



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

Marsha Meytlis¹, Zachary Nichols¹, 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.

© 2012 Elsevier Ltd. All rights reserved.

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 Averbach, 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.

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 $\mathbf{r} = (r_1, \dots, r_n)$ are noise-correlated if and only if

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

where (r_1, \dots, r_n) are the individual neural responses that constitute the population response \mathbf{r} to the stimulus \mathbf{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

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

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 $0.24 \mu\text{W}/\text{cm}^2$ on the retina (in the photopic range); root-mean-squared contrast was $0.087 \mu\text{W}/\text{cm}^2$. The white noise stimulus covered 10×9 squares (with each square corresponding to $160 \times 160 \mu\text{m}$ on the retina); the natural movie stimulus covered 20×18 squares (with each square covering $80 \times 80 \mu\text{m}$ 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 m th neuron’s firing rate was modeled by

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

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

values given by the linear filter output of the models. H_m is the spike history filter for the m th neuron, and τ_m is the sequence of spike times for the m th neuron.

In the coupled model (Fig. 2B) (Pillow et al., 2008), the influence of neuron k on neuron m is modeled by convolution of a linear filter $C_{m,k}$ with the sequence of spikes on neuron k . The outputs of these convolutions for all neighbors k are summed together and added to the stimulus filter output for neuron m , and the sum is then fed into the m th neuron's nonlinearity function, N_m . Formally, the model is

$$\lambda_m^{\text{coupled}}(t) = N_m \left[(X * L_m)(t) + \sum_i \left[H_m(t - \tau_m(i)) + \sum_{k \neq m} C_{m,k}(t - \tau_k(i)) \right] \right], \quad (2)$$

where τ_k is the sequence of spikes times of neuron k .

Each neuron's linear filter L_m was assumed to be a product of a spatial function (extending through 5×5 squares), and a temporal function (extending through 18 time bins, 67 ms each). Dimensionality was further reduced by constraining the temporal function to a linear combination of 10 basis functions (raised cosines), as in Nirenberg, Pandarinath, and Ohiorhenuan (2011) and Nirenberg and Pandarinath (2012), following Pillow et al. (2008). Similarly, coupling filters were parameterized by four raised-cosine basis functions (using additional basis functions was not found to improve fitting), and had a temporal extent of 9 ms (slightly longer than the observed correlogram widths).

Parameters were fit using an expectation maximization procedure, as described by Paninski, Pillow, and Lewi (2007). The quantity maximized is the log likelihood of the observed spike trains under the model (for a given stimulus). Because the models we are using produce an output governed by an inhomogeneous Poisson process, the log likelihood of the observed spike trains can be calculated according to the standard formula:

$$Z = \sum_m \left[\sum_i \log[\lambda_m(\tau_m(i))] - \int_{t=0}^{\text{end}} \lambda_m(t) dt \right]. \quad (3)$$

In this equation, λ_m is the instantaneous firing rate of the m th neuron (predicted by either the independent model (Eq. (1)) or the coupled model (Eq. (2))), and Z is the log-likelihood of the population activity in which the m th neuron spikes at times τ_m . Z is a sum of terms, one for each neuron m . The term for the m th neuron (the expression in the larger square brackets) is composed of a sum and an integral. The sum term boosts the log-likelihood if the neuron produces a spike when the model's instantaneous firing rate (λ_m) is high. The integral is a standard penalty term that prevents models from achieving a high log-likelihood merely by predicting high firing rates overall (Paninski, Pillow, & Lewi, 2007).

To find the parameters that maximize the log likelihood in Eq. (3), we used the following procedure. We began by assuming that the nonlinearity N_m was exponential, since in this case, a global maximum for the log likelihood Z (Eq. (3)) is assured (Paninski, Pillow, & Lewi, 2007). After optimizing the linear filters for an exponential nonlinearity (by gradient ascent), the exponential nonlinearity was replaced by a spline, as in Nirenberg and Pandarinath (2012). Final model parameters were then determined by alternating stages of maximizing the log likelihood with respect to (i) the spline coefficients and (ii) the filter parameters, until a maximum was reached.

Models were validated by observing their spike-train prediction performance for 10 min of novel spiking data, that is, data produced from the retina by out-of sample stimuli (stimuli not used to fit the model).

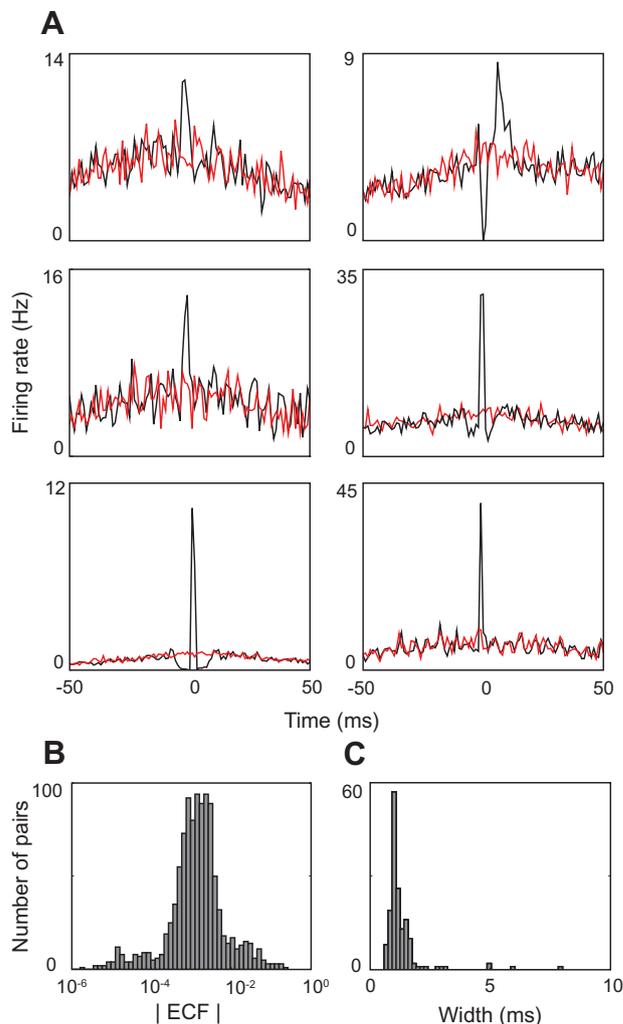


Fig. 1. Degree and timescale of correlated activity for pairs of ganglion cells in the mouse retina. (A) Representative cross-correlograms, showing the range of correlation strengths and timescales. As is standard, each plot shows the firing rate of one cell in the pair relative to the spike times of the cell. *Black* indicates the raw cross-correlogram; *red* indicates the shift-predictor (Perkel, Gerstein, & Moore, 1967). The difference between these two is the noise correlation, the focus of this paper. (B) Distribution of ECFs for all cell pairs in the dataset ($n = 567$ pairs, three retinas). (C) Distribution of half-widths of the central peak of the cross-correlogram for all cell pairs with ECFs $> 0.5\%$ in the dataset ($n = 157$, three retinas). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2.5. Recording

Electrophysiological recordings were obtained in vitro from the isolated mouse retina. Recordings of central retinal ganglion cells (RGCs) were performed on a 64 electrode multi-electrode array using methods described previously (Dedek et al., 2008). Spike waveforms were recorded using a Plexon Instruments Multichannel Neuronal Acquisition Processor (Dallas, TX), and a standard spike sorting method (Fee, Mitra, & Kleinfeld, 1996) was used to identify individual cells. Recordings contained 16–33 cells per retina; cells that had refractory period violations above 5% were excluded.

