


Journal Club

Editor's Note: These short reviews of recent *JNeurosci* articles, written exclusively by students or postdoctoral fellows, summarize the important findings of the paper and provide additional insight and commentary. If the authors of the highlighted article have written a response to the Journal Club, the response can be found by viewing the Journal Club at www.jneurosci.org. For more information on the format, review process, and purpose of Journal Club articles, please see <http://www.jneurosci.org/content/jneurosci-journal-club>.

Information-Limiting Correlations in Neural Populations: The Devil Is in the Details

 **Reeбал Rafeh**^{1,2,3,4,5} and **Geetika Gupta**^{1,3,4,5}

¹Department of Psychology, Western University, London, Ontario N6A 5C2, Canada, ²Department of Physiology and Pharmacology, Western University, London, Ontario N6A 5C1, Canada, ³Brain and Mind Institute, Western University, London, Ontario N6A 5B7, Canada, ⁴Robarts Research Institute, Western University, London, Ontario N6A 5K8, Canada, and ⁵Neuroscience Graduate Program, Western University, London, Ontario N6A 3K7, Canada

Review of [Bartolo et al.](#)

Neurons are noisy. Neuronal activity varies in response to repetitions of the same stimulus. Neuronal noise is ubiquitous in the brain and is often correlated among neighboring neurons, which means that stimulus-independent fluctuations in neuronal activity can affect entire populations of cells ([Schmitz and Duncan, 2018](#)). One might therefore reasonably expect that correlated neuronal noise poses a major challenge for the capacity of the brain to process and store information. How can one brain region reliably interpret a stimulus that was just presented from the noisy output of another region? Yet the brain somehow accomplishes this feat since our perceptions and actions are remarkably stable.

The influence of neuronal noise on population coding has become a focus of intense debate in computational neuroscience over recent decades. Some researchers have found evidence that correlated noise fundamentally constrains the amount of information that can be carried by neuronal population activity ([Cohen and](#)

[Maunsell, 2009](#); [Mitchell et al., 2009](#); [Verhoef and Maunsell, 2017](#)). By contrast, earlier work has shown that correlated noise is inconsequential to the amount of information carried by the population ([Averbeck and Lee, 2003, 2006](#)). In a recent report in *The Journal of Neuroscience*, [Bartolo et al. \(2020\)](#) conducted a series of electrophysiology experiments using multielectrode array recordings in macaques to reconcile these two seemingly conflicting lines of evidence.

[Bartolo et al. \(2020\)](#) recognized a major source of methodological variability in prior monkey electrophysiology experiments examining neuronal noise, namely the total number of neurons recorded in each animal. Using simulations of neuronal activity, previous work had yielded predictions of how neuronal noise impacts information processing at increasing population sizes ([Zohary et al., 1994](#); [Abbott and Dayan, 1999](#)). Based on these predictions, [Bartolo et al. \(2020\)](#) hypothesized that the impact of neuronal noise on population coding becomes more prevalent with increasing population size. If the influence of neuronal noise on population coding scales with the size of the recorded population, this might reconcile why different studies report a different influence of noise correlations on information.

To test their hypothesis, the authors leveraged high-density multielectrode arrays to simultaneously record from hundreds of neurons in two macaque monkeys. They placed the arrays over prefrontal regions known to be involved in the generation of eye movements. Neuronal activity was recorded from these arrays while the monkeys performed a simple task in which they were cued to perform a leftward or rightward saccade.

To estimate the amount of information in the recorded neuronal populations, [Bartolo et al. \(2020\)](#) examined distributions of population activity corresponding to either of the two saccade directions. The greater the overlap between the two response distributions, the less the information carried by the population for distinguishing the two conditions. With this measure of information in hand, [Bartolo et al. \(2020\)](#) next examined how correlated noise influences population coding as the recorded population grows in size.

[Bartolo et al. \(2020\)](#) found that correlated noise limits the amount of information carried by the population. However, this effect became apparent only when the population size surpassed ~350 neurons. Increasing the population size further, up to 700 neurons, revealed still larger effects of noise correlation on population coding. Thus, [Bartolo et al. \(2020\)](#) confirmed their hypothesis that the effect of noise on information encoded by a neuronal population

Received Sep. 2, 2020; revised June 14, 2020; accepted Sep. 2, 2020.

We thank Drs. Marieke Mur and Taylor Schmitz for comments and feedback in the preparation of this manuscript.

Correspondence should be addressed to Reeбал Rafeh rrefeh@uwo.ca.

