractable is the problem?

Sheila H Nirenberg^{1,3} and Jonathan D Victor^{2,3}

¹ Department of Physiology and Biophysics, Weill Medical College of Cornell University, 1300 York Avenue, New York, NY 10021, USA

² Department of Neurology and Neuroscience, Weill Medical College of Cornell University, 1300 York Avenue, New York, NY 10021, USA

³ Institute for Computational Biomedicine, Weill Medical College of Cornell University, 1300 York Avenue, New York, NY 10021, USA

Abstract

Understanding how the brain performs computations requires understanding neuronal firing patterns at successive levels of processing—a daunting and seemingly intractable task. Two recent studies have made dramatic progress on this problem by showing how its dimensionality can be reduced. Using the retina as a model system, they demonstrated that multineuronal firing patterns can be predicted by pairwise interactions.

Review

One of the most challenging problems we face in systems neuroscience is understanding how the brain performs computations. Understanding this means, essentially, understanding how the brain takes a set of inputs and transforms it into a set of outputs. When we study vision, for example, we do this at many levels. We present a visual stimulus (an input) to the photoreceptors and try to determine how it is transformed into a pattern of action potentials (an output) at the level of the ganglion cells, then how the pattern of action potentials at the level of the ganglion cells is transformed into a pattern of action potentials at the level of the lateral geniculate nucleus, then various levels of cortex, until, finally, a behavior is produced.

One of the main reasons this problem has been so difficult is that it requires accurate descriptions of the input and output data at each level. Take the visual system again—at each level (retina, lateral geniculate nucleus, cortex) there are hundreds to thousands of cells in a processing unit, and, at each moment in time, each of these cells is either firing or not firing an action potential. Add to this the fact that the firing of each cell is at least somewhat dependent on the firing of other cells, one can see right away that the problem is very high dimensional.

So how can it be simplified? Two approaches come to mind. One is a top-down approach. For this, one takes the firing patterns produced when an animal performs a task and determines the crucial features (i.e. those that are needed to perform the task) [1]. This provides a way to identify the relevant quantities in the firing patterns (e.g. spike count,

spike timing, temporal correlations), and discard the irrelevant ones. The other way to simplify the problem, the way that is the subject of this review, is to directly parameterize the firing patterns in a low-dimensional way. At first glance, it might seem that any low-dimensional parameterization would be hopelessly inaccurate, but what is exciting, and what we review here, is that it is not. Two recent papers show that a low-dimensional parameterization is dramatically effective—at least at the level of the retina.

The two papers are ‘The structure of multi-neuron firing patterns in primate retina’ by Shlens *et al.* [2**] and ‘Weak pairwise correlations imply strongly correlated network states in a neural population’ by Schneidman *et al.* [3**].

We start with the Shlens *et al.* paper [2**]. In this study, the authors focus on finding the simplest characterization of the response distribution of the output cells of the retina. Previous studies had shown that characterizing the response distribution assuming that the cells are independent would not suffice, that is, would not give the correct frequency of the joint firing events of multiple cells [4-6]. This led them to take the next simplest approach, which was to characterize the response distribution taking pairwise correlations into account. To carry out this, they used a maximum entropy method (see Box 1). Their rationale was that it gives the simplest distribution consistent with the measured pairwise correlations.

Specifically, what Shlens *et al.* [2**] did was analyze the firing patterns of clusters of seven neurons within the retinal mosaic. They tabulated the frequency of each of the possible firing patterns during every 10 ms period. This comes to $2^7 = 128$ firing patterns. (The reason the base is 2 is that ‘firing’ was represented as a binary quantity—that is, each neuron was designated as ‘firing’ or ‘not firing’ during each 10 ms epoch.) The question they asked then was: Can they capture the complete set of firing pattern frequencies using a much smaller number of parameters—the firing frequency of each cell by itself (7 parameters, 1 for each cell) and the frequencies of the pairwise firings (21 parameters, 1 for each pair of cells)? They found, remarkably, that the maximum entropy distribution constructed this way was virtually indistinguishable from the true 128-pattern distribution. Even more remarkably, when they performed the same analysis using only the nearest neighbor pairs, they could also recover the original distribution. (See Box 2 for how they quantified the level of agreement between the predictions of the pairwise model and the observed frequencies of the multineuronal firing events.) For a discussion on why local interactions between pairs of neurons can account for global patterns of activity, see Figure 1.

