

Temporal Precision in the Visual Pathway through the Interplay of Excitation and Stimulus-Driven Suppression

Daniel A. Butts,¹ Chong Weng,² Jianzhong Jin,² Jose-Manuel Alonso,² and Liam Paninski³

¹Department of Biology and Program in Neuroscience and Cognitive Science, University of Maryland, College Park, Maryland 20742, ²Department of Biological Sciences, State University of New York College of Optometry, New York, New York 10036, and ³Department of Statistics and Center for Theoretical Neuroscience, Columbia University, New York, New York 10027

Visual neurons can respond with extremely precise temporal patterning to visual stimuli that change on much slower time scales. Here, we investigate how the precise timing of cat thalamic spike trains—which can have timing as precise as 1 ms—is related to the stimulus, in the context of both artificial noise and natural visual stimuli. Using a nonlinear modeling framework applied to extracellular data, we demonstrate that the precise timing of thalamic spike trains can be explained by the interplay between an excitatory input and a delayed suppressive input that resembles inhibition, such that neuronal responses only occur in brief windows where excitation exceeds suppression. The resulting description of thalamic computation resembles earlier models of contrast adaptation, suggesting a more general role for mechanisms of contrast adaptation in visual processing. Thus, we describe a more complex computation underlying thalamic responses to artificial and natural stimuli that has implications for understanding how visual information is represented in the early stages of visual processing.

Introduction

In the context of a dynamically varying visual input, visual neurons can respond with action potentials timed precisely at millisecond resolution. Such temporal precision in the visual pathway has been observed in the retina (Berry and Meister, 1998; Passaglia and Troy, 2004; Uzzell and Chichilnisky, 2004), lateral geniculate nucleus (LGN) (Reinagel and Reid, 2000; Liu et al., 2001; Butts et al., 2007), and visual cortex (Buracas et al., 1998; Kumbhani et al., 2007). This precision is notable because in many cases the neuronal response has much finer time scales than the stimulus, which might be useful in reconstructing the stimulus in artificial and natural visual stimulus contexts (Butts et al., 2007). As a result, precision is often cited as evidence for a “temporal code” (Theunissen and Miller, 1995; Borst and Theunissen, 1999), which posits additional information about the stimulus represented in the fine temporal features of the spike train.

At the heart of elucidating the role of precision in the neural code is understanding how the timing of visual neuron responses is related to the stimulus. The relationship between the visual stimulus and the neuronal response can be captured, at a coarse level, by predictions based on a neuron’s receptive field (Chichilnisky, 2001; Simoncelli et al., 2004). While studies at finer temporal resolution reveal large discrepancies between predic-

tions of simple receptive field-based models and observed responses in retina and LGN, this is generally attributed to spike-generating machinery of the neuron (Berry and Meister, 1998; Keat et al., 2001; Paninski, 2004; Gaudry and Reinagel, 2007), in contrast to additional computation on the visual stimulus. While more complex processing is known to occur in the LGN (Wang et al., 2007, 2011; Alitto and Usrey, 2008; Sincich et al., 2009), particularly in the context of adaptation (Shapley and Victor, 1978; Mante et al., 2008), such additional processing has not been explicitly linked to the issue of temporal precision.

Here, we investigate the computation underlying the precise timing of LGN neurons in response to both artificial noise stimuli and natural movies by using a new statistical framework that can identify multiple stimulus-driven elements, in addition to spike refractoriness, which drive the neuron’s response. We find that a single receptive field coupled with spike refractoriness cannot explain the observed precision of LGN responses, and instead the precise timing of spikes is modeled parsimoniously as arising from the interplay of excitation and a delayed, stimulus-driven suppression. The coordination between this excitation and suppression allows the relatively slow time courses of stimulus-driven input to cancel each other, except in brief windows where excitation exceeds suppression, resulting in the observed fast response dynamics. While not previously explored in the retina and LGN, such an explanation for precision has been proposed in auditory (Wehr and Zador, 2003; Wu et al., 2008), somatosensory (Gabernet et al., 2005; Wilentz and Contreras, 2005; Okun and Lampl, 2008), and visual cortices (Cardin et al., 2007). Thus, in addition to its direct relevance to understanding processing in the LGN and downstream in the cortex, this work also presents a broader framework to characterize common elements of nonlinear computation in sensory pathways.

Received Dec. 4, 2010; revised June 9, 2011; accepted June 13, 2011.

Author contributions: D.A.B. designed research; D.A.B., C.W., J.J., and J.-M.A. performed research; L.P. contributed unpublished reagents/analytic tools; D.A.B. analyzed data; D.A.B. wrote the paper.

This work was supported by the National Eye Institute Grant EY005253 (J.-M.A.). We thank Stephen David, Mark Goldman, Nadja Schinkel-Bielefeld, and Timm Lochmann for valuable comments on this manuscript. We also thank Kamiar Rahnama Rad for contributing his code to estimate two-dimensional nonlinearities.

Correspondence should be addressed to Daniel A. Butts, Department of Biology, 1210 Biology-Psychology Building 144, University of Maryland, College Park, MD 20742. E-mail: dab@umd.edu.

DOI:10.1523/JNEUROSCI.0434-11.2011

Copyright © 2011 the authors 0270-6474/11/3111313-15\$15.00/0

Materials and Methods

Experimental recordings. All data were recorded from the LGNs of three anesthetized and paralyzed cats (either sex), and surgical and experimental procedures were performed in accordance with United States Department of Agriculture guidelines and approved by the Institutional Animal Care and Use Committee at the State University of New York, State College of Optometry (New York NY). As described in detail by Weng et al. (2005), cats were initially anesthetized with ketamine (10 mg/kg, i.m.) followed by thiopental sodium (20 mg/kg, i.v., supplemented as needed during surgery, and at a continuous rate of 1–2 mg/kg/h, i.v., during recording). A craniotomy and durotomy were made to introduce recording electrodes into the LGN (anterior, 5.5; lateral, 10.5). Animals were paralyzed with atracurium besylate (0.6–1 mg/kg/h, i.v.) to minimize eye movements, and artificially ventilated. LGN responses were recorded extracellularly within layer A; recorded voltage signals were conventionally amplified, filtered, and passed to a computer running the RASPOTIN software package (Plexon). For each cell, spike waveforms were identified initially during the experiment and verified carefully off-line by spike-sorting analysis. Cells were eliminated from this study if they did not have at least 2 Hz mean firing rates in response to all stimulus conditions, if the maximum amplitude of their spike-triggered average (STA) in spatio-temporal white noise was not at least five times greater than the amplitude outside of the receptive field area, or if their firing rate averaged over 1 min changed more than 10% in successive minutes.

Visual stimuli. All stimuli were displayed on a CRT display at a resolution of 0.2° per pixel with a monitor refresh rate of 120 Hz. There are two different types of stimuli used in this study: spatially uniform noise and natural movies. The spatially uniform noise consists of a temporal sequence of luminances displayed uniformly over the entire monitor, with the luminance at each time step selected from a Gaussian distribution with “zero” mean (corresponding to the midway point of the full range of monitor luminance) and an RMS contrast of 0.55, presented at 120 Hz. A two-minute sequence is used to estimate model parameters, and then 64 repeats of a 10 s unique stimulus sequence is used for cross-validation. The natural movie sequence was produced by members of the laboratory of Peter König (Institute of Neuroinformatics, ETH/UNI Zürich, Switzerland) using a removable lightweight CCD camera mounted to the head of a freely roaming cat in natural environments such as grassland and forest (Kayser et al., 2004). A 48 × 48 windowed area was processed to play at 60 Hz and have a constant mean and standard deviation of luminance for each frame (contrast held at 0.4 of maximum) as described by Lesica and Stanley (2006). The natural movie experiment consists of a single 15 min sequence, with the last 10 s of each 90 s interval set aside for cross-validation. Model parameters are determined using the other 800 s.

A different group of neurons was presented with the spatially uniform sequence ($N = 21$) versus the natural movie condition ($N = 20$). Care was taken to ensure that the potential effects of phase locking to the monitor refresh did not affect our results (Butts et al., 2007).

Implementation of the generalized nonlinear modeling framework. The generalized nonlinear modeling (GNM) framework is an extension of the generalized linear modeling (GLM) framework described by McCullagh and Nelder (1989), Brillinger (1992), Paninski (2004), Simoncelli et al. (2004), and Truccolo et al. (2005). Both models predict the instantaneous firing rate of the neuron $r(t)$ at a given time t given the stimulus $\mathbf{s}(t)$, the history of observed spike train $\mathbf{R}(t)$, and potentially other fixed functions of the stimulus such as $\mathbf{W}_{\text{sup}}(t)$ (Eq. 3, below):

$$r(t) = F[\mathbf{k} * \mathbf{s}(t) + \mathbf{h}_{\text{spk}} * \mathbf{R}(t) + \mathbf{h}_{\text{sup}} * \mathbf{W}_{\text{sup}}(t) - \theta]. \quad (1)$$

Here, the parameters of the model are all linear functions inside the spiking nonlinearity F : the linear [excitatory] receptive \mathbf{k} , the spike history term \mathbf{h}_{spk} , a “postsynaptic current” (PSC) term \mathbf{h}_{sup} associated with \mathbf{W}_{sup} , and threshold θ . Note that boldface is used to denote multidimensional parameters (i.e., vectors) such that, for example, $\mathbf{k} = k(\tau) = [k(0) \ k(1) \ k(2) \ \dots]$ and $*$ represents temporal convolution such that $\mathbf{k} * \mathbf{s}(t) = \int k(\tau) \mathbf{s}(t-\tau) d\tau = k[0] \mathbf{s}(t) + k[1] \mathbf{s}[t-1] + k[2] \mathbf{s}[t-2] + \dots$. The linear-nonlinear (LN) model consists of only the linear receptive field (Eq. 1, left) and threshold, the full “GLM” considered adds the spike history

term (middle), and the GNM encompasses all terms in Eq. 1. Schematics of these model structures are diagrammed in Figure 2A.

The spiking nonlinearity $F[\cdot]$ is a fixed function satisfying conditions for efficient optimization given in Paninski (2004): $F[g] = \log(1 + e^g)$. Thus, equation 1 describes a GLM for the parameters \mathbf{k} , \mathbf{h}_{spk} , \mathbf{h}_{sup} (given a fixed \mathbf{W}_{sup}), and its likelihood can be efficiently optimized as follows. The point-process log-likelihood per spike is given by (Paninski, 2004):

$$LL = \frac{1}{N_{\text{spk}}} \left[\sum_{t_s} \log_2 \Pr\{\text{spk}|t_s\} - \sum_t \Pr\{\text{spk}|t\} \right], \quad (2)$$

where $\{t_s\}$ are the observed spike times, and N_{spk} is the number of observed spikes. Parameter values corresponding to the linear terms can be simultaneously optimized by maximizing the likelihood of the model given the data. All model estimation uses gradient ascent of the log-likelihood (LL), which, given the spiking nonlinearity shown above, has no local maxima other than the single global maximum (Paninski, 2004). Also, for a given choice of model parameters, both the LL and its gradient can be directly calculated, resulting in efficiently computed ascents to the globally maximum likelihood.

The GNM framework incorporates optimization of the additional parameters that comprise \mathbf{W}_{sup} , which is an LN model with an associated receptive field \mathbf{k}_{sup} and internal nonlinearity $f_{\text{sup}}(\cdot)$:

$$\mathbf{W}_{\text{sup}}(t) = f_{\text{sup}}[\mathbf{k}_{\text{sup}} * \mathbf{s}(t)]. \quad (3)$$

For a given receptive field and nonlinearity, note that the output is a fixed function of the stimulus, and thus its associated PSC term \mathbf{h}_{sup} can be optimized in the context of the GLM (Eq. 1). The PSC term is constrained to be negative (using a linear constraint during the optimization), making the output, $\mathbf{h}_{\text{sup}} * \mathbf{W}_{\text{sup}}$ “suppressive.” Omitting this negative constraint does not improve model performance and makes for a larger search space with more local maxima that can anchor the full parameter search (described below) to nonoptimal solutions.

The suppressive nonlinearity $f_{\text{sup}}(\cdot)$ can also be efficiently optimized by expressing it as a linear combination of basis functions $f_{\text{sup}}(x) = \sum_n \xi_n[x]$, which are chosen to be overlapping tent-basis functions (Ahrens et al., 2008). Because the coefficients α_n operate linearly on the processed stimulus $\xi_n[\mathbf{k}_{\text{sup}} * \mathbf{s}(t)]$, they—and as a result a piecewise linear approximation of $f_{\text{sup}}(\cdot)$ —can also be optimized in the context of the GLM.

