

## HISTORICAL NEWS AND VIEWS: NEURAL CODING

## Whither variability?

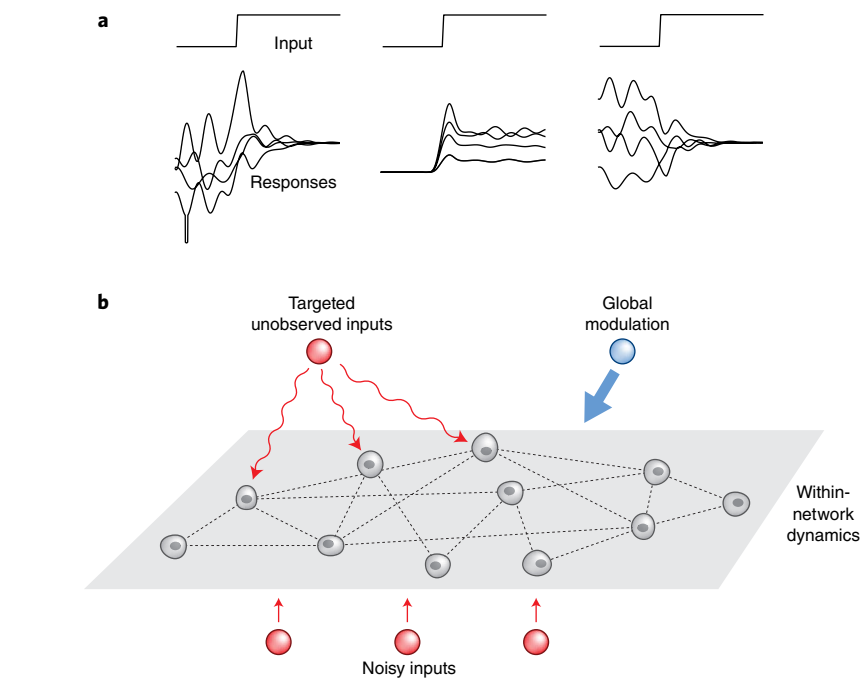
Variability is a ubiquitous aspect of neural recordings. In an influential paper, Churchland et al. (2010) compiled data from many cortical areas to demonstrate that variability generally decreases upon presentation of a stimulus. What are the implications of this finding?

Adrienne L. Fairhall

In a standard artificial neural network trained for image classification, units are active when inputs are presented to it and silent otherwise. In contrast, the brain, right down to the early visual system, displays a swirling mess of ongoing activity, even in the apparent absence of stimuli. The form of this internal activity, how it interacts with incoming information, and what all of this tells us about the nature of neural computation are long-standing puzzles, the answers to which are truly fundamental to our understanding of brain processing.

Alongside the dynamic variations that constitute spontaneous activity, sensory responses are also variable, in that the same input generally leads to a spiking response that differs from trial to trial. The amount of variability varies with the strength of the input and the state of the animal; for instance, whether it is awake or anesthetized. In a now-classic paper from 2010, Churchland and collaborators<sup>1</sup> set out to determine the general characteristics of the relationship between the variability before stimulus onset and the variability during stimulus presentation in multiple brain areas. Given the prevalence of background fluctuation, several outcomes are all quite plausible a priori (Fig. 1): ongoing background activity might be the cause of the variability seen during stimulus presentation, in which case total variation will stay about the same; background variability might sum with additional stimulus-driven variability; there may even be a nonlinear interaction between stimulus and background that amplifies total variability; or the total variability might be reduced as the system engages in its presumed function.

Different outcomes would suggest different interpretations with respect to information processing. If variability is simply noise, active mechanisms may work to suppress it to improve the decodability of sensory responses. Variations may also be signatures of state-dependent processing



**Fig. 1 | Scenarios for network variability.** **a**, Possible trial-by-trial dynamics suggested in ref. <sup>1</sup> at the onset of a stimulus: there may be a stimulus-driven decline in variability, a stimulus-driven increase in variability, or a stimulus-driven decline in variability with little change in mean rate. Image reproduced from ref. <sup>1</sup>, Nature Publishing Group. **b**, Potential sources of variability in cortical firing include noise in inputs, nonlinear dynamics within the network itself, the effects of specific top-down or unobserved inputs, and global state-dependent modulation.

and/or serve computational functions that we have not yet fully grasped.

