International Workshop on Statistical-Mechanical Informatics 2009 (IW-SMI 2009) IOP Publishing
Journal of Physics: Conference Series **197** (2009) 012013 doi:10.1088/1742-6596/197/1/012013

On the presence of high-order interactions among somatosensory neurons and their effect on information transmission

Robin A A Ince[1], Fernando Montani[2], Ehsan Arabzadeh[3], Mathew E Diamond[4] and Stefano Panzeri[2]

[1] Faculty of Life Science, University of Manchester, 3.431 Stopford Building, Oxford Road, Manchester M13 9PL, UK
[2] Robotics, Brain, and Cognitive Sciences Department, Italian Institute of Technology, Via Morego 30, 16163 Genova, Italy
[3] School of Psychology, University of New South Wales, Sydney, New South Wales, Australia
[4] Cognitive Neuroscience Sector, International School for Advanced Studies, Trieste, Italy, and the SISSA Unit, Italian Institute of Technology, Trieste, Italy

E-mail: stefano.panzeri@iit.it

Abstract.
In order to understand how populations of neurons encode information about external correlates, it is important to develop minimal models of the probability of neural population responses which capture all the salient changes of neural responses with stimuli. In this context, it is particularly useful to determine whether interactions among neurons responding to stimuli can be described by a pairwise interaction model, or whether a higher order interaction model is needed. To address this question, we compared real neural population activity obtained from the rat somatosensory cortex to maximum-entropy models which take into account only interaction of up any given order. By performing these comparisons, we found that interactions of order two were sufficient to explain a large amount of observed stimulus-response distributions, but not all of them. Triple-wise interactions were necessary to fully explain the data. We then used Shannon information to compute the impact of high order correlations on the amount of somatosensory information transmitted by the neural population. We found that correlations of order two gave a good approximation of information carried by the neural population, within 4% of the true value. Third order correlations gave an even better approximation, within 2% of the true value. Taken together, these results suggest that higher order interactions exist and shape the dynamics of cortical networks, but play a quantitatively minor role in determining the information capacity of neural populations.

1. Introduction

The activity of sensory neurons provides the information about the external world based upon which animals take decisions. Determining how the activity of populations of neurons represents sensory stimuli is thus of obvious importance for the progress of neuroscience and is a prerequisite for understanding sensory function. Yet, characterizing the details of the neural representations of sensory stimuli has proven to be a daunting challenge so far. This difficulty is due to several reasons. The first reason is that the presentation of even the simplest sensory stimulus elicits the activity of large numbers of neurons, which at any time may either fire

or not fire a spike in response to the stimulus. Thus, neural responses are intrinsically high dimensional [1]. The second problem is that neural responses are variable, and this implies that the stimulus-response relationship must be characterized probabilistically rather than with just a simple tuning curve or transfer function.

However, even despite the above problems, it would be still easy to construct a low dimensional representation of neural populations responses to a stimulus if the responses of each neurons were statistically independent from that of other neurons. In this case, one could measure the response probabilities of each neuron separately, and then construct a population response probability by simply taking the product of the probabilities of the individual neurons. Unfortunately, that's not the case: the response of a neuron does not depend only on the stimulus but also on the activity of the other neurons [2–5]. Theoretical and experimental studies have suggested that such interactions can profoundly affect the information transmitted by neural populations, by increasing the information content of neural populations [6–8], by tagging features to be bound together [9], by stabilizing the temporal relationships between cells against the detrimental effect of trial-to-trial variability [10, 11], by implementing strategies for error correction [12], or by severely limiting the representational capacity of neural populations [13, 14]. As a consequence, neural population responses cannot be described simply using models of independent non-interacting neurons neurons: any model of neural responses must be able to capture all the salient interactions among neurons.

Given that models of the neural responses must include all the salient interactions among neurons, the best hope to find a compact low-dimensional representation of neural responses is if the structure of interactions within neurons in a population could be described only in terms of interactions between pairs of neurons, rather than in terms of interactions between whole large assemblies. A pairwise representation of neural interaction would greatly simplify the number of parameters needed to describe neural representations of stimuli, potentially leading to the development of simple but effective models of decoding neural population activity [1].

The question of whether the structure of correlated activity can be satisfactorily described by considering only pairwise interactions is only recently beginning to be studied systematically both by theory and experiments [15–22] A number of groups [17–19] have addressed this problem demonstrating that in several types of experimental preparations pairwise correlations account for the majority ($> 90\%$) of decrease of response entropy due to all possible interactions (of any order) between neurons. Given that entropy is a measure of variability of responses, these results suggests that pairwise interactions are sufficient to describe reasonably well how interactions constrain neural population activity. However, these results do not clarify fully whether or not high order interactions exist and are statistically significant. Moreover, given that response entropy and information about stimuli are not equivalent concepts, these results do not fully clarify whether high order interactions are quantitatively important to describe how neurons encode information about external stimuli.

