

Spike-triggered neural characterization

Odelia Schwartz

The Salk Institute, La Jolla, CA, USA, &
Howard Hughes Medical Institute, USA



AQ1

Jonathan W. Pillow

Gatsby Computational Neuroscience Unit,
University College London, London, UK



AQ2

Nicole C. Rust

Department of Brain and Cognitive Sciences and
McGovern Institute for Brain Research, MIT,
Cambridge, MA, USA



AQ2

Eero P. Simoncelli

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



AQ3

Neural responses are commonly described using receptive fields. This description may be formalized in a model that operates with a small set of linear filters whose outputs are nonlinearly combined to determine the instantaneous firing rate. Spike-triggered average and covariance analyses can be used to estimate the filters and nonlinear combination rule from extracellular experimental data. We describe this methodology, demonstrating it with simulated model neuron examples that emphasize practical issues that arise in experimental situations.

Keywords: spike-triggered analysis, characterization, neural response

Introduction

A fundamental goal of sensory systems neuroscience is the characterization of the functional relationship between stimuli and neural responses. The purpose of such a characterization is to elucidate the computation being performed by the system. Many electrophysiological studies in sensory areas describe neural firing rates in response to highly restricted sets of stimuli that are parameterized by one or perhaps two stimulus parameters. Although such “tuning curve” measurements have led to considerable understanding of neural coding, they provide only a partial glimpse of the full neural response function. On the other hand, it is not feasible to measure neural responses to *all* stimuli. One way to make progress is to restrict the response function to a particular model (or class of models). In this modeling approach, the problem is reduced to developing a set of stimuli along with a methodology for fitting the model to measurements of neural responses to those stimuli. One wants a model that is flexible enough to provide a good description of neural responses but simple enough that the fitting is both tractable and well constrained under realistic experimental data conditions.

One class of solutions, which we refer to as “spike-triggered analysis,” has received considerable attention in recent years due to a variety of new methodologies, improvements in stimulus generation technology, and

demonstration of physiological results. In these methods, one generally assumes that the probability of a neuron eliciting a spike (i.e., the instantaneous firing rate) is governed only by recent sensory stimuli. More specifically, the response model is assumed to be an inhomogeneous Poisson process whose rate is a function of the stimuli presented during a recent temporal window of fixed duration. In the forward neural response model, the stimuli are mapped to a scalar value that determines the instantaneous firing rate of a Poisson spike generator. Our job in the analysis is to work backward: From the stimuli that elicited spikes, we aim to estimate this firing rate function. The analysis of experimental data is thus reduced to examining the properties of the stimuli within temporal windows preceding each recorded spike, known as the *spike-triggered stimulus ensemble* (Figure 1A).

Understanding how the spike-triggered distribution differs from the raw stimuli is key to determining the firing rate function. It is often useful to visualize the analysis problem geometrically (Figure 1B). Consider input stimuli, which at each time step consist of an array of randomly chosen pixel values (8 pixels in this example). The neural response at any particular moment in time is assumed to be completely determined by the stimulus segment that occurred during a prespecified interval in the past (6 time steps in this example). The overall stimulus dimensionality is high (48 dimensions here), but we can depict a projection of the stimuli onto two space–time vectors. The raw stimulus ensemble and the spike-

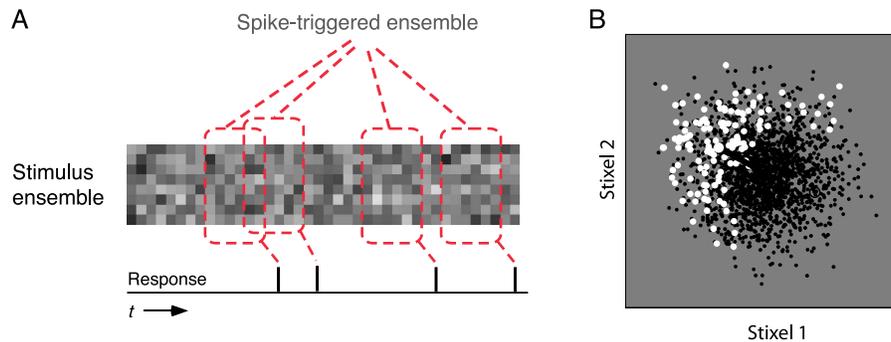


Figure 1. The spike-triggered stimulus ensemble. (A) Discretized stimulus sequence and observed neural response (spike train). On each time step, the stimulus consists of an array of randomly chosen values (eight, for this example). These could represent, for example, the intensities of a fixed set of individual pixels on the screen or the contrast of each of a set of fixed sinusoidal gratings that are additively superimposed. The neural response at any particular moment in time is assumed to be completely determined by the stimulus segment that occurred during a prespecified interval in the past. In this figure, the segment covers six time steps and lags three time steps behind the current time (to account for response latency). The spike-triggered ensemble consists of the set of segments associated with spikes. (B) Geometric (vector space) view of the spike-triggered ensemble. Stimuli (here, 48-dimensional) are projected onto two space–time vectors. In this example, each of the two vectors contained 1 stixel (space–time pixel) set to a value of 1, and the other 47 stixels were set to 0. For these given vectors, the projection is equivalent to the intensity of the corresponding stixel in the stimulus. More generally, one can project the stimuli onto any two 48-dimensional vectors. The spike-triggered stimulus segments (white points) constitute a subset of all stimulus segments presented (black points).

