Is there signal in the noise?

Alexander S Ecker & Andreas S Tolias

A study now shows that variability in neuronal responses in the visual system mainly arises from slow fluctuations in excitability, presumably caused by factors of nonsensory origin, such as arousal, attention or anesthesia.

Responses of cortical neurons appear to be notoriously noisy. Even with repeated presentations of the same visual stimulus, we rarely observe the same spike train twice. This high degree of variability, which is often correlated amongst pairs of neurons, has fascinated neuroscientists for decades. Is it noise, arising from stochastic features of neural architecture, or does it reflect meaningful, yet complicated, signals that we simply do not understand? Thus far, it is still not entirely clear what causes the observed neuronal variability. Moreover, we do not even have simple, parsimonious models to describe it appropriately. A study by Goris, Movshon and Simoncelli in this issue of *Nature Neuroscience* fills this latter gap and proposes a simple conceptual model that provides a new view on cortical response variability.

The authors propose modeling a neuron's firing rate as the product of a sensory drive and a modulatory gain (**Fig. 1a**). Typical models of the early visual system consist of only the first ingredient, the sensory drive (usually called the receptive field or tuning curve), and assume that any variability around this average response is random noise arising in presynaptic neurons or local circuits. Goris *et al.*[1], in contrast, reasoned that signals internal to the brain that are not purely sensory in origin might equally affect neuronal responses and should therefore be explicitly incorporated into a response model. Examples of such signals include arousal, attention or adaptation. As they tend to modulate neural responses multiplicatively[2], the authors modeled the sensory drive as being multiplied by a modulatory gain that fluctuates across trials, subsuming all such internal signals.

But how exactly does one go about fitting such a model? We know very little about all the possible internal signals that might contribute, which ones are most relevant in different contexts, how strongly each of them fluctuates and how each of them affects the response of any given neuron. Thus, one might think we know close to nothing that is relevant and that there is no way of accomplishing this task. Fortunately, none of these needs to be known if we make just one additional assumption: neurons emit spikes according to a Poisson process. This means that, at any given point in time, a neuron either spikes or it doesn't, with a certain probability that is independent of spikes that occurred in the past. In this case, the variance of a neuron's spike count is equal to its mean. Any variability that exceeds the variance predicted by the Poisson process must be a result of fluctuations in the firing rate generated by signals internal to the brain. Because the authors assumed such fluctuations to be driven by a multiplicative gain, they obtained a relatively simple Poisson mixture model, which they termed the modulated Poisson model. Through a clever choice of distribution for the gain (a gamma distribution), they obtained resulting spike counts that follow a well-known distribution, the negative binomial, whose parameters can be readily estimated.

The authors showed that this model nicely captures the variability of single units recorded in several areas along the visual hierarchy of monkeys and substantially outperforms the standard Poisson model for most neurons. Notably, under anesthesia, the classical Poisson model accounts for only 20–30% of the cortical response variability. Gain fluctuations, in contrast, account for the major share of the variance, and their share increases substantially along the visual hierarchy, from ~50% in the lateral geniculate nucleus of the thalamus to ~70% in primary visual cortex and ~80% in cortical area MT. Part of these changes in excitability are quite slow, changing on the order of minutes. Interestingly, fluctuations in excitability are less pronounced and somewhat faster in awake animals performing a visual discrimination task than they are under anesthesia, suggesting that the underlying mechanisms may differ between the two brain states.

The approach of partitioning the variance into firing rate variability and point process

