

Slow gain fluctuations limit benefits of temporal integration in visual cortex

Robbe L. T. Goris

Center for Neural Science, New York University,
New York, NY, USA
Howard Hughes Medical Institute, New York University,
New York, NY, USA
Present address: Center for Perceptual Systems,
The University of Texas at Austin, Austin, TX, USA

Corey M. Ziemba

Center for Neural Science, New York University,
New York, NY, USA
Howard Hughes Medical Institute, New York University,
New York, NY, USA
Present address: Center for Perceptual Systems,
The University of Texas at Austin, Austin, TX, USA

J. Anthony Movshon

Center for Neural Science, New York University,
New York, NY, USA

Eero P. Simoncelli

Center for Neural Science, New York University,
New York, NY, USA
Howard Hughes Medical Institute, New York University,
New York, NY, USA



Sensory neurons represent stimulus information with sequences of action potentials that differ across repeated measurements. This variability limits the information that can be extracted from momentary observations of a neuron's response. It is often assumed that integrating responses over time mitigates this limitation. However, temporal response correlations can reduce the benefits of temporal integration. We examined responses of individual orientation-selective neurons in the primary visual cortex of two macaque monkeys performing an orientation-discrimination task. The signal-to-noise ratio of temporally integrated responses increased for durations up to a few hundred milliseconds but saturated for longer durations. This was true even when cells exhibited little or no adaptation in their response levels. These observations are well explained by a statistical response model in which spikes arise from a Poisson process whose stimulus-dependent rate is modulated by slow, stimulus-independent fluctuations in gain. The response variability arising from the Poisson process is reduced by temporal integration, but the slow modulatory nature of variability due to gain fluctuations is not. Slow gain fluctuations therefore

impose a fundamental limit on the benefits of temporal integration.

Introduction

Many visually driven activities require observers to integrate information over time. Traditional models of perceptual decision-making suggest that the signal-to-noise ratio (SNR) of integrated sensory information should improve in proportion to the square root of integration time (Bloch, 1885; Green & Swets, 1966; Watson, 1979). However, in experimentally controlled tasks, observer performance grows at substantially slower rates, in many cases saturating at relatively short time scales (Barlow, 1958; Grice, 1972; Gorea & Tyler, 1986; Burr & Santoro, 2001). This failure has been ascribed to multiple causes, including rapid sensory adaptation (Barlow, 1958), faulty accumulation of available information (Grice, 1972; McClelland, 1979; Busemeyer, Jessup, Johnson, & Townsend, 2006), and

Citation: Goris, R. L. T., Ziemba, C. M., Movshon, J. A., & Simoncelli, E. P. (2018). Slow gain fluctuations limit benefits of temporal integration in visual cortex. *Journal of Vision*, 18(8):8, 1–13, <https://doi.org/10.1167/18.8.8>.

<https://doi.org/10.1167/18.8.8>

Received February 20, 2018; published August 22, 2018

ISSN 1534-7362 Copyright 2018 The Authors



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

Downloaded From: <https://jov.arvojournals.org/pdfaccess.ashx?url=/data/journals/jov/937434/> on 08/23/2018

incorporation of task-related incentives for rapid decision-making (Kiani, Hanks, & Shadlen, 2008; Drugowitsch, Moreno-Bote, Churchland, Shadlen, & Pouget, 2012).

But temporal integration may be limited by a more elementary factor: the structure of noise in sensory responses. Indeed, a previous study that examined temporal averaging of single-cell responses in anesthetized animals concluded that temporal response correlations reduce the benefits of integration (Osborne, Bialek, & Lisberger, 2004). And measurements of neural activity with a voltage-sensitive dye have suggested that such correlations may also limit temporal integration at the level of neural populations (Chen, Geisler, & Seidemann, 2008). However, the nature and consequences of temporally correlated noise have not been fully explored.

Recently, several groups have shown that responses of visual neurons are often well described as arising from a Poisson process whose stimulus-dependent rate is modulated by slow, stimulus-independent fluctuations in gain (Ecker et al., 2014; Goris, Movshon, & Simoncelli, 2014; Lin, Okun, Carandini, & Harris, 2015; Rabinowitz, Goris, Cohen, & Simoncelli, 2015; Zylberberg, Camaro, Turner, Shea-Brown & Rieke, 2016). Gain fluctuations arise from modulatory signals such as attentional focus (Luck, Chelazzi, Hillyard, & Desimone, 1997; Ecker, Denfield, Bethge, & Tolias, 2015; Rabinowitz et al., 2015), reward expectation (Baruni, Lau, & Salzman, 2015), arousal (Kato, Chu, Isaacson, & Komiyama, 2012), and local gain control circuitry (Carandini & Heeger, 2012). Might these gain fluctuations give rise to a form of temporal correlation that could explain limited temporal integration?

Perceptual decisions are based on the activity of large populations of sensory neurons. Gain fluctuations are often shared across neurons (Lin et al., 2015; Rabinowitz et al., 2015) and can limit the coding capacity of neural populations (Kanitscheider, Coen-Cagli, & Pouget, 2015). It is therefore possible that gain fluctuations limit temporal integration at the behavioral level. In the current study, we examine this by analyzing the responses of single neurons in the context of a matched behavioral task.

We measured responses of orientation-selective V1 neurons in awake macaques performing an orientation-discrimination task and fitted them with a modulated Poisson model. We found that for most neurons, temporal response correlations arising from modulatory fluctuations severely limited the benefits of temporal integration. As predicted by the model, the magnitude of this limitation depended on both the variance of gain fluctuations and the strength of stimulus-driven responses. Rapid sensory adaptation was only weakly associated with limited integration benefits. Analysis of a previously published data set

revealed qualitatively similar effects in extrastriate area MT. Together, these results suggest that slow gain fluctuations impose a fundamental limit on temporal integration throughout the visual cortex.

Methods

The data shown here are reported in full by Goris, Ziemba, Stine, Simoncelli, and Movshon (2017), who used them to examine choice-correlated activity. Here we briefly summarize the behavioral task and unit recording.

Behavioral task

Two male macaque monkeys (one *Macaca mulatta*, one *M. nemestrina*) were trained to perform an orientation-discrimination task. They were seated with their heads stabilized in a dimly lit, sound-isolated room in front of a gamma-corrected CRT monitor. We presented drifting sinusoidal gratings of varying orientation at an eccentricity of approximately 5° of visual angle for a duration of 500 ms. The stimulus was positioned within the neuron's receptive field. Stimulus orientation varied over a range of 30°. The animals judged the orientation of the stimulus relative to a discrimination boundary which was chosen to correspond to the steepest part of the neuron's orientation-tuning function (estimated by visually inspecting the orientation tuning function measured in a fixation task prior to the decision-making task). After the stimulus had disappeared, the animals communicated their decision via a saccadic eye movement toward one of two choice targets. Stimuli were presented in random order and repeated at least 10 times (for the data presented here, mean = 43 repeats). The animals' behavioral orientation acuity approximated that of human observers tested under similar conditions (Goris et al., 2017).