Thus, for a given choice of the suppressive receptive field \mathbf{k}_{sup} , all other terms of the GN model have a single globally optimal solution that can be found efficiently in the GLM framework. [One caveat is that the PSC term \mathbf{h}_{sup} and nonlinearity coefficients $\{\alpha_n\}$ comprise a bilinear optimization problem requiring alternating optimization steps (Ahrens et al., 2008), which in practice converges to a global optimum.] Thus, the difficult task of optimizing the large number of variables associated with the nonlinear terms of the GNM is reduced to determining the best suppressive receptive field \mathbf{k}_{sup} . To limit the number of parameters used to represent \mathbf{k}_{sup} , we use a family of basis functions suggested by Keat et al. (2001):

$$\zeta(t) = \sin\{\pi n[2t/t_F - (t/t_F)^2]\}, \quad (4)$$

which are orthonormalized using the Gram-Schmidt method, with $t_F = 133$ ms for the spatially uniform noise stimulus and $t_F = 217$ ms for the natural movies. This basis set behaves like Fourier decomposition, but with the time axis scaled to have more temporal detail at short latencies. With these basis functions, the observed temporal filters can be represented with relatively few terms: ten is sufficient for the spatially uniform context, and eight sufficient for the natural movies.

With the addition of model parameters inside nonlinear terms (the suppressive receptive field \mathbf{k}_{sup}), the relationship between model parameters and the log-likelihood becomes more complex and no longer has a single global maximum. As a result, the choice of the initial guess of \mathbf{k}_{sup} is quite important and we have tried several different strategies, arriving at the following. First, parameters for the GLM (including the linear receptive field and spike history term) are estimated, and the resulting linear receptive field is then used as an initial guess for the suppressive receptive field \mathbf{k}_{sup} of the GNM. Optimization of \mathbf{k}_{sup} then proceeds by

determining the LL of each choice of \mathbf{k}_{sup} through GLM optimization of the PSC term and internal nonlinearities, and using this as a basis for numerical gradient ascent. We found that it is not necessary to optimize \mathbf{h}_{sup} and $f_{\text{sup}}(\cdot)$ for each choice of \mathbf{k}_{sup} , but rather it is more computationally efficient to alternate between the optimization of \mathbf{k}_{sup} holding \mathbf{h}_{sup} and $f_{\text{sup}}(\cdot)$ fixed, and optimization of \mathbf{h}_{sup} and $f_{\text{sup}}(\cdot)$ holding \mathbf{k}_{sup} fixed. We also found that this optimization is aided by constraining the internal nonlinearities to be monotonic via linear constraints, which prevented the model from getting trapped in some local maxima. An example of the results of optimization of the suppressive term from the initial guess (STA) is shown in Figure 2E–G.

Due to the estimation of model parameters at time resolutions finer than the frame rate, regularization is used to prevent overfitting. A “smoothness” penalty term is added to the LL proportional to the squared slope of the temporal kernels, e.g., $-\sum(t) \lambda_k [k(t) - k(t + \Delta t)]^2$, with a different λ chosen for each temporal function to yield smoothly varying kernels. Note that adding penalty terms of this form to the LL still preserves the efficient properties of the GLM optimization.

After the optimal parameters are obtained, the spiking nonlinearity $F[\cdot]$ is re-estimated using the “histogram method” (Chichilnisky, 2001) by directly measuring the probability of a spike $\text{Pr}\{\text{spk}|g\}$ for each value of the “internal” model output g (i.e., the terms in Eq. 1 inside the spiking nonlinearity). The resulting estimate of $\text{Pr}\{\text{spk}|g\}$ is then fit to the nonlinearity $\text{Pr}\{\text{spk}|g\} = F[g] = \alpha/\beta \log[1 + \exp(\beta(g - \theta))]$, where α is the overall slope of the nonlinearity, θ is the threshold, and β determines how sharp the transition from zero firing rate is around the threshold. Empirically, the resulting estimate of $F[g]$ described the measured spiking nonlinearity $\text{Pr}\{\text{spk}|g\}$ accurately and resulted in better cross-validated performance on average. Note that no saturating function is necessary.

The models applied to the spatially uniform stimulus condition were fit at a time resolution of 1/16 of the frame rate (0.5 ms), and the models in the natural movie condition were fit at a time resolution of 1/4 of the frame rate (4 ms resolution). All parameter estimation was performed using MATLAB (MathWorks).

Spatiotemporal implementation of the GNM. The linear spatiotemporal receptive fields (STRFs) of retinal ganglion cells and LGN neurons are well approximated by separate center and surround components, each individually space–time separable (Dawis et al., 1984; Cai et al., 1997; Allen and Freeman, 2006). The spatiotemporal GLM firing rate can then be expressed as

$$r(t) = F[\mathbf{k}_C * \mathbf{s}_C(t) + \mathbf{k}_S * \mathbf{s}_S(t) - \theta], \quad (5)$$

where $\mathbf{s}_C(t)$ and $\mathbf{s}_S(t)$ represent the spatiotemporal stimulus filtered by the “center” and “surround” spatial kernels, and \mathbf{k}_C and \mathbf{k}_S are the corresponding temporal kernels. Both spatial kernels are circular two-dimensional Gaussian distributions with the same center (x_0, y_0) and different widths (standard deviations), σ_C and σ_S . Because the temporal kernels and resulting model likelihood are uniquely determined by global optimization in the maximum likelihood estimation framework, the spatial parameters can then be optimized by searching in the four-dimensional space $[x_0, y_0, \sigma_C, \sigma_S]$, starting with initial values that describe the location and widths of the peak in the spike-triggered average (Butts et al., 2010). This progresses in a similar way to GNM optimization of the suppressive receptive field described above, because each combination of four spatial parameters can be uniquely associated with an optimal choice of \mathbf{k}_C and \mathbf{k}_S and an associated LL .

The spatiotemporal GN model is optimized for the natural movie condition in much the same way as in the spatially uniform condition, with two main differences described below. As with the purely temporal context, the GLM STRF serves as an initial guess for the suppressive RF, and then the temporal parameters of the suppressive fields are optimized through brute force simultaneously with those of the linear “excitatory” receptive field. The first difference in the spatiotemporal case is the estimation of the additional spatial parameters, which start at values determined by the GLM. Following the first temporal optimization, the spatial parameters for linear excitatory term and suppressive term optimized simultaneously while holding the temporal parameters fixed. These temporal and spatial optimizations are alternated until there is no further improvement in the likelihood.

The second difference between the spatiotemporal model and purely temporal model is that we found that the performance of the GNM could be further improved by making the excitatory term “nonlinear” (i.e., with the same functional form as the suppressive term; Eq. 3) with excitatory receptive field \mathbf{k}_{exc} , associated nonlinearity $f_{\text{exc}}(\cdot)$, and PSC term \mathbf{h}_{exc} . After first optimizing \mathbf{k}_{exc} as a linear receptive field (described above), \mathbf{k}_{exc} was inserted as an initial guess within a nonlinear excitatory term through alternating optimizations of \mathbf{k}_{exc} and \mathbf{k}_{sup} . The optimal values of \mathbf{k}_{exc} and \mathbf{k}_{sup} are generally very close to this initial condition. Thus, while this additional step improved model performance, it did not qualitatively change the other model parameters, other than a rectified excitatory nonlinearity that is similar to the suppressive nonlinearity (data not shown).

More complex nonlinear spatiotemporal models were also tested on this dataset, and while additional complexities could yield even better model performance, the resulting fits were entirely consistent with the relationship between excitation and suppression described here (see Fig. 7). As a result, a more detailed exploration of spatiotemporal processing is outside the scope of this paper.

The relationship between log-likelihood and information in the neuron’s spike train. The log-likelihood (LL) of the data given the model (Eq. 2) is directly comparable to the information that single spikes signal about the stimulus, defined as the single spike information (Brenner et al., 2000): $I_{ss} = \sum_t r(t) \log [r(t)/\bar{r}]$, where \bar{r} is the average firing rate. For example, the “null” model, whose firing rate is not modulated by the stimulus and has a constant firing rate of $r(t) = \bar{r}$, would have $I_{ss} = 0$ and its log-likelihood (Eq. 2) given by $LL_0 = \bar{r} \log \bar{r} - 1$. Conversely, a model that perfectly predicted the firing rate would have its $LL = \sum_t r(t) \log r(t) = I_{ss} + LL_0$. Because LL_0 is just a constant term that depends on the number of observed spikes, maximizing the likelihood is mathematically equivalent to maximizing the information between the stimulus and the observed spike train (Kouh and Sharpee, 2009). To make this comparison explicit, we use log-base-2 to express the units in bits, and report the cross-validated likelihood measures with the constant LL_0 subtracted (see below).

Model cross-validation. Although we applied common prediction performance metrics, including “percentage explained variance” or R -squared (see Fig. 4A, B), the main measure of model performance used in this study is the cross-validated log-likelihood, LL_x , which is given by the log-likelihood of the model (Eq. 2) applied to data that was not used to fit the model parameters and with the log-likelihood of the “null model” LL_0 subtracted (see above, The relationship between log-likelihood and information in the neuron’s spike train). This calibrates the LL_x such that the null model (which just predicts the average firing rate with no stimulus modulation) has $LL_x = 0$, and a perfect model has $LL_x = I_{ss}$.

Using the LL_x as measure of model performance, instead of peristimulus time histogram (PSTH)-based measures like R -squared, is particularly important because the LL_x could be measured using single trials. For the spatially uniform stimulus condition, this allows the state of the neuron from trial-to-trial to be tracked. The nonstationarity that is detected as a result can be attributed to changes in the gain and offset terms (α and Δ , respectively) applied to the parameters of the GNM:

$$r(t) = F\{\alpha_E[\mathbf{k} * \mathbf{s}(t)] + \alpha_S[\mathbf{h}_{\text{sup}} * \mathbf{W}_{\text{sup}}(t)] + \alpha_{\text{RP}}[\mathbf{h}_{\text{spk}} * \mathbf{R}_{\text{spk}}(t)] - (\theta - \Delta)\}. \quad (6)$$

To track these gains and additional offset over the repeated stimulus presentations, the 10 s repeated stimulus is divided into two 5 s segments; the first is used to fit these additional parameters using maximum likelihood estimation (while holding the standard GNM parameters constant), and the second 5 s segment is used for cross-validation. Regularization on the gains and offsets across successive trials is performed with an additional term added to the log-likelihood that penalized changes of the values of each gain and offset from one trial to the next. As with the smoothness penalty described above, such an additional penalty term preserves the efficiency of GLM optimization applied here (Paninski et al., 2010).

To facilitate direct comparisons across models, all reported LL_x values in this condition are measured on the same 5 s segment and are based on

optimally adjusted gains and offsets for each model. For the LN and GL models with fewer terms in Equation 6, correspondingly fewer gains could be adjusted, and furthermore such adjustments have little effect on the models' cross-validated performance. In the meantime, the spike-triggered covariance (STC) model could not be adjusted on a trial-by-trial in a straightforward manner, because its spiking nonlinearity is not parameterized and requires a large amount of data to estimate (Rahnama Rad and Paninski, 2010).

In the natural movie condition, 100 s of the 900 s stimulus sequence is set aside for cross-validation, and not used to determine the model parameters. The LL_x values reported here are thus calculated based on the likelihood of each model for this 100 s stimulus sequence. Adjusting the gains and offsets in this condition does little to improve the model performance, most likely because of the much larger dimensionality of the model parameter space.

The effective GN filters. Because the stimulus-driven suppression of the GN model includes two temporal convolutions, for purposes of display and comparison, we usually show an effective filter that represents the combined output of both stages. The effective filter is the cross-correlation between a Gaussian white noise stimulus (whose mean and variance match that of the true stimulus) and the output of each module driven by this noise, resulting in something akin to a spike-triggered average each for nonlinear element (Eq. 3). These effective filters are pictured in Figure 3E, throughout Figure 6, and in Figure 7B.

Spike-triggered covariance analysis. We analyzed all neurons recorded in the spatially uniform stimulus condition with STC analysis, using standard practices (Schwartz et al., 2006). Briefly, we calculate the average covariance matrix of the stimulus history preceding each spike, projecting out the spike-triggered average (STA) and then subtract off the covariance matrix of the stimulus not triggered on spikes. The eigenvectors of this matrix and the STA represent principal directions in stimulus space that best describe the distribution of spikes. For this data, they are always suppressive (representing decreased variance), and further analysis shows that only the first suppressive STC filter consistently contributes to better cross-validated performance.

To understand the relationship between the "STC filters" (i.e., the STA and first suppressive STC eigenvector) and those found by the GNM, we measure the angles between the STC filters and the effective GN filters by normalizing each filter to have unit magnitude and then taking the arccosine of their dot product. This procedure also is used to project the GNM filters into the space spanned by the STC filters (Fig. 6C). The nonlinearities associated with each filter of the STC model (Fig. 6E) are calculated in the context of the GN model and are consistent with the full two-dimensional nonlinearity associated with a full STC-based LN model.