In this study we aim at contributing to this research by evaluating whether high order interactions among neurons are present in the somatosensory cortical network, and by measuring what is the impact of such high order interactions on the amount of Shannon information that the neural population conveys about somatosensory stimuli.

This paper is organized as follows. We first describe how to use maximum entropy models to quantify whether the probabilities of neural population responses to stimuli can be described with only interactions up to a given order. We then apply these techniques to real neural population activity recorded from the rat somatosensory cortex, and establish whether there are statistically significant interactions among neurons of order higher than two. We then use Shannon information to quantify what is the impact of correlations of a given order on the sensory representation of information.

2. Investigating the order of interaction through the maximum entropy principle

We consider a population of N neurons whose activity is simultaneously observed during a specified short time window following the presentation of a sensory stimulus s taken from a set of S different stimuli. Since the time window is short, we can use a binary representation of the activity of each neuron (with zero or one to respectively denote the presence or absence of spikes from the neuron). Thus, the neuronal population activity is represented by a binary array $\mathbf{x} = \{x_1, ..., x_N\}$ in the space X of all binary arrays of length N, where $x_i = 0$ if neuron i is silent in some time window and $x_i = 1$ if it is firing one or more spikes.

The question that we seek to address is whether we can describe all interactions between the neurons in terms of interactions between up to two neurons only, or whether there are interactions among groups of more than two neurons which cannot be explained in terms of pairwise interactions. A rigorous way to investigate the effects of different orders of interaction is provided by the technique of *maximum entropy*, which was originally introduced in statistical physics [23], and is now beginning to be used in neuroscience [1, 17–19, 24, 25]. In general, the idea of the maximum entropy (ME) principle is to first fix some constraints that are of interest and then seek the simplest, or most random, distribution subject to those constraints. Using entropy as a measure of randomness, asking for the most random distribution corresponds to asking for the distribution with maximal entropy subject to the constraints. This removes all types of correlation or structure in the data that does not result from the constrained features.

The ME formalism can be naturally used to to address the problem of whether we can describe all interactions between neurons in terms of interactions between up to k neurons only, or whether there are higher interactions among more than k neurons which cannot be possibly explained in terms of interactions of order up to k. Measuring all interactions of up to k variables means measuring all the marginal response probabilities involving up to k variables. Therefore any probability matching the observed interactions of up to k elements must obey (apart from the usual non negativity and normalization constraints) the following linear constraints:

$$\begin{aligned} P(x_{i_1} = 1|s) &= \eta_{i_1} \\ P(x_{i_1} = 1, x_{i_2} = 1|s) &= \eta_{i_1,i_2} \\ &\cdots \\ P(x_{i_1} = 1, x_{i_2} = 1, \cdots, x_{i_k} = 1|s) &= \eta_{i_1,i_2,\cdots,i_k} \end{aligned} \tag{1}$$

where (in agreement with the well-established notation of Amari [26]), the set of variables denoted η represent the values of the marginal probabilities, and the subscript indices of η represent the variables involved in the marginal probabilities.

The probability distribution $P^{(k)}_{ME}(\mathbf{x}|s)$ with maximum entropy among those satisfying the above constraints is the one that does not impose the presence of any additional interaction of order higher than k. (The case $k = 1$ is trivial and corresponds to the case in which all neurons fire independently at fixed stimulus). Following [26, 27], it can be shown that there is a unique solution to the constrained maximum entropy problem, which can be written in the following exponential form:

$$P^{(k)}_{ME}(\mathbf{x}|s) = \exp\left\{\sum_i x_i\theta_i + \sum_{i_1<i_2} x_{i_1}x_{i_2}\theta_{i_1 i_2} + \cdots + \sum_{i_1<\cdots<i_k} x_i \cdots x_k \theta_{i_1\cdots i_k} - \psi\right\} \tag{2}$$

To quantify whether interactions of up to k neurons in a population are sufficient to describe the probabilities of neural responses to stimuli, we can quantitatively compare the true distribution $P(\mathbf{x}|s)$of neural responses to the stimulus s consider to the distribution $P^{(k)}_{ME}(\mathbf{x}|s)$. By performing this comparison over a range of values of k, we can empirically determine the minimal k necessary to fit well the empirically measured response probability.

In order to compute the maximum entropy distribution $P^{(k)}_{ME}(\mathbf{x}|s)$ of Eq. (2) from real data, we just need to find the θ coefficients with up to K indices entering the above equation to construct the solution above. These θ coefficients can be determined from the experimentally measured marginal probabilities η of up to k elements through a set of algebraic equations which were derived in the work of Amari [26, 28]. To solve these equations numerically, we used the publicly available `pyentropy` numerical package that we recently developed and published [29]. We refer to Ref [29] for full details and for download of the code implementing the numerical solutions.

3. The role of high order correlations in shaping synchronous discharge in somatosensory cortex

After having described the ME techniques to study the presence of interactions of up to any given order, we apply them to a population of neurons recorded from the whisker representation in the somatosensory cortex of urethane anaesthetized rats. We first describe the dataset; we then compare the ME probabilities at fixed interaction order with the real data; and we finally evaluate the effect of the interaction order on the information about the stimuli carried by the neural responses.