Unit recording

Extracellular recordings were made with dura-penetrating microelectrodes, advanced mechanically into the primary visual cortex (V1). We made recordings from every single unit with a spike waveform that rose sufficiently above noise to be isolated. We first presented suitably vignetted sinusoidal gratings to map each isolated unit's receptive field location in a fixation task. Specifically, the aperture consisted of a raised cosine edge and a flat top which covered three quarters of the window. Thereafter, we

determined preferred stimulus size, spatial frequency, and drift rate by hand and measured neuronal selectivity for orientation by computer. Most analyses reported here—all except the one in Figure 2—include only data from trials in the behavioral decision-making task when the grating was at its most effective orientation (i.e., the orientation which elicited the highest firing rate). When firing rate is high, neural responses tend to be most informative about stimulus identity (Geisler & Albrecht, 1995). This is therefore a favorable condition to study temporal integration.

In addition to our own V1 recordings, we analyzed MT recordings which were previously published by Britten, Shadlen, Newsome, and Movshon (1992) and downloaded from the Neural Signal Archive (www.neuralsignal.org; nsa2004.1). These recordings were made while monkeys judged the direction of a stochastic motion stimulus. We analyze only data from trials in the behavioral task for the most effective motion stimulus.

Analysis

Our analysis focuses on the effects of temporal integration on the reliability of neural responses. We quantified the response reliability of neurons with the squared signal-to-noise ratio (SNR^2), the squared ratio of the mean to the standard deviation of the spike count of a single neuron within a given temporal window. We chose to use SNR^2 because a homogeneous Poisson process predicts linear growth in this quantity with time. We computed the *momentary response* by counting spikes in a sliding 10-ms window and the *integrated response* by counting spikes in a window of 10, 20, 30 ms, etc., beginning at stimulus onset.

We studied both absolute and relative reliability of the integrated response. The latter statistic is defined as the ratio of the measured to the predicted SNR^2 , whereby the prediction is derived under the assumption that response variability is temporally uncorrelated (see Results). We classified relative reliability as statistically significant if it fell outside of the central 95% of the expected null distribution, computed from 1,000 randomly permuted data sets.

We also studied the effects of transient response dynamics (*adaptation*) on temporal integration. We computed the relative strength of the reliability transient by taking the ratio of the maximal SNR^2 during the first 100 ms to the mean SNR^2 over the remaining 400 ms.

For one analysis, we measured neuronal ability to discriminate stimulus orientation by fitting the relationship between stimulus orientation and probability of “clockwise” choice for an ideal observer with a

cumulative Gaussian function. The ideal observer’s choices were obtained by applying a deterministic decision criterion to the responses of each neuron. To minimize bias in the ideal observer’s choices, the criterion was set to the median response to the zero signal stimulus (i.e., the stimulus whose orientation matches the orientation-discrimination boundary). We defined the neuronal discrimination threshold as the standard deviation of the cumulative Gaussian. To investigate how neuronal discrimination threshold depends on temporal integration, we performed this analysis on responses obtained by counting spikes in a window of 50, 73, 108, 232, 341, and 500 ms beginning at the onset of the visual response. For each cell, we determined response latency by maximizing the stimulus-associated response variance (Smith, Majaj, & Movshon, 2005).

Model predictions

We compare the measured effects of temporal integration with the predictions of a previously proposed model for neural response variability: the modulated Poisson model. In this model, spikes arise from a Poisson process whose stimulus-dependent rate is modulated by slow, stimulus-independent fluctuations in gain (Goris et al., 2014). We derive the model’s predictions for temporal integration in the Results.

Results

Limitations on temporal integration in V1

We recorded the activity of individual orientation-selective V1 neurons in two macaque monkeys while they judged the orientation of drifting sinusoidal gratings presented for 500 ms within the neuron’s receptive field. Grating position, size, spatial frequency, and speed were chosen to match the neuron’s preference. Grating orientation varied randomly across trials within a 30° range centered on the steepest part of the orientation tuning curve, and we measured the animals’ perceptual ability to discriminate orientation. A detailed analysis of the behavior and its connection to the neural responses has been previously published (Goris et al., 2017). Here we primarily consider data obtained when the orientation of the grating was optimal for the recorded neuron.

Consider the spikes fired by an example neuron in response to a stimulus that drove it well (Figure 1A, neuron 1). Repeated presentations of the stimulus evoked different responses. This variability limits the

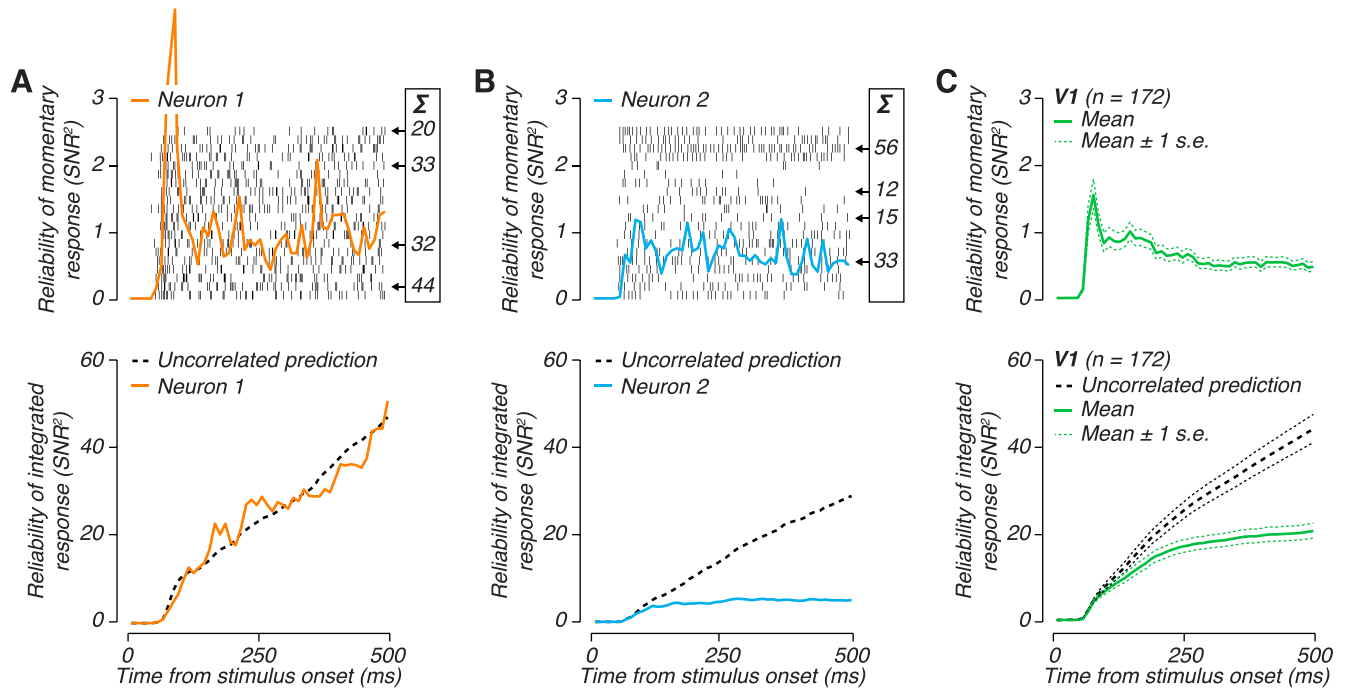


Figure 1. The benefits of temporal integration of V1 activity are limited. (A) Temporal evolution of the reliability (SNR^2 , orange line) of the momentary response (top, 1 moment = 10 ms) and of the integrated response (bottom) for an example neuron. Tick marks in the top panel indicate spikes; the dashed line in the bottom panels illustrates the prediction for temporally uncorrelated response variability. (B) Same as (A) for a different example neuron. (C) Geometric mean of the reliability of the momentary response (top) and the integrated response (bottom) for a population of V1 neurons. The dashed line in the bottom panel is the mean uncorrelated prediction.

