



HHS Public Access

Author manuscript

Nat Neurosci. Author manuscript; available in PMC 2014 December 01.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Published in final edited form as:

Nat Neurosci. 2014 June ; 17(6): 858–865. doi:10.1038/nn.3711.

Partitioning neuronal variability

Robbe L.T. Goris¹, J. Anthony Movshon¹, and Eero P. Simoncelli^{1,2}

¹Center for Neural Science, New York University, 4 Washington Place, Room 809, New York, NY 10003

²Howard Hughes Medical Institute, New York University, 4 Washington Place, Room 809, New York, NY 10003

Abstract

Responses of sensory neurons differ across repeated measurements. This variability is usually treated as stochasticity arising within neurons or neural circuits. However, some portion of the variability arises from fluctuations in excitability due to factors that are not purely sensory, such as arousal, attention, and adaptation. To isolate these fluctuations, we developed a model in which spikes are generated by a Poisson process whose rate is the product of a drive that is sensory in origin, and a gain summarizing stimulus-independent modulatory influences on excitability. This model provides an accurate account of response distributions of visual neurons in macaque LGN, V1, V2, and MT, revealing that variability originates in large part from excitability fluctuations which are correlated over time and between neurons, and which increase in strength along the visual pathway. The model provides a parsimonious explanation for observed systematic dependencies of response variability and covariability on firing rate.

Introduction

Neurons transmit information with sequences of action potentials. These responses are variable – repeated measurements under identical experimental conditions give different spike trains – but the origins of this variability are unknown. If spike generation were variable, it might account for response variability, but *in vitro* measurements indicate that it is highly reliable¹. Variability in synaptic transmission is another possible source², but its magnitude is also believed to be insufficient to account for the observed variability in spiking responses^{3–4}. A more likely explanation is that the variability arises from the accumulation and amplification of small amounts of noise as signals flow through neural circuits⁵. And recent theories propose that the substantial variability in neural responses may arise from the dynamics of recurrent but largely deterministic networks^{6–7}.

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:http://www.nature.com/authors/editorial_policies/license.html#terms

Address for correspondence: Eero P. Simoncelli, Center for Neural Science, New York University, 4 Washington Place, Room 809, New York, NY 10003, Phone: 212-998-3938, FAX: 212-995-4619, eero.simoncelli@nyu.edu.

Author contributions: R.L.T.G., J.A.M., and E.P.S. designed research, R.L.T.G. analyzed data, and R.L.T.G., J.A.M., and E.P.S. wrote the paper.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Regardless of its source, characterizing variability with simple stochastic models has proven useful in understanding the nature of neural coding. The simplest stochastic model is a Poisson process, in which spikes occur independent of one another. A hallmark of the Poisson model is that the variance of the spike count in any time interval is equal to the mean. In visual cortex, the spike count variance typically equals or exceeds the mean, but rarely falls below it^{8,9}. This suggests that Poisson-like behavior is a “floor” state of cortical variability, and raises the question of the origin of the excess variance. Arousal, attention, adaptation and other contextual factors are known to modulate sensory responses^{10–12}. In typical electrophysiological experiments, some of these may be well-controlled, but many are not. The idea that fluctuations in excitability can inflate estimates of neuronal variance has a long history^{8,9}, and we wondered whether a more directed analysis of single-neuron responses might reveal the effect of these factors.

We formalize this hypothesis in a doubly stochastic response model in which spikes arise from a Poisson process whose rate is the product of “drive” and “gain” (the “modulated Poisson model”, Fig. 1). The drive is a reproducible firing rate response to a sensory stimulus; the gain represents modulatory influences on excitability, and can vary across repeated measurements. Under this model, trial-to-trial variability in spike counts can be partitioned into a sum of Poisson point process variance and variance arising from fluctuations in gain. Likewise, spike count covariation can be partitioned into point process covariance, and covariance arising from correlated gain fluctuations.

We found that this model provides an excellent account of single-neuron response distributions in macaque visual thalamus and cortex. Inferred gain fluctuations are correlated over long timescales (minutes), are larger in anesthetized than in awake animals, and (in anesthetized animals) increase in strength along the visual hierarchy. They are also shared across wide areas of primary visual cortex, while point process variance is more localized. Together, these results suggest that much of the response variability of sensory neurons arises from fluctuations in excitability which are correlated over time and between neurons and which increase in strength along the visual pathway.

Portions of these results have been presented in conferences (R.L.T. Goris, J.A. Movshon and E.P. Simoncelli, *Soc. Neurosci. Abstr.* 311.01, 2013; *Cosyne Abstr.*, I-37, 2012), and two other groups have recently presented models similar to the one shown in Fig. 1 (ref. 13, and I.-C. Lin, M. Okun, M. Carandini and K. Harris, *Cosyne Abstr.*, III-11, 2014), though their application and conclusions are somewhat different from our own.

Results

Modulated Poisson framework

Consider the most commonly-used rate model, in which the spike count, N , follows a Poisson distribution:

$$p(N|\mu, \Delta t) = \frac{(\mu \Delta t)^N}{N!} \exp(-\mu \Delta t) \quad (1)$$

where μ is the mean spike rate, and t is the duration of the counting window. Assume that the rate arises from the product of two positive-valued signals:

$$\mu = f(S)G \quad (2)$$

where $f(S)$ is some function of the stimulus, and G is a stimulus-independent gain. In this case, the variance and mean of the spike count in a time interval t are both equal to $f(S)Gt$.

If the gain signal is unobserved, and varies over repeated trials, the observed spike counts over those trials will have a variance larger than predicted by the simple Poisson model. Assume the gain has a mean of one, but fluctuates randomly on a timescale that is slow relative to the interval over which spikes are counted. In this case, the net distribution of spike counts is a mixture of Poisson distributions (Online Methods) whose variance can be decomposed into two components:

$$\text{var}[N|S, \Delta t] = (f(S)\Delta t) + \sigma_G^2(f(S)\Delta t)^2 \quad (3)$$

The first term is simply the variance of a Poisson distribution with mean count $f(S)t$. The second term is the variance of the expected spike count conditioned on the stimulus drive¹⁴, and is proportional to the square of the first term, with a proportionality factor equal to the variance of the gain signal. Thus, in the presence of gain fluctuations, the spike count variance exceeds the mean by an amount that is proportional to the square of the mean. Note that the spike count variance does not depend on the firing rate or duration of the counting window *per se*, but on their product, $f(S)t$.