3.1. Neurophysiological data

The data set (previously published in [30, 31]) consists of 24 simultaneously recorded neural clusters, each sampled with a different electrode with a minimal inter-electrode distance of 400 μm. Spike times from each electrode were determined by a voltage threshold set to a value 2.5 times the root mean square voltage. Since it was not possible to sort well-isolated units from each channel, spikes from the same recording channel were considered together as a single neural cluster. It has been estimated that, under these recoding conditions, each cluster captured the spikes of approximately two to five neurons located near the tip of the electrode (see [32]). Neural activity was recorded in response to stimulation (with a piezoelectric wafer controlled by a voltage generator) consisting of sinusoidal whisker vibrations, each defined by a different value of vibration velocity and delivered for 500 ms (see [31] for full details). Thirteen different values of vibration velocity were tested, ranging between $Af = 0.15$ mm/s and $Af = 47.7$ mm/s. Each value of vibration velocity was treated as a different stimulus s (there were 13 stimulus classes in total). The number of recorded repetitions for each stimulus (called "trials" in neurophysiology), from which the probability of response at fixed stimulus is determined, varied between a minimum of 200 and a maximum of 1400 across stimuli.

It was previously shown [31] that the majority of the information is transmitted very early post stimulus onset (typically between 5 and 30 ms). We therefore concentrated on data taken from these early highly-informative windows.

3.2. Fitting the neurophysiological stimulus-conditional response probabilities to maximum entropy models

We used these data to study the shape of the distribution of neural population responses at fixed stimulus, and the order of neural interactions needed to describe this distribution. We note that some previous seminal studies [17, 18] focused on the probability of response to many different stimuli. However, this has the potential problem that the resulting correlation may arise both from correlations in the stimulus and from actual neural interactions, and it is difficult to separate them [1]. Thus, we considered distributions at fixed stimuli to ensure that we did not investigate neural interactions originating at least in part from correlations in the stimuli.

We decided to consider response distributions of $N = 8$ channels simultaneously recorded (out of the 24 available). The reason was that this population size was big enough to begin observing some effects of high order interactions, while was small enough to be sampled with

the available data and be tested with both quantitative goodness-of-fit tests and the information analysis of the next subsection.

We first considered (figure 1) the distribution of neural responses to one particular stimulus (velocity = 4.3 mm/s; 1200 trials available). Neural responses evolve over time, and thus the distribution of neural responses depends on the particular post-stimulus window considered. Since it has been previously shown [31] that in this system the majority of the information is transmitted very early after stimulus onset (typically between 5 and 30 ms), and since the integration time constant of neurons post-synaptic to the ones analyzed here is likely to be in the range of 10-20 ms, we concentrated on responses taken from short, early, highly informative windows. In particular, we selected the 5-15 and the 5-25 ms post-stimulus windows. The distributions of neural responses in these windows are shown in figure 1 (5-15 ms post-stimulus in figure 1A; and 5-25 ms post-stimulus in figure 1B). In figure 1, the responses were arbitrarily ordered in the x axis to give a monotonically decreasing probability. We compared the real observed distribution with the ME distribution of a given order k ($k = 1, \cdots 3$) obtained with the numerical procedure outlined in the previous section. It is apparent than, for both time windows, the $k = 1$ distribution (which of course corresponds to the case in which all neurons fire independently at fixed stimulus) did not provide a good fit to the data, whereas ME probabilities with $k > 1$ provided a good qualitative agreement. The $k = 3$ distribution provided a small but visible improvement in the goodness of fit with respect to the $k = 2$ distribution.