2.6. Calculating information

For both stimuli (white noise and natural movies) and both models (independent and coupled), we estimated information

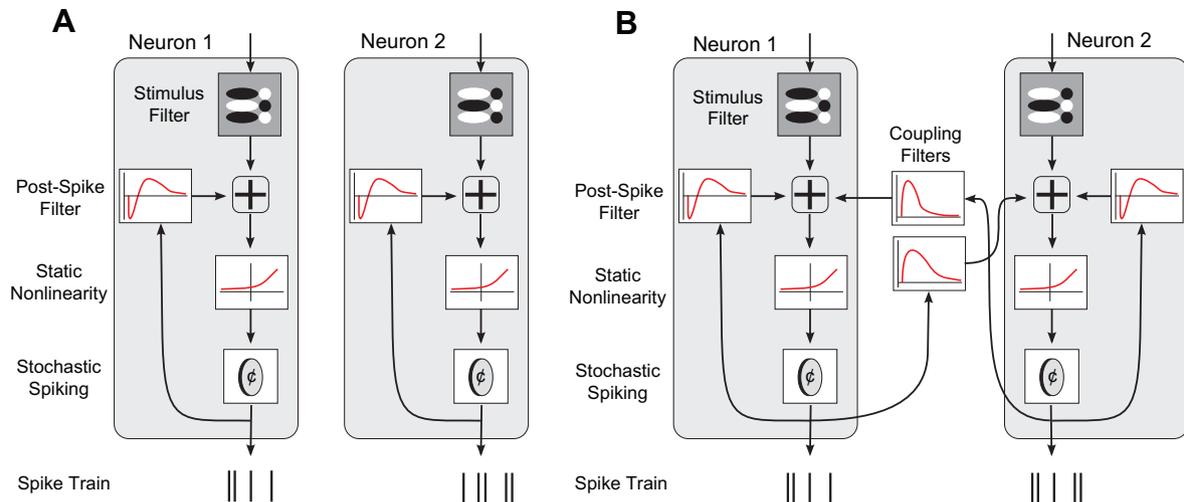


Fig. 2. Schematics of the two input/output models. (A) The independent model. In this model, each neuron within the population is represented by a separate cascade. To generate each neuron's output, the input is convolved with a linear filter. The output of the convolution is then added to the output of a second filter which is convolved with the cells' spiking history. Finally, a static nonlinearity is applied to the combined linear filter output, yielding a rate function for a Poisson spike generator. (B) The coupled model. In the coupled model, the cascades are interconnected – that is, the spiking output of each cell undergoes convolution with neighboring cells' coupling filters, and the result combines additively with the input signal for the neurons. Schematic adapted from Pillow et al., 2008.

rates by decoding the population responses and comparing the accuracy of the decoded stimulus with the actual stimulus. Information calculations were performed from data produced by the test stimuli (i.e., data not used to fit the model parameters).

For white noise, we followed the procedure of Pillow et al. (2008): briefly, we used the Bayes' least squares estimator to decode population responses for both the independent and coupled models. The stimulus variable, x_i , was a 10-element binary vector representing the luminance of a single spatial checker (with spatial location chosen to maximize coverage by the population of recorded cells) over 10 frames. The response variable \mathbf{r} was the set of spike trains recorded across all neurons (r_1, \dots, r_n) for 2 s beginning with the first frame of x_i . For each response \mathbf{r} , we calculated the likelihood that it would have been elicited by each of the 2^{10} possible stimuli x_i . This is the likelihood $\exp(Z)$ given by Eq. (3), where the rate functions on the right-hand side are taken from Eq. (1) for the independent model or from Eq. (2) for the coupled model. It is proportional to $p(\mathbf{r}|x_i)$, the probability that the stimulus x_i would elicit a population response \mathbf{r} . To convert $p(\mathbf{r}|x_i)$ to what we need for decoding, the posterior probability $p(x_i|\mathbf{r})$, we use Bayes' theorem:

$$p(x_i|\mathbf{r}) \propto p(\mathbf{r}|x_i)p(x_i),$$

where in this case $p(x_i)$ is uniform. The Bayes' least squares estimate of the stimulus given the response \mathbf{r} is a probability-weighted sum over all possible stimuli:

$$\hat{x} = \sum p(x_i|\mathbf{r})x_i.$$

Finally, as in Pillow et al. (2008), decoding performance was converted into an estimate of information rate (in bits/s) via

$$H = \log_2 \left[\frac{\det(\langle x_j x_j^T \rangle)}{\det(\langle (\hat{x}_j - x_j)(\hat{x}_j - x_j)^T \rangle)} \right],$$

where $\langle \cdot \rangle$ denotes averaging across all responses. Note that the above equation is an estimate of the log of the signal-to-noise ratio (SNR) between the stimulus x and the decoded estimate \hat{x} . Implicit in this estimate is the assumption that the white noise stimulus can be considered to be approximately Gaussian (see Pillow et al., 2008). There are two reasons that we choose to decode a single spatial pixel: (1) this allows for a direct comparison with Pillow et al. (2008) and (2) it facilitates the information calculations under the

assumptions of a Gaussian channel. We use short snippets for decoding as it allows us to repeat the decoding calculation for many samples of the response, which is a necessary step for computing the SNR.