85 triggered ensemble are then two clouds of points in this
 86 space. Intuitively, the task of estimating the neural
 87 response function corresponds to describing the ways in
 88 which these two clouds differ. In practice, when the input
 89 stimulus space is of high dimensionality, one cannot
 90 estimate the neural response function without further
 91 assumptions.

92 Spike-triggered analysis has been employed to estimate
 93 the terms of a Wiener/Volterra expansion (Korenberg,
 94 Sakai, & Naka, 1989; Marmarelis & Marmarelis, 1978;
 95 Volterra, 1959; Wiener, 1958), in which the mapping from
 96 stimuli to firing rate is described using a low-order
 97 polynomial (see Dayan & Abbott, 2001; Rieke, Warland,
 98 de Ruyter van Steveninck, & Bialek, 1997 for a review).
 99 Although any reasonable function can be approximated as
 100 a polynomial, the firing rate nonlinearities found in the
 101 responses of sensory neurons (e.g., half-wave rectified,
 102 rapidly accelerating and saturating) tend to require a
 103 polynomial with many terms (see, e.g., Rieke et al., 1997).
 104 However, the amount of data needed for accurate
 105 estimation grows rapidly with the number of terms.
 106 Therefore, in an experimental setting where one can
 107 estimate only the first few terms of the expansion, the
 108 polynomial places a strong restriction on the nonlinearity.

109 As an alternative to the polynomial approximation, one
 110 can assume that the response function operates on a low-
 111 dimensional *linear subspace* of the full stimulus space
 112 (Bialek & de Ruyter van Steveninck, 2005; de Ruyter van
 113 Steveninck & Bialek, 1988). That is, the response of a
 114 neuron is modeled with a small set of linear filters whose
 115 outputs are combined nonlinearly to generate the instan-
 116 taneous firing rate. Stated differently, although the stimulus

117 space is high dimensional, it is assumed that the neuron
 118 only cares about a small set of dimensions. This is in
 119 contrast to the Wiener/Volterra approach, which in general
 120 does not restrict the subspace but places a restriction on
 121 the nonlinearity.¹ By concentrating the data into a space of
 122 reduced dimensionality, the neural response can be fit with
 123 less restriction on the form of the nonlinearity.

124 A number of techniques have been developed to estimate
 125 the linear subspace and, subsequently, the nonlinearity. In
 126 the most widely used form of this analysis, the linear front
 127 end is limited to a single filter that serves as an explicit
 128 representation of the “receptive field” of the neuron, but the
 129 nonlinearity is essentially unrestricted. With the right
 130 choice of stimuli, this linear filter may be estimated by
 131 computing the spike-triggered average (STA) stimulus (i.e.,
 132 the mean stimulus that elicited a spike). The STA has been
 133 widely used in studying auditory neurons (e.g., Eggermont,
 134 Johannesma, & Aertsen, 1983). In the visual system, STA
 135 has been used to characterize retinal ganglion cells
 136 (e.g., Meister, Pine, & Baylor, 1994; Sakai & Naka,
 137 1987), lateral geniculate neurons (e.g., Reid & Alonzo,
 138 1995), and simple cells in primary visual cortex (V1;
 139 e.g., DeAngelis, Ohzawa, & Freeman, 1993; Jones &
 140 Palmer, 1987; McLean & Palmer, 1989). Given the STA
 141 filter, one typically has enough experimental data to
 142 construct a nonparametric estimate of the nonlinearity
 143 (i.e., a lookup table; Anzai, Ohzawa, & Freeman, 1999;
 144 Chichilnisky, 2001; deBoer & Kuyper, 1968; Eggermont
 145 et al., 1983). For some classes of nonlinearity, it has also
 146 been shown that one can write down a closed-form
 147 solution for the estimates of the linear filter and non-
 148 linearity in a single step (Nykamp & Ringach, 2002).