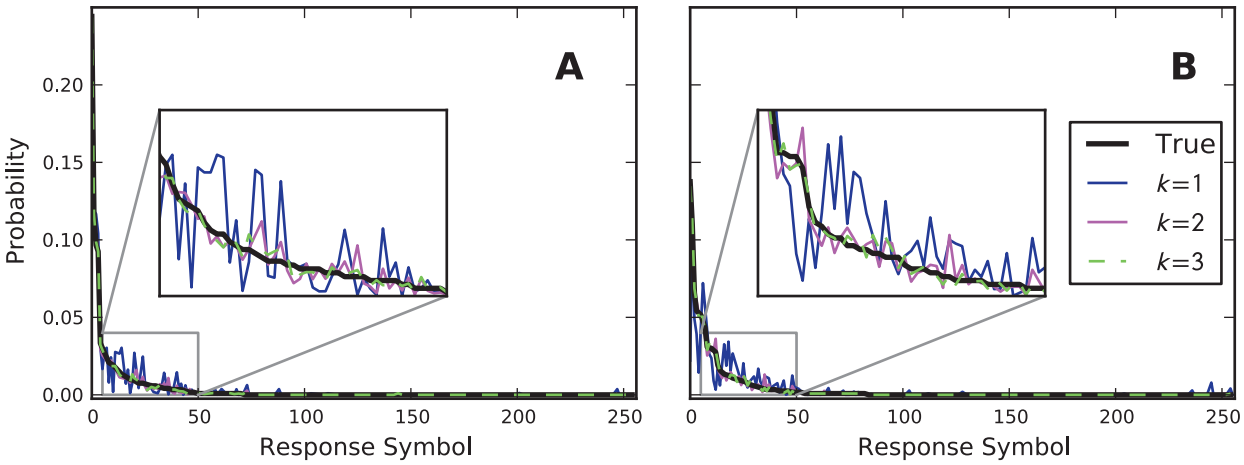

Figure 1. *Maximum entropy models at fixed stimulus.* Here the true distribution (solid black line) and maximum entropy distributions of orders $k = 1, 2, 3$ (solid blue, solid purple, dashed green respectively) are plotted for binary responses to a particular randomly selected group of 8 channels to a 4.3 mm/s stimulus (1200 trials) in the window 5-15 ms (panel **A**) and 5-25 ms (panel **B**) (see text). The response symbols are ordered by decreasing probability value.

To better quantify the goodness of fit of the maximum entropy models across all stimuli and all observed populations of $N = 8$ channels, we repeated the ME calculation of the probability models for each stimulus condition and for 1000 populations made of random combinations of 8 channels from the 24 available. We then computed the probability of rejection of the null hypothesis that the observed distribution comes from the considered ME model using a log-likelihood test (the g-test) [33]. This test produces a chi-square distributed test statistic, similar to the Pearson chi-squared test. Indeed the Pearson chi-squared test was originally developed as an approximation of the log-likelihood ratio test due to the difficulty of calculating logarithms before the advent of computers. The percentage of rejections of the null hypothesis at the

$p = 0.05$ level for the $k = 1$ and $k = 2$ model across all dataset are reported in figure 2 (for the 5-15 ms post-stimulus window) and figure 3 (for the 5-25 ms post-stimulus window). It is apparent that the $k = 1$ model had to be rejected most times, whereas for most stimulus conditions the number of rejections of the $k = 2$ model was below the 5% random level. The $k = 2$ model however failed above chance for longer windows and for the stimulus conditions with more trials available (velocity equal to 1.61, 2.66 and 4.3 mm/s). In contrast, the $k = 3$ model and higher had zero rejections at $p = 0.05$ for both windows and all stimuli.

In sum, the $k = 1$ order model does not explain the response probabiltiies at fixed stimulus in any satisfactory way. The $k = 2$ model explains most (but not all) probability distributions, whereas the $k = 3$ model was sufficient to fully explain the observed response distributions.

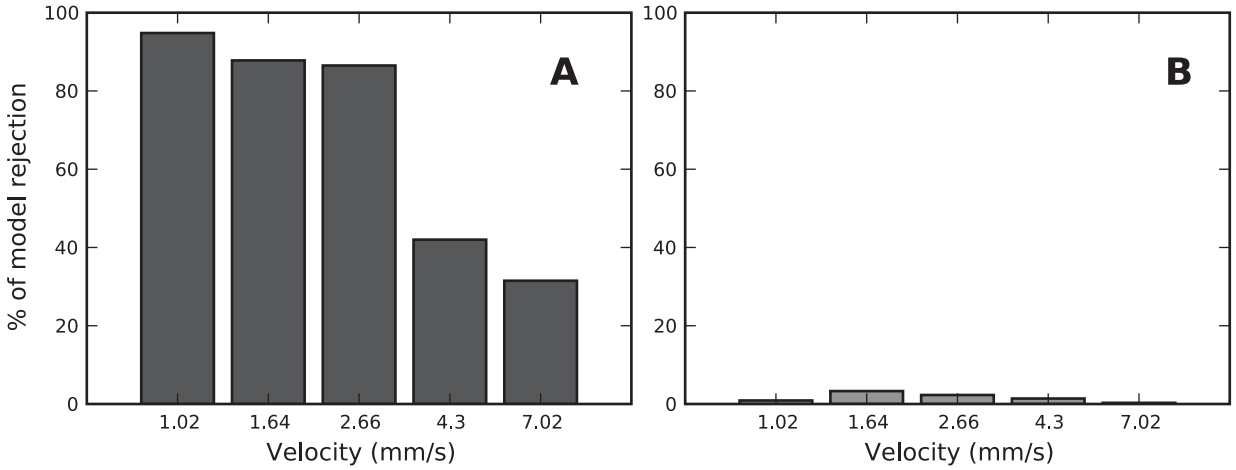

Figure 2. *Goodness of fit measures of maximum entropy distributions for 5-15 ms response window.* 1000 random combinations of 8 channels were chosen from the 24 available and the maximum entropy solutions of different orders were obtained. These were compared to the measured distribution using the log-likelihood g-test (see text) and the percentage of these 1000 trials for which the first order model (panel **A**) and second order model (panel **B**) was rejected (p=0.05) for 5 different fixed stimuli are shown. For the third order model (not shown) no trials were rejected for any stimuli.