Models based on STC analysis are LN cascades with a spiking nonlinearity $F[g_1, g_2, \dots]$ that depends on the combined output of each linear filter, where $g_i(t) = \mathbf{s}(t) * \mathbf{k}_i$. The STC spiking nonlinearity can in principle be determined using the histogram method, just like in the one-dimensional LN model (Chichilnisky, 2001). In this case, the number of spikes in each bin (g_1, g_2) is normalized by the number of times each combination occurred, resulting in the probability of a spike given the pair $\Pr\{\text{spk} | (g_1, g_2)\}$. However, the higher dimensionality of the nonlinearity stretches the number of observed spikes over many more bins, making accurate estimation methods more complicated. As a result, recently developed techniques are used for nonparametric estimation of the two-dimensional nonlinearity (Rahnama Rad and Paninski, 2010), which produces a firing rate prediction.

For several neurons, maximally informative dimension (MID) analysis (Sharpee et al., 2004) was also applied using an MID software package available online from Tatyana O. Sharpee (Salk Institute for Biological Studies, La Jolla, CA), and its results are to be consistent with the STC analysis, as expected (Paninski, 2003; Sharpee et al., 2004).

Measure of neuronal precision. To measure the amount of precision (see Fig. 4D), we use the response time scale as previously defined and implemented in earlier studies (Butts et al., 2007, 2010; Desbordes et al., 2008). This measure is calculated from the autocorrelation of the spike train pooled across multiple repeated trials. A Gaussian distribution is fit to the central peak of the autocorrelation function, and its standard

deviation σ is measured. The response time scale, corresponding to the average width of features in the PSTH of the neuron, is given by $\tau_R = \sigma\sqrt{2}$ (Butts et al., 2007).

Results

We recorded extracellularly from LGN neurons in the context of spatially uniform noise visual stimuli (Fig. 1A), a context in which spike train features are reliable to millisecond precision (Reinagel and Reid, 2000; Liu et al., 2001; Kumbhani et al., 2007). To understand the computation on the visual stimulus underlying the timing of LGN responses, we first used the familiar LN cascade model (Chichilnisky, 2001; Simoncelli et al., 2004), which is based on the neuron's linear receptive field and is typically estimated from the spike-triggered average (Chichilnisky, 2001). In the context of spatially uniform noise, the receptive field is purely temporal (because there is no spatial information in the stimulus), and the receptive field of the example LGN neuron (Fig. 1B) demonstrates both its selectivity to dark-to-light transitions ("ON cell") and the average latency between ON transitions and the neuronal response (Fig. 1A, cyan).

The LN model posits that the neuron's firing rate simply reflects the degree that the stimulus $s(t)$ matches the receptive field \mathbf{k} . The firing rate prediction of the LN model is generated by making this comparison explicit with a linear convolution, resulting in the filtered stimulus $g(t)$ (Fig. 1C, top), which is then processed by the measured spiking nonlinearity (Fig. 2D, red) that maps $g(t)$ to a firing rate (Fig. 1C, bottom).

While the LN model prediction captures the coarse timing and overall magnitude (i.e., area under firing rate curve) of the observed firing rate, it clearly does not capture the fine temporal features of the response: neither the duration of individual events, nor their precise timing (Fig. 1C, bottom). Note that this cannot be repaired by artificially constraining the spiking nonlinearity to be steeper (or even incorporating a deterministic threshold), because while this can result in more precise predicted responses, the predicted firing rate will necessarily have the wrong overall timing. As a result, the output of a model that processes the stimulus with its linear receptive field alone is anchored by the relatively long integration time of the receptive field itself (Butts et al., 2007), and the width of each predicted event (Fig. 1C, bottom, red) reflects the amount of time that a stimulus matching the neuron's receptive field is present. Thus, the observed precision of the LGN response is evidence for more complex nonlinear processing taking place.

Stimulus-driven suppression underlies precision

Although the LN model represents a tremendously simplified description of the processing performed by the underlying neuronal circuitry of the retina and LGN, its success at describing the coarse features of the LGN response suggests that LGN computation on the visual input might still be captured by augmenting the LN model. While nearly all observed firing events are within the LN prediction (Fig. 1C), there is no observed response for a large part of the time that a preferred stimulus is present. This implies that the short duration of the spiking events might stem from an additional processing element that suppresses neuronal firing at certain times when the neurons would otherwise fire.

Because suppression occurs over the time that the stimulus matches the neuron's linear receptive field, this implies that suppression has tuning similar to that of the neuron. Ideally, the stimulus tuning of this suppression would be directly measurable, but such a measurement using only extracellular data is confounded by two problems. First, while spikes unambiguously

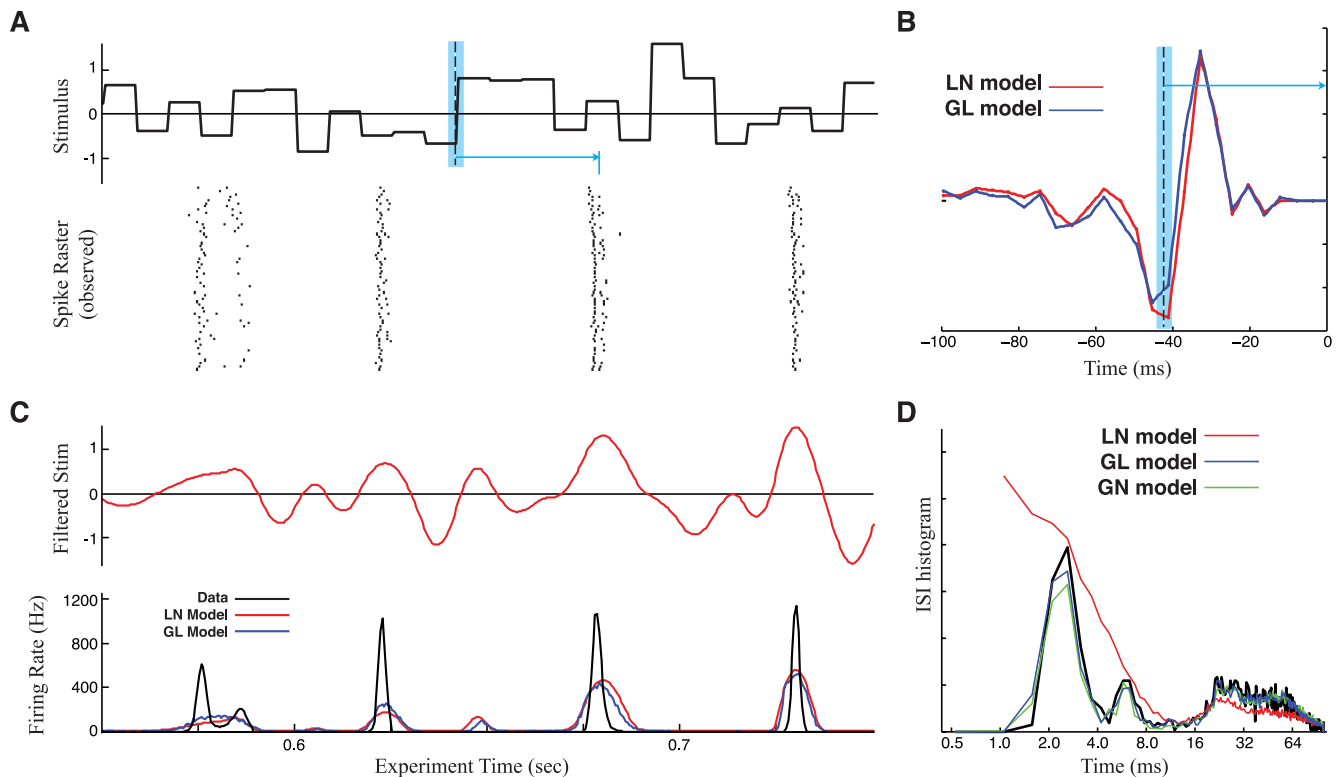


Figure 1. The failure of current models in capturing observed temporal precision. **A**, Spike raster of an LGN Y-cell response to 60 repeated presentations of a full-field temporal noise stimulus (top). **B**, The optimal linear filter describing the temporal processing of the stimulus for the LN model (red) and GL model (blue). This demonstrates the selectivity of the neuron to a dark-to-light transition and the average response latency (cyan), also shown in **A**. **C**, Top, The stimulus filtered by the LN receptive field (Filtered Stim), producing a linear prediction of the neuron's response that is highest when the neuron fires. Bottom, The firing rate predictions of the LN and GL models compared with the PSTH of the neuron's response (black). Both the GL and LN model reflect the time scales of the underlying filtered stimulus, which does not reproduce the precision in the response. **D**, The interspike interval (ISI) distributions predicted by the models compared with that observed (black), demonstrating that, in incorporating a spike history term, the GL model (blue) is able to reproduce this distribution but this is not sufficient to predict the observed precision. The spike history term of the GN model (green, see Fig. 2) is also shown.

signal the presence of some form of underlying excitation, a lack of neuronal response at a given time might be due to suppression, but also can simply be due to a lack of excitation. Thus, a suppressive receptive field must be measured simultaneously with the excitatory receptive field, because the presence of one will affect the estimate of the other. Second, receptive field estimation techniques generally rely on linear modeling, which is limited to identifying a single stimulus-processing element. This is due to the associative property of linear operations, whereby linear processing by two receptive fields (\mathbf{k}_1 and \mathbf{k}_2) is equivalent to processing by a single equivalent receptive field \mathbf{k}' : $[\mathbf{k}_1 * \mathbf{s}(t)] + [\mathbf{k}_2 * \mathbf{s}(t)] = (\mathbf{k}_1 + \mathbf{k}_2) * \mathbf{s}(t) = (\mathbf{k}') * \mathbf{s}(t)$, where the asterisk (*) represents the linear convolution. Thus, a model with more than one element of stimulus tuning must involve embedded nonlinear transforms that act differently on each processing element.

One previously explored possibility is that—because suppression has tuning similar to that of the neuron—it could simply arise from the neuron's own “refractoriness” following spikes (Berry and Meister, 1998). Such a spike response model (Gerstner, 2001) can be efficiently estimated in the context of the generalized linear modeling or GLM framework (Paninski, 2004; Truccolo et al., 2005) because the spike train that results in suppression (i.e., the neuron's own) is known, and thus the linear receptive field and spike refractoriness (or “spike history”) term can be estimated simultaneously (Fig. 2A, middle) (see Materials and Methods). For the neuron in Figure 1, the GL model has a similar receptive field as the LN model (Fig. 2B, blue) and a long suppressive spike history term (Fig. 2C, blue), which is consistent

with previous estimates of refractoriness in this context (Berry and Meister, 1998; Keat et al., 2001; Pillow et al., 2005). The spike history term has two distinguishable components: an “absolute refractory period” at short time scales (up to 1–2 ms) strong enough to veto any spike regardless of the stimulus, and a second “relative refractory period” lasting >15 ms (Berry and Meister, 1998). By including the spike history term, the GL model is able to reproduce the observed interspike distribution, unlike the LN model (Fig. 1D).

However, the GL model is unable to predict the precision of the LGN spike train, and this results in a firing rate prediction with a very similar time course to that of the LN model (Fig. 3A). By comparison, the GLM is successful in capturing the precision of simulated data where precision is generated by spike refractoriness (with a variety of spike-generating mechanisms; data not shown), suggesting that inability of the GLM to explain LGN precision is not simply because of not correctly modeling spike generation. Although we cannot eliminate arbitrary forms of spike refractoriness as an explanation for precision in the LGN, the failure of the GLM in this context suggests a simpler alternative: that its long relative refractory period (Fig. 2C) results from this suppression being correlated with—rather than caused by—the LGN spikes.

Therefore, we developed a new modeling approach to investigate a more general case where suppression is not completely explained by feedback from the neuron's own spike train. In such a situation, the suppression can be thought to have its own stimulus tuning, represented by a “suppressive receptive field” \mathbf{k}_{sup} .

As mentioned above, a second receptive field alone would not be distinguishable from the original excitatory receptive field without some associated nonlinearity. Furthermore, for the output of filtering by the suppressive receptive field to be only suppressive (i.e., negative), it must be [nonlinearly] processed such that it only has a negative output. Thus, we model this suppression with a three stage cascade: a linear receptive field k_{sup} , followed by a static nonlinear function $f_{sup}(\cdot)$, and then by a negatively constrained temporal filter h_{sup} that determines how it influences the neuron's output (Fig. 2A, right). This LNL cascade structure (Korenberg et al., 1989) gives the model flexibility to capture the effects of an inhibitory input, which could be thought of as an LN model (k_{sup} and f_{sup}) filtered by an "inhibitory" (negative) PSC h_{sup} —but the results shown below do not depend on the accuracy of this interpretation. In either case, this model can be thought of as a next order expansion from the LN model, similar in spirit to other nonlinear approaches such as spike-triggered covariance analysis (see below). However, rather than the next order term being based on a mathematical expansion, here the next order term is specifically designed to resemble the type of suppression caused by an inhibitory synaptic input.