To explore this issue, teams from an impressive number of labs contributed data from multiple cortical areas toward a common analysis of variability change during stimulus onset. Results were compiled both from sensory areas during a stimulus presentation, including from visual areas V1 (recorded both extra- and intracellularly) and V4 and from middle temporal visual area, driven both by plaids and dots; and from motor-related areas during the execution of a stereotyped task, including from the lateral intraparietal area, the parietal reach region, dorsal

premotor cortex, and orbitofrontal cortex. A common result emerged in each of these cortical regions: stimulus or motor onset reduces, or ‘quenches’, the level of variability.

The analysis computed measures of variability both at the single-neuron level and across the network. Single-neuron variability was measured using the Fano factor, the ratio of spike-count variance to mean, computed across trials as a function of time in sliding time windows across stimulus (or movement) onset. Considerable care was taken to address potential confounds. It is possible, for example, that at the higher firing rates produced in

response to a stimulus, the spike counts will regularize due to refractory effects. Such effects were ruled out by controlling for mean firing rate. This approach yielded the surprising finding that even when a neuron did not change its mean firing rate, such that it would not have been identified as responsive to the stimulus, its variability nonetheless generally decreased upon stimulus onset.

This result pointed toward the importance of quantifying the variability of stimulus representation at the network level rather than singling out apparently responsive neurons. Multineuronal recordings from V1 and dorsal premotor cortex were analyzed using factor analysis, which parcels variability into 'private' and 'shared', or network, components. While private noise was ascribed to 'spiking process noise'—an attribution which may stand simply as a reasonable definition of what spike process noise is—the method allowed a convincing demonstration that the reductions in variability occur in a correlated way across the network.

What is the origin and meaning of neural variability? With some exceptions, very little variability arises from the process of spiking itself. Neurons fire precisely when driven by a fixed, time-varying input, although more randomly when the input is held stationary<sup>3</sup>, perhaps as a result of gain control amplifying the effects of small input fluctuations. While some variability is introduced through uncertainty in synaptic transmission, it is likely that the majority of variability observed in the brain arises from a combination of the influences of dynamic variables that have not been accounted for, as well as the intrinsic dynamics of networks themselves. These two contributions have quite different interpretations.

Variability due to ongoing brain processes reflects not noise but aspects of state that may be inaccessible to (or at least, unaccounted for by) the experimenter, yet are part of normal embodied computation: effects of respiration and arousal; top-down influences such as attention, motivation, and expectation; efference copies of movements; and additional sensory or cognitive variables<sup>3–5</sup>. It should be possible to peel off the effects, additive or modulatory, of these contributions to activity one by one, as has been the case in recent years as the quality of behavioral and

state monitoring has steadily improved<sup>6</sup>. This will in principle permit the design of experiments or selection of conditions during which these additional identified factors are controlled for.

Much theoretical work has focused on the second issue: generic models of neural networks intrinsically display chaotic dynamics. The amount of variation in a potentially chaotic network is determined by its inputs in ways that may cast mechanistic light on the Churchland et al. results. A steady input can drive an otherwise stable network into a chaotic regime. Yet a highly temporally structured input delivered to a chaotic network can clamp the network's variability enormously, reducing the number of dimensions of variation from on the order of the number of neurons to only a few, albeit dimensions that can vary in time<sup>7</sup>. The degree of variability reduction in chaotic networks also shows a dependence on stimulus characteristics such as frequency<sup>8</sup>, a finding that has some experimental support<sup>9</sup>.

