6. Williams, H. and Mehta, N. (1999) Changes in adult zebra finch song require a forebrain nucleus that is not necessary for song production. *J. Neurobiol.* 39, 14–28
7. Hessler, N.A. and Doupe, A.J. (1999) Singing-related neural activity in a dorsal forebrain-basal ganglia circuit of adult zebra finches. *J. Neurosci.* 19, 10461–10481
8. Leonardo, A. (2004) Experimental test of the birdsong error-correction model. *Proc. Natl. Acad. Sci. U. S. A.* 101, 16935–16940
9. Kao, M.H. *et al.* (2005) Contributions of an avian basal ganglia–forebrain circuit to real-time modulation of song. *Nature* 433, 638–643
10. Olveczky, B.P. *et al.* (2005) Vocal experimentation in the juvenile songbird requires a basal ganglia circuit. *PLoS Biol* 3, e153
11. Tumer, E.C. and Brainard, M.S. (2007) Performance variability enables adaptive plasticity of 'crystallized' adult birdsong. *Nature* 450, 1240–1244
12. Andalman, A.S. and Fee, M.S. (2009) A basal ganglia–forebrain circuit in the songbird biases motor output to avoid vocal errors. *Proc. Natl. Acad. Sci. U. S. A.* 106, 12518–12523
13. Mandelblat-Cerf, Y. *et al.* (2014) A role for descending auditory cortical projections in songbird vocal learning. *Elife* 3, e02152
14. Gadagkar, V. *et al.* (2016) Dopamine neurons encode performance error in singing birds. *Science* 354, 1278–1282
15. Hisey, E. *et al.* A common neural circuit mechanism for internally guided and externally reinforced forms of motor learning. *Nat. Neurosci.* (in press)

Series: Seminal Neuroscience Papers 1978–2017

Science & Society

Neuronal Variability as a Proxy for Network State

Ramon Nogueira,[1,2,*] Sofía Lawrie,[2] and Rubén Moreno-Bote[2,3]

Characterizing how network state modulates cortical dynamics and information processing is an important step for understanding the neural code. In 2010, Churchland *et al.* reported wide experimental evidence showing that spontaneous and stimulus-evoked conditions are two distinct states, as indicated by a marked reduction of neuronal variability after stimulus onset.

The brain is an evolution-shaped dynamical system optimized for guiding the actions that maximize the chances of survival. As an information processing system, one could expect the brain to respond with identical activity patterns when presented with identical inputs. However, a large body of experiments has revealed that cortical responses are variable in both the number of spikes and their precise timings [1,2]. This variability is found even under controlled experimental conditions, for instance when the environmental variables are kept constant, when eye movements are accounted for by fixation, or even when behavior is suppressed by anesthesia [3].

In addition to its consistency across experimental conditions, neuronal variability can be characterized by some distinctive statistical traits. Within trials, the distribution of inter-spike intervals (ISIs) has been typically shown to be broad with a long tail, resembling an exponential distribution [4]. Across trials, variability can be assessed by the Fano factor (FF), which is defined as the ratio between the variance of the spike counts and their mean. A linear relationship with a slope close to one between the spike count variability and its mean has been experimentally shown for single neurons [2], although for high firing rates a supralinear regime has been reported [5]. The statistical trait that the FF is roughly constant over broad ranges of neuronal firing rates is referred to as 'Poisson-like firing' [6] owing to its resemblance to a Poisson process (rate-independent FF = 1).

A natural question arises at this point: can diverse network states be characterized by different statistical traits of spiking variability? And what can these differences in neuronal variability reveal about cortical dynamics and information processing? Mark Churchland and colleagues [7] took these questions in the context of spontaneous (ongoing) versus evoked (stimulus-driven) cortical activity, and analyzed whether they correspond to different states with distinct statistical properties. One possible scenario could be that, before stimulus onset, cortical activity is prepared for action – 'waiting', so to say, within a small region in the abstract, mathematical space of neuronal activity states – and that once the network receives an input, it diverges to a larger region in that space ($H_{1\text{-}1}$ and $H_{1\text{-}2}$; Figure 1A). Another possible scenario could be that spontaneous activity lives within a large region, and that after stimulation the available space of cortical activity shrinks ($H_{2\text{-}1}$ and $H_{2\text{-}2}$; Figure 1B). The first scenario would be characterized by a larger trial-by-trial