Box 1

The maximum entropy principle

The maximum entropy principle [19] (see references [20,21] for pioneering applications to neuroscience) is a method to create a probability distribution from a limited set of measurements. The basic idea is to determine the most random probability distribution consistent with a set of constraints. Randomness is measured by ‘entropy’: if P_j indicates the probability that event j occurs (with $\sum P_j = 1$), then the entropy is defined by $H = -\sum P_j \log P_j$. The maximum entropy distribution is the set of probabilities P_j that maximize H , subject to a set of specified constraints.

Since a maximum entropy distribution is as random as possible given a set of constraints, it provides null hypotheses for quantities not explicitly constrained. Shlens *et al.* [2**] and Schneidman *et al.* [3**] showed that the empirical frequency of multineuronal firing patterns in the retina was consistent with the maximum entropy null hypothesis generated by their pairwise firing frequencies.

Maximum entropy distributions arise in many familiar contexts. A Gaussian distribution is the maximum entropy distribution constrained by a specified mean and variance. A Poisson process is the maximum entropy distribution of events, constrained by an overall firing rate. A product distribution of several variables is the maximum entropy multivariate distribution constrained by the marginal distribution of the individual variables. The Wiener systems-analytic formalism [22] (and, the special cases of spike-triggered averaging and covariance) amount to modeling a system's input-output distribution as a maximum entropy distribution constrained by the input-output correlations and the response variance [23].

Note that this kind of analysis was possible because of the recording techniques Shlens *et al.* [2**] used. With a multielectrode array, they were able to record from nearly all the ON and OFF cells of one cell class in a small region of retina (ON and OFF parasol cells in a 4×8 degree region of peripheral macaque retina). From this, they drew clusters of more or less adjacent cells from either the ON or OFF populations (clusters of 7 cells drawn from 118 ON cells and 175 OFF cells).

The implications of this analysis are most striking when we consider how they apply to a large network of neurons. In a network of N neurons, there are 2^N multineuronal firing patterns whose frequencies must be explained. For $N = 20$, this is over a million; for $N = 50$, it is astronomical. In a pairwise model, $N + N(N - 1)/2$ parameters (the individual neurons' firing frequencies and their pairwise firing frequencies) suffice to account for the multineuronal frequencies; the parameter count has now been reduced substantially (210 for $N = 20$, 1275 for $N = 50$). The further reduction provided by the nearest neighbor model reduces the parameter count to approximately $4N$: Each of the N neurons has approximately 6 nearest neighbors in a retinal mosaic, so there are $6N/2 = 3N$ nearest neighbor interactions and N individual firing rates. The number of parameters that need to be measured is then proportional to the number of neurons. That is, if nearest neighbor interactions determine the full correlational structure, the complexity of the model, and the length of time required to measure its parameters, becomes greatly reduced.

Box 2

How well does a model fit the data?

How can one quantify the agreement of a model for the frequency of global firing patterns (i.e. a probability model) with data? The conceptual problem is that even if one has a correct probability model, one may not be able to predict particular events. For example, even though one knows the correct probability of all poker hands, one cannot predict the results of any particular deal. The basic strategy that Shlens *et al.* [2**] used was as follows. For each reduced probability model (e.g. a maximum entropy model based on frequencies of pairwise interactions), they calculated the likelihood that the model would predict the observed set of global firing patterns. They then determined the ratio of the likelihood for a reduced model to the likelihood for the full empirical model (i.e. one based on the experimentally determined frequency of each firing pattern).

