Vision Research 70 (2012) 44–53

Contents lists available at SciVerse ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Determining the role of correlated firing in large populations of neurons using white noise and natural scene stimuli

Marsha Meytlis[1], Zachary Nichols[1], Sheila Nirenberg*

Department of Physiology and Biophysics, Weill Medical College of Cornell University, New York, NY 10065, United States

ARTICLE INFO

Article history:
Received 19 January 2012
Received in revised form 13 July 2012
Available online 3 August 2012

Keywords:
Correlations
Noise correlation
Retinal ganglion cell
Parallel processing
Information
Vision
Population coding
Decoding
Spike trains

ABSTRACT

The role of correlated firing in representing information has been a subject of much discussion. Several studies in retina, visual cortex, somatosensory cortex, and motor cortex, have suggested that it plays only a minor role, carrying <10% of the total information carried by the neurons (Gawne & Richmond, 1993; Nirenberg et al., 2001; Oram et al., 2001; Petersen, Panzeri, & Diamond, 2001; Rolls et al., 2003). A limiting factor of these studies, however, is that they were carried out using pairs of neurons; how the results extend to large populations was not clear. Recently, new methods for modeling network firing patterns have been developed (Nirenberg & Pandarinath, 2012; Pillow et al., 2008), opening the door to answering this question for more complete populations. One study, Pillow et al. (2008), showed that including correlations increased information by a modest amount, ∼20%; however, this work used only a single retina (primate) and a white noise stimulus. Here we performed the analysis using several retinas (mouse) and both white noise and natural scene stimuli. The results showed that correlations added little information when white noise stimuli were used (∼13%), similar to Pillow et al.'s findings, and essentially no information when natural scene stimuli were used. Further, the results showed that ignoring correlations did not change the quality of the information carried by the population (as measured by comparing the full pattern of decoding errors). These results suggest generalization: the pairwise analysis in several species show that correlations account for very little of the total information. Now, the analysis with large populations in two species show a similar result, that correlations still account for only a small fraction of the total information, and, most significantly, the amount is not statistically significant when natural stimuli are used, making rapid advances in the study of population coding possible.

1. Introduction

In the last several years, there has been a great deal of interest in whether correlations in spiking patterns carry important information (e.g., Meister, Lagnado, & Baylor, 1995; Nirenberg et al., 2001; Schneidman et al., 2006; Wu, Nakahara, & Amari, 2001; also see reviews by Latham and Nirenberg (2005) and Averbeck, Latham, and Pouget (2006)). The question arises frequently because the answer has critical bearing on the research approaches that can be used to understand population coding. If these correlations do carry information, then direct, i.e., brute force, approaches for characterizing population activity cannot be used: one simply cannot find the mapping from stimulus to response, as such a mapping would require estimating response distributions in high dimensions – at least N dimensions for N neurons. For populations of more than 3 or 4 cells, the amount of data needed would be impossibly large, and one would have to turn to indirect approaches, such as estimating the response distributions with the correlational structure modeled parametrically.

In contrast, if correlations do not carry unique information, direct approaches become viable even for large populations. Under these conditions, one can characterize the population response distributions from the single neuron distributions. This latter scenario would allow much more rapid advances in the field of population coding.

Much of the work addressing this issue has focused on pairwise analyses. These studies, which include a broad range of neural areas, showed that correlations carry little information – less than 10% of the total information carried by each pair (Gawne & Richmond, 1993; Nirenberg et al., 2001; Oram et al., 2001; Petersen, Panzeri, & Diamond, 2001; Rolls et al., 2003). How this result scales with population size, however, is still a subject of debate. One possibility is that as population size increases, each pair will continue to contribute about the same amount of information. Since the number of pairs is proportional to the square of the number of neurons, this scaling behavior predicts that for large populations,

---

* Corresponding author. Address: Department of Physiology and Biophysics, Weill Medical College of Cornell University, 1300 York Ave., New York, NY 10065, United States.

E-mail address: shn2010@med.cornell.edu (S. Nirenberg).

[1] These authors contributed equally to the work.

0042-6989/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.visres.2012.07.007

correlations could become a substantial or even dominant carrier of information. Another alternative is that as the size of the population grows, only some pairs contribute, or the contributions of the individual pairs become redundant. In these scenarios, the correlations would remain a small contributor to the total amount of information.

Recently, methods for modeling the firing patterns of retinal ganglion cells to white noise stimuli have been developed and used to address this (Pillow et al., 2008). Analysis of the importance of correlations using these models showed that correlations increased information by a relatively small amount, ~20%; however, the work used only a single retina and a white noise stimulus; no natural stimuli were used.