The variance-to-mean relationship expressed in Eq. (3) is inherent to systems with modulatory effects, and depends on gain only through its variance: all gain distributions with the same variance will produce the same variance-to-mean relationship. To fit the model to data, the form of the gain distribution must be specified. If we assume that the gain follows a Gamma distribution, the overall spike count distribution will follow a negative binomial (Online Methods):

$$p(N|S, \Delta t) = \frac{\Gamma(N+1/\sigma_G^2)}{\Gamma(N+1)\Gamma(1/\sigma_G^2)} \cdot \frac{(\sigma_G^2 f(S)\Delta t)^N}{(\sigma_G^2 f(S)\Delta t + 1)^{(N+1/\sigma_G^2)}} \quad (4)$$

Where $\Gamma(\cdot)$ represents the standard gamma function. This distribution is parameterized by the variance of the gain σ_G^2 and the mean spike count, $f(S)t$, and is readily fit to neural data.

Gain fluctuations increase along the visual hierarchy

We measured responses of neurons in areas LGN, V1, V2 and MT to drifting sinusoidal gratings of the preferred size and speed, varying either in spatial frequency (12 spatial frequencies, ranging from 0 to 10 c/deg) or in drift direction (16 equally spaced directions). Each grating was presented for 1,000 ms and repeated at least five times. Responses were

computed by counting spikes in a 1,000-ms window following response onset. We fit both the Poisson and modulated Poisson models to the responses of each individual neuron.

Consider the actual and model-predicted response distributions for an example V1 cell (Fig. 2a). The mean spike count depends on the direction of motion of the stimulus (center Fig. 2a), and the variance grows with the mean. The variance-to-mean relation is well described by the modulated Poisson model (Fig. 2b). As predicted (eq. 3), this relation does not depend on firing rate or measurement interval (inset Fig. 2b). Comparing histograms of the measured spike counts with the predictions of the models reveals that the modulated Poisson model captures the data much better than the standard Poisson model (Fig. 2a). When the mean spike count is low, both models produce similar response distributions with mode at zero and variance approximately equal to the mean (Fig. 2b). But as the mean spike count grows, model predictions diverge, and the Poisson model fails to account for the shape of the count distributions (Fig. 2a). In contrast, the data are well-described by the modulated Poisson model, suggesting that a substantial part of the response variance arises from fluctuations of modulatory inputs, whose distribution is well-described by a Gamma distribution.

This visual impression is confirmed by statistical analysis of the model fits. To quantify relative goodness of fit of the two models, we performed a cross-validation analysis (CVA), comparing the log-probability of observing the hold-out data under each fitted model (Online Methods). For the example neuron, the value for the modulated Poisson model is much higher than that of the Poisson model. To assess the model fits in absolute terms, we developed a parametric bootstrap test that compares the log-probability of the data with its expectation under the model in question (Online Methods). For the example neuron, the fit of the Poisson model is very poor ($P < 0.001$, absolute goodness of fit test; Fig. 2c), but the modulated Poisson model cannot be rejected ($P = 0.91$; Fig. 2c).

In sum, the variable discharge of this V1 cell is well described as originating from three different sources: the stimulus attributes (i.e., direction of motion), a Poisson point process, and Gamma-distributed fluctuations in excitability. To estimate the relative contribution of each source, we used the modulated Poisson model to partition the spike count variance (Online Methods). Surprisingly, Poisson noise accounts for only a small fraction of the total variance (5.5%). The gain fluctuations account for nearly half of the variance (47.5%), a share comparable to the fraction due to variations in the stimulus drive (47%). The latter is dependent on the set of stimuli and the tuning properties of the neuron. To focus our analysis on the variability across repeated measurements, we consider the portion of within-condition variance that is explained by the excitability fluctuations. For the example neuron in Fig. 2, this fraction is 89.6%.

In our model, stronger gain fluctuations lead to a more rapidly accelerating variance-to-mean relationship, which deviates more and more from the Poisson expectation as spike count grows (Fig. 2d). Note that this accelerating relationship implies that the ratio of variance to mean (Fano factor) can vary within a single spike train, and over short time scales. As such, a single Fano factor provides an incomplete and potentially biased, measure of neuronal variability.

In the modulated Poisson model, the rate arises from the product of two positive-valued signals, one a function of the stimulus and the other arising elsewhere. An alternative model adds these signals rather than multiplying them. Under this additive model, if the two signals are statistically independent the expected spike count variance is given by the sum of the point process variance and a constant equal to the variance of the drive fluctuations, yielding a variance-to-mean relationship (Fig. 2d) which is unlike that in our data.

Fluctuations in excitability modulate neuronal activity throughout visual thalamus and cortex. Neurons in LGN, V1, V2, and MT all exhibit super-Poisson variability (Fig. 3a). Within each area, the variance-to-mean relationship is consistent with the predictions of the modulated Poisson model (compare Fig. 3a to Fig. 2d), but note that the fitted gain variance increases in strength along the visual processing stream. Specifically, the modulated Poisson model systematically outperforms the standard Poisson model for all areas (Fig. 3b), and the proportion of neurons for which this model is selected increases from LGN to V1 ($P < 0.001$, t test), from V1 to V2 ($P = 0.004$) and from V1 to MT ($P = 0.004$). The strength of fluctuations in excitability, as measured by the coefficient of variation of the gain, grows from LGN to V1 ($P < 0.001$, Wilcoxon rank sum test), from V1 to V2 ($P < 0.001$), and from V1 to MT ($P < 0.001$) (Fig. 3c). As information propagates through the visual hierarchy, fluctuations in excitability not only increase in strength, but also account for a larger share of variance. When stimulated with drifting gratings, within-condition variance primarily reflects Poisson-like noise in LGN ($P = 0.02$, Wilcoxon signed rank test), but becomes progressively more dominated by excitability fluctuations in cortex (LGN to V1, $P < 0.001$, Wilcoxon rank sum test; V1 to V2, $P < 0.001$, V1 to MT, $P < 0.001$; Fig. 3d).