amount of information about the stimulus that can be extracted by a downstream decoder. We capture this quantity with SNR^2 , the ratio of the squared mean to the variance of the spike count, which we hereafter term *reliability*. Under our experimental conditions, orientation tuning is stable over time (Mazer, Vinje, McDermott, Schiller, & Gallant, 2002). It is therefore reasonable to harvest the encoded stimulus information by accumulating spikes over the entire duration of the stimulus. In particular, if the mean and variance across trials were identical for all time bins and the spike counts were uncorrelated across time bins, then accumulated reliability would grow linearly with duration. In reality, however, mean and variance usually fluctuate within a trial (Churchland et al., 2010). For neuron 1, reliability measured within 10-ms intervals peaked just after response onset, then quickly dropped and remained low for the rest of the stimulus epoch (Figure 1A, top panel). As a result, if spikes are temporally uncorrelated, reliability of the accumulated response should grow rapidly at first and thereafter at a slower rate (Figure 1A, dashed curve in the bottom panel). For this particular neuron, integrating responses over time yields a reliability that closely follows this prediction (Figure 1A, orange curve in the bottom panel).

Factors other than response dynamics can also limit the rate at which reliability grows with integration time.

Consider example neuron 2 (Figure 1B). Here the epoch of effective temporal integration was brief, even though there was no onset transient. The reliability of accumulated activity saturated after roughly 150 ms, and integration over the full 500-ms interval yielded a reliability that was much smaller than the uncorrelated prediction (Figure 1B, bottom panel). What limits reliability in this neuron? Compared with neuron 1, the spike trains of neuron 2 show substantial variability in overall activity across trials, and firing rates are correlated across time within trials (Figure 1A and B, top panels). Indeed, just as interneuronal correlations limit the benefits of integrating responses of identically tuned neurons (Zohary, Shadlen, & Newsome, 1994), temporal response correlations limit the benefits of integrating activity over time (Osborne et al., 2004; Chen et al., 2008). Specifically, the SNR^2 of the spike count, accumulated over discrete bins, is

$$\text{SNR}^2(T) = \frac{\left(\sum_{i=1}^T E[N_i]\right)^2}{\sum_{i=1}^T \text{var}[N_i] + \sum_{i \neq j}^T \text{cov}[N_i, N_j]}, \quad (1)$$

where N_i is the spike count in the i th time bin, T is the number of time bins being accumulated, and $E[\cdot]$, $\text{var}[\cdot]$, and $\text{cov}[\cdot, \cdot]$ are the mean, variance, and covariance

(respectively) of the binned spike counts. If the counts are uncorrelated, the second term in the denominator is zero and reliability will grow with time. But if responses are temporally correlated, reliability will grow more slowly.

We analyzed responses of 172 orientation-selective V1 neurons (78 from monkey 1, 94 from monkey 2). The population average reveals that the estimated momentary reliability of the neuronal ensemble tended to peak early after response onset, and then fell to a relatively stable level for the remainder of the stimulus epoch (Figure 1C, top panel). Most growth in the reliability of accumulated activity occurred during the first 250 ms, after which it fell short of the uncorrelated prediction (Figure 1C, bottom panel).

The reliability with which neurons can relay stimulus information is oftentimes measured by determining performance in a discrimination task for an ideal observer who only has access to the neural responses (Tolhurst, Movshon, & Dean, 1983). This kind of *neurometric* analysis differs from ours in that it critically relies on a comparison of neuronal responses across stimulus conditions, whereas we have so far only considered responses within a single stimulus condition. We wondered whether the effects of temporal integration would be comparable for both measures of signaling capacity. For each neuron, we computed a family of neurometric functions, using different temporal integration intervals (see Methods). Figure 2A shows three such functions, averaged across all neurons recorded from one animal. Prior to averaging, we converted neurometric performance to d' (Green & Swets, 1966), a statistic whose variance does not depend on performance level. After averaging, we converted d' back into choice proportion. A steeper slope of the neurometric function corresponds to an increase in neuronal sensitivity for stimulus orientation, and is equivalent to a decrease in orientation-discrimination threshold. As can be seen, increasing integration time from 50 to 232 ms systematically increases the slope of the neurometric function (Figure 2A). However, as was the case for SNR^2 (Figure 1C), increasing the integration window beyond 232 ms yielded little further benefit (Figure 2B). This behavior deviates from the traditional prediction that neuronal sensitivity should improve in proportion to the square root of integration time (Figure 2B, dashed curve). Given the similar behavior of both statistics, we hereafter exclusively focus on SNR^2 .

The curves in Figure 1C show that, on average, reliability at the end of a 500-ms trial was about half of the uncorrelated prediction. Because our calculation takes account of response dynamics, this shift must be due to temporal correlations in neuronal firing. How might these correlations arise?

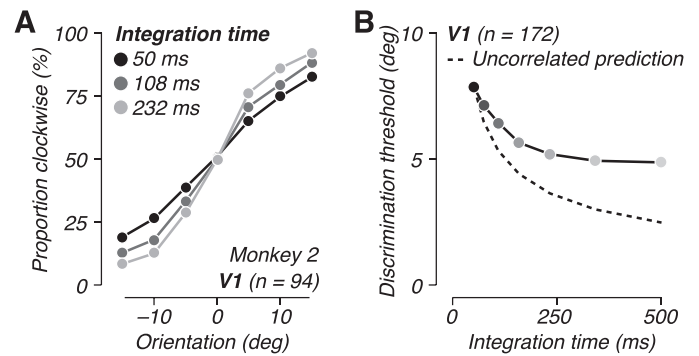


Figure 2. The effects of temporal integration on the neurometric function. (A) Mean neurometric function for a population of V1 neurons recorded from monkey 2, computed for three different integration windows. (B) Temporal evolution of the median orientation-discrimination threshold (estimated from the slope of the neurometric function) for a population of V1 neurons, recorded from two monkeys. The dashed line illustrates the prediction for a model with identical spike-count mean and variance for all time bins but uncorrelated across time. Under this model, thresholds decrease in proportion to the square root of integration time.