To determine the parameters of this model, we developed a generalized nonlinear modeling or GNM framework based on maximum likelihood estimation, which leverages many of the efficiencies of past GLM approaches (Brillinger, 1992; Paninski, 2004; Paninski et al., 2007; Pillow et al., 2008). This framework allows for the simultaneous estimation of the linear excitatory receptive field (Fig. 2B), spike history term (Fig. 2C), and components of the suppressive nonlinear term (Fig. 2E–G). When applied to the same data as the previously considered models, both the linear excitatory and suppressive receptive fields are different from those of the linear models (Fig. 2B,E). Instead, the suppressive receptive field ends up resembling the linear excitatory term of the GNM (Fig. 2, compare B, E, green), which is also similarly different from those of the linear models. The suppressive output, however, is delayed by the PSC term h_{sup} (Fig. 2F), which has nearly the same time course as the relative refractory period of the previously measured spike history term (Fig. 2C). With the suppressive term included in the model, the spike history term now almost completely lacks suppression for intervals greater than 5 ms. Because the spike history and PSC are fit simultaneously in the GNM (see Materials and Methods), the attenuation of the spike history term and corresponding appearance of the suppressive PSC reflects the best description of the data, rather than any bias in the modeling framework.

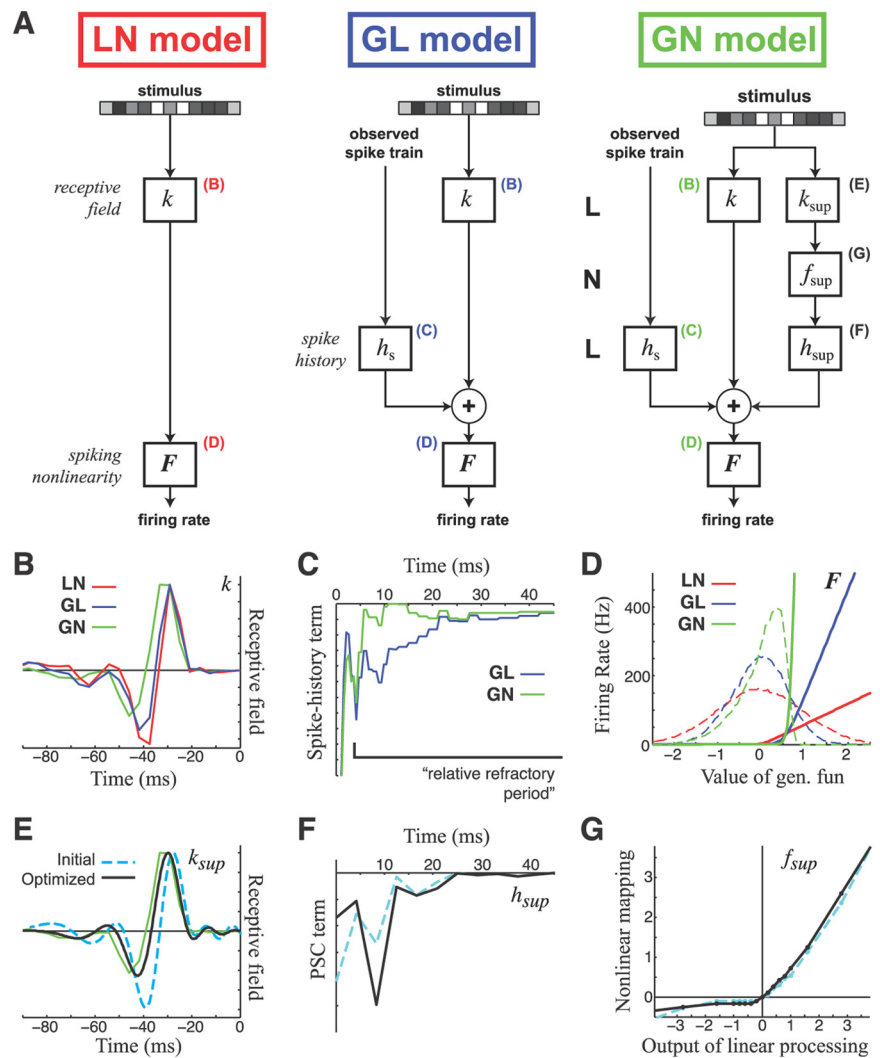


Figure 2. Augmentation of the LN model to explain suppression underlying precision. **A**, Model schematics demonstrating the evolution from the linear-nonlinear or LN model (left), to the generalized linear or GL model with the addition of a spike history term (center), to the generalized nonlinear or GN model with the addition of a suppressive LNL cascade (right). **B–G**, The model components for the neuron considered in Figure 1. **B**, The linear [excitatory] receptive field for all three models; note the slower time course of the GN excitation (green). **C**, The spike history term of the GL model (blue), demonstrating an absolute and relative refractory period. The relative refractory period disappears when the suppressive term is added in the GNM (green). **D**, The spiking nonlinearities of the three models (solid lines) compared with the distributions of the generating functions (gen. fun) (dashed lines, i.e., sum of terms of model inside the spiking nonlinearity). The spiking nonlinearity of the GNM (green) is nearly vertical, demonstrating that it predicts a high firing rate for a small fraction of the stimuli, compared with the spiking nonlinearities from the other models. **E–G**, Components of the suppressive term demonstrating the effect of optimization from the initial guess (cyan) to the optimized answers (black). **E**, The suppressive receptive field (black) has tuning similar to that of the excitatory tuning (green) reproduced from **B**. **F**, The PSC term delays the output of the suppressive term. **G**, The internal nonlinearity is rectifying, as expected for a purely suppressive contribution.

Because the GN excitatory receptive field is notably slower and less biphasic than those of the LN and GL models (Fig. 2B), this suggests that the standard stimulus tuning of the neuron represented by the linear receptive field results from a combination of a slower, less biphasic excitatory tuning and the delayed suppression. Thus, this analysis reveals a different view of the stimulus selectivity of the neuron.

Precision is explained by the interplay of excitatory and suppressive receptive fields

The GN model is able to capture the precise timing of the LGN neuron response (Fig. 3A): both the duration of individual firing events (Fig. 3B) and their precise onset times (Fig. 3C). The im-

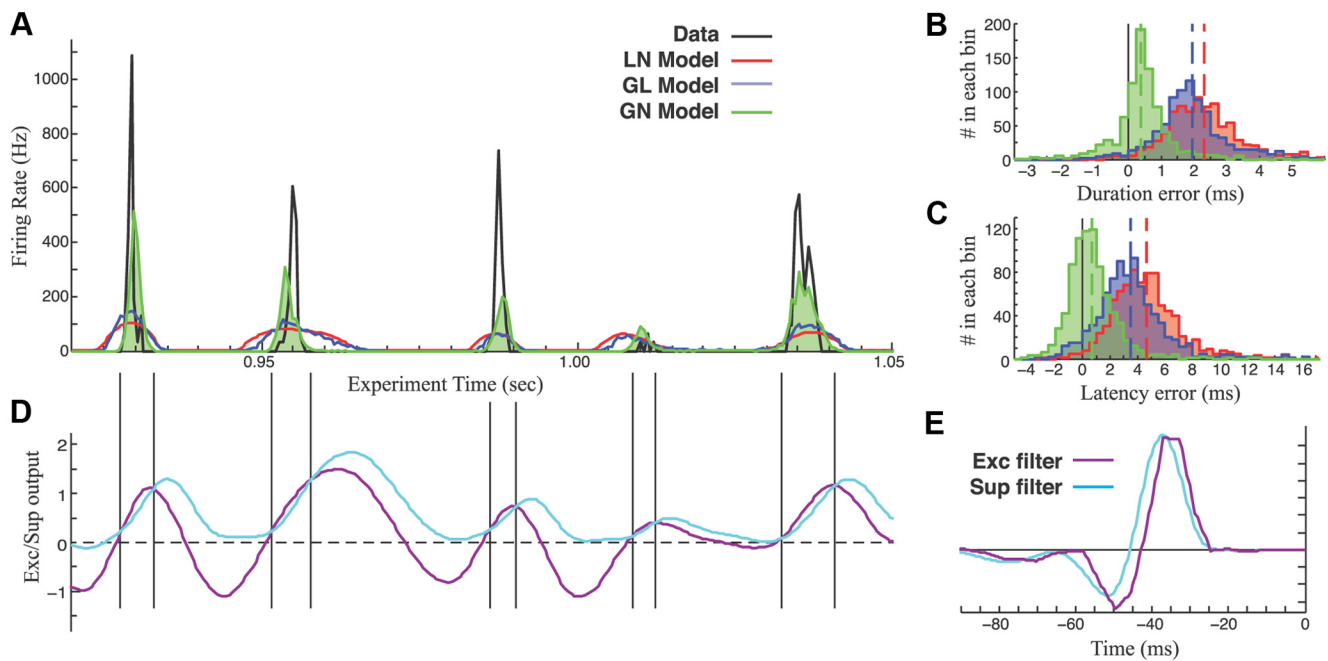


Figure 3. Fine temporal processing in the LGN is explained through the interplay of excitation and suppression. **A**, The predictions of the GNM (green) on a cross-validation stimulus sequence closely match the observed firing rate (black) compared with the LN (red) and GL (blue) predictions. **B**, The ability of each model to predict the duration of each firing event on an event-by-event basis, showing that the GN model successfully predicts more than just the averages (dashed lines). Firing events are simply parsed using >5 ms gaps in the firing rate, and duration is defined at $2 \times$ the standard deviation of spike times in each event. The distributions represent histograms over the duration across a 10 s cross-validation stimulus sequence. **C**, The GN model also captures the precise timing of the onset or latency of each event compared with the other models. **D**, The precise firing of the GN model occurs when the output of the linear [excitatory term] (purple) is greater than the output of the suppressive term (cyan), shown over the same stimulus sequence as the firing rates (**A**). Thus, although temporal processing of each term is relatively slow, the neuron only fires for a fraction of the time (vertical black bars) that a stimulus that matches the excitatory receptive field is present. **E**, The effects of suppression (Sup) are delayed relative to excitation (Exc) as shown by comparing the excitatory receptive field (purple) to the effective suppressive receptive field, calculated from the cross-correlations between the stimulus and the suppressive term output. Suppressing timing arises through the combined effects of the suppressive receptive field (Fig. 2E) and the PSC term (Fig. 2F).

proved performance of the GN model is also reflected in conventional measures of model performance, such as the fraction of explained variance, or R-squared (R^2) (Fig. 4A,B) and the cross-validated log-likelihood (Fig. 4C,D). For the neuron shown, the GNM explains more than double the variance of the data compared with the LN and GL models (Fig. 4B). This improved cross-validated performance of the model is particularly significant given the increased number of parameters of the GN model, which implicitly lead to worse cross-validation unless they are capturing real elements of stimulus processing.

The precise timing of the model is generated by the interplay of the two receptive fields (Fig. 3D), with the output of the excitatory receptive field preceding that of the suppressive output. To picture how the precise firing of the model arises from the interplay of excitation and suppression, the output of the suppressive element is inverted and shifted by the spiking threshold, such that a nonzero firing rate occurs when the amount of excitation exceeds the amount of suppression. This comparison demonstrates that the neuron's response occurs in brief temporal windows where excitation exceeds suppression, allowing for precise response timing to emerge from relatively slow integration and processing of the stimulus (i.e., the vertical lines demarcate only a fraction of the features of the more slowly evolving excitatory and suppressive terms). This delay is evident comparing the excitatory filter with the effective suppressive filter (Fig. 3E), calculated by the cross-correlation between the stimulus and output of the suppressive term.

By comparison, the time courses of the LN and GL model predictions are anchored to the stimulus filtering of their single linear receptive field, which is fixed by the frequency content of

the stimulus combined with the filtering properties of the temporal receptive fields (Butts et al., 2007). The resulting time course of this linear stimulus filtering is comparable to that of the excitatory GNM filter alone (Fig. 3D, magenta), although it is shifted earlier in time such that the output peaks with the data, resulting in a latency difference in the filters themselves (Fig. 2B). While one could construct a more precise model by using a linear filter that was more high-pass, the resulting output of this filter would not describe the data, which constrains such a linear model to these slower time scales. From this perspective, the presence of the suppression in the GNM allows for the observed precision with the appropriate stimulus selectivity.

A more subtle but equally important element of the GNM processing is suppressive nonlinearity (Fig. 2G), without which these two filters would linearly sum into a single linear filter (Fig. 2B, red). Without the nonlinear rectification of the suppressive term—the effects of which can be seen in Figure 3D (cyan)—the suppressive term would result in excitation by the opposite stimulus type, resulting in spurious excitations that are not in the data.