3.3. Effect of Interactions on somatosensory information encoding

The above finding that there are significant high order interactions suggests that they should not be neglected in models of information transmission, but it does not tell how much these correlations are important. To quantify this, we next computed the information between the stimulus and the population activity, and we compared it to that derived from the maximum entropy models.

The mutual information between the stimuli and the neural population activity is defined as follows:

$$I(S;X) = H(X) - H(X|S) \tag{3}$$

where $H(X)$ and $H(X|S)$ are the response entropy and noise entropy respectively:

$$H(X) = -\sum_{\mathbf{x}} P(\mathbf{x}) \log_2 P(\mathbf{x}) \tag{4}$$

$$H(X|S) = -\sum_{\mathbf{x},s} P(s)P(\mathbf{x}|s) \log_2 P(\mathbf{x}|s) \tag{5}$$

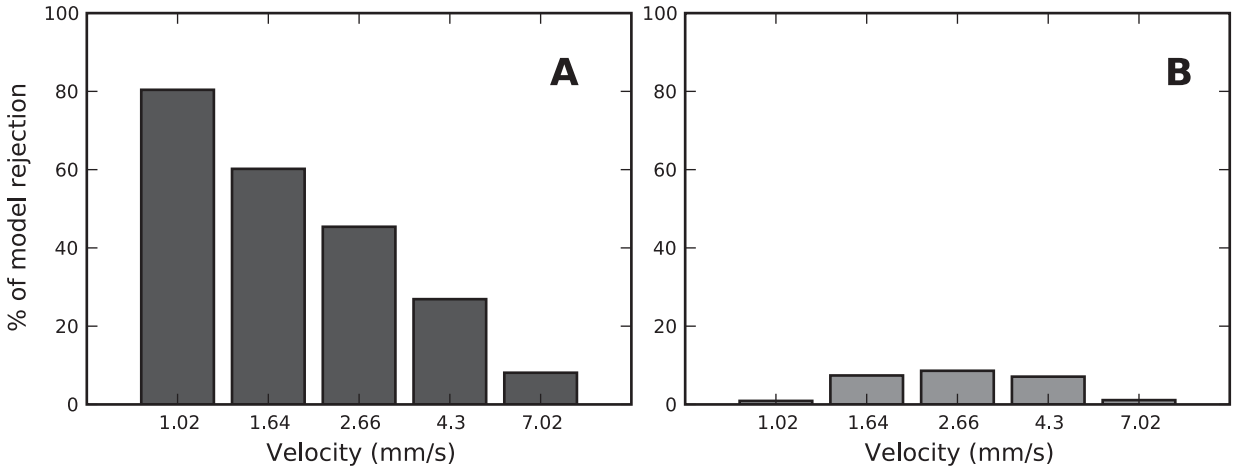

Figure 3. *Goodness of fit measures of maximum entropy distributions for 5-25 ms response window.* 1000 random combinations of 8 channels were chosen from the 24 available and the maximum entropy solutions of different orders were obtained. These were compared to the measured distribution using the log-likelihood g-test (see text) and the percentage of these 1000 trials for which the first order model $k = 1$ (panel **A**) and second order model $k = 2$ (panel **B**) was rejected (p=0.05) for 5 different fixed stimuli are shown. For the third order model (not shown) no trials were rejected for any stimuli.

where in the above $P(\mathbf{x}) = \sum_s P(\mathbf{x}|s)P(s)$. We computed $I(S;X)$ from the the population activity as follows. First, to improve the sampling properties of this data set we pooled the two lowest velocity stimuli (0.15 mm/s and 2.3 mm/s) and the two highest velocity stimuli (29.5 mm/s and 47.7 mm/s) resulting in 11 stimulus classes with a minimum of 600 trials and a maximum of 1400 trials. Then, we empirically measured $P(\mathbf{x}|s)$, $P(s)$ and $P(\mathbf{x})$ and plugged them into the above equations. The limited sampling bias [34] was corrected for using a combination of the shuffled estimator $I_{\rm sh}$ [25, 34] with the Panzeri-Treves analytic correction [35] from the Pyentropy library [29].

We investigated the impact of interactions at a given order k by calculating the mutual information that would result from a system exhibiting the probability distributions obtained from the maximum entropy solution, as follows:

$$I^{(k)}(S;X) = H^{(k)}(X) - H^{(k)}(X|S) \tag{6}$$

where $H^{(k)}(X)$ and $H^{(k)}(X|S)$ are the response and noise entropies respectively of the k-th order maximum entropy model. These entropies are obtained by replacing $P(\mathbf{x}|s)$ and $P(\mathbf{x})$ with $P^{(k)}_{ME}(\mathbf{x}|s)$ and $P^{(k)}_{ME}(\mathbf{x})$ in Eqs. (4,5), where $P^{(k)}_{ME}(\mathbf{x}) = \sum_s P^{(k)}_{ME}(\mathbf{x}|s)P(s)$. Then

$$I^{(k)}(S;X) = \sum_{\mathbf{x},s} P(s)P^{(k)}_{ME}(\mathbf{x}|s)\log_2 \frac{P^{(k)}_{ME}(\mathbf{x}|s)}{P^{(k)}_{ME}(\mathbf{x})} \tag{7}$$