Temporal integration in the modulated Poisson model

We recently showed that spike counts of neurons in many areas of visual cortex can be described by a Poisson process whose stimulus-dependent rate is modulated by slow, stimulus-independent fluctuations in gain (Goris et al., 2014):

$$N_i \sim \text{Poiss}(\mu_i g_i),$$

where N_i is spike count in the i th time bin, μ_i is the mean stimulus-driven spike count in that bin, and g_i is a stochastic gain signal for that bin, with mean 1 and variance σ_g^2 . Under this model, we can express the statistics of the binned counts as (Goris et al., 2014)

$$\begin{aligned} E[N_i] &= \mu_i \\ \text{var}[N_i] &= \mu_i + \mu_i^2 \sigma_g^2 \\ \text{cov}[N_i, N_j] &= \mu_i \mu_j \text{cov}[g_i, g_j] \end{aligned} \quad (2)$$

The modulated Poisson model has direct implications for response integration. Consider the special case in which both the mean response and gain signal are constant over a trial (see Appendix for the general case). In this case, $\text{cov}[g_i, g_j] = \sigma_g^2$, and substituting Equation 2 into Equation 1 yields

$$\text{SNR}^2(T) = \frac{(T\mu)^2}{T\mu + T^2\mu^2\sigma_g^2} = \frac{T\mu}{1 + T\mu\sigma_g^2}. \quad (3)$$

In the absence of gain fluctuations ($\sigma_g^2 = 0$), the model reduces to the standard Poisson, and reliability grows

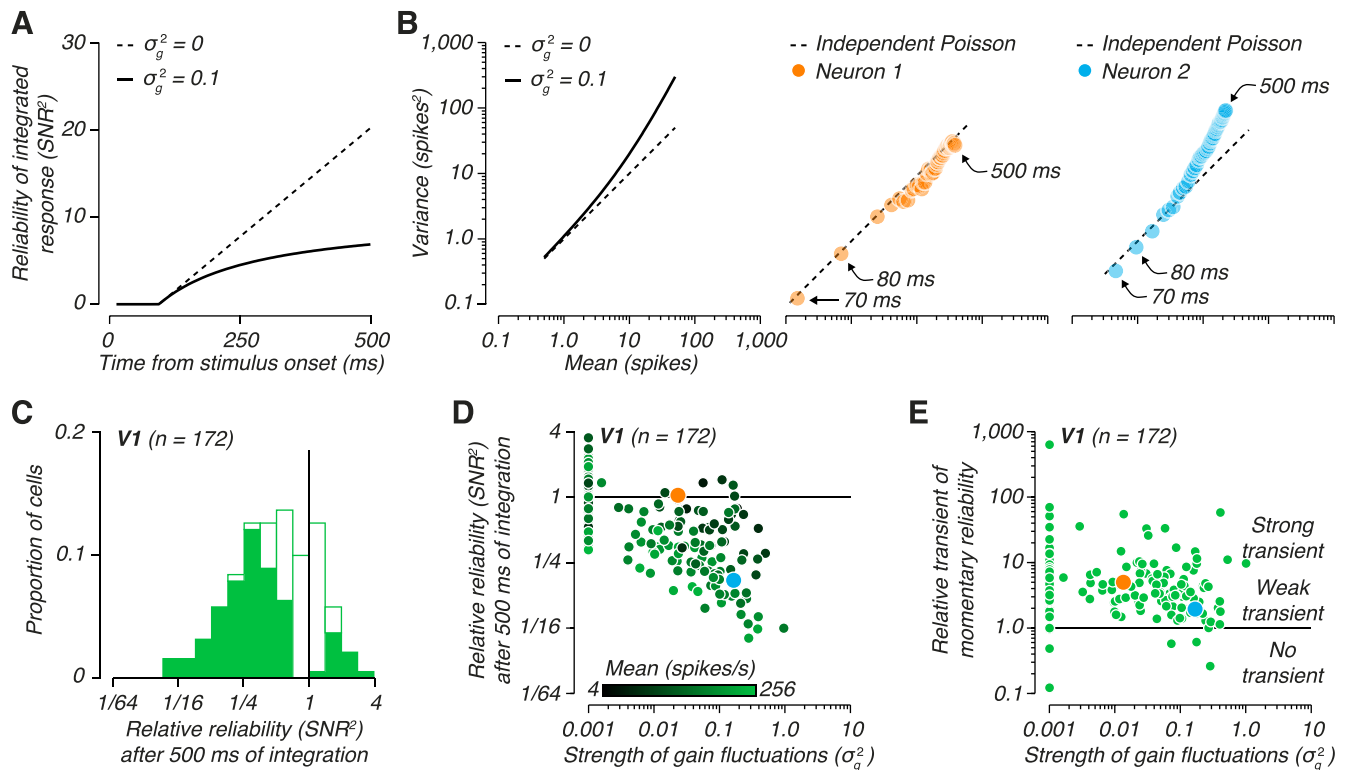


Figure 3. The modulated Poisson model accounts for the effects of temporal integration. (A) Predicted temporal evolution of the SNR^2 of integrated activity for two levels of gain fluctuations, for a model neuron with constant momentary reliability. (B) Left: Predicted variance-to-mean relationship for two levels of gain fluctuations, for the same simulated neuron. Middle: Variance-to-mean relationship of example neuron 1 (see Figure 1), measured with variable duration windows. Right: Variance-to-mean relationship of example neuron 2 (see Figure 1). (C) Distribution of the reliability after 500 ms of integration, relative to the uncorrelated prediction, for a population of V1 neurons. Filled entries in the distribution indicate neurons for which SNR^2 differed significantly from the uncorrelated prediction (1,000-fold bootstrap, $p < 0.05$). (D) Relative reliability after 500 ms of integration plotted against strength of gain fluctuations, estimated for a population of V1 neurons. Points are shaded according to the mean firing rate (over 500 ms). Orange and blue points correspond to neurons 1 and 2, respectively. (E) Relative amplitude of transient of the momentary response plotted against strength of gain fluctuations.

linearly with time (Figure 3A), as for neuron 1 (Figure 1A). But in the presence of substantial gain fluctuations, response variability is temporally correlated and reliability saturates with time (Figure 3A), as was the case for neuron 2 (Figure 1B).

In the modulated Poisson model, saturation of reliability arises from the growth of variance with accumulated activity. Specifically, when fluctuations in gain are weak, response variance is approximately equal to the mean spike count, as predicted by a Poisson spiking model (Figure 3B, left panel) and seen in neuron 1 (Figure 3B, middle panel). When fluctuations in gain are strong, response variance grows as the squared mean (Figure 3B, left panel), as seen in neuron 2 (Figure 3B, right panel).

Under the modulated Poisson model, the limits on temporal integration imposed by gain fluctuations differ across neurons. This can be seen easily by considering the relative reliability of integrated activity (i.e., the ratio of the reliability under the modulated

Poisson model to the reliability under an independent Poisson model):

$$\text{SNR}_{\text{Rel}}^2(T) = \frac{1}{1 + T\mu\sigma_g^2}. \quad (4)$$

Reliability will deviate more from the uncorrelated prediction for neurons with stronger gain fluctuations.