This likelihood ratio indicates how difficult it is to distinguish the predictions of the reduced model from that of the full empirical model. For example, Shlens *et al.* [2**] determined that it would take about a minute of data before the likelihood of the pairwise model and the likelihood of the full empirical model would differ by a factor of two.

The logarithm of the likelihood ratio is closely related to the Kullback-Leibler divergence [24], which is a principled way to measure how difficult it is to distinguish two

probability distributions. Schneidman *et al.* [3**] used a conceptually similar strategy based on the related Jensen–Shannon divergence [24].

Schneidman *et al.* [3**] performed a similar analysis and also found that maximum entropy models based on pairwise interactions could account for retinal ganglion cell firing patterns. Their paper, though, uses the analysis for an additional purpose—to make inferences about the neural code. Here we briefly discuss what the inferences are and whether or not they are justified.

The main proposal the authors put forth is that the pairwise interactions present in the firing patterns imply that the retina uses an error-correcting code. Briefly, an error-correcting code is one in which the signals in a system are correlated so that messages can be correctly decoded in the face of noise (see Box 3 for an example).

The possibility of error correction occurring in the retina is reasonable. Ganglion cell firing is noisy, so one could imagine that some correlation in the system might be useful for ensuring reliable information transmission. (The benefits, though, must be weighed against a potential loss of efficiency [7].) What the authors propose, though, is that there is much more correlation than one would have thought, and from this they conclude that error correction is a dominating coding mechanism.

There are two problems with this proposal. The first concerns the method they used to determine the extent of the correlations, and the second concerns the source of the correlations.

Box 3

Error-correcting codes

An error-correcting code is a means for obtaining reliable signaling from unreliable components. To see how an error-correcting code works, we consider a toy sensory system and environment. The environment consists of a stimulus that is present half of the time. The sensory system is built out of neurons that signal the presence or absence of the stimulus by the presence or absence of a spike—but do so incorrectly with probability P . If there is only one neuron in the sensory system, then the probability of an incorrect message is P . The error probability can be reduced by adding identical neurons that observe the same stimulus and by decoding the message by a majority vote. For example, with three neurons, errors will occur only when all neurons signal incorrectly (probability P^3) or when two neurons are in error and one is not (probability $3P^2(1 - P)$). The total chance of an error ($P^3 + 3P^2(1 - P)$) is less than P . With sufficiently many neurons, this majority-rule code can achieve any desired degree of accuracy.

To determine the extent of the correlations, Schneidman *et al.* [3**] used an information-theoretic measure. They applied the measure to analyses of subsets of N observed neurons within the retinal mosaic. For each N -neuron subset, they compared an estimate of the entropy of the observed pattern of responses, S_N , with the entropy of a network of independent neurons, S_1 . The difference between these two quantities, $I_N = S_1 - S_N$, is a measure of the total amount of correlation. Correlations are considered to dominate network behavior when I_N approaches S_1 because at this point, the correlations are so strong so as to reduce the entropy of observed firing patterns, S_N to 0.

Using the pairwise maximum entropy model, Schneidman *et al.* [3**] determined these quantities for small N and extrapolated to the whole network. In their extrapolation, I_N

approaches S_1 for $N \sim 200$, that is, correlations dominate. However, the data that form the basis for the extrapolation stop at $N = 15$. At this point, I_N is only one-tenth of S_1 , which means that correlations are far from dominating. Thus, their conclusion that error correction dominates is built on an extrapolation that extends an order of magnitude beyond the limits of the data. What is troubling is that the rationale for the form of the extrapolation they used is unclear [25*]. Many other extrapolations would also have fit the data but would have led to different conclusions.