For natural scene stimuli, we calculated information by an approach that did not require the Gaussian approximation. The stimulus variable x consisted of the starting point of the 150-frame natural movie sequence, that is, x_0 was the movie that started at frame 0, x_1 was the movie that started at frame 1, etc. Each response \mathbf{r} was taken as the spike trains recorded over a single 67 ms segment of the movie. The decoded stimulus was computed as the element from the stimulus set that maximized the posterior probability:

$$\hat{x} = \operatorname{argmax}[p(x_i|\mathbf{r})].$$

Mutual information was then calculated between x and \hat{x} with the standard plug-in estimator of entropy (Antos & Kontoyiannis, 2001). In order to ensure that the stimulus entropy was not a critical factor in our calculations, we also performed this analysis using half of the stimuli (x_0, x_2, x_4 , etc.) and a third of the stimuli (x_0, x_3, x_6 , etc.). As a further check for robustness, we repeated the calculation with the response variable consisting of longer segments of the spike trains, at lengths of 133 ms and 200 ms (the same was done for the decoding used for confusion matrices described below).

For information estimates, 95% confidence intervals were determined by taking 1000 bootstrap resamples of 200 decoded segments for the white noise stimulus conditions, and by taking 200 bootstrap resamples of 60 decoded segments (for each of the stimuli) for natural scenes.

Finally, we used the natural scene stimuli and responses to construct confusion matrices, as in Pandarinath, Victor, and Nirenberg (2010). Briefly, a confusion matrix gives the probability that a neural response to a presented stimulus will be decoded as that stimulus (Hand, 1981). To find these probabilities, we used the same Bayesian approach as for calculating information. Each recorded response \mathbf{r} was decoded by choosing the stimulus x most likely to have produced it, that is, the stimulus for which $p(x_i|\mathbf{r})$ was maximized. Here again, $p(x_i|\mathbf{r})$ was calculated from $p(\mathbf{r}|x_i)$ using Bayes' theorem, where the prior $p(x_i)$ was uniform. For each presentation of stimulus x_i that resulted in a response \mathbf{r} that was decoded as

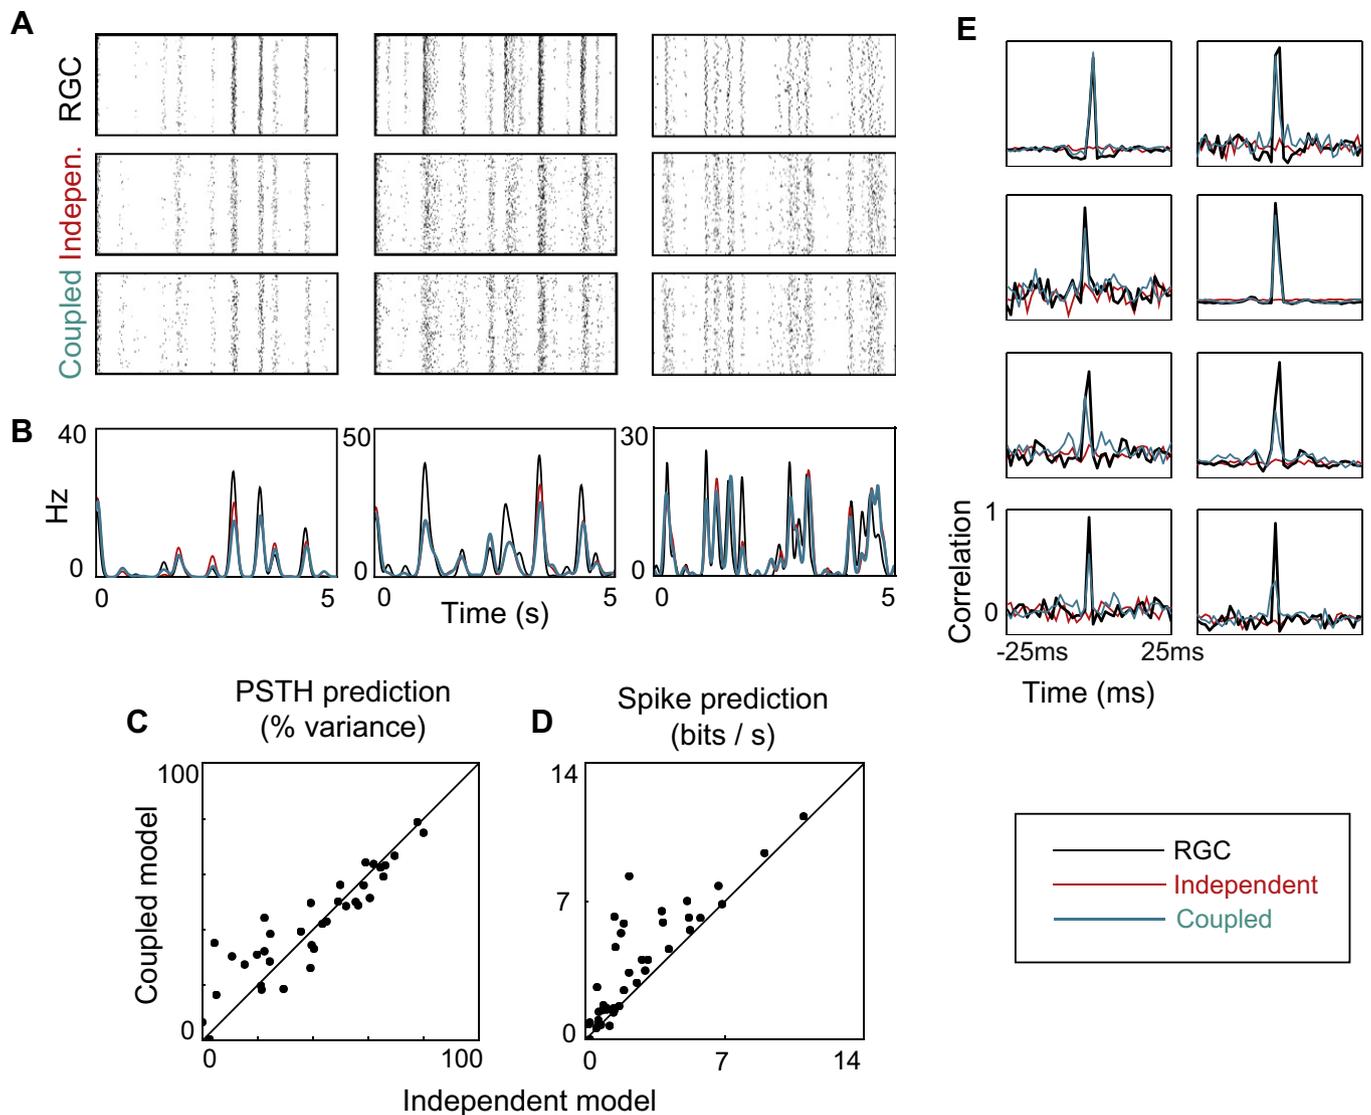


Fig. 3. Model performance for white noise stimuli. (A) Recorded responses from three representative retinal ganglion cells (RGCs) to repeats of an out-of-sample binary white noise stimulus (*top*), along with responses derived from the independent model (*middle*) and the coupled model (*bottom*). (B) Corresponding PSTH's of the ganglion cell responses (*black*), the independent model output (*red*), and the coupled model output (*blue*). (C) Quality of PSTH prediction by independent and coupled models across the three retinas. Coupling yields no advantage in predicting the average response. (D) Log-likelihoods of ganglion cell responses for the independent and coupled models. The coupled model is better able to explain individual spike trains. (E) Representative shift-corrected cross-correlograms for pairs of ganglion cells (*black*) and for the independent (*red*) and coupled models (*blue*). The coupled model is able to capture the correlated activity, while the independent model is not.

stimulus x_j , the entry at position (i, j) in the confusion matrix was incremented.