To test this prediction, we estimated σ_g^2 and relative reliability for each neuron. For the latter statistic we computed the ratio of measured to predicted reliability at 500 ms. For the former statistic, we used the method described by Goris et al. (2014). Specifically, we assumed that the gain is constant within a trial and distributed across trials according to a gamma distribution with mean 1 and variance σ_g^2 . We estimated this parameter by maximizing the likelihood of the full set of observed spike counts in the orientation-discrimination task under a gamma mixture of Poisson distributions. As can be seen in Figure 3C, the entire distribution of relative reliability is shifted below 1

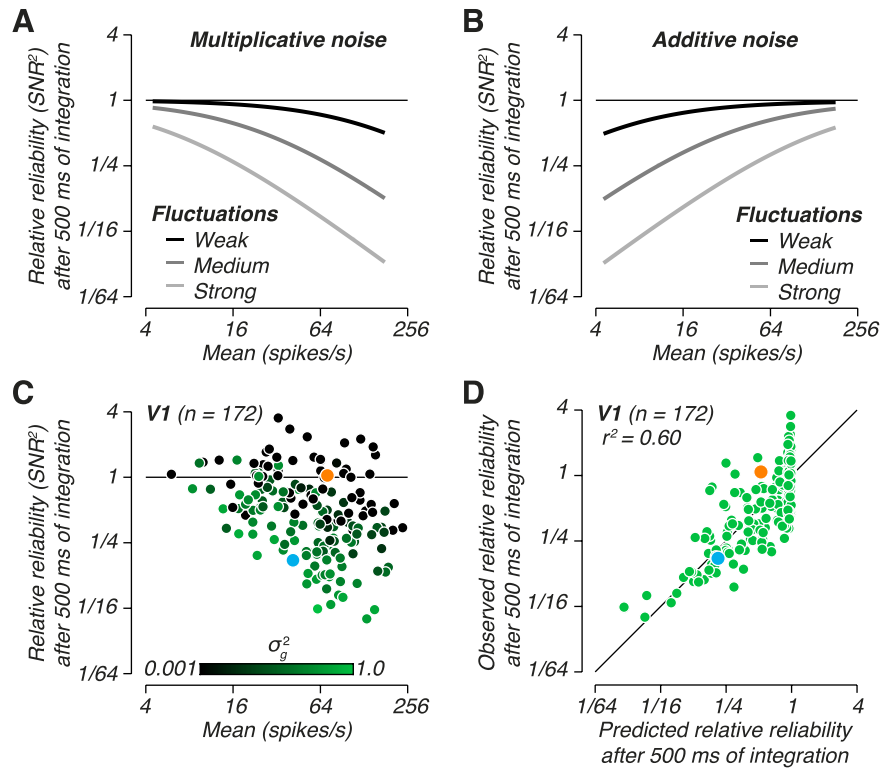


Figure 4. Comparison of multiplicative and additive noise models. (A) Simulated relative reliability after 500 ms of integration plotted against mean firing rate for three levels of multiplicative noise. (B) Simulated relative reliability after 500 ms of integration plotted against mean firing rate for three levels of additive noise. (C) Relative reliability after 500 ms of integration plotted against mean firing rate, estimated for a population of V1 neurons. Points are shaded according to the strength of gain fluctuations. Orange and blue points correspond to neurons 1 and 2, respectively. (D) Relative reliability after 500 ms of integration plotted against predictions derived from the modulated Poisson model using Equation 4.

(geometric mean = 0.47, $p < 0.001$, t test), but the effects of temporal integration are quite diverse. We found the account of the modulated Poisson model for this diversity to be consistent with the data: Relative reliability decreases with gain fluctuations (Figure 3D, $r_s = -0.61$, $p < 0.001$, Spearman correlation).

We have shown that the growth of reliability with integration time can slow either because of a drop in the momentary reliability (i.e., after an initial transient) or because of gain fluctuations. We wondered whether these factors might be related. We measured the strength of the reliability transient by taking the ratio of the maximal SNR^2 during the first 100 ms to the mean SNR^2 over the remaining 400 ms. As can be seen in Figure 3E, the association between both factors was weak ($r_s = -0.22$, $p < 0.01$, Spearman correlation), probably too weak to be meaningful.

Equation 4 reveals a second prediction of the modulated Poisson model: Reliability will deviate more from the uncorrelated prediction for neurons with higher firing rates (Figure 4A). This is a direct consequence of the multiplicative nature of gain fluctuations. If slow fluctuations in firing rate were additive instead, as proposed by some authors (Arieli,

Sterkin, Grinvald, & Aertsen, 1996; Lin et al., 2015), relative reliability would increase with firing rate (Figure 4B; see Appendix). Again, we found the prediction of the modulated Poisson model to be consistent with the data: Across the population of V1 neurons, relative reliability decreases with firing rate (Figure 4C, $r_s = -0.37$, $p < 0.001$, Spearman correlation).

Together, the influence of firing rate and gain fluctuations on relative reliability suggest that the modulated Poisson model should make reasonable predictions on a cell-by-cell basis. We found that predictions derived from Equation 4 accounted for a substantial portion of the variance in observed relative reliability (Figure 4D, $r^2 = 0.60$, $p < 0.001$).

Determinants of encoded information

Equation 3 reveals that different factors control reliability at different moments during the trial. Early in the trial, when T is small, reliability grows with T at a rate of μ . But when T is large, the denominator is dominated by the gain fluctuations term, and reliability

