



Functional imaging and neuronal information processing

Angel Nevado^a, Malcolm P. Young^b, Stefano Panzeri^{b,*}

^a*Departamento de Física Teórica, Universidad Autónoma de Madrid, 28049 Cantoblanco, Madrid, Spain*

^b*Department of Psychology, University of Newcastle upon Tyne, Henry Wellcome Building for Neuroecology, Medical School, Framlington Place, Newcastle upon Tyne, NE2 4HH, UK*

Abstract

Since the functional magnetic resonance imaging (fMRI) signal is likely to reflect a spatial average of the activity of neurons with partly dissimilar response properties, its interpretation in terms of neuronal information processing may be difficult. We explored this problem by means of a simple computational model based on known tuning properties of individual neurons and on the assumption that the fMRI BOLD signal is proportional to the local mean activity. We investigated the relationship between information transmitted by the population, BOLD signal and information conveyed by the BOLD signal. We found that the relationships between these variables were complex and dependent on voxel-size. © 2002 Elsevier Science B.V. All rights reserved.

Keywords: BOLD signal; fMRI; Spikes; Information; Spatial averaging

1. Introduction

Functional magnetic resonance imaging (fMRI) provides a non-invasive method to access indirectly neuronal activity in the human brain, by measuring blood-oxygen level-dependent (BOLD) changes in different brain regions and in different conditions [7]. Although many aspects of the relationship between neuronal activity and the BOLD fMRI signal are still to be clarified [4], it is believed that fMRI signals reflect blood flow and oxygenation, that blood flow changes reflect local metabolism, and that metabolism reflects neuronal population activation [7,5]. The relatively good signal-to-noise ratio of fMRI measures makes it possible to measure the magnitude of

* Corresponding author. Tel.: +44-191-2226174; fax: +44-191-2225622.

E-mail addresses: anevado@delta.ft.uam.es (A. Nevado), m.p.young@ncl.ac.uk (M.P. Young), stefano.panzeri@ncl.ac.uk (S. Panzeri).

BOLD signal changes, not only its presence or absence in certain brain regions [7]. This allows a quantitative characterization of responses of brain areas to parametric variations of stimulus variables of interest, and hence it may lead to better understanding of how sensory information is represented in the human brain [7,5]. However, there is a potential problem in using the parametric (or ‘computational’ [7]) design of the fMRI experiment to understand the nature of neuronal representations. It is clear that, whatever its exact relationship with the underlying neuronal activity, the BOLD fMRI signal, due to its vascular origin, can only reflect a spatial average of the neuronal signal with a resolution that at best is of the order of fractions of millimeter. Information is transmitted in the brain by the spiking electrical activity of populations of individual neurons. Characterizing neuronal information processing from a spatially averaged signal may lead to distortions if the region of the cortex over which the average is performed contains neurons with dissimilar stimulus–response profiles. In this paper, we try to shed some light on this issue by investigating how spatially averaging the neuronal signal may influence the reconstruction of sensory representations in the cortex from computational fMRI experiments. As stressed above, the precise link between neuronal activity and BOLD has still to be fully understood. However, comparisons of neuronal responses measured in monkeys and BOLD signal measured in analogous brain regions of humans indicate that the fMRI signal may be proportional to the mean local firing rate (i.e. population activation in the voxel) [5,3]. For simplicity, we hence assume in this study that there is a linear relationship between the BOLD signal and the mean firing rate in a given fMRI voxel. Though simplified, this assumption may be precise enough in several circumstances [5,3], and it has the advantage of allowing us to focus on the effect of the spatially averaged nature of the BOLD signal in a clearer way. We make use of a computational model to address quantitatively the relationships between the information carried by neuronal population firing, the changes in the average of the population activity over a voxel, and the information carried by the spatially averaged population activity, at different voxel resolutions. We then use these relationships to investigate how the stimulus encoding properties of the underlying neuronal population can be inferred from the measurement of the BOLD signal.