$I_k(S;X)$ was computed from the data as follows. First, and as for $I(S;X)$, we pooled the two lowest velocity stimuli (0.15 mm/s and 2.3 mm/s) and the two highest velocity stimuli (29.5 mm/s and 47.7 mm/s) resulting in 11 stimulus classes. Then, for each stimulus class, we obtained the maximum entropy solution, $P^{(k)}_{ME}(\mathbf{x}|s)$, for each order of interest and for each stimulus conditional response. Then, from each of these stimulus-conditional maximum entropy

solutions, we simulated data with the same number of trials as available in the experimental data set (this was different for each stimulus). These trials were generated using inverse transform sampling. This is done to ensure a fair comparison between the measured data and the generated data; any bias effects should affect both equally. The values obtained were averaged over 1000 repetitions to remove any trial to trial variation from the inverse transform sampling step.

Figure 4 shows the effect of including higher order interactions on information. The first order maximum entropy models conveyed significantly higher information than the true probabilities. This means that interactions in this system have mostly a limiting effect on information. The second and third order information was significantly lower than the first order one and was very close to to the true one. We found that correlations of order two gave a good approximation of information carried by the neural population, as they were within 3% and 4% of the true value for windows of 5-15 ms (figure 4A) and 5-25 ms (figure 4B) post-stimulus respectively. Third order correlations gave an even better approximation, within 2% of the true value for both post-stimulus windows (figure 4). This suggests that interactions of order two or higher are sufficient to describe well how neural interactions limit the rate of information transmission.

It is interesting that the mutual information of the system was already well approximated by models containing interactions of up to order 2 or 3. This is significant, since it greatly reduces the parameters required to describe the system. While it is still challenging to sample up to third order marginals reliably in an experiment, it is a much more tractable problem than the case where all orders of interaction must be accurately determined.

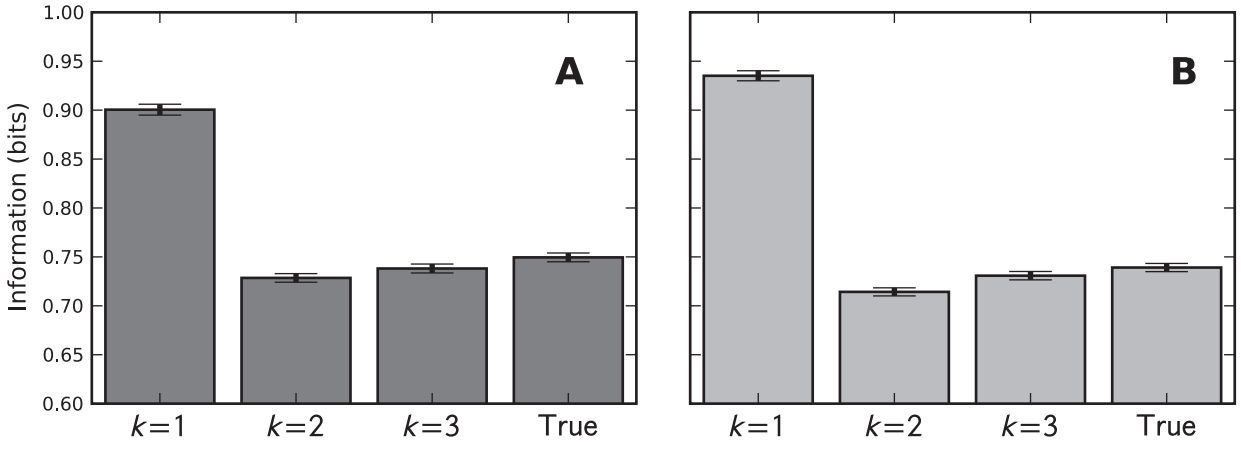

Figure 4. *Effect of interactions on mutual information.* 1000 random combinations of 8 channels were chosen from the 24 available and the maximum entropy solutions of different orders were obtained. These were used to compute the mutual information in a system with only interactions of up to order $k = 1, 2, 3$. The means of these values over the different channel combinations are compared to the mean of the true information for response windows of 5-15 ms (panel **A**) and 5-25 ms (panel **B**). Errorbars show ± 1 s.e.m.

4. Discussion

In this paper, we used maximum-entropy techniques to investigate the presence of interactions of order higher than two in the rat somatosensory cortex, and to evaluate their effect on the information that somatosensory cortical populations carry about the velocity of whisker stimulation. Considering population responses made of spiking activity collected from up to 8 simultaneously recorded locations, we found that interactions of order two were sufficient to

explain a large amount of observed stimulus-response distributions, but not all of them. Triple-wise interactions were necessary to fully explain the data. We then used Shannon information to compute the impact of high order correlations on the amount of somatosensory information transmitted by the neural population. We found that correlations of order two gave a good approximation of information carried by the neural population, within 4% of the true value. Third order correlations did even better, proving an approximation within 2% of the true value of information carried by neural activity. In the following, we briefly discuss the implication of our results, as well as differences and similarities with respect to previous work.