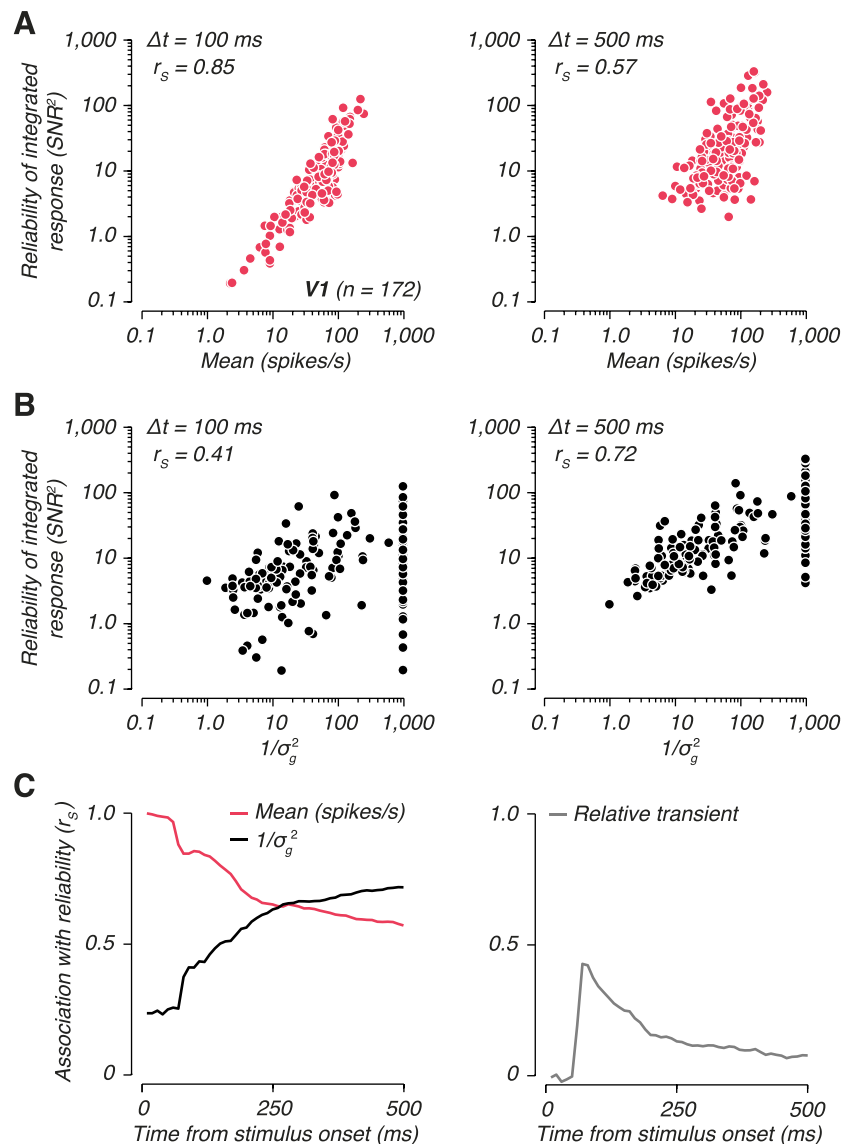


Figure 5. The primary determinant of encoded information changes with time. (A) The relationship between reliability of integrated activity and mean firing rate after 100 ms (left) and 500 ms (right) of stimulus exposure for a population of V1 neurons. (B) The relationship between reliability of integrated activity and estimates of $1/\sigma_g^2$ after 100 and 500 ms of stimulus exposure for the same population of V1 neurons. (C) Temporal evolution of the Spearman correlation (r_s) between the SNR^2 of integrated activity and mean firing rate (left, red), and $1/\sigma_g^2$ (left, black); and between the SNR^2 of integrated activity and the relative amplitude of the transient of the momentary response (right).

saturates at a level of $1/\sigma_g^2$. To test these predictions, we examined the correlation of μ and σ_g^2 with the reliability of the accumulated response. Early in the trial (100 ms after stimulus onset), neuron-to-neuron variations in reliability were strongly associated with differences in mean firing rate and only weakly associated with differences in the estimated strength of gain fluctuations (Figure 5A and B, left panels). But by the end of the stimulus epoch, the importance of mean firing rate had diminished, while the importance of gain fluctuations had grown (Figure 5A and B, right panels). After 250 ms of integration, gain fluctuations were a better predictor of reliability than mean firing rate

(Figure 5C, left). In contrast, variations in strength of the initial transient were associated with variations in reliability only around the time of response onset (Figure 5C, right).

Temporal integration in extrastriate area MT

Is imperfect temporal integration a general property of the visual cortex? In previous work, we have found that the modulated Poisson model could account for spiking variability throughout the visual cortex, and

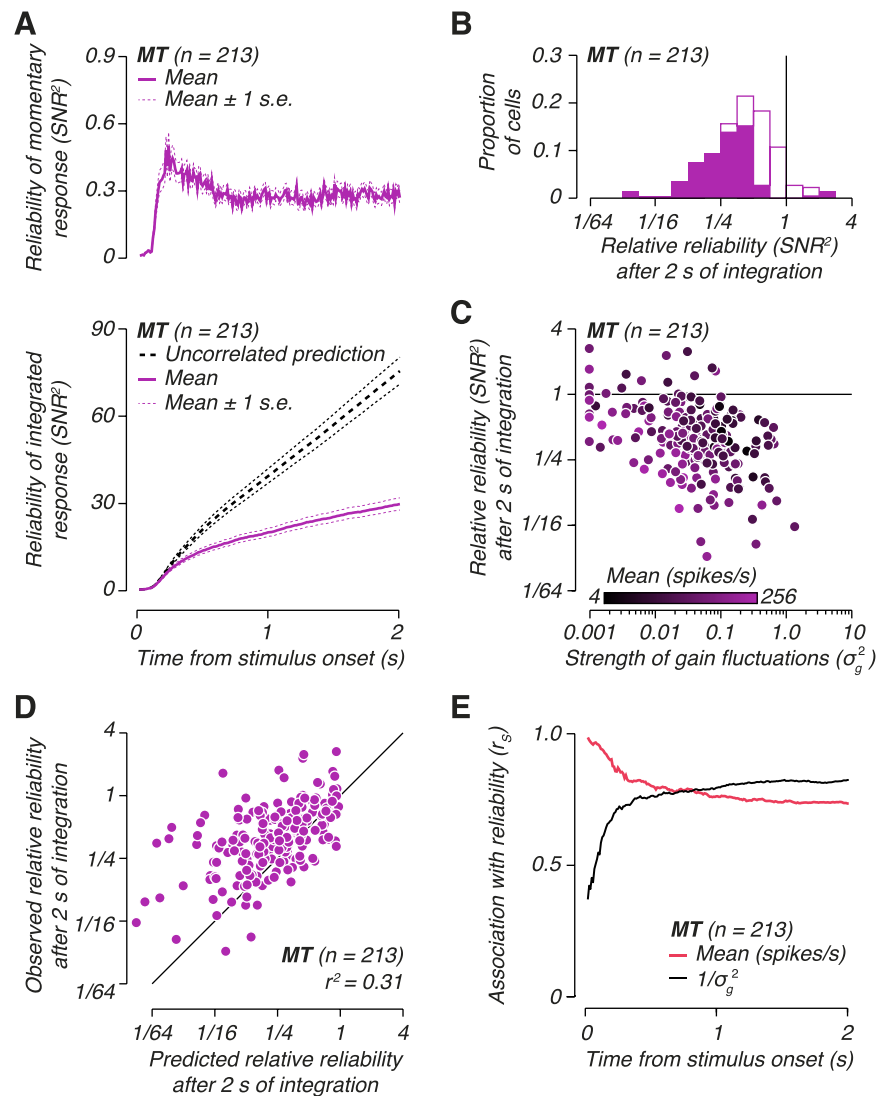


Figure 6. Limited benefits of temporal integration are also found in neurons from cortical area MT. (A) Data in the same format as Figure 1C for the population of MT neurons reported by Britten et al. (1992). (B) Analysis as in Figure 3C for the population of MT neurons. (C) Analysis as in Figure 3D for the population of MT neurons. (D) Analysis as in Figure 4D for the population of MT neurons. (E) Analysis as in Figure 5C for the population of MT neurons.

that the estimated contribution of gain fluctuations was more substantial in later stages of the visual hierarchy (Goris et al., 2014). Therefore, we expect that temporal integration will be imperfect across the visual cortex, and more so in later stages. As a specific example, we examine a published data set of responses in extrastriate area MT, downstream of V1, that were recorded while monkeys judged the direction of a stochastic motion stimulus (Britten et al., 1992). Consistent with previous observations in anesthetized animals (Osborne et al., 2004), the benefits of temporal integration were significantly less than expected under the uncorrelated model (Figure 6A). As was the case in V1, this phenomenon was observed to varying degrees in most individual neurons (Figure 6B). This diversity could in part be accounted for by differences in the strength of

gain fluctuations (Figure 6C, $r_s = -0.37$, $p < 0.001$, Spearman correlation). Predictions from the modulated Poisson model also accounted for a substantial proportion of the variance in relative reliability in MT (Figure 6D, $r^2 = 0.31$, $p < 0.001$). However, model performance for MT was worse than for V1. For the majority of MT neurons, the model underestimated the benefits of temporal integration (Figure 6D), presumably because the saturation of reliability with integration time was less complete than for neurons in V1 (Figure 6A). This might be due to the fact that the moving dot stimuli were not repeated exactly across trials (and thus some response variability arises from temporally independent stimulus fluctuations), or it might signify that the gain varied somewhat within each 2-s trial. Despite this, many features of the MT

data were qualitatively consistent with both the V1 data and the modulated Poisson model, including a shift in the factors controlling reliability over the course of the trial (Figure 6E).