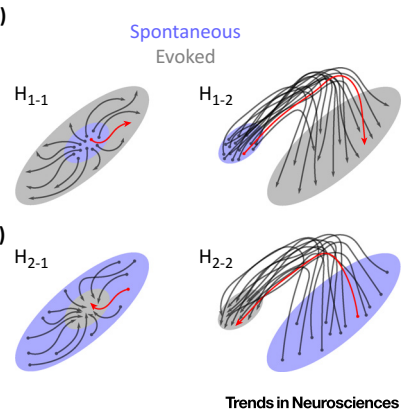

Figure 1. Hypothetical Possible Changes in Neuronal Variability during the Transition from Spontaneous to Stimulus-Evoked Conditions. (A) In hypothesis 1 (H_1), cortical activity in the spontaneous condition (blue) 'lives' within a small region in the neural activity space (defined as a multidimensional space where each dimension is the activity of a single neuron). Once the network receives an input, it diverges to a larger set of available space (grey). In this scenario, trial-by-trial variability would be larger at evoked versus spontaneous activity. Hypothesis 1 can be further refined into two sub-categories: in $H_{1\text{-}1}$, the spontaneous activity space resides within the evoked activity space, whereas in $H_{1\text{-}2}$ it resides outside it. (B) In hypothesis 2 (H_2), cortical activity in the spontaneous condition (blue) encompasses a larger region of the neuronal activity space, and the available space shrinks (grey) after stimulus presentation. In this scenario, trial-by-trial variability would decrease after stimulus onset. In $H_{2\text{-}1}$ the evoked activity space resides within the spontaneous space, whereas in $H_{2\text{-}2}$ it resides outside it.