3. Results

Our goal is to determine the extent to which information transmission by a neuronal population – the output cells of the retina – depends on correlations, specifically noise correlations (see Section 2). Our strategy is to decode the population responses with two decoders: one that ignores correlations and one that exploits them. The decoding is carried out via a Bayesian framework, and therefore requires a model of the relationship between stimulus and response. The model determines how we treat the correlations: when we decode using an “independent” model, we ignore correlations; when we decode using a “coupled” model, we take them into account (see Section 2 for descriptions of the models).

To approach this goal, we proceeded in four steps. First, we measured the degree of correlation in our system, the ganglion cells of the mouse retina. Second, we constructed the two input/

output models and verified that they perform as expected – that is, that the coupled model predicts correlations while the independent model does not. Third, using each model, we decoded out-of-sample stimuli (i.e., the test stimuli), both white noise and natural scenes. Finally, using the decoding results, we compared the information rates for the two models and assessed the quality of the decoding using confusion matrices.

We begin by characterizing the correlations in our model system, the mouse retina (Fig. 1). Fig. 1A shows six representative cross-correlograms, covering the range of strengths and timescales. The lower panels summarize the results for the population: Fig. 1B shows the distribution of correlation strengths, quantified by the excess correlated fraction, ECF (Nirenberg et al., 2001; see Section 2), and Fig. 1C shows the distribution of correlation timescales. The range of ECF's encountered (up to 33%) agrees with that previously reported for mouse retina (up to 34% in Nirenberg et al. (2001), and 33% in Jacobs et al. (2009)) and other mammalian species (up to 27% in cat (Mastrorarde, 1983) and 28% in rabbit (DeVries, 1999)). Similarly, the range of correlation timescales (from <1