2. The model

We defined the neuronal response as the number of spikes fired by the neuron in a given time window. We hence ignored non-stationarities in the neuronal response which may play a role in the dynamics of blood oxygenation changes [2]. The neuronal response to a stimulus θ was assumed to follow a Poisson distribution whose mean followed a Gaussian tuning function f :

$$f(|\theta - \theta_p|) = m \exp\left(-\frac{(\theta - \theta_p)^2}{2\sigma_f^2}\right) + b. \quad (1)$$

The tuning function (1) was characterized by the following parameters: preferred stimulus (θ_p), stimulus modulation (m), spontaneous firing rate (b), and width of the tuning function (σ_f). This Gaussian tuning curve model is a good description of the tuning

properties of e.g. MT neurons to angular variables, such as motion direction. In addition, we assumed that the distribution of preferred stimuli of the neurons within a voxel was a Gaussian of width σ_p and centered at a certain preferred stimulus $\hat{\theta}_p$:

$$p(\theta_p) = \frac{1}{\sqrt{2\pi\sigma_p^2}} \exp\left(-\frac{(\theta_p - \hat{\theta}_p)^2}{2\sigma_p^2}\right). \quad (2)$$

The preferred stimulus in the voxel $\hat{\theta}_p$ was then set to zero without loss of generality. This means that the stimulus value θ has to be interpreted as the difference between the physical value of the presented stimulus and the preferred stimulus in the voxel. The width of the distribution of the preferred stimulus is obviously related to the spatial resolution of the voxel. The larger is the voxel size, the larger is σ_p .

We then investigated the relationship between the information about stimuli that is conveyed by the neuronal population through the firing rates of each individual neuron and the stimulus coding properties that can be extracted in various ways from the spatially averaged signal.

3. Quantities

In order to quantify how information about a particular stimulus value θ is transmitted by the activity of the neuronal population, we used Fisher information $J(\theta)$. Fisher information is a good measure on the encoding accuracy of a particular stimulus θ because its inverse is the (Cramer–Rao) lower bound on the mean squared stimulus reconstruction error, and it is defined as follows:

$$J(\theta) = \sum_{\mathbf{r}} p(\mathbf{r}|\theta) \left(\frac{\partial}{\partial \theta} \log p(\mathbf{r}|\theta) \right)^2, \quad (3)$$

where \mathbf{r} is the neuronal population response vector, each element of the vector being the number of spikes emitted by a neuron in the population, and $p(\mathbf{r}|\theta)$ is the stimulus-conditional probability of a neuronal response \mathbf{r} .

Since the neurons fire independently of each other, the Fisher information (Eq. (3)) conveyed by the responses of all the single neurons in the population is the sum of the Fisher information from each neuron [1]:

$$J(\theta) = \int d\theta_p p(\theta_p) J_{\text{neuron}}(|\theta - \theta_p|). \quad (4)$$

The information J_{neuron} conveyed by single-neuron firing according to a Poisson process has the following expression [1]:

$$J_{\text{neuron}}(|\theta - \theta_p|) = \frac{((\partial/\partial\theta)f(|\theta - \theta_p|))^2}{f(|\theta - \theta_p|)} T, \quad (5)$$

where T is the time window from which the spikes are counted.

To understand the effect of the spatial average of the neuronal signal intrinsic in fMRI measures, it is useful to compare the information carried by the ensemble of neurons in the voxel, Eq. (4), to quantities constructed from the measures of the fMRI

signal. Two quantities are of obvious interest. The first quantity is the change of the averaged activity with stimulus. It is defined as the fractional change with stimulus of the averaged neuronal activity in the voxel, $\Delta r(\theta) = (f_{\text{voxel}}(\theta) - b)/b$, where $f_{\text{voxel}}(\theta)$ is the averaged activity in the voxel when stimulus θ is presented

$$f_{\text{voxel}}(\theta) = \int d\theta_p p(\theta_p) f(|\theta - \theta_p|). \quad (6)$$

Under the assumptions of our model, Δr is proportional to the percentage changes in the BOLD signal, and hence corresponds to the traditional way to quantify the fMRI response to different stimuli.

