Reliable recall of spontaneous activity patterns in chaotic cortical networks

Olivier Marre, Pierre Yger, Andrew Davison, Yves Frégnac

RELIABLE RECALL OF SPONTANEOUS ACTIVITY PATTERNS IN CHAOTIC CORTICAL NETWORKS

O. Marre & P. Yger, A.P. Davison and Y. Frégnac
Unit des Neurosciences Intégratives et Computationnelles
C.N.R.S, 1 Avenue de la Terrasse
91198, Gif sur Yvette, France
{name}@unic.cnrs-gif.fr

ABSTRACT

Irregular ongoing activity in cortical networks is often modelled as a result of recurrent connectivity. Yet it remains unclear how its presence corrupts the signal transmission evoked by the sensory drive. Here we show that reproducible responses in a generic recurrent cortical-like network can be obtained if the imposed external drive is consistent with patterns previously seen in the spontaneous, self-sustained activity. A subset of neurons in the network, constrained to replay a spiking pattern previously recorded during spontaneous activity, reliably drives the remaining, free-running neurons to reproduce the rest of the pattern. Comparison with surrogate Poisson patterns indicates that such transmission of input patterns is optimal for inputs with the statistical properties of the spontaneous activity. We propose that the similarity between spontaneous and evoked activity in sensory cortical areas could be the signature of learned efficiency of transmission across cortical networks.

KEY WORDS

Neuronal network, ongoing activity, reproducibility, information transmission, chaos

1 Introduction

Understanding the efficiency of signal transmission across primary sensory cortical areas remains an open challenge. On the one hand, a high level of trial-to-trial variability in evoked sensory responses is frequently reported, imposing an averaging of neural activity over time and cell assemblies to extract information, in a classical rate code scheme [1]. On the other hand, low trial-to-trial variability has been reported in both sub-cortical (retina, LGN, the trigeminal system) and cortical areas (V1, A1), implying a neural code based on the precise, absolute timing of spikes, such as is predicted by a number of theories (review in [2]), including the synfire chain theory [3]. Resolving the apparent discrepancy between experimental reports will require an understanding of the neuronal and network mechanisms underlying the measured variability.

Since neuronal integration and spike generation are highly reliable for fluctuating current inputs injected at the soma [4], much of the response variability observed *in vivo* must originate from the connectivity of the neural circuits. Even in the absence of external drive, ongoing neuronal activity is highly irregular, characterized by a coefficient of variation above 0.8 [5], although the discharge statistics are still a matter of debate. As an alternative to modelling ongoing activity as an external noise source uncorrelated with the signal [6, 7], models of recurrent networks have been proposed that explain the irregular activity as being self-generated [8, 9, 10]. Such models, characterized by large excitatory and inhibitory populations sparsely connected by relatively strong synapses, are able to display a dynamic but stable regime called Asynchronous Irregular (AI) [11] which closely resembles the spontaneous activity observed *in vivo*. To achieve reliable signal propagation within such a chaotic recurrent network, some studies introduced a change in the network structure, either by selectively and substantially increasing synaptic weights for a rate-based input [10] or by adding connections along a predetermined propagation path. Even in this latter case, the synfire chain stimulation can induce "synfire explosions", which can subsequently silence the network activity. Several questions remain to be answered: i) Does response reliability in these networks depend on stimulus features ? ii) Is there a specific class of stimuli that maximizes the reliability of signal transmission ?

While previous studies with AI networks tried to adapt the connectivity matrix and the path selection rules in such a way as to optimize the transmission of the chosen (*a priori*) input patterns, we decided to explore the converse approach: leave the network structure unchanged and find the activity patterns that are best propagated and sustained by the network. Since the patterns found in on-going activity are by definition highly compatible with the recurrent architecture of the network, we propose to drive the network with temporal pattern segments extracted from the recorded ongoing activity of the exact same network. We have designed a new stimulation paradigm mimicking the spontaneous activity of the network model, which we refer to as the "frozen paradigm". We consider an active network containing two interconnected populations of neurons, only one of which receives external input. The selection of which neuron belongs to which population is made at random and the connectivity of two neurons is independent of which populations they belong to. We first record a spontaneous pattern across the whole network (both populations) and then, while the spontaneous activity is ongoing, we force the neurons of the input-recipient population to replay the sequence of spontaneous activity previously recorded. We then measure the extent to which these clamped or frozen cells drive the free-running neurons to replay the spontaneous pattern previously recorded. Replay of the recorded pattern by the free-running neurons is then equivalent to the propagation and transmission of that pattern through the network.