The interplay of excitation and suppression as an explanation for precision is fundamentally different than the source being the dynamics of spike generation and results in a different estimate of the underlying tuning to stimuli of the neuron (Fig. 2B). The model's ability to predict the onset and offset times of individual firing events likely stems from the individually tuned excitatory and suppressive receptive fields and implies the existence of a richer computation that can be accomplished in the LGN than if its responses were governed by a single receptive field—in particular for spatiotemporal stimuli (see below). Supporting this role

in more general computation, the improved performance of the GN model extends to coarse time resolutions where precision is not a factor (Fig. 4A), with the GNM explaining over 80% of the variance compared with 60% of the LN model.

To demonstrate the performance gain across the population, we use the cross-validated log-likelihood LL_x , which is much more sensitive than R^2 and can be measured using single trials to reveal much more detail about the state of the neuron over time (see below). The LL_x is in units of bits per spike and is measured relative to the null model, which assumes a mean firing rate with no stimulus-tuned elements, making the LL_x comparable to the single-spike information (Brenner et al., 2000; Kouh and Sharpee, 2009) (see Materials and Methods). For the neuron considered above, the GN model gives a 94% improvement in the LL_x , with $LL_x(\text{LN}) = 1.64$ bits/spk and $LL_x(\text{GN}) = 3.18$ bits/spk. Across the neurons in this study, the GN model yields an improvement over both the LN and GL models in every case (Fig. 4C), with mean 47% improvement compared to a 15% improvement of the GL model. These results are summarized in a box plot (Fig. 4D) that compares LL_x of the GN model to all other models considered in this study, including a model based on spike-triggered covariance (described below). The GN model provides a better explanation for the data for every neuron in the study, and in many cases the improved performance represents a significant proportion of the information in the spike train.

Rather than being uniform across all neurons, the improvement gained from the GN model is highest for the most precise neurons (Fig. 4E). Neurons with the strongest suppression have a narrower window over which excitation exceeds suppression, resulting in greater precision. In fact, the example neuron considered above is the most precise neuron in the study and correspondingly had the largest improvement in the GN framework (Fig. 4E, arrow). Choosing it as an example highlights that this approach is most successful with the neurons that are the least well described by the LN model, although its qualitative features are typical of all other neurons in this study.

Dynamic adaptation of precise timing

Although the overall performance of the GNM is significantly better on average than that of other models, this is often not the case for the first few trials of the repeated cross-validation stimulus (Fig. 5A). For the first few stimulus repeats, the performance of the GNM is below that of the LN and GL models before stabilizing at a much better performance for the duration of the repeats. Because blocks of repeated stimulus presentations used for cross-validation always followed pauses in the experiments when no visual stimulus was presented, we hypothesized that this shift might be due to adaptation in the state of the neuron during the few trials. However, the fact that this nonstationarity is not evi-

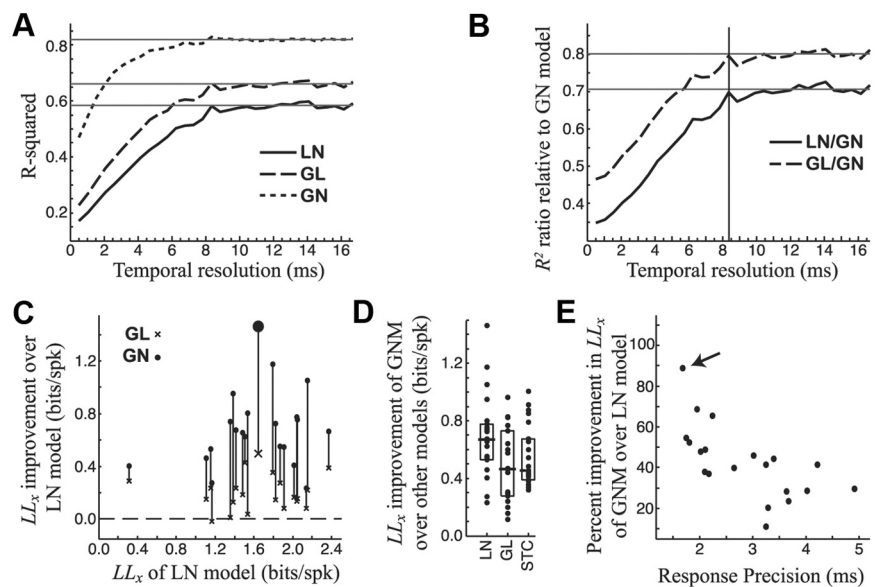


Figure 4. Performance of the GN model demonstrated by two cross-validated measures. **A**, The fraction of variance (R^2) of the observed firing rate (from the neuron in Figs. 1–3) explained by each model as a function of bin size, ranging from the temporal resolution of the fits (0.5 ms) up to double the frame duration. Although the GN model was constructed to capture performance at fine time scales (where it explains almost half the variance compared with other models that drop below 0.2), it also outperforms the LN and GL models at low time resolution (horizontal gray lines). **B**, The ratio of the R^2 performance of the GNM to the other two models, showing how the performance of other models begins to decrease for time scales smaller than the frame rate (vertical line) down to the finest time resolutions where the GNM explains more than twice the variance than that explained by the other models. Note that even at coarse time resolutions, where precision is not a factor, the GN model still outperforms the GL and LN models, implying that the interplay of excitation and suppression accounts for more than just timing predictions. **C**, The improvement in cross-validated log-likelihood LL_x of the GL (x symbols) and GN (circles) models, plotted as a function of the LL_x of the LN model, for the 21 LGN neurons in this study. The GN model yields a better performance (than both other models) for every neuron. The degree of improvement has little relationship to the initial LL_x of the LN model, but the relative magnitude of performance increase can be seen by comparing the horizontal axis to the vertical axis. The example neuron considered in previous figures is shown with a larger circle. bits/spk, Bits per spike. **D**, The increase in LL_x of the GNM compared with the other three models explicitly considered in this study. Each neuron is shown as a point ($N = 21$), with the median increase shown as a horizontal line and the box showing the 25 and 75% levels. Note that the GNM improved significantly over all other models considered for every neuron in the study, with a median increase around 0.5 bits/spk. **E**, The percentage improvement of the GN model is related to the precision of the neuron, defined as the time scale of the spike train autocorrelation function (see Materials and Methods). The example neuron in previous figures (arrow) is the most precise neuron in the study but exhibits the trends of all other neurons in this study.

dent in the performance of the less precise models suggests that this adaptation primarily affects the timing of the response rather than gross changes in firing rate.

To investigate this, we allow the gains of each model component and the overall spike threshold to dynamically change from trial to trial. We fitted these extra parameters using the first 5 s of each trial and performed cross-validation using the second 5 s (see Materials and Methods). This reveals a consistent pattern across most of the population of LGN neurons in this study ($N = 16$ of 21); the gains of the excitation and suppression are significantly smaller initially and quickly recover to their equilibrium values (Fig. 5B, top). At the same time, the spike threshold is initially much lower and increases with the same time course (Fig. 5B, bottom). The result of these opposing shifts is a fairly constant firing rate, where initially lower gains are paired with a higher “offset” such that it takes less excitation to get to threshold. These trial-to-trial adjustments of the gains and offset lead to a much more consistent cross-validated performance for the initial trials (Fig. 5C) and also stabilize slower fluctuations in the LL_x observed in some recordings (data not shown).

Although these simultaneous adjustments of gain and threshold have only small effect on the overall firing rate, they have

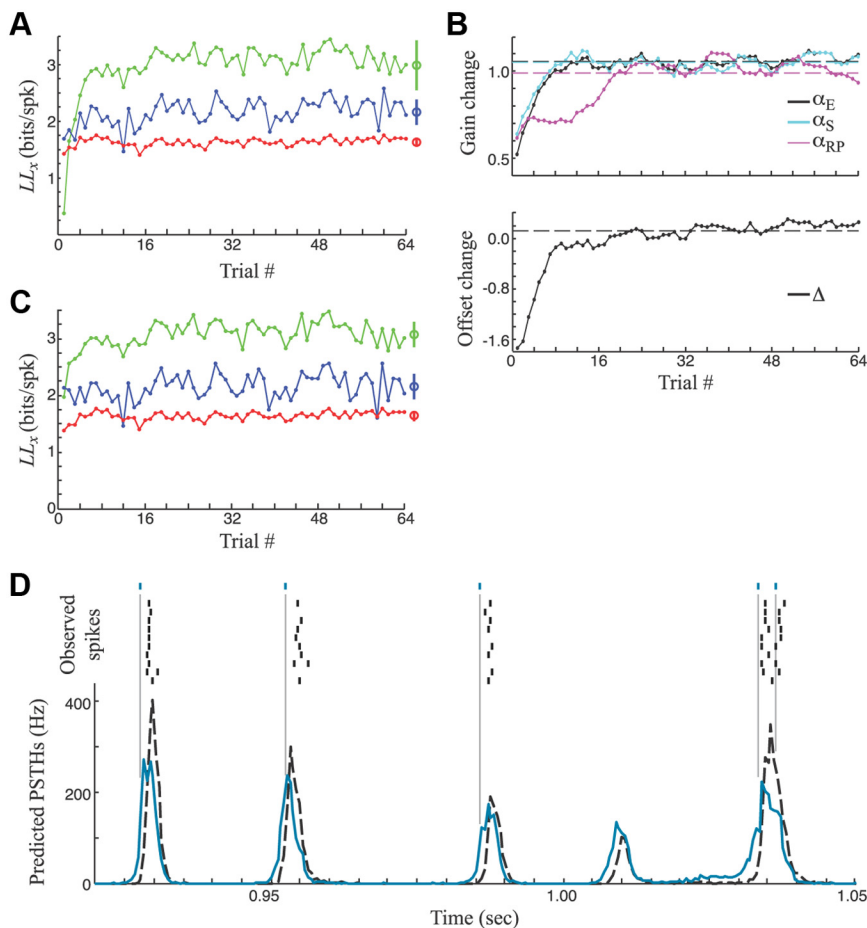


Figure 5. Adaptation of spike timing through measured changes in gain and spike-threshold. **A**, Cross-validated log-likelihoods (LL_x) of the LN (red), GL (blue), and GN (green) models on successive repeated trials of the same stimulus, demonstrating that the first trials of the GNM are very poorly predicted, but this quickly changes and stabilizes over several trials. bits/spk, Bits per spike. **B**, Top, Gains of the excitatory (α_E), suppressive (α_S), and spike history (α_{RP}) terms of the GN model that best predict the spike train of the first 5 s of each trial. Bottom, The simultaneously fit offset Δ of the spike threshold. Together, this suggests that the gains of the excitatory and suppressive terms are initially much smaller at the beginning of the trial, but the neuron is closer to spike threshold, and these values stabilize over time. **C**, Applying these gains and offsets dramatically stabilizes the quality of the prediction of the GNM, as demonstrated by the LL_x of the second 5 s of the repeated stimulus for the GNM (green). Similar adjustments to the gains and offsets of the other models (LN, red; GL, blue) have little effect on the LL_x . **D**, The difference in predictions of the GNM from the first trial (cyan) and at stabilized values (black) predict that spike times in general should be earlier initially. This agrees with observed spike trains on the corresponding section of cross-validation stimulus (top), which shows the observed spike trains of the first trial (cyan) compared with later trials (black). The small differences in spike timing between the first and later trials explain why this was only detected in the LL_x of the GNM (**A**).

significant effects on the predicted timing of firing events. This is demonstrated by comparing the predicted firing rates from the first few trials versus at equilibrium (Fig. 5D). While the observed PSTH cannot be directly compared (because it takes many repeated trials to estimate), the observed spike times themselves confirm this timing relationship (Fig. 5D, top).

Thus, this trial-based analysis underscores the ability of the GNM to predict the precise timing of the neuronal spike train. Such subtle differences in timing are not detectable using previous models and furthermore cannot be observed in the PSTH, which must be estimated from many repeated trials. Both the excitatory and suppressive gains, as well as the relative change in spiking threshold, are in principle observable experimentally, and thus these modeling results make specific predictions about the effects of adaptation on response timing, as well as their underlying mechanisms.

The relationship of the GNM to more general nonlinear characterizations

In explicitly modeling the nonlinear processing in LGN neurons as the sum of separate excitatory and suppressive elements, the GNM assumes a much more specific structure for the nonlinear processing of LGN neurons compared with other nonlinear modeling approaches, such as spike-triggered covariance or STC (de Ruyter van Steveninck and Bialek, 1988; Schwartz et al., 2006) and Maximally Informative Dimensions analysis or MID (Sharpee et al., 2004). Such approaches have been recently applied to both retinal ganglion cells and LGN neurons (Fairhall et al., 2006; Sincich et al., 2009), but the same relationship between the multiple temporal filters was not found. To understand the differences between solutions detected by these different methods, as well as to provide a larger context to interpret models found through nonlinear analysis, here we compare the GNM to models based on STC.

