Diverse weighting of shared input noise prevents information saturation in a population code

Pratik S. Sachdeva, Michael R. DeWeese
Redwood Center for Theoretical Neuroscience
University of California, Berkeley

Noise is a prominent feature of neural systems: neural responses will vary trial-to-trial despite constant experimental stimuli. In sensory cortex, response variability is often correlated - these *noise correlations* are of theoretical interest because their structure can strongly influence the fidelity of a population code. Previous work has demonstrated that shared input noise can induce specific noise correlations, called *differential correlations*, that cause information to saturate in a neural population. We present a network of linear-nonlinear neurons in which we induce differential correlations by injecting noise to model, for instance, shared synaptic noise from irrelevant upstream action potentials. We show that by applying a diverse set of synaptic weights to the injected noise, the network can prevent information saturation and further improve the accuracy of its population code, despite an overall increase of noise in the system. This improvement results because the noise correlations are restructured in a way that is beneficial for decoding. Thus, by diversifying synaptic weights, a population of neurons can remove the harmful effects imposed by afferents that are uninformative about a stimulus. Interestingly, we also find cases where an abundance of weight diversity is harmful to the network, implying that there is some balanced regime where diversifying synaptic weight is optimal.

Significance

Population coding is hypothesized to be one mechanism of neural coding in sensory cortex. Noise correlations can be beneficial or harmful to a population code, depending on the relationship between the geometry of noise correlations and the tuning curve of neurons. In particular, Moreno-Bote et al. (2014) found that differential correlations, which can arise from shared input noise, are detrimental to a population code because they cause the population response distribution to lie parallel to the tuning curve. Thus, differential correlations result in information saturation in a neural population.

Sensory cortex must overcome the adverse effects of differential correlations. To explore possible strategies, we induced differential correlations in a linear-nonlinear network by introducing shared input noise. We find that the network can remove the differential correlations by applying a diverse range of weights to the input noise. Thus, diverse weighting of synapses presents an effective strategy to improve population coding.

Model

Our network consists of N neurons, each with a linear filter followed by a nonlinearity. The neurons accept two inputs: (1) a stimulus s and (2) injected Gaussian noise ξ_I with mean zero and variance σ_I^2. The linear filter of neuron i computes the linear combination $\ell_i = v_i s + w_i \xi_I$ and passes ℓ_i to its nonlinearity $g_i(x)$. The output of the nonlinearity, $r_i = g_i(\ell_i)$, can be thought of as the average firing rate for the ith neuron.

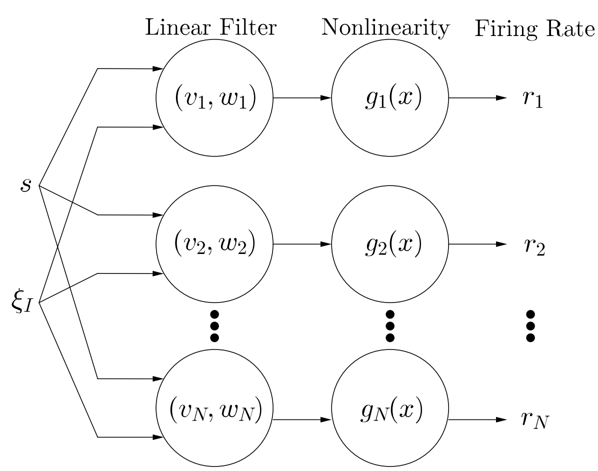

Figure 1

Then, for a specific trial, the firing rate is drawn from a distribution with mean r_i. We consider two cases. First, in the *Gaussian* case, the average firing rate r_i is set as the mean of a Gaussian distribution with variance of unity. In the *Poisson* case, r_i is set as the mean of a Poisson distribution.