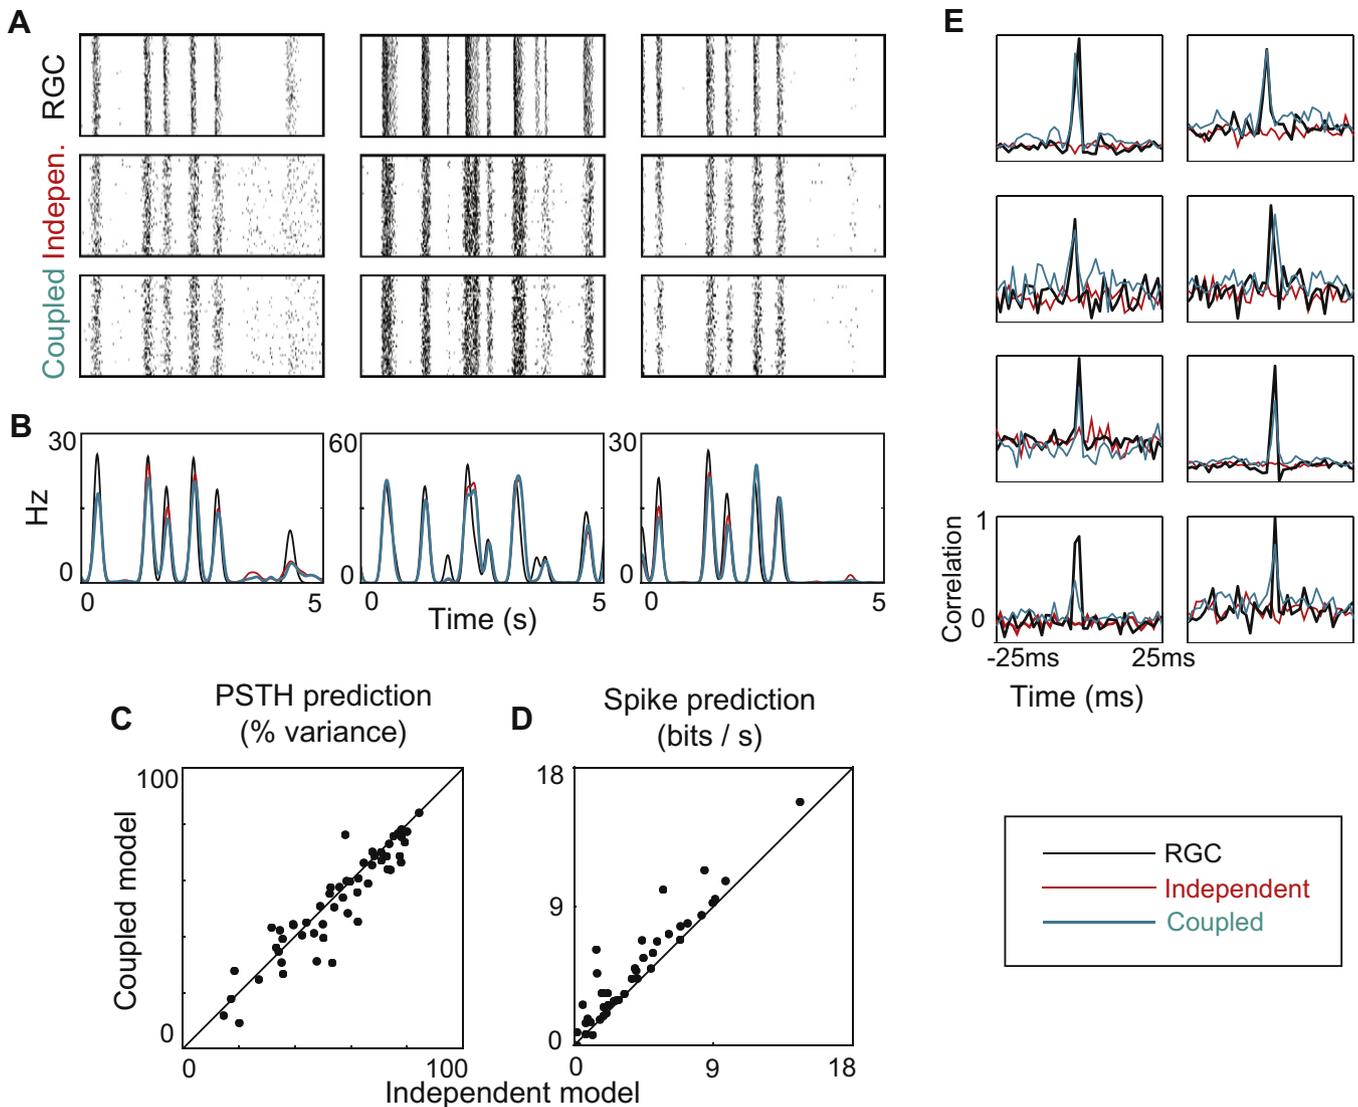


Fig. 4. Model performance for natural scene stimuli. (A) Recorded responses of three representative ganglion cells (RGCs) to repeats of an out-of-sample natural scene stimulus (top), along with responses derived from the independent model (middle) and the coupled model (bottom). (B) Corresponding PSTH's of the ganglion cell responses (black), the independent model output (red), and the coupled model output (blue). (C) Quality of the PSTH prediction by independent and coupled models for the three retinas. Coupling yields no advantage in predicting the average response. (D) Log-likelihoods of ganglion cell responses for the independent and coupled models. The coupled model is better able to explain individual spike trains. (E) Representative shift-corrected cross-correlograms for pairs of ganglion cells (black) and for the independent (red) and coupled models (blue). The coupled model is able to capture the correlated activity, while the independent model is not. All of these findings parallel those of Fig. 3, obtained for white noise stimulation.

to 5 ms, with a tail out to 11 ms) is also consistent with results reported for other mammalian species, including cat, rabbit, and mouse (DeVries, 1999; Mastronarde, 1983; Nirenberg et al., 2001). Note that the range of ECFs in the retina is broad, and their potential impact, as discussed in the literature, is proposed to occur in two ways: through the cumulative effect of the many small correlations that make up the majority of the correlations, or through groups of highly correlated cells, which, though more rare, could also carry significant information (see Chatterjee et al., 2007; Meister, Lagnado, & Baylor, 1995; Nirenberg et al., 2001; Pillow et al., 2008; Schneidman et al., 2006; Wu, Nakahara, & Amari, 2001; also discussed in Latham and Nirenberg (2005) and Averbeck, Latham, and Pouget (2006)).

Having established that our dataset is representative of the typical levels of correlated activity in the retina, we constructed the two retinal input/output models mentioned above – the one that ignores correlations (the independent model), and the one that takes them into account (the coupled model). The two models

are shown schematically in Fig. 2. The independent model (Fig. 2A) consists of a linear filter, a static nonlinearity, and Poisson spike generation for each neuron (i.e., an LNP cascade) and a post-spike filter for modeling refractoriness. The coupled model (Fig. 2B) consists of the same with an added dependence on the firing of other neurons in the population.

In Figs. 5 and 6, we will show the results of using the models to decode stimuli, but first we show that they have the necessary properties to allow us to test the hypothesis that correlations matter. Specifically, we show that both models predict the average responses well, indicating that they are capturing signal correlations, but only the coupled model predicts the detailed firing patterns and their correlations, indicating that it is capturing both signal and noise correlations, the latter being the correlations that are the subject of so much debate.

Fig. 3 shows the evaluation of the two models using the white noise stimulus. Fig. 3A shows rasters from several representative cells produced by the out-of-sample white noise stimulus (the

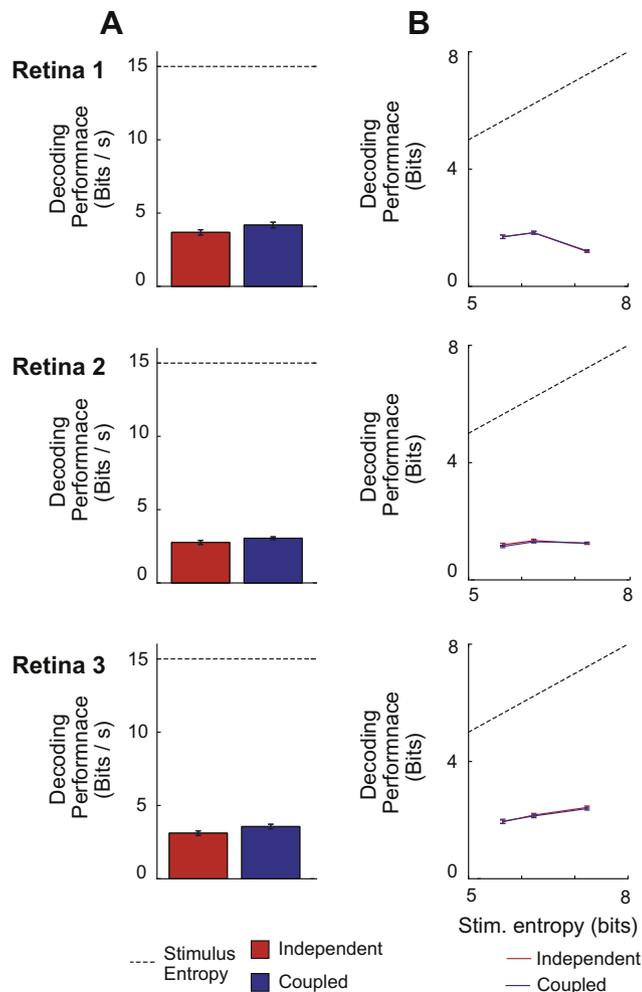


Fig. 5. The contribution of correlations (noise correlations) to information for white noise and natural stimuli. As shown in Figs. 3 and 4, taking correlations into account (via the coupled model) substantially improves the prediction of individual spike trains. However, as shown here, this addition leads to only a modest increase in the amount of information carried. (A) Information decoded from populations using white noise stimuli with an independent (red) and coupled (blue) model. The coupled model (which takes noise correlations into account) adds a small amount of information (13% on average). (B) Information decoded for populations using natural scene stimuli with an independent (red) and coupled (blue) model. To control for effects due to stimulus entropy, we use three levels of task difficulty (corresponding to the number of stimuli to decode). We find that the coupled model does not add significant information for any of the three difficulties. Error bars show 95% confidence intervals. Population size for retina 1 was 17 cells, for retina 2 was 13 cells, and for retina 3 was 26 cells. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

testing set), along with predictions from the independent and coupled models, and Fig. 3B shows the corresponding average peristimulus time histograms (PSTHs).