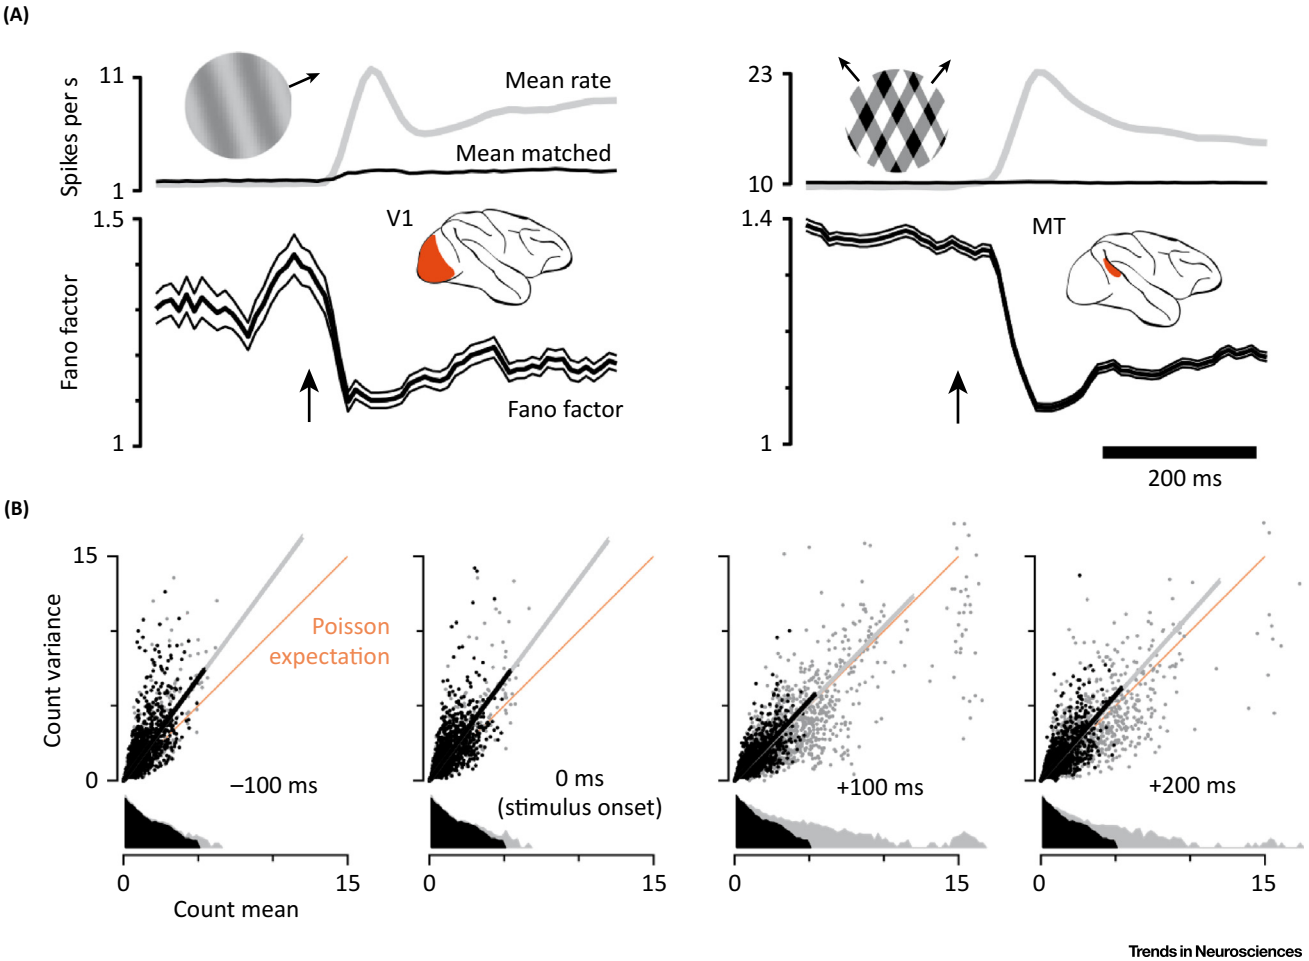

Figure 2. Stimulus Onset Quenches Neuronal Variability Regardless of Stimulus-Evoked Activity Modulations. (A) Fano factor (FF), as a proxy for neuronal variability, is reduced after stimulus onset in monkey cortical areas V1 (primary visual area, left panel) and MT (middle temporal area, right panel) (datasets encompassing other brain areas, stimuli, and brain states are given in [7]). Data are aligned to stimulus onset (arrow). The mean rate (top row, grey), the 'mean matched' rate (top row, black) and the FF (bottom row, black with flanking SEM) were computed using a 50 ms sliding window. (B) The 'mean matching' procedure explained graphically. The raw (grey line) and the mean-matched (black line) FF correspond to the slope obtained after regressing spike count variance versus the mean spike count (orange line, FF = 1). Each datapoint in each time-window (panels) corresponds to the mean spike count and variance of a given neuron and experimental condition. To control for the stimulus-evoked increase in activity when calculating the FF, points from the original distribution of mean spike counts for each time-window (grey histogram) were randomly excluded until the distribution matched the greatest common distribution across time-windows (black histogram). Controlling for mean activity did not modify quantitatively the reported results. Figure adapted, with permission, from [7].

variability during the evoked than the ongoing regime, whereas the latter would be characterized by the opposite situation. Interestingly, experimental evidence [8] and theoretical predictions [9] seemed to indicate that the reduction of response variability during evoked activity could be a general property of cortical activity.

Churchland and colleagues explored the effects of stimulus onset upon the trial-by-trial activity variability in a large set of experimental datasets, spanning different brain areas, stimuli, and brain states – representing a type of comprehensive and comparative analysis that should be more common in neuroscience. The analyses were conducted using a variety of statistical metrics, but the main result was based on the FF as a proxy for spiking variability and its time modulation through the course of the trial. Consistently across 14 datasets, the authors reported a significant decrease in single-cell variability during evoked activity compared to the variability during the

spontaneous condition (Figure 2A). Because a decrease in the FF could simply be a consequence of an increase in the mean firing rate at stimulation, which is typically observed, a novel method (mean matching) was introduced to rule out this possibility. This method allowed constant firing-rate distributions to be created from the originally recorded neuronal populations before and after stimulation by randomly discarding datapoints until all the corresponding histograms matched across time (Figure 2B). Even after controlling for firing-rate differences in spontaneous and evoked conditions by this method, the reported FF decreased during stimulation.

Once the effect on the variability of single neurons was shown to be firing-rate-independent and consistent across all datasets, Churchland *et al*. explored the variability at the population level using simultaneously recorded neurons. The covariance matrix of the neuronal spike counts was decomposed as the sum of two matrices, one representing the variance due to single-neuron spiking variability, and another accounting for shared variability, common to all neurons in the population. This decomposition showed that the decrease in variability originated predominantly from factors that were common to the entire population. This observation led the authors to an interpretation where shared variability can be seen as a variable that is modulated by network state.

Altogether, the study by Churchland and colleagues represented a milestone in the field for several reasons. This work has been an inspiration for a set of theoretical studies that aimed to reproduce the main experimental results using neuronal network models [10–12]. These studies recapitulated the experimentally observed decrease of the FF after stimulus onset, but two different mechanisms underlie the decrease in variability in these models. In one of these studies [10] the spontaneous trial-by-trial variability arises because of the chaotic dynamics that characterize large recurrent networks. When this network is stimulated, it undergoes a phase transition to a non-chaotic regime, reflected in a decrease in trial-by-trial variability. In the other account [11,12], the network has access to a different number of stable attractors during ongoing and evoked activity. During spontaneous activity, the network is found in a multi-attractor regime; therefore, it 'visits' different stable states by means of its inherent chaotic behavior, which is translated into high variability. After stimulus onset, the set of available and visited attractors is reduced, and variability across trials decreases. It remains an open question which of the interpretations is closer to the physiological reality underlying the findings by Churchland *et al*., and other accounts might also be possible. It is also plausible that, despite the consistency of the reduction of variability phenomenon as observed experimentally across brain regions, the underlying physiological mechanism differs from one cortical circuit to another.