<https://doi.org/10.1523/JNEUROSCI.0917-20.2020>

Copyright © 2020 the authors

increases with population size. They also provide a lower bound on the population size needed for detecting information-limiting correlations in prefrontal cortex during a simple saccade task. However, their results leave open how noise correlations arise, what their potential neurobiological function is, and what mechanisms the brain is endowed with for altering the dynamics of correlated noise in population coding. Below, we discuss each of these three issues through the lens of recent theoretical and empirical neuroscience research.

We first consider the issue of how noise correlations arise from population activity. To do so, we use the metaphor of a choir for our neuronal population, where each singer's voice is the spiking output of a neuron, and the driving stimulus is the choir conductor. What tune a singer sings (the spiking output of the neuron) is determined by both the input from the conductor (the stimulus) and their harmonization with neighboring singers' voices (the neuronal population). Now consider that the choir is composed of groups of neighboring singers, with each group specializing in a different pitch: the sopranos, altos, tenors, and basses. So too, cortical populations are composed of ensembles of neurons, with each ensemble preferring different stimulus features: a specific orientation, color, or eye movement. For a high note in the song, the conductor signals the sopranos, whose voices then synchronize and dominate the other groups in the choir. Similarly, an ensemble of neurons strongly activated by their preferred stimulus will synchronize and dominate other ensembles of neurons. They dominate not only because they are strongly excited by the stimulus, but also because they inhibit neighboring neurons with similar preferences (Verhoef and Maunsell, 2017). Thus, the outputs of the neighboring neurons are the joint result of excitatory stimulus inputs and inhibitory synchronized fluctuations originating from the dominant population. These fluctuations are correlated noise.

The noisy correlated fluctuations appear to represent an important mechanism by which neurons maintain accurate coding of information—they are not an accident but rather an important component of neurobiological function. Specifically, shared noise allows neurons to maintain an excitatory–inhibitory balance (Reynolds and Heeger, 2009; Carandini and Heeger, 2012; Verhoef and Maunsell, 2017). Consider a neuronal ensemble maximally excited by its preferred stimulus,

leading to the injection of shared noise into neighboring similarly tuned ensembles. These noisy fluctuations reduce the amount of information carried by other “competing” ensembles and, thus, inhibit their contribution to the population code. In this way, populations can dynamically code stimuli through continuous perturbation of this excitatory–inhibitory balance (Schmitz and Duncan, 2018).

This coding scheme works well in most cases, but it also leaves the brain open to mistakes. Consider a noisy fluctuation received by multiple similarly tuned neuronal ensembles at the same time. This wave of noise could momentarily “wobble” the distribution of outputs in an entire population such that it mimics a stimulus change, leading downstream populations to misinterpret what stimulus actually occurred. Recent theoretical and empirical work has demonstrated that this type of effect does indeed occur, though it accounts for a relatively small proportion of shared noise in the cortical population activity (Moreno-Bote et al., 2014; van Bergen et al., 2015; Nogueira et al., 2020). These forms of noise correlation are thus termed information-limiting noise because they distort the stimulus information carried by the population. It is this form of information-limiting noise that Bartolo et al. (2020) were interested in isolating. Their findings demonstrate that the detectable effect of information-limiting noise scales with the number of neurons recorded.

The information-limiting effects of correlated noise only become apparent at larger population sizes because the correlated noise then starts to dominate the population response (Zohary et al., 1994; Shadlen and Newsome, 1998; Averbeck et al., 2006). Let us consider a population consisting of five neurons versus a population consisting of a hundred neurons. Five neurons can make 20 pairs while a hundred neurons can make 9900 pairs. The number of possible pairs thus grows exponentially with population size. Shared fluctuations because of correlated noise between neuron pairs contribute to the fluctuations of the population response. Therefore, the relative contribution of correlated noise to the population response grows exponentially with population size. Now consider adding a neuron to an existing large population. The unique contribution of the neuron will be dwarfed by the correlated noise. Correlated noise thus places an upper bound on the stimulus information that can be gained by including additional neurons (Zohary et al., 1994; Shadlen and Newsome, 1998; Averbeck et

al., 2006). The upper bound is lower and reached more quickly with increasing strengths of correlated noise (Zohary et al., 1994). Because electrophysiology studies of monkey visual cortex often report analyses on similarly tuned populations of neurons (Mitchell et al., 2009; Smith and Sommer, 2013; Ruff and Cohen, 2014), these studies likely captured the effects of information-limiting correlations despite recording from smaller populations.