Fig. 3C and D summarize the results for the whole dataset. As can be seen in the figures, both models predict the average responses equally well (Fig. 3C). Fig. 3D then shows that, in addition, the coupled model predicts the detailed firing patterns better than the independent model. The increase was quantified by calculating the log likelihood of the observed spike trains, given each model (Section 2, Eq. (3)). As shown in Fig. 3D, the coupled model consistently achieves the better performance.

Finally, to confirm that the coupled model was achieving the better performance for the right reason, that is, because it was capturing the noise correlations in the spike trains, we calculated cross-correlograms. Each plot in Fig. 3E shows the shift-corrected cross-correlogram of the true responses of a pair of neurons, along

with the predictions from the independent and coupled models (the shift-corrected correlograms show only the noise correlations, not the full correlograms which also contain signal correlations as in Pillow et al. (2008) and Oizumi et al. (2010)). As shown in the figure, the coupled model provides a good fit to the cross-correlograms, whereas the independent model fails to predict any peaks at all.

Fig. 4 provides a parallel analysis of the two models for natural scene stimuli. As was the case for the responses to white noise, the independent and coupled models provide equally good predictions of the average responses. This is seen in representative rasters (Fig. 4A) and PSTHs (Fig. 4B), and in the summary across all cells in all retinas (Fig. 4C). Also in agreement with the findings for white noise, the coupled model is better able to predict individual spike trains (Fig. 4D) and cross-correlograms (Fig. 4E).

Having shown that the coupled model captures the noise correlations, we now use a Bayesian decoding framework (given in Section 2) to determine how much information these correlations carry.

Fig. 5A shows the results of the decoding analysis for the white noise stimulus in the three retinas. As shown in the figure, capturing the correlations led to only a small increase in decoding performance, about 1 bit/s. On average this amounted to a 13% increase in information.

Fig. 5B extends this analysis to the responses to natural scenes and shows that for these stimuli, the contribution of noise correlations is even less, effectively undetectable. As was the case for white noise, taking noise correlations into account via the coupled model yields superior predictions of spike trains on a trial-by-trial basis (Fig. 4D). But, as shown in Fig. 5B, despite this, taking noise correlations into account yielded no detectable increase in the amount of information that can be decoded about the stimulus. Similar results were obtained when other response lengths were used for decoding (see Section 2). Thus, the information in the correlations is essentially redundant to the information in the independent firing patterns, and including them in the decoding does not add information.

The Supplement shows two additional analyses that further support these conclusions: Fig. S1 shows that the results are robust to goodness of fit (as measured by R^2), and Fig. S2 shows that the results hold when the analysis is restricted to subpopulations of cells with particularly high mutual correlations.

At this point, we have shown that taking noise correlations into account has minimal effect on the amount of information carried by the population. We now take this one step further, and ask whether they change the *kind* of information that is carried. To do this, we constructed confusion matrices corresponding to the information calculations for the natural scene stimuli shown in Fig. 5B. Two confusion matrices are shown: one is constructed with 50 stimuli drawn from the natural scene movies, and the other with 75 stimuli, also drawn from the natural scene movies. The stimuli were 67 ms movie segments. These matrices, shown in Fig. 6, indicate which stimuli are correctly decoded, and which ones are mistaken for another. Rows correspond to the presented stimulus; columns correspond to the decoded stimulus. The intensity of a pixel in row i and column j corresponds to the probability that the neural response to stimulus i is decoded as stimulus j . Thus, perfect decoding would result in a bright diagonal line, and off-diagonal elements represent the pattern of decoding errors.

As is clear from Fig. 6, the pattern of decoding for the independent model is the same as the pattern for the coupled model – that is, the pattern of correct classifications (the on-diagonal elements), and the pattern of errors (the off-diagonal elements) are the same for both models. This finding also held for decoding among 150 stimuli and for decoding with longer response durations (133 ms and 200 ms). We conclude that under natural scene stimulation,

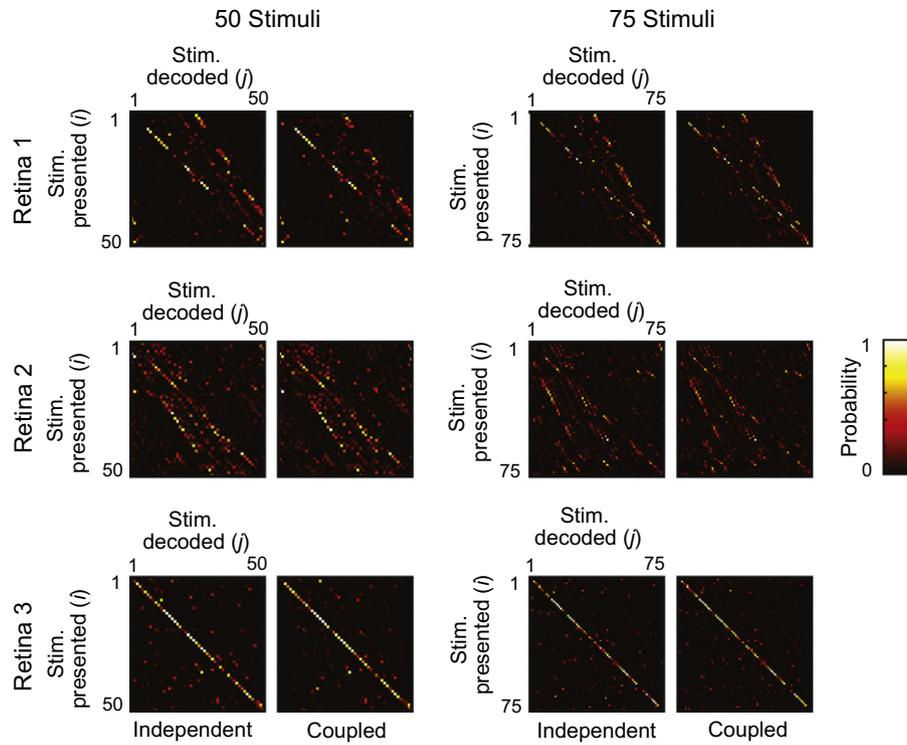


Fig. 6. Characterization of decoding errors on natural scenes via confusion matrices. Two confusion matrices are shown. One was constructed with 50 stimuli drawn from the natural scene movies, and the other with 75 stimuli, also drawn from the natural scene movies. The stimuli were 67 ms movie segments. In each matrix, the rows correspond to each presented stimulus, and the columns correspond to the decoded stimulus. The intensity of a pixel in row i and column j corresponds to the probability that the neural response to stimulus i is decoded as stimulus j . On-diagonal elements indicate correct decodings; off-diagonal elements indicate the pattern of errors. Extending the findings of the information calculation (Fig. 5), the confusion matrices are nearly indistinguishable for the independent and coupled models.