The second issue is the source of the correlation. The authors use natural scenes as their stimuli. This makes sense in that the goal is to determine if error correction occurs with behaviorally relevant stimuli. The problem is that natural scene stimuli themselves have correlations [8], and this is not controlled for. As a result, it is not clear whether correlations in the ganglion cell output reflect an error correcting mechanism in the retina, or merely correlations in the stimulus. As the authors agree (Elad Schneidman, personal communication), this could be addressed experimentally.

In sum, two recent studies show that the problem of characterizing the activity of large populations of neurons, at least in the retina, can be dramatically reduced. In the first study, Shlens *et al.* [2**] found, using a maximum entropy model, that the activity could be accounted for by pairwise interactions. Moreover, nearest neighbor interactions sufficed. Schneidman *et al.* [3**] used the same approach and also found that pairwise interactions sufficed. They used their analysis, though, for a different purpose: to assert that the retina uses an error-correcting code. The jury is still out on whether their data support this conclusion.

The success of maximum entropy methods in achieving a striking dimensional reduction opens the door to further applications of this approach, in retina and beyond (Ohiorhenuan *et al.*, abstract III-90, CoSyNe, Salt Lake City, UT, February 2007; Tang *et al.*, abstract II-62, CoSyNe, Salt Lake City, UT, February 2007; Yu *et al.*, abstract I-16, CoSyNe, Salt Lake City, UT, February 2007). Further applications could also consider firing patterns distributed over time [9-16] and how firing patterns might be contingent on the stimulus [4,17,18]. Both kinds of extensions are straightforward, at least in principle.

Acknowledgments

The authors thank Peter Latham for his very helpful comments. This work is supported partly by National Eye Institute Grants 1RO1 EY12978 to S Nirenberg and 2RO1 EY9314 to J Victor.

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
 - of outstanding interest
1. Nirenberg S, Jacobs A, Fridman G, Latham P, Douglas R, Alam N, Prusky G. Ruling out and ruling in neural codes. *J Vis* 2006;6:889.
 - 2••. Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, Litke AM, Chichilnisky EJ. The structure of multi-neuron firing patterns in primate retina. *J Neurosci* 2006;26:8254–8266. The authors recorded from parasol cell networks in primate retina and showed that a maximum entropy model based solely on nearest neighbor interactions could account for multineuronal firing patterns. [PubMed: 16899720]