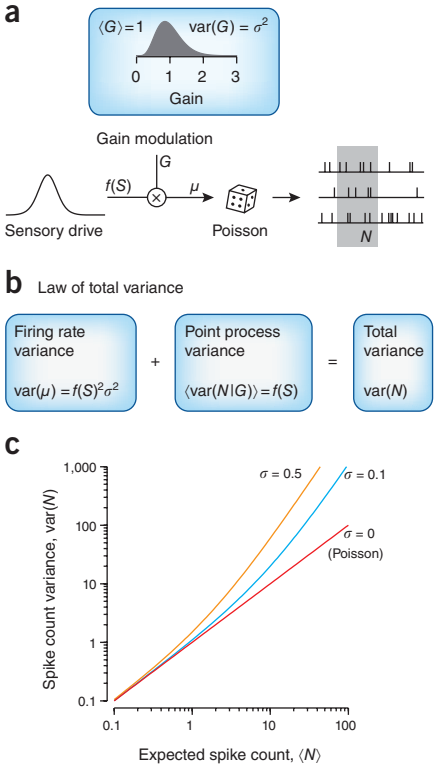

Figure 1 The modulated Poisson model and the law of total variance. (**a**) The neuron's firing rate, μ, is the product of a sensory drive (the tuning curve, $f(S)$) and a modulatory gain (G). Spikes are generated according to a Poisson process with rate μ. The gain is assumed to be constant within one trial and distributed across trials according to a gamma distribution with mean of 1 and variance of σ^2 (top). (**b**) The law of total variance states that the total variance can be decomposed into the variance of the conditional expectation plus the expected conditional variance: $\mathrm{var}(N) = \mathrm{var}(\langle N \mid G \rangle) + \langle \mathrm{var}(N \mid G) \rangle$. This decomposition forms the basis for separating the spike count variance (right) into firing rate variance (left) and point process variance (middle). (**c**) A supralinear mean-variance relationship characterizes the modulated Poisson model. Because of the multiplicative nature of the gain, the firing rate variance grows quadratically with the average firing rate.

Alexander S. Ecker and Andreas S. Tolias are in the Department of Neuroscience, Baylor College of Medicine, Houston, Texas, USA, and the Bernstein Centre for Computational Neuroscience, Tübingen, Germany. Alexander S. Ecker is also at the Werner Reichardt Center for Integrative Neuroscience and Institute of Theoretical Physics, University of Tübingen and the Max Planck Institute for Biological Cybernetics, Tübingen, Germany. Andreas S. Tolias is also in the Department of Computational and Applied Mathematics, Rice University, Houston, Texas, USA.
e-mail: alexander.ecker@uni-tuebingen.de or astolias@bcm.edu

(residual) variance through application of the law of total variance (**Fig. 1b**) has recently been applied by several other groups[3,4]. In fact, it is very common in statistical applications such as analysis of variance or, more generally, generalized linear models. What makes the model by Goris *et al.*[1] stand out is the multiplicative nature of the modulation. In addition to being supported by experimental data on, for instance, attentional modulation[2], this model makes a very specific prediction: if the fluctuations in the gain are independent of the neuron's firing rate, then the spike count variance induced by the gain modulation should scale with the square of the average stimulus-evoked response[5,6], leading to a relatively simple formula (**Fig. 1c**)

$$\mathrm{var}(N) = \langle N\rangle + \sigma^2\langle N\rangle^2 \qquad (1)$$

Here $\langle N\rangle = f(S)\Delta t$ is the expected spike count of a neuron, which depends on both the firing rate $f(S)$ and the time window Δt used to count spikes. Such an expanding mean-variance relation is indeed common for cortical neurons[7], but so far there has been no principled account for this phenomenon. For instance, under an additive model, the variances would simply add and the relative effect of the gain modulation would be weaker when stimulus drive is strong.

The modulated Poisson model is not limited to analyzing the variability of single neurons. It can also be very useful for analyzing correlations in the variability of pairs of neurons (sometimes called noise correlations). Correlated variability is of great interest to many researchers because the correlation structure should depend on how neurons are connected to each other or whether (and to what extent) they receive common input. Moreover, the correlation structure places important constraints on the fidelity of a population code[8]. In our view, one of the most fundamental contributions of the study by Goris *et al.*[1] is that it provides a powerful framework for understanding the origin of such noise correlations. Analogously to the approach taken for the variances, they describe the correlations as being generated by two components: point process correlations, which arise as a result of shared noise in the sensory afferent pathways[9], and gain correlations, which are a result of gain modulation by unobserved internal signals[10] (a partitioning that we proposed independently using a simplified version of the model[4,5]). Similarly to equation (1) above, the gain-induced covariance depends quadratically on the expected spike count, whereas the dependence is linear for the point process–induced covariance. This relationship allows separation of the two components. Notably, the modulated Poisson framework predicts that one will observe a wide range of spike count correlations, even when point process correlations and gain correlations are stable properties of neurons that do not change across stimulus conditions. For example, in the presence of strong gain fluctuations, there will be higher spike count correlations for stimuli that drive the cells strongly compared with stimuli that drive the cells suboptimally.