Response correlations in the modulated Poisson framework

Trial-to-trial response fluctuations are often correlated among simultaneously recorded neurons¹⁵. Pairwise response correlations can arise when neurons receive shared sensory input⁵, but also when they are subject to correlated modulatory influences that are not sensory in origin^{16,17}. Our model provides a vehicle for separating the effects of these two contributions. Specifically, for doubly stochastic processes, the spike count covariance can be decomposed in a manner that is analogous to the variance decomposition introduced in Eq. (3) (Online Methods):

$$\text{cov}[N_i, N_j | S, \Delta t] = r_{P_{ij}} \sqrt{f_i(S) f_j(S) \Delta t^2} + r_{G_{ij}} \sigma_{G_i} \sigma_{G_j} f_i(S) f_j(S) \Delta t^2 \quad (5)$$

where $r_{P_{ij}}$ indicates the point process correlation (assumed to be independent of stimulus S), and $r_{G_{ij}}$ the gain correlation. The first term is the covariance expected for the spike counts of neuron i and j with a constant gain of one. The second component is the covariance of the conditional expectations that arises from correlated gain fluctuations, and can generate spike count correlations even when the two point processes are independent. Note that Eq. (5) reduces to the expression for spike count variance (Eq. (3)) when $i=j$. The spike count correlation is obtained by dividing this equation by the square root of the product of the spike count variances of neurons i and j , as expressed in Eq. (3), yielding a complicated

dependence of spike count correlation on stimulus drive, and on the variance of the gain signals in the two neurons.

Studies of response correlation typically combine normalized responses across conditions to estimate a single spike count correlation for a pair of neurons^{18,19}. But equation (5) implies that even if the point process and gain correlations are both stable properties of a given cell pair, measured spike count correlations can vary dramatically with stimulus drive.

Specifically, when the geometric mean of the stimulus-driven response of the two neurons is low, the spike count correlation will approximate the point process correlation r_{Pij} . In contrast, when this geometric mean is high, the spike count correlation will be dominated by the gain correlation r_{Gij} . If r_{Pij} and r_{Gij} differ, the spike count correlation will depend on the stimulus conditions, and any single estimate will provide an incomplete and potentially biased measure of neuronal covariability.

We analyzed the responses of four populations of simultaneously recorded neurons in the superficial layers of macaque primary visual cortex²⁰. The recordings were made over 2.5 h, during which gratings drifting in 72 equally spaced directions were presented for 1,280 ms each, interleaved with a 1,280 ms blank screen and repeated 50 times in random order. We analyzed 379 well-isolated units (62, 94, 87, and 136 from the four data sets). We fit the modulated Poisson model to the responses of each neuron, and then fit the model to the spike count correlations of pairwise combinations of neurons in each data set (Online Methods).

Consider the actual and predicted response correlations for three example pairs of V1 cells (Fig. 4). For each neuron, the mean spike count depended on the direction of motion of the stimulus (Fig. 4a–c), and for each pair of neurons, the spike count correlation estimated from the combined normalized responses (r_{SC}) was small and positive. But estimates of spike count correlation computed separately for different response levels reveal a variety of different behaviors, including a decrease with mean response strength (Fig. 4d), independence of response strength (Fig. 4e), or an increase with response strength (Fig. 4f). The modulated Poisson model can mimic each of these behaviors using particular choices of the underlying point process and excitability correlations, r_{Pij} and r_{Gij} . We estimated both statistics for each cell pair, and Fig. 4g–i illustrate the full spike count correlation surface predicted by the fitted models.

The structure of correlations in primary visual cortex

V1 spike count correlations have been shown to depend on cortical distance and tuning similarity¹⁹, but as revealed by the examples in Fig. 4, this measure mixes two underlying sources of correlation. We wondered whether these two sources might be differently structured across the neuronal population. In particular, common point process variance might be local in space and time, while joint excitability fluctuations could reflect network changes that affect larger populations of neurons on a slower time scale.

We examined how point process correlation and excitability correlation depend on electrode separation and on the similarity of tuning (Fig. 5, Online Methods). Point process correlations were on average smaller than excitability correlations ($P < 0.001$, $F_{1,34} = 24.16$,

ANCOVA) and fell more rapidly with electrode distance ($P < 0.001, F_{1,34} = 18.05$). Note that neurons recorded on the same electrode (the leftmost datum in Fig. 5a–b) might have inflated the significance of the latter difference since they are vulnerable to spike sorting errors that can alter the measured correlation^{15,17,21}. Nevertheless, we find that excluding this datum does not alter the conclusion ($P < 0.001, F_{1,32} = 41.17$). Analysis of tuning similarity effects (Fig. 5c) revealed that the average point process and excitability correlation had a largely similar (linear) dependency on tuning correlation ($P = 0.6, F_{1,14} = 0.28$), differing only in their means (Fig. 5d).

In summary, we find that the relationship between response correlations and mean responses in V1 depends on both inter-neural distance and tuning similarity: Spike count correlations are most likely to decrease with response strength for nearby neurons with similar tuning (Fig. 4d). As inter-neural distance or tuning dissimilarity grows, response correlations will initially tend to be independent of the mean responses (Fig. 4e), and eventually exhibit an increasing relation (Fig. 4f).

The temporal structure of gain fluctuations

We assume that the Poisson point process variance is independent across trials, so all temporal structure in neural responses that is not explained by the stimulus must arise from structure in the gain signal. Consider how normalized spiking activity evolves over 2.5 h for three simultaneously recorded neurons (Fig. 6a). Response strength rises and falls over minutes, suggesting that excitability varies slowly. This is consistent with the slow decay of the autocorrelation of the gain (Fig. 6b), computed from temporal correlations in the trial-by-trial responses (Online Methods). At a time lag of 5 sec, 307/379 neurons had positive autocorrelation (t -test, $P < 0.05$). For longer time lags, this fraction gradually decreased (at a time lag of 2 minutes, $r > 0$ for 185/379 neurons; at a time lag of 10 minutes, this fraction drops to 81/379). Although the average profiles for the four data sets differ in detail, they share a slow falloff, indicating that changes in excitability persist for many minutes (Fig. 6c).

Trial-to-trial changes in excitability are often correlated across neurons (Fig. 5). It might therefore be expected that the slow component of excitability fluctuations is also shared across neurons. However, across neurons, excitability correlations decrease dramatically at all time lags exceeding 0 sec (Fig. 6d). The slow component of excitability fluctuations thus appears to be local.

Effects of anesthesia

We have shown that neural response variability in visual cortex of anesthetized macaques originates largely from excitability fluctuations that are correlated over time. To what degree does this arise from fluctuations in the state of anesthesia? To address this, we analyzed data recorded in area MT of awake monkeys performing a psychophysical task involving random dot kinematogram stimuli^{18,22}. We fit both the Poisson and modulated Poisson model to the responses of 307 individual neurons. As in anesthetized animals, MT neurons in awake monkey typically exhibit super-Poisson variability (Fig. 7a), and the modulated Poisson framework successfully accounts for this variability. The inclusion of gain fluctuations

improves goodness of fit for 224/307 neurons (CVA analysis). At a significance level of 5%, the Poisson model is accepted for only 28/307 of neurons, while the modulated Poisson model is accepted for 294/307 of cells (absolute goodness of fit test). Fluctuations in excitability were the main source of within-condition variance ($P < 0.001$, Wilcoxon signed rank test; Fig. 7b). Finally, the average autocorrelation function reveals that gain fluctuations in awake animals exhibit a slow temporal structure similar in timescale to that seen under anesthesia (Fig. 7c).