Here we performed the analysis using several retinas, large populations of cells, and both white noise and natural stimuli. The results showed that correlations added little information when white noise stimuli were used (~13%), similar to the results of Pillow et al. (2008), and essentially no information when natural stimuli were used. Thus, the second alternative is more likely the correct one: correlations are a relatively minor contributor to the information carried by populations of neurons – not only for neuronal pairs, but also for whole populations. This is also consistent with a smaller population study with natural scenes in salamander (Oizumi et al., 2010). Thus, correlations likely play a different role in network functioning (e.g., reinforcing network learning (reviewed in Feldman (2009)) or shaping network development (reviewed in Blankenship and Feller (2010))).

2. Methods

2.1. Defining correlations

Two types of correlations are commonly referred to in the literature. One is called "noise correlation" (Gawne & Richmond, 1993) and is the focus of this paper. Neural responses $\boldsymbol{r} = (r_1, \ldots, r_n)$ are noise-correlated if and only if

$$p(\boldsymbol{r}|x) \neq \prod_{i=1}^{n} p(r_i|x),$$

where $(r_1, \ldots, r_n)$ are the individual neural responses that constitute the population response $\boldsymbol{r}$ to the stimulus x.

The second type is called "signal correlation" (Gawne & Richmond, 1993) and differs from noise correlation in that it takes the average over all stimuli. Neural responses are signal-correlated if and only if

$$p(\boldsymbol{r}) \neq \prod_{i=1}^{n} p(r_i).$$

To provide intuition for what these two types of correlations are, we give an example, following from Nirenberg and Latham (2003). Suppose one presents a flash of light while recording from two ON-type ganglion cells that lie far apart on the retina (such that their receptive fields do not overlap). Because the cells are both ON cells, they will both fire at the onset of the flash. The similarity in their response is an example of signal correlations, and its role in neural coding is clear and not disputed. If, though, the two cells are close enough to receive common input from presynaptic cells (e.g., common photoreceptors, amacrine cells, etc.), then they would show correlations above and beyond the signal correlations. These extra correlations are the noise correlations; their contribution to the information carried by the cells has become the subject of much debate and is the focus of this paper.

2.2. Stimuli

The retinas were stimulated with two photopic, grayscale stimuli of identical luminance and contrast: binary spatio-temporal white noise (WN) and a grayscale natural scene movie (NS). The natural scene movie was recorded in New York City's Central Park, and had a temporal power spectrum of $1/f^{2.04}$, where f is temporal frequency, and a spatial power spectrum of $1/\omega^{2.09}$, where ω is spatial frequency. Both were presented at 15 Hz, using an LCD projector driven by a computer running custom software on a real-time version of Red Hat Linux. Luminance was 0.24 $\mu\text{W/cm}^2$ on the retina (in the photopic range); root-mean-squared contrast was 0.087 $\mu\text{W/cm}^2$. The white noise stimulus covered 10×9 squares (with each square corresponding to 160×160 μm on the retina); the natural movie stimulus covered 20×18 squares (with each square covering 80×80 μm on the retina). For each stimulus, we had a training set, which was used to fit model parameters, and a testing set, which was used for evaluating the models and making calculations; the latter were referred to as the "out-of-sample stimuli".

2.3. Measuring degree of correlation

The degree of correlation was measured using the excess correlated fraction (ECF) following Nirenberg et al. (2001). For each pair of cells, the ECF was calculated as follows: first, the "raw fraction" of correlated spikes was determined. This was the number of spikes that occurred within 1 ms of each other divided by the total number of spikes produced by the pair. A second quantity, the "shifted fraction", was then determined. It was obtained by pairing responses from the two cells when they were presented with the stimulus at different times, i.e., when their responses were shifted by one repeat relative to each other (Perkel, Gerstein, & Moore, 1967). The shifted fraction was then calculated by counting the number of spikes in the shifted pair that occurred within 1 ms of each other and dividing this by the total number of spikes for the pair. The ECF is then the difference between the raw and shifted fraction.

Shift-corrected cross-correlograms were generated in a manner similar to that described above for obtaining the ECF. Briefly, for each pair of neurons, the raw cross-correlogram was first determined from the two cells' simultaneously-recorded responses. The "shift predictor" was calculated from their responses recorded on separate repeats, and was then subtracted from the raw cross-correlogram to yield the shift-corrected cross-correlogram (Perkel, Gerstein, & Moore, 1967).

2.4. Independent and coupled models

Two models were constructed from the neural responses: one in which the neurons were treated as independent, and one in which coupling among the neurons was included. Each model consisted of a set of parameters that were fit by maximizing the log-likelihood of observed spiking data collected for 10 min using the training stimulus set (see Section 2.2).

In the independent model (Fig. 2A), the mth neuron's firing rate was modeled by

$$\lambda_m^{independent}(t) = N_m\left[(X * L_m(t)) + \sum_i H_m(t - \tau_m(i))\right], \tag{1}$$

where X is the stimulus, * denotes spatiotemporal convolution, L_m is the spatiotemporal impulse response corresponding to the linear filter for the mth neuron, and N_m is a function that describes its nonlinearity. The nonlinearities N_m were parameterized as cubic spline functions with six knots. Knots were spaced to cover the range of