STC analysis is a widely applied nonlinear approach that can identify multiple “directions in stimulus space” to which a neuron is sensitive (de Ruyter van Steveninck and Bialek, 1988; Schwartz et al., 2006). Furthermore, given the [Gaussian noise] stimulus used in this study, STC and MID solutions should theoretically obtain the same estimates given the model form (Paninski, 2003; Sharpee et al., 2004), and this was verified for several neurons (data not shown). When applied to the example neuron considered in previous figures, STC analysis reveals two relevant directions in stimulus space for most LGN neurons in this study (Fig. 6A, top): an excitatory STA filter and a suppressive filter corresponding to the most significant dimension identified by STC. Together, these will be referred to as the STC filters. These two filters have a qualitatively similar relationship as the excitatory and suppressive receptive fields of

the GNM, with suppression delayed relative to excitation. However, they are clearly quite different in detail (Fig. 6A, bottom), with the STC filters at more spread latencies and significantly more biphasic relative to the GN filters.

The discrepancies between the temporal filters found by these different methods can be understood by considering the “subspace” of stimuli spanned by the two STC filters (Pillow and Simoncelli, 2006; Schwartz et al., 2006). The two effective GN filters are largely contained in this “STC subspace” and thus can be represented as linear combinations of the two STC filters (Fig. 6B) with only small differences with regard to their rebound. This relationship can be pictured geometrically by projecting all four filters into the plane defined by the STC filters (Fig. 6C). This shows that the effective GNM filters are in between the STA and suppressive STC filter, having more intermediate latencies and more closely resembling each other. Similar relationships be-

tween the STC and GN models are found with every neuron in this study, although in some cases the GNM filters did not project quite as well into the plane defined by STC (data not shown).

The different descriptions of the stimulus selectivity of these methods can be understood by picturing the locations of the stimuli associated with each spike projected into the subspace described (Fig. 6D). If the particular stimulus directions of this subspace had no relationship to the neuron's spikes, then the spike cloud projected into this subspace would have a width (i.e., standard deviation) shown by the black circle (Fig. 6D). However, the "spike cloud" does not follow this circle and is extended along the horizontal axis and compressed along the vertical axis (Fig. 6D). Because the STA does indeed point to the center of the spike cloud (by definition), the decreased variance along the vertical axis is explained by a symmetric nonlinear suppression (Fig. 6E, right), which is found when mapping the spiking nonlinearity associated with this dimension (see Materials and Methods). Such symmetric suppression is common in models of contrast adaptation (Shapley and Victor, 1978; Schwartz and Simoncelli, 2001; Schwartz et al., 2006; Mante et al., 2008), where neuronal processing is based on the STA and adaptation to contrast augments this processing. The bowl-like shape of the associated nonlinearity shows that stimuli that either match or are opposite to the suppressive STC filter will cause suppression, which is often described as divisive normalization (Schwartz and Simoncelli, 2001).

In contrast, the GN model suggests that the STA represents the combined effects of excitation and suppression, and that the true excitation of the neuron points one side of the spike cloud (Fig. 6D, blue). Because of asymmetric suppression that is rectifying (not "squaring" or symmetric), only spikes on one side of the excitatory filter are suppressed (Fig. 6D, green), resulting in the same overall location of spikes along the horizontal axis. The GN and STC models thus offer different explanations for processing of excitation and suppression within the same "stimulus subspace" in a way that produces the same spike cloud (Fig. 6F).

However, as mentioned above, these descriptions are not equivalent in their ability to describe the data, and the cross-validated performance of the GN model is significantly higher compared with models based on STC for every neuron in this study (Fig. 4D). Furthermore, the GN model structure is more readily interpreted in physiological terms—potentially as separate excitatory and inhibitory processing (see Discussion)—in part because of the monotonic nonlinearity (Fig. 2G) and PSC term (Fig. 2F) associated with the suppressive filter. In fact,

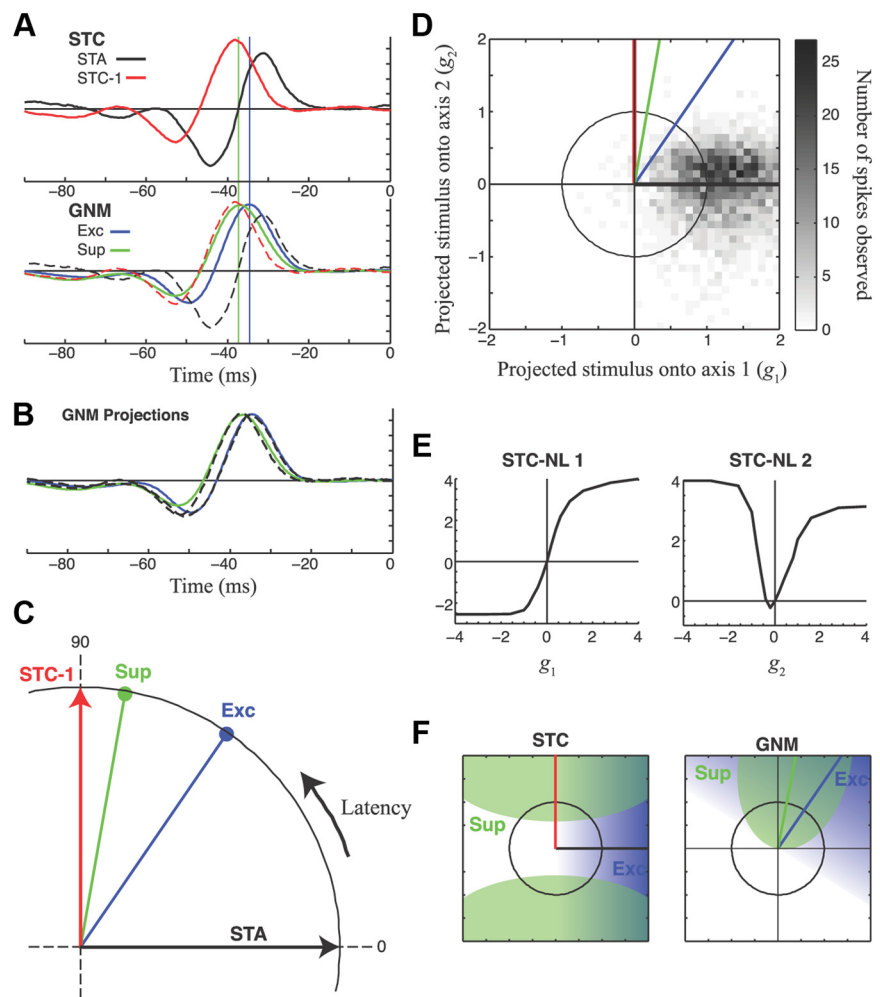


Figure 6. The geometrical relationship of excitation and suppression. **A**, STC analysis also detects an excitatory (black) and a delayed suppressive (red) filter, although they do not match the effective excitatory (Exc) and suppressive (Sup) filters of the GN model (bottom). **B**, Instead, the GN filters can be projected into the subspace spanned by the STC filters (black), suggesting that STC correctly identifies the relevant stimulus subspace of the neuron without identifying particular directions. **C**, The geometrical relationship between STC and GN filters can be shown in the plane spanned by the STA (black) and suppressive STC filter (red). The GN filters (blue, green) are expressed as unit vectors in this plane and are slightly shorter than unity because they do not fully project into this plane. **D**, A two-dimensional histogram of the stimuli associated with each spike, projected into the plane spanned by the STC filters and pictured relative to the filters of both models. Note that the STA (black) labels the center-of-mass of this spike cloud, and the suppressive STC (red) identifies the direction where this cloud is narrowed relative to the variance of other stimulus dimensions (black circle). The excitatory and suppressive GNM filter projections (blue, green) are also shown, as in **C**. **E**, A projection of the STC nonlinearities showing a saturating nonlinearity for the excitatory (STA) direction (left) and a bowl-shaped nonlinearity for the suppressive direction (right). **F**, A diagram showing how these two different models predict similar geometry of the spike cloud as the interplay between excitation and suppression. The STC-based model (left) predicts excitation centered at the middle of the spike cloud (blue) and suppression on both sides (green), resulting in the observed narrowing. In contrast, the GNM (right) predicts excitation (blue) at a slightly earlier latency than the suppressive STC filter and suppression at longer latency (green), which attenuates the longer-latency spike cloud, resulting in a different explanation for the narrowing of the spike cloud in the direction of the suppressive-STC filter.

monotonicity of the nonlinearity and nonpositivity of the PSC term is explicitly enforced in the model parameter search, which speeds the fitting procedure and eliminates local maxima in the likelihood. Not doing so can often result in the GNM finding a solution much more similar to the STC filters, but this occurs at a local maximum of the likelihood rather than a global maximum. The structure of the GNM also allows for the incorporation of spike refractoriness (Fig. 2C), which can only be taken into account in the STC model as an additional filter acting on the stimulus (data not shown). In this sense, the GN model is structured to account for a biologically implementable solution

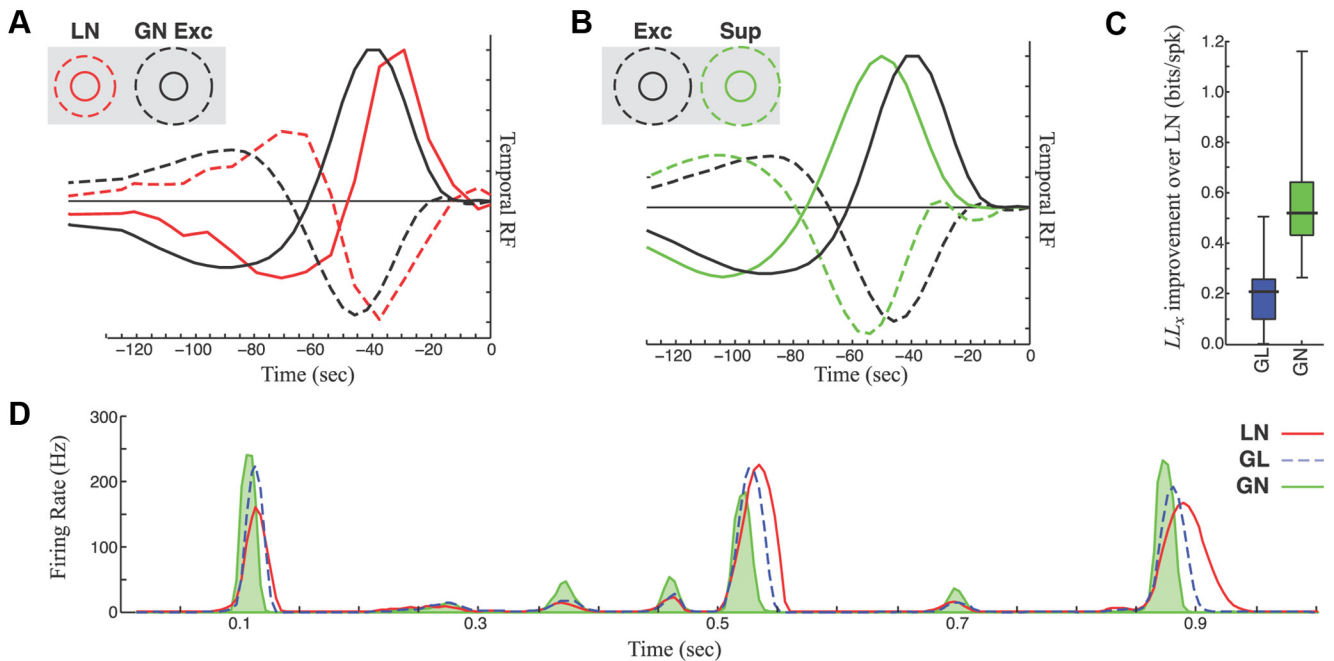


Figure 7. Precision generated by the interplay of excitation and suppression during natural movies **A**, The relationship between the linear (LN) receptive field (RF) (red) and the excitatory (Exc) receptive field of the GNM (black) measured in the context of a natural movie stimulus, revealing the same relationships as for the purely temporal case (Fig. 2). Center (solid) and surround (dashed) components of each spatiotemporal receptive field, STRF, were assumed to be individually space–time separable, given by a two-dimensional circular Gaussian (inset) paired with each temporal kernel. Inset shows the standard deviation of each Gaussian, relative to the vertical dimension of the gray box, which was 1.64° . **B**, A comparison between the excitatory (black) and suppressive (green) STRFs of the GNM, demonstrating the delay between excitation and suppression responsible for precision in this stimulus context. **C**, Cross-validated likelihoods LL_x of the GL and GN models compared with the LN model. bits/spk, bits per spike. **D**, Firing rate predictions of each model demonstrating that the GN prediction is the most precise. Note that these recordings did not include appropriate repeated trials to directly compare with an observed PSTH.

that improves model predictions and leads to specific physiological predictions of the origins of each model component (see Discussion).