The second quantity of interest is the information $J_{\text{mr}}(\theta)$ that the changes of the fMRI signal convey about the identity of the stimuli. The Fisher information J_{mr} conveyed by the fMRI signal (which we assume proportional to the mean activity in the population) has the same expression as Eq. (5), with the difference that it now depends on the averaged activity in the voxel, rather than on the tuning curve of the individual neuron:

$$J_{\text{mr}}(\theta) = \frac{((\partial/\partial\theta)f_{\text{voxel}}(\theta))^2}{f_{\text{voxel}}(\theta)} T. \quad (7)$$

Obviously, the comparison of the spatially averaged fMRI quantities with the information transmitted by the neuronal population is going to tell us how the quantities derived from fMRI experiments may be used to draw conclusions about the information transmitted by the neurons generating the BOLD signal. Since the relationship between mean firing rates and tuning parameters is complex [6], and the relationship between tuning parameters and information is complex [1], we expected that the relationship between mean firing rates and information processing could be very complex as well, and worth a systematic investigation.

4. Results and conclusions

We used the formulae above to investigate whether stimuli that elicit higher activity in the voxel also corresponded to stimuli that are better represented by the neurons in information theoretic terms, and to investigate whether the relationship between the BOLD signal and the neuronal information depends on the voxel size. In Fig. 1, we plot some examples of our results in this investigation (a more complete account of these results will be presented in a paper in preparation). $J(\theta)$, $\Delta r(\theta)$ and $J_{\text{mr}}(\theta)$ are plotted as a function of θ , the distance between the presented stimulus and the preferred stimulus in the voxel. For coarse spatial resolutions (large σ_p ; Fig. 1a) the neuronal population information $J(\theta)$ depends on the stimulus in a way similar to the mean firing rate $\Delta r(\theta)$, but the information available in the rate changes $J_{\text{mr}}(\theta)$ is very different from its neuronal counterpart $J(\theta)$. For finer spatial resolutions (small σ_p ; Fig. 1b), the opposite is true. The neuronal population information $J(\theta)$ is well predicted by the information in the mean activity, $J_{\text{mr}}(\theta)$. However, the mean activity $\Delta r(\theta)$ is not a good predictor of the underlying information processing. For example, dips in information processing $J(\theta)$ correspond to maxima in the mean activity $\Delta r(\theta)$. The implications of this result are twofold. First, it implies that, depending on the




Fig. 1. A comparison between $J(\theta)$, $\Delta r(\theta)$ and $J_{mr}(\theta)$ as a function of θ , the difference between the presented stimulus and the stimulus preferred by neurons in the voxel. To allow for a better comparison of the different curves the y -axis is expressed in arbitrary units: (a) low spatial resolution case: $m = 4, b = 1, \sigma_f = 1, \sigma_p = 2$; (b) high spatial resolution case: $m = 4, b = 1, \sigma_f = 1, \sigma_p = 0.5$.

spatial resolution of the experiment, different ways of quantifying the BOLD signal changes with stimulus should be used to characterize neuronal information processing. Second, it shows that the stimulus that activates the voxel most is not necessarily the one best represented by the neurons in the voxel.

Acknowledgements

We thank K. Friston for many useful discussions. This research was supported by the Wellcome Trust, by the Comunidad Autónoma de Madrid and by an MRC Research Fellowship to SP.

References

- [1] N. Brunel, J.-P. Nadal, Mutual information, Fisher information and population coding, *Neural Comput.* 10 (1998) 1731–1757.
- [2] K.J. Friston, A. Mechelli, R. Turner, C.J. Price, Nonlinear responses in fMRI: Balloon models, Volterra kernels and other hemodynamics, *Neuroimage* 12 (2000) 466–477.
- [3] D.J. Heeger, A.C. Huck, W.S. Geisler, D.G. Albrecht, Spikes versus BOLD: what does neuroimaging tell us about neuronal activity? *Nature Neurosci.* 3 (2000) 631–633.
- [4] N.K. Logothetis, J. Pauls, M. Augath, T. Trinath, A. Oeltermann, Neurophysiological investigation of the basis of the fMRI signal, *Nature* 412 (2001) 150–157.
- [5] G. Rees, K. Friston, C. Koch, A direct quantitative relationship between the functional properties of human and macaque V5, *Nature Neurosci.* 3 (2000) 716–723.
- [6] J.W. Scannell, M.P. Young, Neuronal population activity and functional imaging, *Proc. Roy. Soc. B* 266 (1999) 875–881.
- [7] B.A. Wandell, Computational neuroimaging of human visual cortex, *Ann. Rev. Neurosci.* 22 (1999) 145–173.