Many observations suggest, however, that spontaneous activity is not purely chaotic, but rather is quite structured. Spontaneous activity in visual cortex has been seen to reflect the structure of natural activity patterns<sup>10</sup> and to do so increasingly throughout development<sup>11</sup>. It also has characteristics that recapitulate recently experienced inputs<sup>12</sup>. Researchers have thus speculated that spontaneous activity may reflect the carving into cortical networks of an internal model of the statistics of its inputs and, further, that instantaneous manifestations of activity may reflect random sampling from this learned distribution<sup>13</sup>.

A potentially alternate view, or at least one that requires reconciliation with previous findings, comes from recent widefield imaging studies that provide access to measures of neural activation across large regions of cortex. These images show that large-scale activity patterns propagate across cortical areas in waves<sup>14</sup>. Such spatiotemporally structured activity presumably also contributes to correlated variability. A study of stimulus-driven reduction of variability in V4 using widefield imaging found that, at these scales, the structure of spontaneous and driven variability is quite distinct and suggests that global activity patterns may underlie a significant fraction of

spontaneous correlated variability<sup>15</sup>. Analogous to the Churchland et al. findings, these shifts in noise correlation upon stimulus onset were observed even for stimulus contrasts that barely changed mean firing rates. The possible computational roles of such global activity patterns are as yet unclear<sup>14</sup>.

In sum, the questions both addressed and raised by the Churchland et al. study are profound ones whose resolution is still in progress, aided by emerging techniques and the kind of data analysis this group has helped to pioneer. In particular, the work drew early attention to the importance of higher-order statistics of neural activity, the study of which is now well supported by current recording methods. Finally, the work stands as an example of the type of conceptual question that can be addressed in a comparative way by applying a common analysis to multiple diverse datasets, serving as a motivation for the open sharing of data. □

Adrienne L. Fairhall

Dept. of Physiology and Biophysics and UW Institute for Neuroengineering, University of Washington, Seattle, Washington, USA.

e-mail: fairhall@u.washington.edu

Published online: 11 February 2019  
<https://doi.org/10.1038/s41593-019-0344-0>

#### References

- Churchland, M. M. et al. *Nat. Neurosci.* **13**, 369–378 (2010).
- Mainen, Z. F. & Sejnowski, T. J. *Science* **268**, 1503–1506 (1995).
- Niell, C. M. & Stryker, M. P. *Neuron* **65**, 472–479 (2010).
- McGinley, M. J. et al. *Neuron* **87**, 1143–1161 (2015).
- Hawrylycz, M. et al. *Proc. Natl Acad. Sci. USA* **113**, 7337–7344 (2016).
- Musall, S., Kaufman, M. T., Gluf, S. & Churchland, A. Preprint at *bioRxiv* <https://doi.org/10.1101/308288>.
- Lajoie, G., Lin, K. K. & Shea-Brown, E. *Phys. Rev. E* **87**, 052901 (2013).
- Rajan, K., Abbott, L. F. & Sompolinsky, H. *Phys. Rev. E* **82**, 011903 (2010).
- White, B., Abbott, L. F. & Fiser, J. *J. Neurophysiol.* **108**, 2383–2392 (2012).
- Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A. & Arieli, A. *Nature* **425**, 954–956 (2003).
- Berkes, P., Orbán, G., Lengyel, M. & Fiser, J. *Science* **331**, 83–87 (2011).
- Carrillo-Reid, L., Miller, J. E., Hamm, J. P., Jackson, J. & Yuste, R. *J. Neurosci.* **35**, 8813–8828 (2015).
- Orbán, G., Berkes, P., Fiser, J. & Lengyel, M. *Neuron* **92**, 530–543 (2016).
- Deneux, T. et al. *Neurophotonics* **4**, 031222 (2017).
- Muller, L., Chavane, F., Reynolds, J. & Sejnowski, T. J. *Nat. Rev. Neurosci.* **19**, 255–268 (2018).

#### Competing interests

The authors declare no competing interests.