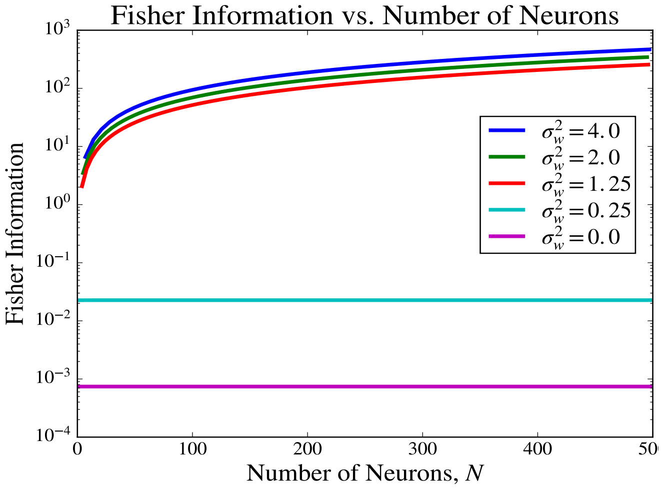

Figure 2

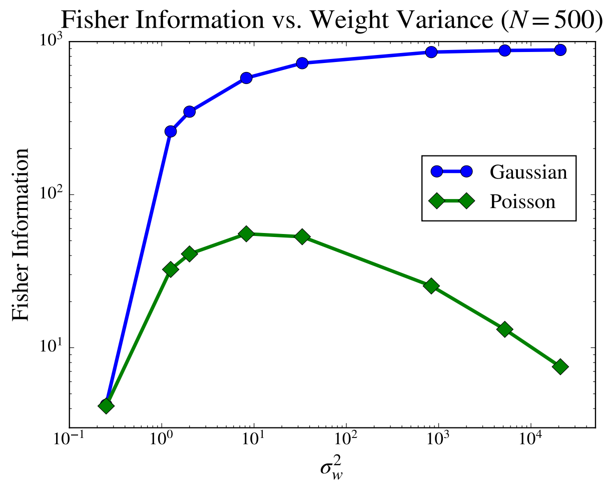

Figure 3

We assess the strength of the population code with the *linear Fisher information*. Fisher information provides a lower bound to the variance of the optimal unbiased estimator and is a standard measure of a population code's ability to discriminate a stimulus. The linear Fisher information, an approximation to the Fisher information, is given by

$$I_F(s) = \mathbf{f}'(s)^T \boldsymbol{\Sigma}^{-1}(s)\mathbf{f}'(s)$$

where $\mathbf{f}(s)$ is the tuning curve of the neurons, $\boldsymbol{\Sigma}(s)$ is the covariance between the firing rates of the neurons, and derivatives are taken with respect to s. Thus, for a choice of nonlinearity $\mathbf{g}(x)$, the Fisher information is dictated by the correlational structure established by a choice of weights $(\mathbf{v}, \mathbf{w})$.

Results

We find that increasing the weight diversity of noise (quantified by the variance of the weights) increases the Fisher information. A particular case is shown in Figure 2, where we plot the Fisher information (derived analytically) as a function of N for a Gaussian network and nonlinearity $g(x) = x^2$, while varying only $\mathbf{w}$. For low weight variance, the Fisher information saturates, implying the existence of differential correlations. Once the weights become more diverse, however, the Fisher information increases without bound.

In Figure 3, we compare the Fisher information between the Gaussian and Poisson networks (both with size $N = 500$) against the variance of the noise weights σ_w^2. Initially, the Fisher information increases as the variance of weights. The Poisson network, however, exhibits decreasing Fisher information for large σ_w^2. Hence, there is some balanced regime in which diversifying noise is optimal. We also find that the Fisher information increases with increasing stimulus weight diversity σ_v^2, as has been observed previously (Ecker et al., 2011)

Weight diversification is an effective strategy because the weights redistribute the noise correlations so that the population response distribution lies more orthogonal to the tuning curves $\mathbf{f}(s)$ of the neurons. Thus, the differential correlations disappear, allowing the Fisher information to increase without bound.

We have observed similar enhancement of Fisher information with increasing weight diversity for other choices of $g(x)$, such as a Gaussian nonlinearity and soft threshold. These results imply that a viable strategy by which neurons can improve their population coding is to apply a diverse weighting to synaptic inputs. Thus, the harmful effects of an upstream neuron whose action potentials are unimportant for the current stimulus can be ameliorated, even if the weighting increases the magnitude of noise for specific neurons.