4.1. Interactions of order higher than two are present in cortical populations
A first result is that pairwise interactions were not enough to fully explain the observed stimulus-response distributions, thereby confirming that higher order interactions between groups of neurons are indeed a part of cortical dynamics, as suggested by earlier studies demonstrating the existence of multineuronal firing patterns in a variety of experimental preparations [4, 24, 36–40]. An important and still open question is what may be the potential function of high order interactions. Modeling work suggests that higher order interactions may have a role in regularizing the overall rate of discharge of the population, for example by avoiding concentration of firing on the maximally firing state [41]. It would be interesting for future research to understand whether high order interaction serve this purpose in real local cortical networks.

4.2. Different measures of the impact of interactions at a given order
A series of recent studies [17–19] quantified the quality of pairwise models by using a measure called the fraction of full network information. The fraction of full network information that is captured by a second order ME model is defined [18, 42] as the the ratio between the reduction of total entropy accounted for by all interactions of order > 2 (computed as the difference between the pairwise maximum entropy and the the total entropy taking into account all correlations) and the reduction of entropy due to all interactions of any order (computed as the difference between the entropy of the independent probability model and the the total entropy taking into account all correlations). This entropy-based metric has many interesting properties and advantages, discussed in [42]. In particular, given that entropy is a measure of variability of responses, this metric is useful to investigate whether pairwise interactions are enough to describe how interactions shape the variability of neural responses. A novelty of our work, with respect to several previous studies, is that we introduced and used some new metrics to assess the quality of pairwise models in describing the response probabilities. These metrics are potentially useful because they allow addressing some questions which are in principle different from the ones that can be addressed using the fraction of network information, as we elaborate next.

The fraction of full network information cannot be used to clarify whether or not high order interactions exist and are statistically significant. This is because, as discussed in [22], this metric can be relatively small even the pairwise and the true probabilities have very different shapes. To overcome this potential problem, here we introduced a different procedure to assess the goodness of the pairwise model. This simple procedure, based on fitting ME distributions to the data numerically using Amari's formalism [26] and by assessing the difference between ME and true distribution using a g-test statistic, revealed that interactions of order higher than two are significant in the analyzed dataset.

Moreover, in this study we were interested in whether pairwise interactions were sufficient to describe how the population of neurons encodes information about the stimulus. This question cannot possibly be addressed by quantifying the fraction of network information. This is because of two reasons. The first is that the fraction of network information measures the reduction of network variability specifically attributable to correlations up to a given order, and the variability of the population response is not equivalent to the information about the stimuli

carried by the population. The second is that typically mutual information is smaller than both the response and the noise entropy: an impact that may be proportionally small for entropy may be proportionally much larger when considering information. Thus, we had to introduce a new metric based on comparing the amount of Shannon information about the stimuli which is obtained when considering either the full response model or the one taking into account only correlations up to a given order k. The novel finding obtained using this metric was that interaction of order two gave a a very good (but not perfect) approximation of information carried by the neural population. These results extended a previous attempt from our group [20] to measure Shannon information from pairwise models which made the simplifying assumptions that all neurons and interactions are perfectly homogenous. This simplifying assumption was no longer used or necessary in the present work.

Another potentially interesting metric is to consider the information lost by a downstream decoder if it was to ignore higher order correlation when decoding. This metric has recently been investigated by Oizumi and colleagues [43].

4.3. Scaling of the effect of higher order with the population size

In this work, we decided to analysis populations responses made of spiking activity collected from up to 8 simultaneously recorded locations. Fixing the population size to 8 was somewhat arbitrary, but it was empirically the largest population size that allowed us to compute reliably, with the present methodology, both the goodness of fit of ME models and the information carried by the neuronal population. Analyzing 8 simultaneous recording sites was a significant improvement with respect to previous analyses based on Shannon information (which were usually limited to two or three cells), but is of course still far from enough to understand the behavior of the cortical network. A crucial question for the interpretation of these results is how the significance of interaction of order higher than two, and the information they carry, scales when increasing the population size. In fact, it would be unwise to assume that the behavior of these measures extrapolates smoothly from small to large population sizes [22]. Measuring information from large populations will require to further improve the techniques for estimating information from limited sampling [34], which is a topic of current research for neurostatisticians.

Acknowledgements

We are grateful to R Senatore, R Petersen and MA Montemurro for many useful discussions and interactions. This research was conducted under the framework of the BMI Project at the Department of Robotics, Brain and Cognitive Sciences of the Italian Institute of Technology. RAAI was supported by the EPSRC CARMEN grant.