Qualitatively, there is thus an excellent agreement between these data sets. However, comparing the coefficient of variation of the excitability signal reveals that its fluctuations are significantly stronger in the anesthetized cortex (compare Fig. 3c with Fig. 7b). It is also notable that the time course of this correlation is substantially shorter in the awake data set, suggesting that the fluctuations in gain seen under the two conditions may have different origins.

Discussion

The distinction between inputs that drive the responses of a neuron and those that modulate its response is well known²³. But while both factors influence the mean response, response variance is usually treated as if it arises solely from noise in driving inputs. Our analysis suggests instead that variability of modulatory signals may underlie much of the response variability in visual cortex. We have shown that a Poisson spiking model whose rate is multiplied by a fluctuating gain signal can explain why the variance of spike count generally grows faster than the mean (Figs. 2 and 3), and that an additive noise model cannot account for this basic behavior (Fig. 2d). The modulated Poisson model also explains why the covariance of spike counts in pairs of neurons can exhibit a diversity of behaviors depending on their means (Fig. 4). Fitting the model to individual neurons from a variety of visual areas reveals that gain fluctuations account for a substantial fraction of spike count variability, and that in anesthetized animals this fraction increases as one ascends the visual hierarchy (Fig. 3).

Poisson spiking models have been common in neuroscience since the 1950's^{24–26}, and are implicit in all analyses in which responses are summarized solely with mean spike counts. They provide the simplest statistical description of the data, are readily fit to data, and capture the basic fact that spike count variances grow with the mean. Poisson processes can mimic the spiking behaviors of integrate-and-fire models⁴, as well as the responses of model neurons embedded in cortical networks with balanced configurations of excitatory and inhibitory inputs^{5–7}. But neuronal responses generally exhibit super-Poisson variability^{5,7–9,27–28}, and the model presented here provides a simple but effective means of explaining this behavior. Some authors have reported sub-Poisson variability, especially in brief time bins immediately after a sudden stimulus onset^{29–32}. This is likely a consequence of neuronal refractoriness^{33–34}, which could be naturally incorporated into our modeling framework through modulatory spike-feedback terms in a generalized linear model^{35–36}.

Our work complements recent work on the temporal dynamics of firing rate variability^{14,28}. One study used the mean-matched Fano factor as a proxy for firing rate variability, but did

not explore the dependence of this value on the mean²⁸. Another estimated firing rate variability by subtracting an estimate of the point process variance from the measured spike count variance¹⁴. The estimate of point process variance is taken from the minimum of an ensemble of experimental measurements, and is therefore vulnerable to the well-known problems associated with estimating extreme values of sample distributions. In contrast, we found it useful to make explicit distributional assumptions: specifically, that gain is drawn from a Gamma distribution, and that the spikes arise from a Poisson process. The resulting negative binomial spike count distribution (a particular continuous mixture of Poisson distributions) is easily fit to data and validated (Fig. 2d, see also ref. 37), and its parameterization as a modulated Poisson distribution provides a natural interpretation in terms of excitability fluctuations. The spike count distribution bears some resemblance to the discrete mixture-of-Poissons model that has been successfully used to account for super-Poisson variability³⁸, but has the added advantage that it has far fewer parameters, and provides an interpretation in terms of modulatory input. It is important to note, on the other hand, that our analyses of joint behaviors are based solely on correlations. A full model for joint neural responses would allow better fitting and validation, but requires an explicit description of both the joint point process distribution, as well as the joint distribution of gain signals.

Our analyses indicate that fluctuations in gain account for a substantial portion of spike count correlation. This has been long-suspected, but it has been difficult to segregate the effects of shared sensory inputs from correlated modulatory influences¹⁶. Our model provides an explicit method for achieving this (Eq. 5). In V1, this analysis reveals that point process correlation and excitability correlation have different structure. Point process correlations decrease rapidly with cortical distance, while excitability correlations change less. Both point process and excitability correlations increase with tuning similarity (Fig. 5). Together, these patterns suggests that point process correlations are caused by inputs that are shared within local functional circuits, while excitability correlations likely arise from fluctuations in modulatory signals that affect larger populations of neurons (e.g., sensory adaptation, metabolic resource availability, attentional signals, reward signals).

The analysis in Fig. 3 indicates that the portion of spike count variability attributable to modulatory fluctuations increases as one ascends the visual hierarchy. This presumably depends on the context in which responses are measured, including the choice of stimuli, the presence or absence of anesthesia, and the cognitive state of the animal. For example, attentional mechanisms are widely believed to act by increasing the gain of neurons involved in a task. Relative to conditions of uncontrolled attention, we might expect that this would lead to a net increase in the mean, but a decrease in the variance, of the modulatory input. Under the modulated Poisson model, this would result in an increase in spike count, accompanied by a decrease in Fano Factor, as has been reported in area V4 (ref. 39). The corresponding predictions for spike count correlation are more complicated, since the relative contributions of the point process and gain correlations depends on the stimulus drive to each cell, as well as the variance of the gain of each cell (Eq. 5). But if the gain correlation is larger than the point process correlation and the two cells are receiving similar stimulus drive, our model predicts that a reduced modulatory variance would decrease spike count correlation, as has also been observed³⁹.

The modulated Poisson model has broad implications for the characterization of neuronal response variability. It is common practice to summarize variability with the Fano factor (the ratio of the spike count variance to its mean). As clearly revealed in Figs. 2–3, this measure is highly dependent on the conditions under which it is estimated. For example, the Fano factor for the neuron in Fig. 2b would be near 1 if estimated from trials in which the response was low (less than 1 spike per time bin), but would be significantly higher if estimated from trials with larger responses. Thus, the Fano factor, by itself, does not provide a reliable summary of neural variability. In contrast, spike count variability of the modulated Poisson model may be decomposed into a sum of two distinct values, one corresponding to the Poisson component (which is always equal to the mean) and the second arising from unobserved modulatory influences, which grows as the square of the mean (Eq. 2). The amplitude of this second term, which corresponds to the variance of the modulated gain, provides a compact summary of variability beyond that expected from the basic Poisson model.