This insight may help us to reconcile what appeared to be contradictory experimental findings on correlated variability in the monkey visual system. The original view was that correlated variability arises as a result of shared noise in the sensory afferent pathway[9]. Although this hypothesis seemed to be confirmed by a number of experimental measurements[11,12], we observed substantially lower levels of correlations in V1 of awake, fixating monkeys[5], and other studies pointed to a top-down source of correlated variability during decision making[10,13], questioning the sensory noise hypothesis. We recently used an approach very similar in spirit to that of Goris *et al.*[1] to analyze the correlation structure under anesthesia compared with that under passive fixation[4]. We partitioned the variability into two components: a network state, which is shared among all cells in the local circuit, and a residual variance. This analysis revealed that, under anesthesia, noise correlations were dominated by a common fluctuating factor, slowly modulating all cells together. These fluctuations resembled a common gain and accounted for most of the elevated correlations observed in comparison with that under passive fixation, where spike count correlations were low and variances were close to the mean (even for cells with high firing rates). Adopting the authors' hypothesis that the Poisson process represents a 'floor' state of cortical variability[1], our data suggest that gain fluctuations are nearly absent under our experimental conditions with awake, fixating monkeys and chronic recordings. This interpretation raises the question of why gain fluctuations should be weak during passive fixation, where the subject's internal state is not well controlled, but stronger during difficult behavioral tasks, which provide a means of controlling and assessing cognitive state[14]. We suspect that when an animal performs a difficult task at perceptual threshold, task-related signals such as attention or prior expectations fluctuate between the two alternatives offered by the task (for example, the two stimuli in an attention task[10] or the possible choices in a forced-choice task[13,15]), introducing a low-dimensional axis of variability (namely, the attention axis[10]) and comodulating cells with similar stimulus selectivity with respect to the feature relevant for the task (for example, receptive field location, direction of motion). In contrast, in a passive fixation experiment, there is no task, so there is no reason for any particular top-down signal to be particularly strong.

Coming back to the question posed in the beginning, it seems that under many conditions there is actually a signal in the noise. The new framework developed by Goris *et al.*[1] moves us much closer to characterizing this signal, in terms of both its spatial and its temporal correlation structure. However, what the model cannot do (yet) is tell us the exact value of the gain on a single-trial basis. Fortunately, this limitation may soon be overcome. As many internal signals affect a large number of neurons at the same time, we can infer them from jointly recorded population activity[4,10]. Thus, extending the modulated Poisson framework to a full state-space model for the joint activity of neuronal populations is an exciting and promising avenue for future research. Used in combination with modern multielectrode recording or imaging techniques, such a model would allow us to read out signals internal to the brain and may greatly advance our understanding of the computations underlying decision making and cognition.

COMPETING FINANCIAL INTERESTS
The authors declare no competing financial interests.

1. Goris, R.L.T., Movshon, J.A. & Simoncelli, E.P. *Nat. Neurosci.* **17**, 858–865 (2014).
2. Maunsell, J.H.R. & Treue, S. *Trends Neurosci.* **29**, 317–322 (2006).
3. Churchland, A.K. *et al. Neuron* **69**, 818–831 (2011).
4. Ecker, A.S. *et al. Neuron* **82**, 235–248 (2014).
5. Ecker, A.S. *et al. Science* **327**, 584–587 (2010).
6. Pillow, J.W. & Scott, J.G. *Adv. Neural Inf. Process. Syst.* **25**, 1907–1915 (2012).
7. Britten, K.H., Shadlen, M.N., Newsome, W.T. & Movshon, J.A. *Vis. Neurosci.* **10**, 1157–1169 (1993).
8. Abbott, L.F. & Dayan, P. *Neural Comput.* **11**, 91–101 (1999).
9. Shadlen, M.N. & Newsome, W.T. *J. Neurosci.* **18**, 3870–3896 (1998).
10. Cohen, M.R. & Maunsell, J.H.R. *J. Neurosci.* **30**, 15241–15253 (2010).
11. Zohary, E., Shadlen, M.N. & Newsome, W.T. *Nature* **370**, 140–143 (1994).
12. Smith, M.A. & Kohn, A. *J. Neurosci.* **28**, 12591–12603 (2008).
13. Cohen, M.R. & Newsome, W.T. *Neuron* **60**, 162–173 (2008).
14. Cohen, M.R. & Kohn, A. *Nat. Neurosci.* **14**, 811–819 (2011).
15. Britten, K.H., Shadlen, M.N., Newsome, W.T. & Movshon, J.A. *J. Neurosci.* **12**, 4745–4765 (1992).