References

[1] Nirenberg S H and Victor J D 2007 *Curr. Opin. Neurobiol.* **17** 397–400
[2] Li C L 1959 *Science* **129** 783–4
[3] Mastronarde D 1983 *J. Neurophysiol.* **49** 303–24
[4] Abeles M, Bergman H, Margalit E and Vaadia E 1993 *J. Neurophysiol.* **70** 1629–38
[5] Luczak A, Bartho P and Harris K D 2009 *Neuron* **62** 413–25
[6] Abbott L and Dayan P 1999 *Neural Comput.* **11** 91–101
[7] Oram M W, Földiák P, Perrett D I, Oram M W and Sengpiel F 1998 *Trends Neurosci.* **21** 259–65
[8] Pola G, Thiele A, Hoffmann K P and Panzeri S 2003 *Network: Comp. Neural.* **14** 35–60
[9] von der Malsburg C 1999 *Neuron* **24** 95–104
[10] Chase S and Young E 2007 *P. Natl. Acad. Sci. USA* **104** 5175–80
[11] Gollisch T and Meister M 2008 *Science* **319** 1108–11
[12] Schneidman E, Bialek W and Berry M J 2003 *J. Neurosci.* **23** 11539–53

[13] Zohary E, Shadlen M N and Newsome W T 1994 *Nature* **370** 140–43

[14] Mazurek M and Shadlen M 2002 *Nat. Neurosci.* **5** 463–71

[15] Bohte S, Spekreijse H and Roelfsema P 2000 *Neural Comput.* **12** 153–79

[16] Nakahara H and Amari S I 2002 *Neural Comput.* **14** 2269–316

[17] Shlens J, Field G D, Gauthier J L, Grivich M I, Petrusca D, Sher A, Litke A M and Chichilnisky E J 2006 *J. Neurosci.* **26** 8254–66

[18] Schneidman E, Berry II M J, Segev R and Bialek W 2006 *Nature* **440** 1007–12

[19] Tang A, Jackson D, Hobbs J, Chen W, Smith J L, Patel H, Prieto A, Petrusca D, Grivich M I, Sher A, Hottowy P, Dabrowski W, Litke A M and Beggs J M 2008 *J. Neurosci.* **28** 505–18

[20] Montani F, Ince R A A, Senatore R, Arabzadeh E, Diamond M and Panzeri S 2009 *Phil. Trans. R. Soc.* A **367** 3297–310

[21] Roudi Y, Tyrcha J and Hertz J 2009 *Phys. Rev.* E **79** 051915

[22] Roudi Y, Nirenberg S and Latham P E 2009 *PLoS Comput. Biol.* **15** e1000380

[23] Jaynes E 1957 *Phys. Rev.* **106** 620–30

[24] Martignon L, Deco G, Laskey K, Diamond M, Freiwald W and Vaadia E 2000 *Neural Comput.* **12** 2621–53

[25] Montemurro M A, Senatore R and Panzeri S 2007 *Neural Comput.* **19** 2913–57

[26] Amari S I 2001 *IEEE Trans. Inform. Theory* **47** 1701–11

[27] Cover T M and Thomas J A 2006 *Elements of Information Theory, 2nd Ed.* (John Wiley & sons)

[28] Amari S I and Nagaoka H 2000 *Methods of Information Geometry* (Oxford University Press)

[29] Ince R A A, Petersen R S, Swan D C and Panzeri S 2009 *Front. Neuroinformatics* **3** 4

[30] Arabzadeh E, Petersen R S and Diamond M E 2003 *J. Neurosci.* **23** 9146–54

[31] Arabzadeh E, Panzeri S and Diamond M E 2004 *J. Neurosci.* **24** 6011–20

[32] Petersen R S and Diamond M E 2000 *J. Neurosci.* **20** 6135–43

[33] Woolf B 1972 *Ann. Hum. Genet.* **35** 397–409

[34] Panzeri S, Senatore R, Montemurro M A and Petersen R S 2007 *J. Neurophysiol.* **98** 1064–72

[35] Panzeri S, Schultz S R, Treves A and Rolls E T 1999 *Proc. R. Soc.* B **266** 1001–12

[36] Schnitzer M J and Meister M 2003 *Neuron* **37** 499–511

[37] Schrader S, Gruen S, Diesmann M and Gerstein G 2008 *J. Neurophysiol.* **100** 2165–76

[38] Harris K D 2005 *Nat. Rev. Neurosci.* **6** 399–407

[39] Riehle A, Gruen S, Diesmann M and Aertsen A 1997 *Science* **278** 1950–53

[40] Gruen S, Diesmann M and Aertsen A 2002 *Neural Comput.* **14** 43–80

[41] Amari S I, Nakahara H, Wu S and Sakai Y 2003 *Neural Comput.* **15** 127–42

[42] Schneidman E, Still S, Berry M J and Bialek W 2003 *Phys. Rev. Lett.* **91** 238701

[43] Oizumi M, Ishii T, Ishibashi T, Hosoya T and Okada M 2009 *Advances in Neural Information Processing Systems* vol 21 ed Koller D, Schuurmans D, Bengio Y and Bottou L (Cambridge: MIT press) pp 1225–32