An analogous issue arises with the common practice of summarizing the covariation of pairs of neurons with spike count correlation. This measure can depend critically on the conditions under which it is measured (Fig. 4), and thus provides a potentially biased summary of covariability. Indeed, previous studies have noted that neural correlations can increase with firing rate¹⁵, and that this behavior can be explained by models in which spikes arise when a Gaussian-distributed membrane voltage crosses a threshold^{15,40–41}. Our model is quite different but accounts for this phenomenon, as well as the more diverse range of behaviors seen across different cell pairs (Fig. 4), by summarizing covariation with two distinct values. One value represents the point process correlation (which could arise from noise in common sensory inputs⁴²), and the other the correlation of the modulatory influences. At low firing rates, response correlations primarily reflect the former, whereas at high firing rates they reflect the latter. This insight suggests a resolution for a recent controversy in the literature regarding the nature and magnitude of cortical correlations. Despite decades of experimental evidence that cortical cells exhibit modest but significant spike count correlations on the order of 0.1–0.3 (ref. 15), a recent study in macaque V1 found correlations statistically indistinguishable from zero, and concluded that previously reported values arose from experimental confounds¹⁷. However, average firing rates were unusually low in this study, and measured correlations were significantly positive for the small subset of neuronal pairs with high responses. The interpretation arising from our model and data (Fig. 4) is that the correlation values reported in ref. 17 primarily reflect point process correlations, which are dominant at low firing rates and fall rapidly with cortical distance, whereas most previously reported values reflect correlated modulation, which is only evident at high firing rates, but is generally more substantial and falls more slowly with cortical distance.

Our analysis suggests that fluctuations in gain are correlated over long time scales within but not across neurons (Fig. 6). This implies that the mechanisms underlying slow drifts in response gain differ from the mechanisms that give rise to instantaneous gain correlations. The local nature of these drifts rules out that they result simply from global state changes induced by factors like anesthesia or arousal. That said, comparison of responses in area MT of behaving and anesthetized macaque revealed that cortex is less stable under anesthesia:

Fluctuations in gain are significantly stronger (Fig. 3c and 7b) and slower (Fig. 6c and 7c) in the anesthetized cortex.

What are the implications of our findings for understanding the representation of sensory information in the brain? We believe that the gradual increase in the strength of modulatory fluctuations along the visual pathway (Fig. 3) reflects the gradual transformation and combination of visual signals with information from other sources, such as other sensory inputs, top-down signals representing attention, arousal, metabolic state, reward expectations, emotional state, and so forth. Regardless of the source of modulatory variability, its increasing weight as information ascends cortical pathways raises the question of why it does not overwhelm the sensory information encoded in higher-level visual areas⁵, with deleterious effects on the accuracy of behaviors that arise from decoding those neural responses^{43–44}. It is perhaps worth noting that information encoded in the relative responses of groups of neurons will not be affected by common modulatory changes, such as those resulting from attentional focus¹¹. Resolving this mystery is a critical step in understanding how stable perception of the visual world can arise from apparently volatile neural activity, a step made easier by the principled and proven analysis that we have presented here.

Methods

Surgical preparation

We recorded from 18 anesthetized, paralyzed, adult macaque monkeys of either sex (3 *Macaca nemestrina*, 1 *Macaca mulatta*, and 14 *Macaca cynomolgus*). Our standard procedures for the surgical preparation of animals and single-unit recordings have been reported in detail previously⁴⁵. Briefly, experiments typically lasted 5–6 d, during which we maintained anesthesia with infusion of sufentanil citrate (6–30 µg kg⁻¹ h⁻¹) and paralysis with infusion of vecuronium bromide (Norcuron; 0.1 mg kg⁻¹ h⁻¹) in isotonic dextrose-Normosol solution. We monitored vital signs (heart rate, lung pressure, end-tidal pCO₂, EEG, body temperature, urine flow, and osmolarity) and maintained them within appropriate physiological ranges. Pupils were dilated with topical atropine. The eyes were protected with gas-permeable contact lenses, and refracted with supplementary lenses chosen through direct ophthalmoscopy. At the conclusion of data collection, the animal was killed with an overdose of sodium pentobarbital. All procedures were conducted in compliance with the National Institute of Health Guide for the Care and Use of Laboratory Animals, and with the approval of the New York University Animal Welfare Committee.

Unit recording

Extracellular recordings were made with quartz-platinum-tungsten microelectrodes (Thomas Recording), advanced mechanically into the brain through a craniotomy and small durotomy. Electrode insertion angle and location varied across experiments, depending on the targeted area. We distinguished V1 from V2 on the basis of depth from the cortical surface and changes in the receptive field location of the recorded units. Area MT was identified from the brisk direction-selective responses of isolated neurons. We made recordings from every single unit with a spike waveform that rose sufficiently above noise

to be isolated. Stimuli were presented in random order. Data are reported from every unit for which we completed at least 5 repetitions.

Visual stimulation

We presented visual stimuli on a gamma-corrected CRT monitor (Eizo T966; mean luminance, 33 cd/m²) at a resolution of 1,280 × 960 with a refresh rate of 120 Hz. Stimuli were presented using Expo software (<http://corevision.cns.nyu.edu>) on an Apple Macintosh computer. For each isolated unit, we first determined its ocular dominance and occluded the non-preferred eye. We presented circularly windowed sinusoidal grating stimuli to map each cell's receptive field, determined its preferred size and speed, and then measured selectivity for orientation or spatial frequency.

Analysis of single electrode recordings

Responses were computed by counting spikes in a 1,000-ms window following response onset. We estimated latency for each cell by choosing the latency that maximized the variance of the tuning curve⁴⁶. The modulated Poisson model describes a doubly-stochastic process (also known as a “Cox process”⁴⁷), in which spiking responses are generated by a Poisson process whose rate is the product of drive and gain (Eqs. 1 and 2). To make the model identifiable, we assume that the gain, G , is constant within trials, and is distributed across trials according to a Gamma distribution with a mean of one and variance of σ_G^2 :

$$p(G;r,s) = \frac{G^{r-1} \exp(-G/s)}{s^r \Gamma(r)}$$

with shape parameter $r=1/\sigma_G^2$, and scale parameter $s=\sigma_G^2$. The mean spike count of the modulated Poisson model is the product of the gain, G , the stimulus driven rate, $f(S)$, and the bin duration, t , and is thus also distributed according to a Gamma distribution, with parameters $r=1/\sigma_G^2$ and $s=\sigma_G^2 f(S) \Delta t$. The spike count distribution is a Gamma mixture of Poisson distributions, and marginalizing over the gain variable yields a negative binomial distribution (Eq. 4) for the spike count⁴⁸:

$$p(N;r,s) = \frac{\Gamma(N+r)}{\Gamma(N+1)\Gamma(r)} \left(\frac{1}{1+s}\right)^r \left(\frac{s}{1+s}\right)^N$$

With parameters r and s described above, this distribution has a mean of $rs = f(S) t$ and a variance of $rs+rs^2 = f(S)\Delta t + \sigma_G^2 f(S)^2 \Delta t^2$ as expressed in Eq. (3). The negative binomial can also be derived as a Poly-gamma mixture of Gaussian distributions⁴⁹, a form that allows efficient inference³⁷, but lacks the modulatory interpretation provided here.