Discussion

The ability to integrate different sources of sensory information over time is a cornerstone of our perceptual and cognitive capabilities (Yang & Shadlen, 2007; Ganmor, Landy, & Simoncelli, 2015). It is often assumed that a benefit of temporal integration is the mitigation of limitations arising from noisy sensory representations. However, this crucially depends on the temporal structure of neural response variability (Osborne et al., 2004; Chen et al., 2008). We have shown that fluctuations in the activity of most V1 and MT neurons are correlated over time, and that this correlation severely limits the benefits of temporal integration. Indeed, for the V1 and MT data that we have analyzed, temporal correlations play a much more important role than rapid sensory adaptation. Our analysis suggests that temporal correlations arise from slow fluctuations in neuronal excitability: A Poisson spiking model with a fluctuating gain signal explains why the reliability of integrated activity saturates with time, and why the relative benefits of temporal integration are smaller for neurons that fire more vigorously. These phenomena likely extend beyond V1 and MT, as gain fluctuations appear to be a general property of the visual cortex (Goris et al., 2014).

We tailored the orientation-discrimination task to the tuning properties of the neurons we recorded from, so as to maximize the relevance of their responses for behavior. Ultimately, however, we do not know which sensory signals support perceptual behavior. Our observations do not rule out the possibility that downstream areas manage to isolate the inputs of the subset of sensory neurons whose responses are not impaired by significant gain fluctuations. Or perhaps these areas implement a temporal read-out strategy which suffers less from gain fluctuations than simple accumulation (Chen et al., 2008), for example, by incorporating knowledge of the modulatory signal driving the fluctuations (Rabinowitz et al., 2015). While these hypotheses cannot be dismissed, to be consistent with the limited behavioral benefits of temporal integration (Barlow, 1958; Grice, 1972; Gorea & Tyler, 1986; Burr & Santoro, 2001) they would need to be combined with a different loss-of-information mechanism. The simpler interpretation is that imperfect temporal integration at the behavioral level arises from imperfect temporal integration at the level of sensory neurons.

The execution of perceptual tasks such as detection, discrimination, and identification relies on combining responses across both neurons and time. It is well known that pooling responses across identically tuned neurons yields limited benefits for signaling capacity when these responses are correlated (Zohary et al., 1994). And if the pooling rule is not optimal, across-neuron correlations can also limit the capacity of diversely tuned populations (Shadlen, Britten, Newsome, & Movshon, 1996; Graf, Kohn, Jazayeri, & Movshon, 2011; Goris, Putzeys, Wagemans, & Wichmann, 2013; Moreno-Bote et al., 2014). While the origin of between-neuron correlations is generally not known, recent work suggests they may in large part arise from shared gain fluctuations (Ecker et al., 2014; Goris et al., 2014; Lin et al., 2015; Rabinowitz et al., 2015; Zylberberg et al., 2016), which cannot be averaged out by pooling across neurons. Here we have demonstrated that gain fluctuations also affect temporal integration at the level of single cells. Thus, the variability that limits the benefits of cross-neuron pooling and cross-time integration might arise from a common source.

We found that the benefits of temporally integrating sensory responses are rather short-lived and diminish substantially after a few hundred milliseconds. While some previous reports have documented related observations (Uka & DeAngelis, 2003; Osborne et al., 2004; Chen et al., 2008), others have not. In particular, one study of V1 neurons has concluded that neural discriminability improves significantly with integration time (Vogels & Orban, 1990). The apparent discrepancy with our results may be due to the fact that that study considered only two temporal intervals (175 ms and 450 ms)—for these intervals, our data also reveal a mild improvement in signaling capacity (Figure 2B). A second reason may be that we optimized the stimulus for the tuning properties of each neuron, whereas Vogels and Orban did not. Optimized stimuli yield stronger neural responses, which we have shown are associated with less efficient temporal integration (Figure 4C).

What are the implications of our findings for understanding the role of temporal integration in perception? First, under natural viewing conditions observers explore visual scenes by rapidly moving their eyes from one location to another. Fixations between successive eye movements usually last between 200 and 400 ms (Yarbus, 1967), a duration approximately matched to that at which we have found integration of sensory responses to yield diminishing returns (Figures 1C and 6A). Second, we have shown that slow fluctuations in sensory gain impose a fundamental limit on the perceptual benefits that can be obtained from temporal integration by a decoder unaware of these fluctuations. Indeed, gain fluctuations affect the *encoding* of sensory information. Previous studies have suggested that the limited benefits of temporal integration in behavioral tasks may arise during the

decoding of sensory signals, either through imperfect (“leaky”) integration (Grice, 1972; McClelland, 1979; Busemeyer et al., 2006) or through a policy preference to minimize response latency (Kiani et al., 2008; Drugowitsch et al., 2012). In contrast, our analysis provides evidence for the idea that elementary limitations in sensory encoding may be critical—it may not matter whether integration is leaky or organisms have an urgent need to respond fast, if sensory integration is intrinsically limited.

Keywords: visual cortex, computational modeling, temporal integration, decoding, temporal modulation

Acknowledgments

This work was supported by National Institutes of Health Grants EY04440 and EY022428, Howard Hughes Medical Institute, National Science Foundation Graduate Research Fellowship to C.M.Z., and postdoctoral fellowships from the Fund for Scientific Research of Flanders and the Belgian American Educational Foundation to R.L.T.G. The authors declare no competing financial interests.

Commercial relationships: none.

Corresponding author: Eero P. Simoncelli.

Email: eero.simoncelli@nyu.edu.

Address: Center for Neural Science, New York University, New York, NY, USA.