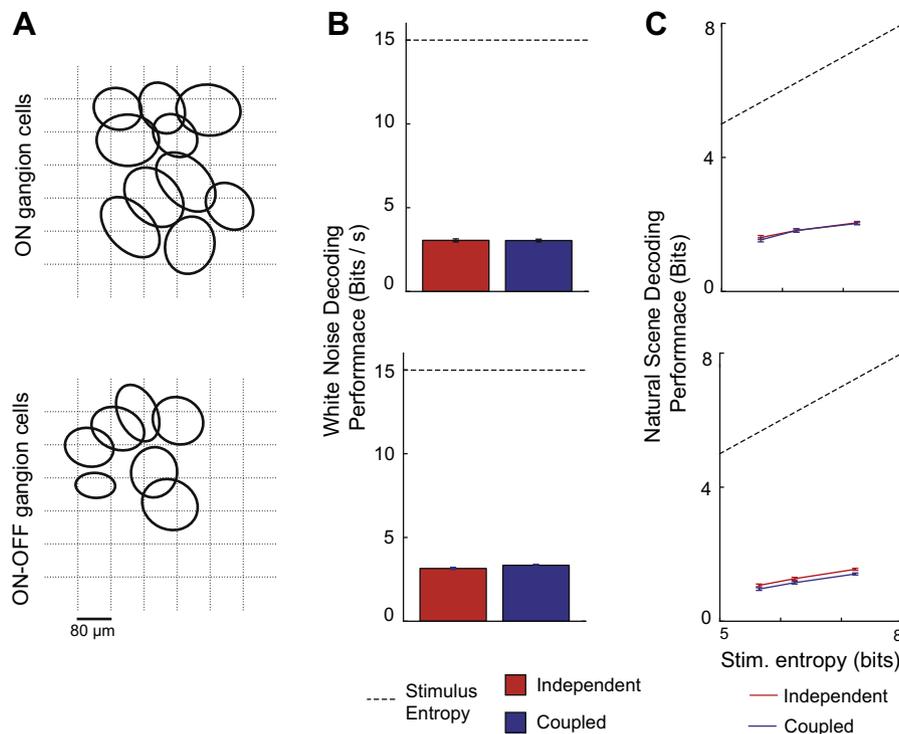


Fig. 7. Information carried by correlations within single-cell mosaics. (A) The receptive field mosaics for a patch of ON cells (top) and a patch of ON-OFF cells (bottom). The receptive fields were calculated by fitting 2-D Gaussian functions to the spatial kernels of the models, then plotting the contour at height $1/e$ relative to the peak. (B) Information decoded for the white noise stimulus. Calculations were made using only the ON cell patch (top) and only the ON-OFF cell patch (bottom). For each of the patches here, the inclusion of correlations led to a very small increase in information (smaller than for the full retinal populations). (C) Information decoded for the natural scene stimulus. For each of the patches here, correlations did not add any significant amount of information, consistent with the results for full populations.

taking noise correlations into account has little impact not only on the amount of information that is present, but also on the kind of information – that is, it has almost no effect on which stimuli are properly decoded, and which are confused.

4. Discussion

The question of whether neural responses can be understood or decoded without taking noise correlations into account is crucial to the study of population coding. If cells cannot be treated as independent units, then a recording from one neuron cannot be understood without taking into account the neurons interacting with it. Determining the correlational structure for large populations of neurons is nontrivial: using brute force methods on populations with more than a few neurons is not possible, since the amount of data needed to characterize the population response scales exponentially with the number of cells in the population. Therefore, alternative approaches, such as characterizing the correlational structure parametrically, would be needed. For this reason, the answer to the question of whether noise correlations carry information has direct and significant impact on the direction of research efforts in the field (Averbeck & Lee, 2004; Nirenberg et al., 2001; Latham & Nirenberg, 2005; Averbeck, Latham, & Pouget, 2006; see Brown, Kass, & Mitra, 2004 in addition for further discussion of the problems of studying large neuronal populations).

Though this question has already been studied for pairs of neurons, finding that noise correlations add little or negligible amounts of information (Gawne & Richmond, 1993; Levine, 2004; Nirenberg et al., 2001; Oram et al., 2001; Petersen, Panzeri, & Diamond, 2001; Rolls et al., 2003), the possibility remained that they could contribute in larger amounts in more complete populations. The concern remained in part because of the matter of scaling: as the number N of neurons in a population grows, the number of pairs of neurons grows in proportion to N^2 . Thus, if the noise correlations from each pair were to contribute non-redundantly to the total information, their contribution would eventually become very large even though the contribution from individual pairs is small. Furthermore, higher-order (non-pairwise) interactions could also contribute in large populations. This study however confirms that neither of these was the case: even in the large populations we used (as high as 26 neurons), noise correlations continued to account for only a small fraction of the total information.

4.1. Correlations do not carry significant information in complete subpopulations

Unlike the two previous population studies (Oizumi et al., 2010; Pillow et al., 2008), our analysis included more than one cell class; this was done to allow correlations among different cell types to be included, such as those between ON cells and ON-OFF cells or OFF cells and ON-OFF cells. To relate our results to previous studies, though, we also performed the analysis on two cell type-specific patches (Fig. 7) – a patch of 10 ON cells and a patch of 7 ON-OFF cells. We found that our results still hold for both of these patches, and, in fact, within these single cell type patches, correlations account for even less of the information than in the full mixed populations, possibly because single cell type mosaics do not include the strong correlations that exist between different cell types with overlapping response properties (DeVries, 1999; Greschner et al., 2011) (e.g., between ON and ON-OFF cells).

4.2. Ignoring correlations does not affect the quality of the information

We emphasize that our analysis went a step beyond determining the contribution of noise correlations to the amount of infor-

mation: we showed that taking noise correlations into account also does not change the *kind* of information carried. We showed this using confusion matrices, which delineate which stimuli were correctly decoded, and exactly which errors were made. As shown in Fig. 6, the confusion matrices that emerge when correlations are taken into account are very nearly identical to the confusion matrices that emerge when they are ignored. Thus, taking noise correlations into account does not change either the amount or the quality of the information carried by retinal spike trains.

5. Conclusions

We have shown that for multiple types of stimuli (white noise and natural scenes), cells in a retinal population can be treated as independent with little or no loss of information. This greatly simplifies, and, therefore, facilitates the study of retinal population coding. The fact that the result holds for natural stimuli is notable because natural stimuli have both local and long range correlations that are not found in white noise, and, as a result, the noise correlations they induce may differ. Hence, generalization of the result from white noise to natural scenes is nontrivial, and, since natural stimuli are the biologically relevant ones, this generalization is critical to know.