- 3••. Schneidman E, Berry MJ II, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 2006;440:1007–1012. The authors recorded from ganglion cell networks in tiger salamander and guinea pig retina and showed that a maximum entropy model in which all pairs of neurons interact could account for multineuronal firing patterns. They extrapolated this model to infer that error correction dominates network behavior during stimulation by natural scenes. See main text for a discussion of concerns about this conclusion. [PubMed: 16625187]
4. Nirenberg S, Carcieri SM, Jacobs AL, Latham PE. Retinal ganglion cells act largely as independent encoders. *Nature* 2001;411:698–701. [PubMed: 11395773]
5. Mastronarde DN. Correlated firing of cat retinal ganglion cells. I. Spontaneously active inputs to X- and Y-cells. *J Neurophysiol* 1983;49:303–324. [PubMed: 6300340]
6. Mastronarde DN. Correlated firing of cat retinal ganglion cells. II. Responses of X- and Y-cells to single quantal events. *J Neurophysiol* 1983;49:325–349. [PubMed: 6300341]
7. Atick JJ, Redlich AN. Towards a theory of early visual processing. *Neural Comput* 1990;2:308–320.
8. Field DJ. Relations between the statistics of natural images and the response properties of cortical cells. *J Opt Soc Am [A]* 1987;4:2379–2394.
9. Abeles, M. *Corticonics: Neural Circuits of the Cerebral Cortex*. Cambridge UP; Cambridge: 1991.
10. Oram MW, Wiener MC, Lestienne R, Richmond BJ. Stochastic nature of precisely timed spike patterns in visual system neuronal responses. *J Neurophysiol* 1999;81:3021–3033. [PubMed: 10368417]
11. Mao BQ, Hamzei-Sichani F, Aronov D, Froemke RC, Yuste R. Dynamics of spontaneous activity in neocortical slices. *Neuron* 2001;32:883–898. [PubMed: 11738033]
12. Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, Yuste R. Synfire chains and cortical songs: temporal modules of cortical activity. *Science* 2004;304:559–564. [PubMed: 15105494]
13. Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsaki G. Organization of cell assemblies in the hippocampus. *Nature* 2003;424:552–556. [PubMed: 12891358]
14. Ji D, Wilson MA. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nat Neurosci* 2007;10:100–107. [PubMed: 17173043]
15. Dragoi G, Buzsaki G. Temporal encoding of place sequences by hippocampal cell assemblies. *Neuron* 2006;50:145–157. [PubMed: 16600862]
16. Robbe D, Montgomery SM, Thome A, Rueda-Orozco PE, McNaughton BL, Buzsaki G. Cannabinoids reveal importance of spike timing coordination in hippocampal function. *Nat Neurosci* 2006;9:1526–1533. [PubMed: 17115043]
17. Samonds JM, Allison JD, Brown HA, Bonds AB. Cooperative synchronized assemblies enhance orientation discrimination. *Proc Natl Acad Sci USA* 2004;101:6722–6727. [PubMed: 15096595]
18. Samonds JM, Zhou Z, Bernard MR, Bonds AB. Synchronous activity in cat visual cortex encodes collinear and cocircular contours. *J Neurophysiol* 2006;95:2602–2616. [PubMed: 16354730]
19. Jaynes ET. Information theory and statistical mechanics. *Phys Rev* 1957;106:620–630.
20. Bohte SM, Spekreijse H, Roelfsema PR. The effects of pair-wise and higher order correlations on the firing rate of a post-synaptic neuron. *Neural Comput* 2000;12:153–179. [PubMed: 10636937]
21. Nakahara H, Amari S. Information-geometric measure for neural spikes. *Neural Comput* 2002;14:2269–2316. [PubMed: 12396564]
22. Wiener, N. *Nonlinear Problems in Random Theory*. Wiley; New York: 1958.
23. Victor JD, Johannesma P. Maximum-entropy approximations of stochastic nonlinear transductions: an extension of the Wiener theory. *Biol Cybern* 1986;54:289–300. [PubMed: 3741900]
24. Cover, TM.; Thomas, JA.; Schilling, DL. *Elements of Information Theory*. Wiley; New York: 1991.
- 25•. Tkacik G, Schneidman E, Berry MJ II, Bialek W. Ising models for networks of real neurons. *arXiv:q-bio.NC* 2006;0611072v1:1–4. This study showed that pairwise maximum entropy models account for multineuronal correlations in samples of up to 40 retinal ganglion cells. In addition, they simulated behavior of networks of up to 120 retinal ganglion cells. Note that they did not determine whether the crucial extrapolation of Schneidman *et al.* [3••] remained valid.




Figure 1.

Why can local interactions between pairs of neurons account for global patterns of activity among many neurons? The basic reason is that local interactions induce long-range correlations. In the simple example shown here, each neuron is connected to its nearest neighbor by a gap junction. The local interactions between neurons A and B and between neurons B and C induce a longer range correlation between neuron A and C. The connection strength is such that approximately one-third of the spikes of adjacent neurons are synchronous. Consequently, even though there is no direct interaction between neurons A and C, approximately one-ninth of the time that neuron A spikes, so does neuron C. The strength of the long-range correlations induced by nearest neighbor interactions decays exponentially with distance. Note also that even though the physical connections are only pairwise, the interactions often induce joint firing events that involve three (green) or more (blue) neurons. The same principle applies to neural networks with inhibitory local interactions, or mixtures of excitatory and inhibitory interactions, or with a more complex connection topology.