Our study investigates the recall efficiency and reliability and the noise resistance achieved during frozen pattern stimulation. Simulations show that our paradigm produces efficient convergence, even in a mathematically-defined chaotic network, and that this is preserved over a broad range of parameters. Finally, we discuss biological implications of the possible cooperation between stimulus-evoked and ongoing activity in sensory cortical networks.

2 Results

2.1 Convergence of the network activity to a target activity pattern

We applied our paradigm to a generic, sparsely connected (2%), recurrent network of 8000 excitatory and 2000 inhibitory integrate and fire neurons. In order to make the present study comparable with previously published recurrent models, the set of parameter values was that used in [10] to put the network in an asynchronous irregular regime in which it generates self-sustained activity with a mean rate of 10.2 Hz and a mean coefficient of variation of the inter-spike interval (ISI CV) of 1.55. When 50% of all neurons are forced to replay a spontaneous pattern recorded previously, the spiking activity of the free-running neurons converges reliably to the target activity : repetitions of the same stimulation elicit temporally precise and reproducible spikes characterized by their temporal alignment across repetitions. The sub-threshold membrane potentials V_{m} of the free-running neurons closely follow to the target activity waveformsas soon as the input-recipient population is frozen.

2.2 Influence of input temporal structure on reliability

We next studied the influence of the high-order temporal structure of the stimulation patterns on the reliability of signal transmission. It could be argued that the reliability of the responses could have been obtained with any other imposed stimulation. We thus compared the frozen stimulation with surrogate stimulations having the same number of spikes, but a shuffled temporal structure. Several different surrogates were compared: a first series consisting of temporally "jittered" patterns with standard deviation of the jitter ranging from 5 to 25 ms; and a second series consisting of Poisson spike trains with the mean firing rate for each neuron equal either to (i) that of the same neuron during the reference, frozen pattern (called "local Poisson" patterns (LP), equivalent to an infinite jitter) or to (ii) the mean firing rate of the whole network during the reference pattern ("global Poisson" pattern (GP)). Note that all these controls have approximately the same number of spikes as the reference spontaneous pattern. These different surrogate patterns represent a progressively increasing deviation from spontaneous activity statistics, while keeping first-order statistics unchanged.

For this comparison, we chose a level of frozen neurons (50%) at which each free neuron receives equal numbers of connections from frozen and free neurons. The more the stimulation statistics deviate from those of spontaneous activity the lower is the reliability. This difference is significant for jitter of 15 ms and greater ($p < 0.001$) and for both Poisson surrogates ($p < 0.0001$), being largest for the global Poisson process. These changes in reliability are accompanied by a reduction in evoked activity in the free running network: the mean firing rate drops significantly ($p < 0.0001$) for local and global Poisson surrogate patterns. Taken together, these measures imply a larger degradation in signal transmission, the more the input statistics deviate from the spontaneous statistics. Thus, for the same number of spikes, the forced replay of spontaneous ongoing patterns seems to induce a local optimality in signal transmission.

To extend further our exploration of the network response to Poisson stimulation, we explored the effect on mean activity and reliability of independently varying the mean excitatory and inhibitory firing rates of the global Poisson stimulation (note that local Poisson stimulation gives qualitatively similar results). The combination of high excitatory and low inhibitory stimulation rates produces a high mean firing rate in the free-running units, but a low trial-to-trial reliability. In contrast, low excitatory but high inhibitory stimulation rates induce a highly reliable response, but with a very low mean firing rate. We empirically fitted this inverse relationship between reliability and mean activity of the free running units by the following power law:

$$\mathrm{Rate}^{\alpha} \cdot \mathrm{Reliability}^{\beta} = K \cdot (\mathrm{Spike\ count}) \qquad (1)$$

($r^2 = 0.95$) where the spike count is the total number of spikes in the imposed stimulation pattern, summed over all the frozen neurons, and K is a constant ($\alpha = 0.81$, $\beta = 1.88$, $K = 11.02$).