References

- Arieli, A., Sterkin, A., Grinvald, A., & Aertsen, A. (1996, September 27). Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science*, *273*, 1868–1871.
- Barlow, H. B. (1958). Temporal and spatial summation in human vision at different background intensities. *The Journal of Physiology*, *141*, 337–350.
- Baruni, J. K., Lau, B., & Salzman, C. D. (2015). Reward expectation differentially modulates attentional behavior and activity in visual area V4. *Nature Neuroscience*, *18*, 1656–1663.
- Bloch, M. A.-M. (1885). Expériences sur la vision: Essai d’Optique sur la gradation de la lumière *Comptes Rendus de Séances de La Société de Biologie*, *37*, 493–495.
- Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, *12*, 4745–4765.
- Burr, D. C., & Santoro, L. (2001). Temporal integration of optic flow, measured by contrast and coherence thresholds. *Vision Research*, *41*, 1891–1899.
- Busemeyer, J. R., Jessup, R. K., Johnson, J. G., & Townsend, J. T. (2006). Building bridges between neural models and complex decision making behaviour. *Neural Network*, *19*, 1047–1058.
- Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, *13*, 51–62.
- Chen, Y., Geisler, W. S., & Seidemann, E. (2008). Optimal temporal decoding of neural population responses in a reaction-time visual detection task. *Journal of Neurophysiology*, *99*, 1366–1379.
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S. . . . Shenoy, K. V. (2010). Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nature Neuroscience*, *13*, 369–378.
- Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *The Journal of Neuroscience*, *11*, 3612–3628.
- Ecker, A. S., Berens, P., Cotton, R. J., Subramanian, M., Denfield, G. H., Caldwell, C. R. . . . Tolias, A. S. (2014). State dependence of noise correlations in macaque primary visual cortex. *Neuron*, *82*, 235–248.
- Ecker, A. S., Denfield, G. H., Bethge, M., & Tolias, A. S. (2015). On the structure of neuronal population activity under fluctuations in attentional state. *The Journal of Neuroscience*, *36*, 1775–1789.
- Ganmor, E., Landy, M. S., & Simoncelli, E. P. (2015). Near-optimal integration of orientation information across saccades. *Journal of Vision*, *15*(16):8, 1–12, <https://doi.org/10.1167/15.16.8>. [PubMed] [Article]
- Geisler, W. S., & Albrecht, D. G. (1995). Bayesian analysis of identification performance in monkey visual cortex: Nonlinear mechanisms and stimulus certainty. *Vision Research*, *35*, 2723–2730.
- Gorea, A., & Tyler, C. W. (1986). New look at Bloch’s law for contrast. *Journal of the Optical Society of America A*, *3*, 52–61.
- Goris, R. L., Movshon, J. A., & Simoncelli, E. P. (2014). Partitioning neuronal variability. *Nature Neuroscience*, *17*, 858–865.
- Goris, R. L. T., Putzeys, T., Wagemans, J., & Wichmann, F. A. (2013). A neural population

- model for visual pattern detection. *Psychological Review*, *120*, 472–496.
- Goris, R. L. T., Ziemba, C. M., Stine, G. M., Simoncelli, E. P., & Movshon, J. A. (2017). Dissociation of choice formation and choice-correlated activity in macaque visual cortex. *The Journal of Neuroscience*, *37*, 5195–5203.
- Graf, A. B., Kohn, A., Jazayeri, M., & Movshon, J. A. (2011). Decoding the activity of neuronal populations in macaque primary visual cortex. *Nature Neuroscience*, *14*, 239–245.
- Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Grice, G. R. (1972). Application of a variable criterion model on auditory reaction time as a function of the type of catch trial. *Perception & Psychophysics*, *12*, 103–107.
- Kanitscheider, I., Coen-Cagli, R., & Pouget, A. (2015). Origin of information-limiting noise correlations. *Proceedings of the National Academy of Sciences, USA*, *112*, E6973–E6982.
- Kato, H. K., Chu, M. W., Isaacson, J. S., & Komiyama, T. (2012). Dynamic sensory representations in the olfactory bulb: Modulation by wakefulness and experience. *Neuron*, *76*, 962–975.
- Kiani, R., Hanks, T. D., & Shadlen, M. N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *The Journal of Neuroscience*, *28*, 3017–3029.
- Lin, I. C., Okun, M., Carandini, M., & Harris, K. D. (2015). The nature of shared cortical variability. *Neuron*, *87*, 644–656.
- Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, *77*, 24–42.
- Mazer, J. A., Vinje, W. E., McDermott, J., Schiller, P. H., & Gallant, J. L. (2002). Spatial frequency and orientation tuning dynamics in area V1. *Proceedings of the National Academy of Sciences, USA*, *99*, 1645–1650.
- McClelland, J. L. (1979). On the time relations of mental processes: An examination of systems of processes in cascade. *Psychological Review*, *86*, 287–330.
- Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., & Pouget, A. (2014). Information-limiting correlations. *Nature Neuroscience*, *17*, 1410–1417.
- Osborne, L. C., Bialek, W., & Lisberger, S. G. (2004). Time course of information about motion direction in visual area MT. *The Journal of Neuroscience*, *24*, 3210–3222.
- Rabinowitz, N. C., Goris, R. L. T., Cohen, M., & Simoncelli, E. P. (2015). Attention stabilizes the shared gain of V4 populations. *eLife*, *4*, e08998.
- Shadlen, M. N., Britten, K. H., Newsome, W. T., & Movshon, J. A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *The Journal of Neuroscience*, *16*, 1486–1510.
- Smith, M. A., Majaj, N. J., & Movshon, J. A. (2005). Dynamics of motion signaling by neurons in macaque area MT. *Nature Neuroscience*, *8*, 220–228.
- Tolhurst, D. J., Movshon, J. A., & Dean, A. F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research*, *23*, 775–785.
- Uka, T., & DeAngelis, G. C. (2003). Contribution of middle temporal area to coarse depth discrimination: Comparison of neuronal and psychophysical sensitivity. *The Journal of Neuroscience*, *23*, 3515–3530.
- Vogels, R., & Orban, G. A. (1990). How well do response changes of striate neurons signal differences in orientation: A study in the discriminating monkey. *The Journal of Neuroscience*, *10*, 3543–3558.
- Watson, A. B. (1979). Probability summation over time. *Vision Research*, *19*, 515–522.
- Yang, T., & Shadlen, M. N. (2007, June 28). Probabilistic reasoning by neurons. *Nature*, *447*, 1075–1080.
- Yarbus, A. L. (1967). *Eye movements and vision*. New York: Plenum Press.
- Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994, July 14). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, *370*, 140–143.
- Zylberberg, J., Camaro, J., Turner, M. H., Shear-Brown, E., & Rieke, F. (2016). Direction-selective circuits shape noise to ensure a precise population code. *Neuron*, *89*, 369–383.

Appendix

Under the modulated Poisson model, the growth of SNR^2 as a function of integration time may be

written (by substituting Equation 2 into Equation 1) as

$$\text{SNR}^2(T) = \frac{\left(\sum_{i=1}^T \mu_i\right)^2}{\sum_{i=1}^T \mu_i + \sum_{i=1}^T \mu_i^2 \sigma_g^2 + \sum_{i \neq j}^T \mu_i \mu_j \text{cov}[g_i, g_j]},$$

where T is the number of time bins being accumulated, μ_i is the mean stimulus-driven spike count in the i th time bin, and g_i is a stochastic gain signal for that bin, with mean 1.0 and variance σ_g^2 .

SNR^2 is an absolute measure of reliability. We also considered relative reliability of integrated activity for two models of neural response statistics (see Figure 4A). We define relative reliability as the ratio of the

reliability in the presence of rate fluctuations to the reliability with no such fluctuations. For the modulated Poisson model, in which rate fluctuations arise from a multiplicative source, relative reliability is

$$\text{SNR}_{\text{Rel}}^2(T) = \frac{1}{1 + T\mu\sigma_g^2}.$$

Under an alternative model in which rate fluctuations arise from an additive rather than a multiplicative source, relative reliability is

$$\text{SNR}_{\text{Rel}}^2(T) = \frac{1}{1 + \sigma_g^2/T\mu}.$$