One question that naturally arises is whether the conclusions – that population coding can be studied by treating the neurons as independent – will generalize to other regions of the CNS. At the level of pairs of neurons, the contribution of noise correlations is similar between the retina and many cortical areas, including visual, motor, and somatosensory (Oram et al., 2001; Petersen, Panzeri, & Diamond, 2001; Rolls et al., 2003), and this contribution is small. While this suggests generalization, it does not assure it. Good parametric models for capturing global input–output relationships will open the door to direct tests along the lines of the ones reported here.

Acknowledgments

We thank I. Bomash for helpful discussion and J. Pillow for both helpful discussion and for generously contributing software, particularly at the initiation of the project. This work was supported by NIH EY012978 to S.N., NIH EY07977 and the Tri-Institutional Training Program in Vision Research for M.M. and Z.N., and the Tri-Institutional Training Program in Computational Biology for Z.N.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.visres.2012.07.007>.

References

- Antos, A., & Kontoyiannis, I. (2001). Convergent properties of functional estimates for discrete distributions. *Random Structures and Algorithms*, 19, 163–193.
- Averbeck, B. B., Latham, Peter E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5), 358–366.
- Averbeck, B. B., & Lee, D. (2004). Coding and transmission of information by neural ensembles. *Trends in Neurosciences*, 27(4), 225–230.
- Blankenship, A. G., & Feller, M. B. (2010). Mechanisms underlying spontaneous patterned activity in developing neural circuits. *Nature Reviews Neuroscience*, 11, 18–29.
- Brown, E. N., Kass, R. E., & Mitra, P. P. (2004). Multiple neural spike train data analysis: State-of-the-art and future challenges. *Nature Neuroscience*, 7(5), 456–461.
- Chatterjee, S., Merwine, D. K., Amthor, F. R., & Grzywacz, N. M. (2007). Properties of stimulus-dependent synchrony in retinal ganglion cells. *Visual Neuroscience*, 24, 827–843.

- Dedek, K., Pandarinath, C., Alam, N. M., Wellershaus, K., Schubert, T., Willecke, K., et al. (2008). Ganglion cell adaptability: Does the coupling of horizontal cells play a role? *PLoS ONE*, 3(3), e1714.
- DeVries, S. H. (1999). Correlated firing in rabbit retinal ganglion cells. *Journal of Neurophysiology*, 81(2), 908–920.
- Fee, M. S., Mitra, Partha P., & Kleinfeld, D. (1996). Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability. *Journal of Neuroscience Methods*, 69(2), 175–188.
- Feldman, D. E. (2009). Synaptic mechanisms for plasticity in neocortex. *Annual Review of Neuroscience*, 32, 33–35.
- Gawne, T. J., & Richmond, B. J. (1993). How independent are the messages carried by adjacent inferior temporal cortical neurons? *Journal of Neuroscience*, 13(7), 2758–2771.
- Greschner, M., Shlens, J., Bakolitsa, C., Field, G. D., Gauthier, J. L., Jepson, L. H., et al. (2011). Correlated firing among major ganglion cell types in the primate retina. *Journal of Physiology*, 589(1), 75–86.
- Hand, D. J. (1981). *Discrimination and classification*. Chichester, UK: Wiley.
- Jacobs, A. L., Fridman, G., Douglas, R. M., Alam, N. M., Latham, P. E., Prusky, G. T., et al. (2009). Ruling out and ruling in neural codes. *PNAS*, 106(14), 5936–5941.
- Latham, P. E., & Nirenberg, S. (2005). Synergy, redundancy, and independence in population codes, revisited. *Neuroscience*, 25(21), 5195–5206.
- Levine, M. W. (2004). The potential coding utility of intercell cross-correlations in the retina. *Biological Cybernetics*, 91(3), 182–187.
- Mastrorarde, D. N. (1983). Correlated firing of cat retinal ganglion cells. I. Spontaneously active inputs to X- and Y-cells. *Journal of Neurophysiology*, 49(2), 303–324.
- Meister, M., Lagnado, L., & Baylor, D. A. (1995). Concerted signaling by retinal ganglion cells. *Science*, 270(5239), 1207–1210.
- Nirenberg, S., Carcieri, S. M., Jacobs, A. L., & Latham, P. E. (2001). Retinal ganglion cells act largely as independent encoders. *Nature*, 411(6838), 698–701.
- Nirenberg, S., & Latham, P. E. (2003). Decoding neuronal spike trains: How important are correlations? *PNAS*, 100(12), 7348–7353.
- Nirenberg, S., & Pandarinath, C. (2012). Retinal prosthetic strategy with the capacity to restore normal vision. *PNAS*. <http://dx.doi.org/10.1073/pnas.1207035109>.
- Nirenberg, S., Pandarinath, C., & Ohiohenuan, I. (2011). *Retina Prosthesis*. International Patent WO2011106783.
- Oizumi, M., Ishii, T., Ishibashi, K., Hosoya, T., & Okada, M. (2010). Mismatched decoding in the brain. *Journal of Neuroscience*, 30(13), 4815–4826.
- Oram, M. W., Hatsopoulos, N. G., Richmond, B. J., & Donoghue, J. P. (2001). Excess synchrony in motor cortical neurons provides redundant direction information with that from coarse temporal measures. *Journal of Neurophysiology*, 86(4), 1700–1716.
- Pandarinath, C., Victor, J. D., & Nirenberg, S. (2010). Symmetry breakdown in the ON and OFF pathways of the retina at night: Functional implications. *Journal of Neuroscience*, 30(30), 10006–10014.
- Paninski, L., Pillow, J., & Lewi, J. (2007). Statistical models for neural encoding, decoding, and optimal stimulus design. *Progress in Brain Research*, 165, 493–507.
- Perkel, D. H., Gerstein, G. L., & Moore, G. P. (1967). Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophysical Journal*, 7, 419–440.
- Petersen, R. S., Panzeri, S., & Diamond, M. E. (2001). Population coding of stimulus location in rat somatosensory cortex. *Neuron*, 32(3), 503–514.
- Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., et al. (2008). Spatio-temporal correlations and visual signaling in a complete neuronal population. *Nature*, 454, 995–999.
- Rolls, E. T., Franco, L., Aggelopoulos, N. C., & Reece, S. (2003). An information theoretic approach to the contributions of the firing rates and the correlations between the firing of neurons. *Journal of Neurophysiology*, 89(5), 2810–2822.
- 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.
- Wu, S., Nakahara, H., & Amari, S. (2001). Population coding with correlation and an unfaithful model. *Neural Computation*, 13, 775–797.