We fit both the ordinary and modulated Poisson models to the responses of each individual neuron. The maximum likelihood estimator of the parameter μ that characterizes a Poisson distribution is the sample mean, which we estimate separately for each stimulus condition.

The modulated Poisson model has two parameters (r and s) that are related to the stimulus drive and the variance of the gain. We used a simplex algorithm (the Matlab function ‘fminsearch’) to search for the value of $\{r, s\}$ that minimized the negative log likelihood (NLL) of the full set of observed responses.

To assess the models' goodness-of-fit, we performed a 100-fold cross-validation. We fitted both models to a data-set consisting of all trials except for one randomly chosen trial per stimulus condition and compared the average log-probability of the hold-out data under each model (Fig. 3b). To evaluate the models' absolute goodness-of-fit, we compared the log-probability of the data with that of a large number of simulated data sets drawn from the fitted model. Each simulated data set has the same number of repeats as the observed data. We consider a model fit acceptable if the log-probability of the real data lies within the central 95% of the distribution of log-probability of the simulated data. This test can detect both under-dispersion and over-dispersion.

We used the modulated Poisson model to estimate the fraction of spike count variance that arises from the gain signal. As is standard in ANOVA, one can partition the sum of squares into components arising from variations in the stimulus (S_{stim}), the gain signal (S_{gain}), and the point process (S_{pp}):

$$\begin{aligned} \sum_k (N_k - \bar{N})^2 &= \sum_k (N_k - \bar{N}_k)^2 + \sum_k (\bar{N}_k - \bar{N})^2 \\ &= \sum_k \bar{N}_k + \sum_k \sigma_G^2 \bar{N}_k^2 + \sum_k (\bar{N}_k - \bar{N})^2 \\ &= S_{pp} + S_{gain} + S_{stim} \end{aligned}$$

where N_k is the spike count on the k th trial, \bar{N}_k is the spike count averaged over those trials in which the presented stimulus was the same as that of the k th trial, and \bar{N} is the spike count averaged over all trials. The second line follows from the first given the spike count variance of the modulated Poisson model, as expressed in Eq. (3). The fraction of within condition variance that arises from excitability fluctuations is given by $S_{gain} / (S_{gain} + S_{pp})$.

Analysis of multielectrode array recordings

Full details regarding these data sets may be found in ref. 20. Briefly, an array containing 96 fixed electrodes was used to record from single units in the superficial layers of macaque primary visual cortex. The animal was anesthetized and stimulated with sinusoidal gratings. We included all units that could be tracked throughout the entire experiment and whose mean response exceeded 2 spikes/sec for at least one stimulus condition. We opted to leave out one set of array recordings (data set 3 in ref. 20) because the responses exhibited uncharacteristically strong correlations across time and neurons; we suspect the cortical surface was still recovering from array insertion. In our modeling framework, pairwise response correlations can result from both correlated point processes, as well as correlated gain fluctuations. To separate these two sources, we used the covariance decomposition formula in Eq. (5) and searched for the point process correlation and gain correlation that maximized the likelihood of the observed response relations. Rather than fitting response

correlations directly, we performed the optimization on Z-transformed spike count correlations:

$$z = \frac{1}{2} \ln \left(\frac{1+r}{1-r} \right)$$

where r is the spike count correlation, and \ln the natural logarithm. The advantages of this Fisher transformation include a quicker convergence to normality and variance stabilization. We used a simplex algorithm to find the $[r_{Pij}, r_{Gij}]$ values that maximized the likelihood of the observed response correlations under a normal residual distribution. This approach works well for most neuronal pairs (Fig. 4), but the model parameters are not always well constrained. For instance, when the modulation of the direction tuning curve is weak for both neurons, many parameter combinations will yield a similar likelihood. For the population analysis (Fig. 5), we therefore took the uncertainty associated with the parameter estimates into account. We modeled the population distribution of the mean parameter values with a bivariate normal distribution, and searched for the parameter values that maximized the likelihood across all neuronal pairs. Poorly constrained pairs of neurons have a flatter likelihood function and are therefore down-weighted in this analysis. Although we consider this approach preferable, simple averaging of the parameter estimates yielded similar results.

To study the temporal structure of the gain signal, we estimated its autocorrelation function. We first removed the stimulus-induced temporal structure by subtracting the appropriate stimulus-elicited mean response from the observed responses. We then estimated the remaining covariance between pairs of responses separated by different time lags. To obtain the autocorrelation of the excitability signal at a given time lag, this covariance is normalized by the variance of the excitability signal, $S_{gain}/(N - 1)$.

Acknowledgments

This work was supported by NIH grants EY04440, EY022428, the Howard Hughes Medical Institute, and postdoctoral fellowships from the Fund for Scientific Research of Flanders and the Belgian American Educational Foundation awarded to RLTG. We are grateful to Romesh Kumbhani and Neil Rabinowitz for helpful discussions, and to members of the Movshon Laboratory for sharing their data.

Reference List

1. Mainen ZF, Sejnowski TJ. Reliability of spike timing in neocortical neurons. *Science*. 1995; 268:1503–1506. [PubMed: 7770778]
2. Allen C, Stevens CF. An evaluation of causes for unreliability of synaptic transmission. *Proc Natl Acad Sci USA*. 1994; 91(22):10380–10383. [PubMed: 7937958]
3. Softky W, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of small EPSPs. *J Neurosci*. 1993; 13:334–350. [PubMed: 8423479]
4. Stevens, CF.; Zador, A. When is an integrate-and-fire neuron like a Poisson neuron?. In: Mozer, M.; Touretzky, DS.; Hasselmo, M., editors. *Advances in Neural Information Processing Systems*. Vol. 8. 1996. p. 103-109.

5. Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci*. 1998; 18(10):3870–3896. [PubMed: 9570816]
6. van Vreeswijk CA, Sompolinsky H. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*. 1996; 274:1724–1726. [PubMed: 8939866]
7. Vogels TP, Abbott LF. Signal Propagation in Networks of Integrate-and-Fire Neurons. *J Neurosci*. 2005; 25:10786–10795. [PubMed: 16291952]
8. Tomko GJ, Crapper DR. Neuronal variability: non-stationary responses to identical visual stimuli. *Brain Res*. 1974; 79:405–418. [PubMed: 4422918]
9. Tolhurst DJ, Movshon JA, Thompson ID. The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Exp Brain Res*. 1981; 41(3–4):414–419. [PubMed: 7215502]
10. Kato HK, Chu MW, Isaacson JS, Komiyama T. Dynamic Sensory Representations in the Olfactory Bulb: Modulation by Wakefulness and Experience. *Neuron*. 2012; 76(5):962–975. [PubMed: 23217744]
11. Luck SJ, Chelazzi L, Hillyard SA, Desimone R. Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J Neurophysiol*. 1997; 77:24–42. [PubMed: 9120566]
12. Benucci A, Saleem AB, Carandini M. Adaptation maintains population homeostasis in primary visual cortex. *Nat Neurosci*. 2013; 16(6):724–729. [PubMed: 23603708]
13. Ecker AS, Berens P, Cotton RJ, Subramanyan M, Denfield GH, Cadwell CR, Smirnakis SM, Bethge M, Tolias AS. State dependence of noise correlations in macaque primary visual cortex. *Neuron*. 2014; 82(1):xxx–xxx.
14. Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN. Variance as a signature of neural computations during decision-making. *Neuron*. 2011; 69(4):818–831. [PubMed: 21338889]
15. Cohen MR, Kohn A. Measuring and interpreting neuronal correlations. *Nat Neurosci*. 2011; 14(7): 811–819. [PubMed: 21709677]
16. Brody CD. Correlations without synchrony. *Neural Computation*. 1999; 11:1537–1551. [PubMed: 10490937]
17. Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS. Decorrelated Neuronal Firing in Cortical Microcircuits. *Science*. 2010; 327:584–587. [PubMed: 20110506]
18. Zohary E, Shadlen MN, Newsome WT. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*. 1994; 370:140–143. [PubMed: 8022482]
19. Smith MA, Kohn A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci*. 2008; 28:12591–12603. [PubMed: 19036953]
20. Graf ABA, Kohn A, Jazayeri M, Movshon JA. Decoding the activity of neuronal populations in macaque primary visual cortex. *Nat Neurosci*. 2011; 14:239–245. [PubMed: 21217762]
21. Pillow JW, Shlens J, Chichilnisky EJ, Simoncelli EP. A model-based spike sorting algorithm for removing correlation artifacts in multi-neuron recordings. *PLoS One*. 2013; 8(5):1–14.
22. Britten K, Shadlen MN, Newsome WT, Movshon JA. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci*. 1992; 12:4745–4765. [PubMed: 1464765]
23. Sherman SM, Guillory RW. On the actions that one nerve cell can have on another: Distinguishing “drivers” from “modulators”. *Proc Natl Acad Sci USA*. 1998; 95:7121–7126. [PubMed: 9618549]
24. Kuffler SW, Fitzhugh R, Barlow HB. Maintained activity in the cat's retina in light and darkness. *J Gen Physiol*. 1957; 40:683–702. [PubMed: 13428983]
25. Grossman RG, Viernstein LJ. Discharge Patterns of Neurons in Cochlear Nucleus. *Science*. 1961; 134:99–101. [PubMed: 17807391]
26. Siebert WM. Frequency discrimination in auditory system – place or periodicity mechanisms? *Proc IEEE*. 1970; 58:723–730.

27. Geisler WS, Albrecht DG. Bayesian analysis of identification performance in monkey visual cortex: Nonlinear mechanisms and stimulus certainty. *Vision Res.* 1995; 35:2723–2730. [PubMed: 7483312]

28. Churchland MM, Yu BM, Cunningham JP, et al. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci.* 2010; 13:369–378. [PubMed: 20173745]

29. Gur M, Beylin A, Snodderly DM. Response variability in primary visual cortex (V1) of alert monkey. *J Neurosci.* 1997; 17:2914–2920. [PubMed: 9092612]

30. Gershon ED, Wiener MC, Latham PE, Richmond BJ. Coding strategies in monkey V1 and inferior temporal cortices. *J Neurophysiol.* 1998; 79:1135–1144. [PubMed: 9497396]

31. Oram MW, Wiener MC, Lestienne R, Richmond BJ. The stochastic nature of precisely timed spike patterns in visual system neural responses. *J Neurophysiol.* 1999; 81:3021–3033. [PubMed: 10368417]

32. Barbieri R, Quirk MC, Frank LM, Wilson MA, Brown EN. Construction and analysis of non-Poisson stimulus-response models of neural spike train activity. *J Neurosci Methods.* 2001; 105(1):25–37. [PubMed: 11166363]

33. Kara P, Reinagel P, Reid RC. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron.* 2000; 27:635–646. [PubMed: 11055444]

34. Amarasingham A, Chen TL, Geman S, Harrison M, Sheinberg D. Spike count reliability and the Poisson hypothesis. *J Neurosci.* 2006; 26:801–809. [PubMed: 16421300]

35. Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate Effects. *J Neurophysiol.* 2005; 93:1074–1089. [PubMed: 15356183]

36. Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP. Spatiotemporal correlations and visual signaling in a complete neuronal population. *Nature.* 2008; 454:995–999. [PubMed: 18650810]

37. Pillow, JW.; Scott, JG. Fully Bayesian inference for neural models with negative-binomial spiking. In: Bartlett, P.; Pereira, FCN.; Burges, CJC.; Bottou, L.; Weinberger, KQ., editors. *Advances in Neural Information Processing Systems.* Vol. 25. 2012. p. 1907–1915.