Excitation and suppression generate precision in spatiotemporal contexts

The time scales of LGN responses become significantly longer in the context of natural visual stimuli, but LGN responses to natural movies are still significantly more precise than would be expected by linear processing alone (Butts et al., 2007, 2010). To investigate whether this precision could arise from a similar interplay of excitatory and suppressive elements, we developed a spatiotemporal version of the GNM described by replacing the purely temporal receptive fields with spatiotemporal receptive fields or STRFs. Spatial processing was performed by the common difference-of-Gaussians model (Cai et al., 1997) consisting of overlapping center and surround circular Gaussians, each with an associated temporal kernel (see Fig. 7A). This incorporates the observation that spatiotemporal processing is not space–time separable in LGN neurons, although when considered separately the center and surround themselves are (Dawis et al., 1984; Cai et al., 1997; Allen and Freeman, 2006). As with the GL and GN models described for purely temporal data, all parameters of each STRF, including the position and width of each Gaussian and their associated temporal kernels, are optimized through maximum likelihood estimation (see Materials and Methods). Note that we expect nonlinear spatiotemporal processing in LGN neurons to be more complicated than this basic phenomenological model (Mante et al., 2008), but are simply aiming here for the next order model above linear that is consistent with previous studies of spatiotemporal processing.

GL and GN models were applied to LGN recordings in the context of natural movies recorded from a camera on the head of a cat walking in a forest (Kayser et al., 2004). Unlike spike-triggered average approaches (Theunissen et al., 2001), maximum likelihood estimation automatically adjusts for stimulus correlations, making parameter estimation in the context of natural stimuli no more difficult than with uncorrelated “noise” stimuli.

The spatiotemporal GNM in the context of natural movies results in a very similar picture of the generation of precision as compared with the spatially uniform cases considered above. For all neurons ($N = 20$), there was a strong suppressive term with very similar spatiotemporal properties as the excitation (Fig. 7B), although delayed through the PSC term. Note that the delay between excitation and suppression is in addition to the differences in temporal processing between center and surround of each (Fig. 7A,B). The spatiotemporal tuning of excitation was also significantly different from that found in the context of the LN model (using comparable maximum likelihood estimation) (Fig. 7A), which presumably combines the effects of excitation and suppression into a single linear term.

The resulting GNM predictions are a better description of the recorded spike trains than the LN model for every neuron recorded (Fig. 7C), in large part because its ability to predict the precise responses of LGN neurons (Fig. 7D). In conclusion, the interplay of excitation and suppression observed in the spatially uniform stimulus condition is also present and relevant in natural stimulus contexts and is necessary to explain the observed precision (Butts et al., 2007), as well as more subtle features of LGN responses in these conditions that are only evident through the statistical analysis described.

Discussion

Here we provide evidence that the precise timing of LGN neuron responses reflects the interplay of excitation and stimulus-tuned suppression that is detectable in a variety of stimulus contexts. Suppression is delayed relative to excitation, such that LGN responses occur in a brief window where the amount of excitation exceeds the amount of suppression. The interplay between these two stimulus processing elements appears to underlie the precise timing of LGN neuron responses, which has been observed both in spatially uniform noise (Reinagel and Reid, 2000; Liu et al., 2001; Kumbhani et al., 2007) and natural movies (Butts et al., 2007). Such interplay might have its origins in the retina (see below) and subsequently could also underlie the temporal precision observed in retinal ganglion cell responses (Berry and Meister, 1998; Liu et al., 2001; Passaglia and Troy, 2004; Uzzell and Chichilnisky, 2004). Likewise, such a means of generating precise neuronal responses may be quite general in sensory processing and is similar to a proposed function of feedforward inhibition in the somatosensory cortex (Gabernet et al., 2005; Wilent and Contreras, 2005; Okun and Lampl, 2008), auditory cortex (Wehr and Zador, 2003; Wu et al., 2008), and visual cortex (Cardin et al., 2007).

We derived this explanation of visual processing in the LGN using only extracellular [single-unit] data and statistical validation, similar in spirit (and/or methodology) to much previous work in the visual pathway (Berry and Meister, 1998; Wu et al., 2006; Mante et al., 2008; Pillow et al., 2008; Sincich et al., 2009). In this sense, observed spike times are leveraged to evaluate different hypotheses of how the visual circuit processes stimuli, and the large increase in the performance of the model proposed here over that of previous models of retina and LGN processing (Fig. 4D) corresponds to its ability to explain aspects of physiological data that have not been otherwise explained. Furthermore, the resemblance of the identified computation to properties of feedforward inhibition and the model's success at accurately describing the precise timing of LGN spike trains suggest new mechanistic hypotheses (see below) that might be validated by targeted intracellular experiments, as is possible in the retina (Murphy and Rieke, 2006; Manookin et al., 2008) and LGN (Wang et al., 2007, 2011).

The possible sources of suppression

The delayed suppression detected by our model most likely arises from inhibitory interneurons, with the delay due to a disynaptic pathway. Inhibitory interneurons exist in both the retina and LGN, and the effects we report may come from either or both. Most directly, interneurons in the LGN receive direct retinal input and form inhibitory synapses onto both LGN neurons and their presynaptic retinal ganglion cell terminals (Dubin and Cleland, 1977). Studies of LGN neurons *in vitro* found evidence for some LGN neurons that receive excitatory and inhibitory input from the same retinal ganglion cell (Blitz and Regehr, 2005), although such “push” inhibition has not been observed *in vivo* (Wang et al., 2007, 2011).

There is also evidence that precision is at least partially generated in the retina (Berry and Meister, 1998; Passaglia and Troy, 2004; Uzzell and Chichilnisky, 2004; Murphy and Rieke, 2006) and enhanced across the retinogeniculate synapse (Carandini et al., 2007; Casti et al., 2008; Rathbun et al., 2010; Wang et al., 2010). The LNL(N) cascade structure of the GNM (Fig. 2A) is conceptually similar to previously proposed models of retinal processing (Korenberg et al., 1989), including the effects of contrast adaptation (Meister and Berry, 1999; Shapley and Victor,

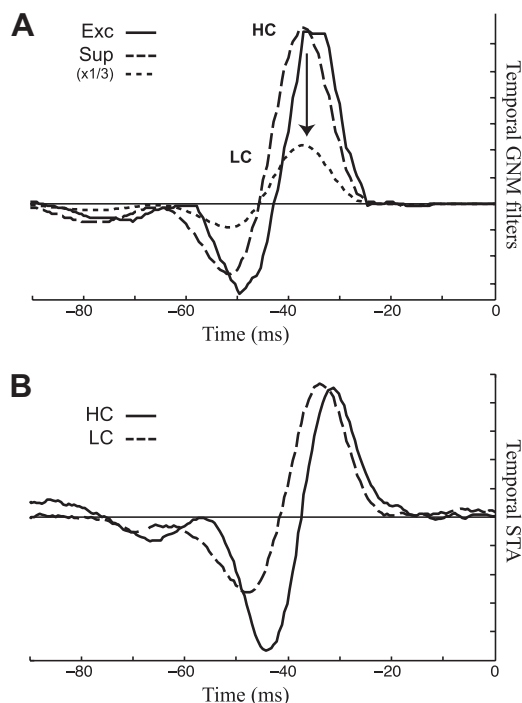


Figure 8. Elements of contrast gain control explained through the modulation of suppression relative to excitation. **A**, Simulated data are generated by the GNM using the excitatory (Exc) and suppressive (Sup) filters considered in previous figures (solid and dashed lines, respectively) for the simulated “high contrast” condition. Low contrast is simulated by reducing the magnitude of the stimulus by a factor of 2 while scaling the magnitude of the output of the suppressive term by a factor of 3 (dotted line). Note that the kernel itself is not scaled (rather the magnitude of the signal after the suppressive nonlinearity, but this approximates the effect). **B**, The spike-triggered average, STA, of the simulated data in high contrast (HC, solid line) and low contrast (LC, dashed line) conditions, demonstrating how the STA in high contrast becomes shorter latency and more biphasic—hallmarks of contrast gain control. The STA is scaled to have comparable variance for direct comparison of the shapes.

1978). This possibility is also consistent with our findings, because such retinal processing would be inherited by the LGN and then contribute to the observed computation underlying LGN response properties.

In fact, the GNM can replicate the dominant effects of contrast adaptation if the suppression is more strongly modulated by contrast than by excitation (Shapley and Victor, 1981; Butts and Casti, 2009). Higher contrasts would evoke stronger suppression, and the effects of the delayed suppressive filter would impart sensitivity to higher frequencies and decreased response latency (Shapley and Victor, 1978; Zaghoul et al., 2005). Also, more suppression at high contrasts would result in lower “gain” overall. These effects can be explicitly reproduced by the GNM by decreasing the relative strength of suppression in low contrast and can be seen in the resulting STA of simulated data (Fig. 8), yielding results consistent with previous contrast studies (Chander and Chichilnisky, 2001; Zaghoul et al., 2005). While similar in structure, the GNM results suggest that “contrast normalization” may not in fact be symmetric or sign invariant, as in past models of adaptation (Meister and Berry, 1999; Schwartz and Simoncelli, 2001; Mante et al., 2008). Rather, the GN model finds an asymmetric solution with monotonic nonlinearities as an alternative explanation for “bowl-shaped” suppression (Fig. 6), which achieves similar computational effects but has much better performance in reproducing the data in this study.

While contrast adaptation plays a large role in shaping the LGN response to natural movies (Mante et al., 2008), the nonlinear effects of contrast adaptation are generally not considered outside of the context of contrast-varying stimuli, except in the anticipation of moving stimuli (Berry et al., 1999). Our study thus suggests a larger role for contrast gain control mechanisms in determining the precise timing of LGN responses. In doing so, it also implies that the impact of such mechanisms can be characterized with high fidelity outside of the context of contrast adaptation. This interpretation is thus far consistent with preliminary results applying the GNM to data at multiple contrasts from cat retina (Casti et al., 2008) and would be directly verifiable through intracellular experiments in the retina targeted to the contexts that this study suggests.

Nonlinear modeling of extracellular data

Rather than using intracellular recording, we measured putative excitatory and inhibitory tuning underlying LGN neuron responses by using extracellular recordings (Butts et al., 2007). An internal nonlinear stage in the GNM (Fig. 2A) allows for multiple stimulus-tuned elements that contribute nonlinearly to the LGN response, and efficient maximum likelihood estimation (Paninski, 2004) can be used to find both the receptive field filters and their associated nonlinearities. This approach thus joins several other nonlinear approaches that can identify multiple receptive fields, including spike-triggered covariance (de Ruyter van Steveninck and Bialek, 1988; Pillow and Simoncelli, 2006; Schwartz et al., 2006), maximally informative dimensions (Sharpe et al., 2004), information-theoretic spike-triggered average and covariance (Pillow and Simoncelli, 2006), and neural network approaches (Lau et al., 2002; Prenger et al., 2004).

Other techniques have previously identified the presence of multiple stimulus-tuned elements in the retina (Fairhall et al., 2006) and LGN (Sincich et al., 2009) (but see Wang et al., 2010). However, the filters found in these studies do not directly reveal the interplay of excitation and suppression found in this study. Likewise, in this study, STC analysis did not find the same filters as the GNM, although it did find roughly the same relevant stimulus subspace (Fig. 6). These differences likely stem from our focus on temporal precision in this study, as well as the particular elements of the GNM structure and optimization framework. First, the LGN neuron response to spatially uniform noise is one of the most precise examples of visual processing, which likely enhanced our ability to detect solutions relating specifically to explaining precision and also precipitated a high temporal resolution of our analyses (0.5 ms). By contrast, in a previous study of precision in the primate retina (Pillow et al., 2005), the LN model largely captured the observed duration of firing events, and that study focused instead on explaining the exact spike count, variability, and intra-event timing by using a model with spike refractoriness. Because the duration of firing events in our LGN data is 2–4 times shorter than the predictions of LN and GL models (Fig. 3B), we focused instead on explaining this aspect of precision, which enabled the effects of spike refractoriness and stimulus-driven suppressive terms to be distinguished.

Other differences between our findings and previous nonlinear studies are also likely due to particular aspects of the GNM framework and how it is optimized, including: (1) the ability of the GNM to simultaneously optimize multiple filters and their associated nonlinearities with other model parameters; (2) the incorporation of other elements directly into the GN model, such as the spike history term; and (3) the LNL(N) cascade structure

of the nonlinear elements, allowing for an efficient parameter search for biologically plausible solutions. In particular, the LNL structure of the GN model mirrors observations comparing intracellular and extracellular processing in the retina (Korenberg et al., 1989) and is also similar to contrast gain control models. The LNL structure purposefully emulates the computation resulting from input from other neurons processing the stimulus, with the “PSC term” reflecting the temporal integration of external inputs that are themselves defined by LN processing. By constraining the nonlinearity to be monotonic and the PSC term to be positive (excitatory) or negative (suppressive), the GN model can thus search a space describing realistic inputs, allowing biological constraints to be implemented and resulting in a more direct resemblance between the resulting model and underlying mechanisms of interest.