If noise correlations represent an important mechanism by which neuronal ensembles compete for dominance in the population code, and can also limit the amount of information carried in populations of neurons, the brain must be endowed with mechanisms for adjusting noise correlations. Over the past decade, multielectrode monkey electrophysiology research has shown that attention represents one such mechanism (Cohen and Maunsell, 2009; Mitchell et al., 2009; Rabinowitz et al., 2015; Verhoef and Maunsell, 2017). Attention is critical for biasing competition among similarly tuned neuronal ensembles to enable fine-grained discrimination of behaviorally relevant stimuli. This biasing mechanism appears to be accomplished in part by reducing shared inhibitory inputs arriving in the dominant (preferred) ensemble. Another mechanism the brain may use to deal with correlated noise is to read out information from neuronal populations that are diversely tuned (Ecker et al., 2011; Tremblay et al., 2015; Leavitt et al., 2017). If the sampled population has sufficiently heterogeneous tuning functions, information processing remains unaffected by correlated noise (Ecker et al., 2011).

Bartolo et al. (2020) demonstrated that population size plays a crucial role in our ability to detect information-limiting noise correlations in population activity. Their study further underscores the multiple theoretical and methodological considerations needed to systematize this body of research: neuronal population size, tuning similarity, and stimulus competition. The devil is in the details. Each of these variables can dramatically influence the relationship of noise correlation to information coding in population activity. Future electrophysiology work with more extensive coverage of the brain in multiple different patches of cortex and a broader array of stimuli and tasks will likely lead to further advances in our understanding of the functional basis of noise correlations in cognition and potentially will pave the way for translating these discoveries to humans.

References

- Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101.
- Averbeck BB, Lee D (2003) Neural noise and movement-related codes in the macaque supplementary motor area. *J Neurosci* 23:7630–7641.
- Averbeck BB, Lee D (2006) Effects of noise correlations on information encoding and decoding. *J Neurophysiol* 95:3633–3644.
- Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366.
- Bartolo R, Saunders RC, Mitz AR, Averbeck BB (2020) Information limiting correlations in large neural populations. *J Neurosci* 40:1668–1678.
- Carandini M, Heeger DJ (2012) Normalization as a canonical neural computation. *Nature Reviews Neuroscience* 13:51–62.
- Cohen MR, Maunsell JH (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12:1594–1600.
- Ecker A, Berens P, Tolias A, Bethge M (2011) The effect of noise correlations in populations of diversely tuned neurons. *J Neurosci* 31:14272–14283.
- Leavitt ML, Mendoza-Halliday D, Martinez-Trujillo JC (2017) Sustained activity encoding working memories: not fully distributed. *Trends Neurosci* 40:328–346.
- Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63:879–888.
- Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A (2014) Information-limiting correlations. *Nat Neurosci* 17:1410–1417.
- Nogueira R, Peltier NE, Anzai A, DeAngelis GC, Martínez-Trujillo J, Moreno-Bote R (2020) The effects of population tuning and trial-by-trial variability on information encoding and behavior. *J Neurosci* 40:1066–1083.
- Rabinowitz NC, Goris RL, Cohen M, Simoncelli EP (2015) Attention stabilizes the shared gain of V4 populations. *Elife* 4:e08998.
- Reynolds JH, Heeger DJ (2009) The normalization model of attention. *Neuron* 61:168–185.
- Ruff DA, Cohen MR (2014) Attention can either increase or decrease spike count correlations in visual cortex. *Nat Neurosci* 17:1591–1597.
- Schmitz TW, Duncan J (2018) Normalization and the cholinergic microcircuit: a unified basis for attention. *Trends Cogn Sci* 22:422–437.
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.
- Smith MA, Sommer MA (2013) Spatial and temporal scales of neuronal correlation in visual area V4. *J Neurosci* 33:5422–5432.
- Tremblay S, Pieper F, Sachs A, Martínez-Trujillo J (2015) Attentional filtering of visual information by neuronal ensembles in the primate lateral prefrontal cortex. *Neuron* 85:202–215.
- van Bergen RS, Ma WJ, Pratte MS, Jehee JF (2015) Sensory uncertainty decoded from visual cortex predicts behavior. *Nat Neurosci* 18:1728–1730.
- Verhoef BE, Maunsell JH (2017) Attention-related changes in correlated neuronal activity arise from normalization mechanisms. *Nat Neurosci* 20:969–977.
- Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143.