38. Wiener MC, Richmond BJ. Decoding Spike Trains Instant by Instant Using Order Statistics and the Mixture-of-Poissons Model. *J Neurosci.* 2003; 23(6):2394–2406. [PubMed: 12657699]

39. Cohen MR, Maunsell JHR. Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci.* 2009; 12:1594–1601. [PubMed: 19915566]

40. Dorn JD, Ringach DL. Estimating membrane voltage correlations from extracellular spike trains. *J Neurophysiol.* 2003; 89:2271–2278. [PubMed: 12686584]

41. de la Rocha J, Doiron B, Shea-Brown E, Josi K, Reyes A. Correlation between neural spike trains increases with firing rate. *Nature.* 2007; 448:802–806. [PubMed: 17700699]

42. Vidne M, Ahmadian Y, Shlens J, Pillow JW, Kulkarni J, Litke AM, Chichilnisky EJ, Simoncelli EP, Paninski L. Modeling the impact of common noise inputs on the network activity of retinal ganglion cells. *J Comput Neurosci.* 2012; 33(1):97–121. [PubMed: 22203465]

43. Goris RLT, Putzeys T, Wagemans J, Wichmann FA. A neural population model for visual pattern detection. *Psychol Rev.* 2013; 120(3):472–496. [PubMed: 23915083]

44. van den Berg R, Shin H, Chou WC, George R, Ma WJ. Variability in encoding precision accounts for visual short-term memory limitations. *Proc Natl Acad Sci USA.* 2012; 109(22):8780–8785. [PubMed: 22582168]

45. Cavanaugh JR, Bair W, Movshon JA. Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *J Neurophysiol.* 2002; 88:2530–2546. [PubMed: 12424292]

46. Smith MA, Majaj NJ, Movshon JA. Dynamics of motion signaling by neurons in macaque area MT. *Nat Neurosci.* 2005; 8:220–228. [PubMed: 15657600]

47. Cox DR. Some Statistical Methods Connected with Series of Events. *J R Stat Soc.* 1955; 17(2): 129–164.

48. Greenwood M, Yule GU. An Inquiry into the Nature of Frequency Distributions of Multiple Happenings, with Particular Reference to the Occurrence of Multiple Attacks of Disease or Repeated Accidents. *J R Stat Soc A.* 1920; 83:255–279.

49. Polson NG, Scott JG, Windle J. Bayesian inference for logistic models using Polya-Gamma latent variables. *J Am Stat Soc*. 2013; 108:1339–1349.






Figure 1.


The modulated Poisson model. Spikes are generated by a Poisson process whose rate is the product of two signals: a stimulus-dependent drive, $f(S)$, that is under experimental control, and a gain signal, G , that summarizes the net effect of stimulus-independent modulatory inputs that are assumed to fluctuate slowly relative to the duration of experimental trials.

**Figure 2.**


Neural response variability originates in substantial part from gain fluctuations. **(a)** Actual and model-predicted response distributions for a V1 neuron, stimulated with gratings drifting in different directions. Responses were computed by counting spikes in a 1,000-ms window following response onset. The mean response varies with drift direction (center). Spike count histograms (outer ring) were calculated from 125 stimulus repetitions. Response distributions are superimposed on the best fitting probability densities of the Poisson (red) and gamma-modulated Poisson (blue) models. **(b)** Variance-to-mean relationship of the neural responses (grey dots, one per direction of motion), compared with predictions of the Poisson model (red line) and the modulated Poisson model (blue line). The inset shows this relation for three directions of motion (red, black, and green), where each data point is obtained from a randomly selected epoch with duration drawn uniformly from the range 1–1,000 ms in the corresponding spike raster (the red data are taken from the inset raster). Mean and variance are computed over all trials. **(c)** Log-probability of the cell responses under the Poisson model (red triangle) and the modulated Poisson model (blue triangle). Histograms illustrate the expected range of the log-probability statistic (computed with a 1,000 run parametric bootstrap) for the Poisson model (red) and the modulated Poisson model (blue). **(d)** Variance-to-mean relationships predicted by the modulated Poisson model and an additive model for weak (orange) to strong (green) fluctuations in gain.

**Figure 3.**


Comparison of neural response variability for cells in different visual areas. (a) Variance-to-mean relationship for 63 LGN cells (orange), 396 V1 cells (green), 189 V2 cells (blue) and 137 MT cells (violet). Each data point illustrates the mean and variance of the spike count in a 1,000-ms window of one cell for one stimulus condition. (b) Comparison of the predictive accuracy of the Poisson and modulated Poisson models. Log-likelihood is computed for a set of hold-out data and expressed per spike (Online Methods). (c) Distribution of stimulus-independent fluctuations in gain, summarized by the coefficient of variation of the gain. Triangles indicate the median value for each area. (d) Fraction of within-condition variance explained by gain fluctuations. Asterisks (****) indicate $P < 0.0001$.

**Figure 4.**


Response correlation analysis for three example pairs of simultaneously recorded V1 neurons. (a–c) Mean response to drifting sinusoidal gratings, as a function of direction (72 stimulus conditions, 50 repeats, 1,280-ms count window). (d–f) Spike count correlation as a function of the geometric mean of the mean spike counts of the two neurons. Each data point corresponds to a different stimulus condition. The blue line shows the correlations predicted by the modulated Poisson model, and the surrounding light blue region indicates \pm one standard deviation of the distribution of estimates computed from 50 repeats. (g–i) Spike count correlation as a function of the mean response of the two neurons, as predicted by the modulated Poisson model (color indicates correlation, points indicate response means for different stimulus conditions, as depicted in the two tuning curves shown in the first column).

**Figure 5.**

Model-based decomposition of measured spike count correlations into gain and point process correlations. (a–b) Measured spike count correlation (a), and inferred point process and gain correlations (b), as a function electrode distance. Thickness of lines indicates the 95% confidence interval. (c–d) Measured and inferred correlations plotted as a function of the correlation in mean responses (i.e., tuning curves) of the two neurons.

**Figure 6.**

Gain fluctuations are correlated over time. (a) Normalized responses as a function of time for three simultaneously recorded V1 neurons. (b) The autocorrelation function of the inferred gain for the example neurons. (c) The autocorrelation function of the gain, averaged across units for each data set. (d) The cross-correlation function of the gain, averaged across pairs for each data set.

**Figure 7.**

Analysis of spike count variance for a population of MT neurons recorded in awake, behaving macaques^{18,22}. (a) Variance-to-mean relationship for 307 MT cells. Each data point illustrates the mean and variance of the spike count in a 2,000-ms window of one cell for one stimulus condition. (b) Distribution of stimulus-independent fluctuations in gain, summarized by the coefficient of variation of the gain (top) and fraction of within-condition variance explained by gain fluctuations (bottom). (c) The autocorrelation function of the gain, averaged across units (trials are assumed to be separated by 5 sec).