Thus, application of an appropriately structured model reveals significant nonlinear processing in LGN neurons, which—outside of the context of luminance and contrast adaptation—have been generally thought of as well described by linear models (Carandini et al., 2005; Shapley, 2009). Strong nonlinear elements, including separately tuned spatiotemporal elements of excitation and suppression, have a significant impact on the LGN response in both artificial and natural stimulus contexts. Nonlinear processing in the LGN can have significant ramifications for cortical processing (Priebe and Ferster, 2006), and in that sense these results lend to a stronger foundation on which to build an understanding of computation across the visual pathway.

Notes

Supplemental material for this article is available at <http://www.cfls.umd.edu/biology/ntlab/NTlab/GNM.html>. Exploration of alternative models for LGN processing. Part 1: Additional analysis of spike-triggered covariance models. Part 2: Model performance on simulated spike trains. This material has not been peer reviewed.

References

- Ahrens MB, Paninski L, Sahani M (2008) Inferring input nonlinearities in neural encoding models. *Network* 19:35–67.
- Alitto HJ, Usrey WM (2008) Origin and dynamics of extraclassical suppression in the lateral geniculate nucleus of the macaque monkey. *Neuron* 57:135–146.
- Allen EA, Freeman RD (2006) Dynamic spatial processing originates in early visual pathways. *J Neurosci* 26:11763–11774.
- Berry MJ 2nd, Meister M (1998) Refractoriness and neural precision. *J Neurosci* 18:2200–2211.
- Berry MJ 2nd, Brivanlou IH, Jordan TA, Meister M (1999) Anticipation of moving stimuli by the retina. *Nature* 398:334–338.
- Blitz DM, Regehr WG (2005) Timing and specificity of feed-forward inhibition within the LGN. *Neuron* 45:917–928.
- Borst A, Theunissen FE (1999) Information theory and neural coding. *Nat Neurosci* 2:947–957.
- Brenner N, Strong SP, Koberle R, Bialek W, de Ruyter van Steveninck RR (2000) Synergy in a neural code. *Neural Comput* 12:1531–1552.
- Brillinger D (1992) Nerve cell spike train data analysis: a progression of technique. *J Am Stat Assoc* 87:260–271.
- Buracas GT, Zador AM, DeWeese MR, Albright TD (1998) Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 20:959–969.
- Butts DA, Casti ARR (2009) Nonlinear computation performed at the retinogeniculate synapse. *Soc Neurosci Abstr* 35:849.5.
- Butts DA, Weng C, Jin J, Yeh CI, Lesica NA, Alonso JM, Stanley GB (2007) Temporal precision in the neural code and the timescales of natural vision. *Nature* 449:92–95.
- Butts DA, Desbordes G, Weng C, Jin J, Alonso JM, Stanley GB (2010) The episodic nature of spike trains in the early visual pathway. *J Neurophysiol* 104:3371–3387.
- Cai D, DeAngelis GC, Freeman RD (1997) Spatiotemporal receptive field

- organization in the lateral geniculate nucleus of cats and kittens. *J Neurophysiol* 78:1045–1061.
- Carandini M, Demb JB, Mante V, Tolhurst DJ, Dan Y, Olshausen BA, Gallant JL, Rust NC (2005) Do we know what the early visual system does? *J Neurosci* 25:10577–10597.
- Carandini M, Horton JC, Sincich LC (2007) Thalamic filtering of retinal spike trains by postsynaptic summation. *J Vis* 7:20.1–2.11.
- Cardin JA, Palmer LA, Contreras D (2007) Stimulus feature selectivity in excitatory and inhibitory neurons in primary visual cortex. *J Neurosci* 27:10333–10344.
- Casti A, Hayot F, Xiao Y, Kaplan E (2008) A simple model of retina-LGN transmission. *J Comput Neurosci* 24:235–252.
- Chander D, Chichilnisky EJ (2001) Adaptation to temporal contrast in primate and salamander retina. *J Neurosci* 21:9904–9916.
- Chichilnisky EJ (2001) A simple white noise analysis of neuronal light responses. *Network* 12:199–213.
- Dawis S, Shapley R, Kaplan E, Tranchina D (1984) The receptive field organization of X-cells in the cat: spatiotemporal coupling and asymmetry. *Vision Res* 24:549–564.
- de Ruyter van Steveninck R, Bialek W (1988) Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc R Soc Lond B Biol Sci* 234:379–414.
- Desbordes G, Jin J, Weng C, Lesica NA, Stanley GB, Alonso JM (2008) Timing precision in population coding of natural scenes in the early visual system. *PLoS Biol* 6:e324.
- Dubin MW, Cleland BG (1977) Organization of visual inputs to interneurons of lateral geniculate nucleus of the cat. *J Neurophysiol* 40:410–427.
- Fairhall AL, Burlingame CA, Narasimhan R, Harris RA, Puchalla JL, Berry MJ 2nd (2006) Selectivity for multiple stimulus features in retinal ganglion cells. *J Neurophysiol* 96:2724–2738.
- Gabernet L, Jadhav SP, Feldman DE, Carandini M, Scanziani M (2005) Somatosensory integration controlled by dynamic thalamocortical feed-forward inhibition. *Neuron* 48:315–327.
- Gaudry KS, Reinagel P (2007) Contrast adaptation in a nonadapting LGN model. *J Neurophysiol* 98:1287–1296.
- Gerstner W (2001) A framework for spiking neuron models: the spike response model. In: *The handbook of biological physics*, Vol 4 (Moss F, Gielen S, eds), pp 469–516. Amsterdam: Elsevier Sciences.
- Kayser C, Kording KP, König P (2004) Processing of complex stimuli and natural scenes in the visual cortex. *Curr Opin Neurobiol* 14:468–473.
- Keat J, Reinagel P, Reid RC, Meister M (2001) Predicting every spike: a model for the responses of visual neurons. *Neuron* 30:803–817.
- Korenberg MJ, Sakai HM, Naka K (1989) Dissection of the neuron network in the catfish inner retina. III. Interpretation of spike kernels. *J Neurophysiol* 61:1110–1120.
- Kouh M, Sharpee TO (2009) Estimating linear-nonlinear models using Rényi divergences. *Network* 20:49–68.
- Kumbhani RD, Nolt MJ, Palmer LA (2007) Precision, reliability, and information-theoretic analysis of visual thalamocortical neurons. *J Neurophysiol* 98:2647–2663.
- Lau B, Stanley GB, Dan Y (2002) Computational subunits of visual cortical neurons revealed by artificial neural networks. *Proc Natl Acad Sci U S A* 99:8974–8979.
- Lesica NA, Stanley GB (2006) Decoupling functional mechanisms of adaptive encoding. *Network* 17:43–60.
- Liu RC, Tzovev S, Rebrik S, Miller KD (2001) Variability and information in a neural code of the cat lateral geniculate nucleus. *J Neurophysiol* 86:2789–2806.
- Manookin MB, Beaudoin DL, Ernst ZR, Flagel LJ, Demb JB (2008) Disinhibition combines with excitation to extend the operating range of the OFF visual pathway in daylight. *J Neurosci* 28:4136–4150.
- Mante V, Bonin V, Carandini M (2008) Functional mechanisms shaping lateral geniculate responses to artificial and natural stimuli. *Neuron* 58:625–638.
- McCullagh P, Nelder J (1989) *Generalized linear models*. London: Chapman and Hall.
- Meister M, Berry MJ 2nd (1999) The neural code of the retina. *Neuron* 22:435–450.
- Murphy GJ, Rieke F (2006) Network variability limits stimulus-evoked spike timing precision in retinal ganglion cells. *Neuron* 52:511–524.
- Okun M, Lampl I (2008) Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nat Neurosci* 11:535–537.
- Paninski L (2003) Convergence properties of three spike-triggered analysis techniques. *Network* 14:437–464.
- Paninski L (2004) Maximum likelihood estimation of cascade point-process neural encoding models. *Network* 15:243–262.
- Paninski L, Pillow JW, Lewi J (2007) Statistical models for neural encoding, decoding, and optimal stimulus design. In: *Computational neuroscience: theoretical insights into brain function*, Vol 165 (Cisek P, Drew T, Kalaska J, eds), pp 493–507. Boston: Elsevier Science.
- Paninski L, Ahmadian Y, Ferreira DG, Koyama S, Rahnema Rad K, Vidne M, Vogelstein J, Wu W (2010) A new look at state-space models for neural data. *J Comput Neurosci* 29:107–126.
- Passaglia CL, Troy JB (2004) Information transmission rates of cat retinal ganglion cells. *J Neurophysiol* 91:1217–1229.
- Pillow JW, Simoncelli EP (2006) Dimensionality reduction in neural models: an information-theoretic generalization of spike-triggered average and covariance analysis. *J Vis* 6:414–428.
- Pillow JW, Paninski L, Uzzell VJ, Simoncelli EP, Chichilnisky EJ (2005) Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *J Neurosci* 25:11003–11013.
- Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454:995–999.
- Prenger R, Wu MC, David SV, Gallant JL (2004) Nonlinear V1 responses to natural scenes revealed by neural network analysis. *Neural Netw* 17:663–679.
- Priebe NJ, Ferster D (2006) Mechanisms underlying cross-orientation suppression in cat visual cortex. *Nat Neurosci* 9:552–561.
- Rahnema Rad KR, Paninski L (2010) Efficient, adaptive estimation of two-dimensional firing rate surfaces via Gaussian process methods. *Network* 21:142–168.
- Rathbun DL, Warland DK, Usrey WM (2010) Spike timing and information transmission at retinogeniculate synapses. *J Neurosci* 30:13558–13566.
- Reinagel P, Reid RC (2000) Temporal coding of visual information in the thalamus. *J Neurosci* 20:5392–5400.
- Schwartz O, Simoncelli EP (2001) Natural signal statistics and sensory gain control. *Nat Neurosci* 4:819–825.
- Schwartz O, Pillow JW, Rust NC, Simoncelli EP (2006) Spike-triggered neural characterization. *J Vis* 6:484–507.
- Shapley R (2009) Linear and nonlinear systems analysis of the visual system: why does it seem so linear? A review dedicated to the memory of Henk Spekreijse. *Vision Res* 49:907–921.
- Shapley RM, Victor JD (1978) The effect of contrast on the transfer properties of cat retinal ganglion cells. *J Physiol* 285:275–298.
- Shapley RM, Victor JD (1981) How the contrast gain control modifies the frequency responses of cat retinal ganglion cells. *J Physiol* 318:161–179.
- Sharpee T, Rust NC, Bialek W (2004) Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput* 16:223–250.
- Simoncelli EP, Pillow J, Paninski L, Schwartz O (2004) Characterization of neural responses with stochastic stimuli. In: *The cognitive neurosciences* (Gazzaniga M, ed), pp 327–338. Cambridge, MA: MIT.
- Sincich LC, Horton JC, Sharpee TO (2009) Preserving information in neural transmission. *J Neurosci* 29:6207–6216.
- Theunissen F, Miller JP (1995) Temporal encoding in nervous systems: a rigorous definition. *J Comput Neurosci* 2:149–162.
- Theunissen FE, David SV, Singh NC, Hsu A, Vinje WE, Gallant JL (2001) Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network* 12:289–316.
- Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN (2005) A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J Neurophysiol* 93:1074–1089.
- Uzzell VJ, Chichilnisky EJ (2004) Precision of spike trains in primate retinal ganglion cells. *J Neurophysiol* 92:780–789.
- Wang X, Wei Y, Vaingankar V, Wang Q, Koepsell K, Sommer FT, Hirsch JA (2007) Feedforward excitation and inhibition evoke dual modes

- of firing in the cat's visual thalamus during naturalistic viewing. *Neuron* 55:465–478.
- Wang X, Hirsch JA, Sommer FT (2010) Recoding of sensory information across the retinthalamic synapse. *J Neurosci* 30:13567–13577.
- Wang X, Vaingankar V, Sanchez CS, Sommer FT, Hirsch JA (2011) Thalamic interneurons and relay cells use complementary synaptic mechanisms for visual processing. *Nat Neurosci* 14:224–231.
- Wehr M, Zador AM (2003) Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* 426:442–446.
- Weng C, Yeh CI, Staelzel CR, Alonso JM (2005) Receptive field size and response latency are correlated within the cat visual thalamus. *J Neurophysiol* 93:3537–3547.
- Wilentz WB, Contreras D (2005) Dynamics of excitation and inhibition underlying stimulus selectivity in rat somatosensory cortex. *Nat Neurosci* 8:1364–1370.
- Wu GK, Arbuckle R, Liu BH, Tao HW, Zhang LI (2008) Lateral sharpening of cortical frequency tuning by approximately balanced inhibition. *Neuron* 58:132–143.
- Wu MC, David SV, Gallant JL (2006) Complete functional characterization of sensory neurons by system identification. *Annu Rev Neurosci* 29:477–505.
- Zaghloul KA, Boahen K, Demb JB (2005) Contrast adaptation in subthreshold and spiking responses of mammalian Y-type retinal ganglion cells. *J Neurosci* 25:860–868.