When compared to the Poisson surrogate stimulations, the particular structure of the imposed spontaneous pattern seems to be best adapted to the network connectivity: in order to reach similar levels of both reliability and response strength as observed for spontaneous statistics, the use of Poisson stimulation would require a large increase (160% of the spontaneous rate) in the spike frequency imposed on the frozen neurons. Furthermore, no spontaneous activity persists after the end of the Poisson surrogate stimulations: in 75% of trials, and independent of stimulation firing rate, the network can no longer generate stable, self-sustained spontaneous activity and falls silent. Taken together, these results indicate that the structure of ongoing activity enables more efficient and stable signal transmission than uncorrelated stimulations, whose mismatch with the underlying connectivity tends to impair further self-sustained activity in the same network.

3 Conclusion

In this article we have shown that reproducible activity patterns in a generic, recurrent cortical-like network can be obtained if the statistics of the imposed external drive match those of spontaneous, self-sustained activity. When a sufficiently large fraction of neurons within the network is constrained to replay a pattern generated during spontaneous activity, this "frozen" sub-population reliably drives the remaining, free-running neurons to reproduce the target activity pattern, a process that can be considered, equivalently, as the recall of a previously experienced pattern, and the completion of the full pattern. In exploring further the requirements for optimal recall, we showed that those input patterns which share high-order statistical similarities with spontaneously generated activity are the ones which elicit the most reliable recall.

Our results, obtained in generic neural networks, could be a general framework for pattern propagation and recall in cortex. This dynamic convergence towards reproducible neural assembly patterns may support pattern completion within associative memory networks and, more generally, efficient message transmission within any asynchronous irregular activity regime. The theoretical novelty of our model is to demonstrate the feasibility of a highly temporally precise encoding of stimuli within chaotic background activity.

Previous modelling studies of self-sustained recurrent neuronal networks have used ad-hoc stimulation patterns to drive the trajectory of the network away from its spontaneous attractor. Conversely, in our case, we let the activity stay on the spontaneous attractor and make it converge to one of its previously explored sub-trajectories. First, reliably driving the activity into a spontaneous trajectory needs fewer spikes than driving it reliably out of the attractor. This latter point is demonstrated by our surrogate Poisson stimulations which can induce efficient transmission when using many more spikes, but cannot induce a self-sustained pattern persisting after the end of the stimulation. Second, in our paradigm, the interest of the high dimensionality of the spontaneous AI regime is to give access to a huge diversity of different spontaneous patterns, enabling the transmission of many different messages across the same network. The AI regime may thus be an optimal state in terms of representation.

4 Acknowledgements

We thank Sami El Boustani, Dr Alain Destexhe and Dr Cyril Monier for helpful discussions and critical comments. This work was supported by CNRS, ANR (Natstats) and European Commission-funded (Bio-I3: Facets FP6-2004-IST-FETPI 15879) grants to Y.F.

References

[1] M. N. Shadlen and W. T. Newsome. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci*, 18(10):3870–3896, 1998.

[2] P. König, A. K. Engel, and W. Singer. Integrator or coincidence detector? the role of the cortical neuron revisited. *Trends Neurosci*, 19(4):130–137, 1996.

[3] M. Abeles. *Local cortical circuits: an electrophysiological study*. Berlin Springer-Verlag, 1982.

[4] Z. F. Mainen and T. J. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.

[5] I. Timofeev, F. Grenier, M. Bazhenov, T. J. Sejnowski, and M. Steriade. Origin of slow cortical oscillations in deafferented cortical slabs. *Cereb Cortex*, 10(12):1185–1199, 2000.

[6] M. C. W. van Rossum, G. G. Turrigiano, and S. B. Nelson. Fast propagation of firing rates through layered networks of noisy neurons. *J Neurosci*, 22(5):1956–1966, 2002.

[7] M. Diesmann, M. O. Gewaltig, and A. Aertsen. Stable propagation of synchronous spiking in cortical neural networks. *Nature*, 402(6761):529–533, 1999.

[8] H. Sompolinsky, A. Crisanti, and H. J. Sommers. Chaos in random neural networks. *Phys Rev Lett*, 61(3):259–262, 1988.

[9] C. van Vreeswijk and H. Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726, 1996.

[10] T. P. Vogels and L. F. Abbott. Signal propagation and logic gating in networks of integrate-and-fire neurons. *J Neurosci*, 25(46):10786–10795, 2005.

[11] N. Brunel. Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *J Comput Neurosci*, 8(3):183–208, 2000.