The work of Churchland *et al*. has also inspired follow-up experimental studies that characterized the link between variability and network properties, or even between variability and cognitive state. Hussar and Pasternak [13], for instance, examined the link between variability and the engagement of the animal in the task, and reported a reduction of the FF at relevant task periods compared to spontaneous activity, as well as an inverse relationship between variability and behavioral performance. Based on the rationale that low variability indicates that the network more closely represents the variables that are important for the task, instead of wandering around other irrelevant states, FF modulations are nowadays taken as a popular indicator of network and animal engagement with the task at hand.

The consistency of the reduction of variability during evoked compared to spontaneous activity in the work of Churchland *et al*. is striking. The precise implications, however, in terms of the possible scenarios depicted in Figure 1 (or other accounts) remain unclear. Does the reduction in variability imply that the available activity space is larger during spontaneous activity [14]? Do evoked responses occupy a different region of activity space, or are they embedded in the spontaneous region ($H_{2\text{-}2}$ vs $H_{2\text{-}1}$, respectively, in Figure 1B)? Is the reduction in variability during evoked activity only a byproduct of statistical conditioning? In the latter case, during spontaneous activity there could be larger variability simply because of lack of conditioning and control of, for instance, previous history variables [15]. In sum, it is crucial, we think, to further clarify the relationship between variability and available activity space (as a proxy of network state). Nailing down these questions would require novel metrics to better quantify the size of 'visited' and 'available' sections of the neuronal activity space. We hope that future work will help to resolve these questions and advance our understanding of the roles of neural variability.

Acknowledgments

We thank Ashok Litwin-Kumar for comments on the manuscript. R.M.B. is supported by BFU2017-85936-P and FLAGERA-PCIN-2015-162-C02-02 from the Ministry of Economy and Competitiveness (MINECO, Spain) and an International Scholar Award from the Howard Hughes Medical Institute (ref. 55008742).

[1]Center for Theoretical Neuroscience, Mortimer B. Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY, USA
[2]Center for Brain and Cognition, and Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain
[3]Serra Húnter Fellow Programme, Universitat Pompeu Fabra, Barcelona, Spain

*Correspondence: rn2446@columbia.edu (R. Nogueira).
https://doi.org/10.1016/j.tins.2018.02.003

References

1. Tolhurst, D.J. *et al.* (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* 23, 775–785
2. Shadlen, M.N. and Newsome, W.T. (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* 18, 3870–3896
3. Kohn, A. and Smith, M.A. (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J. Neurosci.* 25, 3661–3673
4. Softky, W.R. and Koch, C. (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* 13, 334–350
5. Goris, R.L.T. *et al.* (2014) Partitioning neuronal variability. *Nat. Neurosci.* 17, 858–865
6. Moreno-Bote, R. (2014) Poisson-like spiking in circuits with probabilistic synapses. *PLoS Comput. Biol.* 10, e1003522
7. Churchland, M.M. *et al.* (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* 13, 369–378
8. Monier, C. *et al.* (2003) Orientation and direction selectivity of synaptic inputs in visual cortical neurons: a diversity of combinations produces spike tuning. *Neuron* 37, 663–680
9. Sussillo, D. and Abbott, L.F. (2009) Generating coherent patterns of activity from chaotic neural networks. *Neuron* 63, 544–557
10. Abbott, L.F. *et al.* (2011) Interactions between intrinsic and stimulus-evoked activity in recurrent neural networks. In *The Dynamic Brain: An Exploration of Neuronal Variability and Its Functional Significance* (Ding, M. and Glanzman, D. L., eds), pp. 65–82, Oxford University Press
11. Deco, G. and Hugues, E. (2012) Neural network mechanisms underlying stimulus driven variability reduction. *PLoS Comput. Biol.* 8, e1002395
12. Litwin-Kumar, A. and Doiron, B. (2012) Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.* 15, 1498–1505
13. Hussar, C. and Pasternak, T. (2010) Trial-to-trial variability of the prefrontal neurons reveals the nature of their engagement in a motion discrimination task. *Proc. Natl. Acad. Sci. U. S. A.* 107, 21842–21847
14. Luczak, A. *et al.* (2009) Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* 62, 413–425
15. Nogueira, R. *et al.* (2017) Lateral orbitofrontal cortex anticipates choices and integrates prior with current information. *Nat. Commun.* 8, 14823