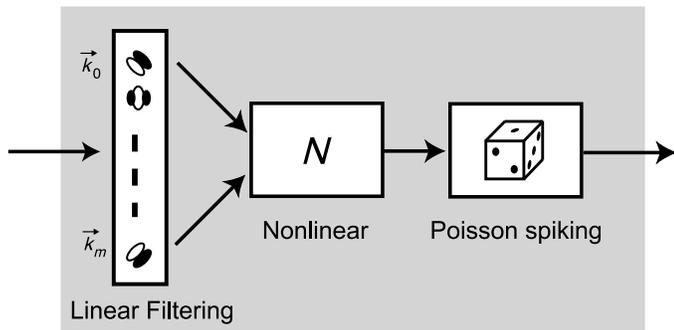


Figure 2. Block diagram of the LNP model. On each time step, the components of the stimulus vector are linearly filtered by $\vec{k}_0 \dots \vec{k}_m$. The responses of the linear filters are then passed through a nonlinear function $N(\cdot)$, whose output determines the instantaneous firing rate of a Poisson spike generator.

This methodology may be extended to the recovery of multiple filters (i.e., a low-dimensional subspace) and the nonlinear combination rule. One approach to finding a low-dimensional subspace is the spike-triggered covariance (STC; Bialek & de Ruyter van Steveninck, 2005; de Ruyter van Steveninck & Bialek, 1988). STC has been used to characterize multidimensional models and a nonlinear combination rule in systems ranging from the invertebrate motion system (Bialek & de Ruyter van Steveninck, 2005; Brenner, Bialek & de Ruyter van Steveninck, 2000; de Ruyter van Steveninck & Bialek, 1988) to songbird forebrain auditory neurons (Sen, Wright, Doupe, & Bialek, 2000) to the vertebrate retina cells (Pillow, Simoncelli, & Chichilnisky, 2003; Schwartz, Chichilnisky, & Simoncelli, 2002) and mammalian cortex (Horwitz, Chichilnisky, & Albright, 2005; Rust, Schwartz, Movshon, & Simoncelli, 2004, 2005; Touryan, Lau, & Dan, 2002). In addition, several authors have recently developed subspace estimation methods that use higher order statistical measures (Paninski, 2003; Sharpee, Rust, & Bialek, 2003, 2004). A review of spike-triggered subspace approaches may also be found in Ringach (2004) and Simoncelli, Pillow, Paninski, & Schwartz (2004).

Despite the theoretical elegance and experimental applicability of the subspace methods, there are a host of issues that an experimentalist is likely to confront when attempting to use them: How should one choose the stimulus space? How many spikes does one need to collect? How does one know if the recovered filters are significant? How should one interpret the filters? How do the filter responses relate to the nonlinear firing rate function? and so on. In this article, we describe the family of spike-triggered subspace methods in some detail, placing emphasis on practical experimental issues, and demonstrating these (where possible) with simulations. We focus our discussion on the STA and STC analyses, which have become quite

widely used experimentally. A software implementation of the methods described is available on the Internet at <http://www.cns.nyu.edu/~lcv/stc/>.

The linear–nonlinear Poisson (LNP) model

Experimental approaches to characterizing neurons are generally based on an underlying response model. Here, we assume a model constructed from a cascade of three operations:

1. a set of linear filters, $\{\vec{k}_1 \dots \vec{k}_m\}$,
2. a nonlinear transformation that maps the instantaneous responses of these filters to a scalar firing rate, and
3. a Poisson spike generation process, whose instantaneous firing rate is determined by the output of the nonlinear stage.

This LNP cascade is illustrated in Figure 2. The third stage, which essentially amounts to an assumption that the generation of spikes depends only on the recent stimulus (and not on the history of previous spike times), is often not stated explicitly but is critical to the analysis.

If we assume a discretized stimulus space, we can express the instantaneous firing rate of the model as:

$$r(t) = N(\vec{k}_1 \cdot \vec{s}(t), \vec{k}_2 \cdot \vec{s}(t), \dots, \vec{k}_m \cdot \vec{s}(t)), \quad (1)$$

where $\vec{s}(t)$ is a vector containing the stimuli over an appropriate temporal window preceding the time t . Here, the linear response of filter i (i.e., the projection or dot product of the filter \vec{k}_i with the stimuli $\vec{s}(t)$) is given by $\vec{k}_i \cdot \vec{s}(t)$. The nonlinear transformation $N(\cdot)$ operates over the linear filter responses.

Spike-triggered analysis

We aim to characterize the LNP model by analyzing the spike-triggered stimulus ensemble. The spike-triggered analysis techniques proceed as follows:

1. Estimate the low-dimensional linear subspace (set of filters). This effectively projects the high-dimensional stimulus into a low-dimensional subspace that the neuron cares about.
2. Compute the filter responses for the stimulus, and estimate the nonlinear firing rate function based on these responses. As noted earlier, typical physiological data sets allow nonparametric estimates of the nonlinearity for one or two filters but require more model restrictions as the number of filters increases.

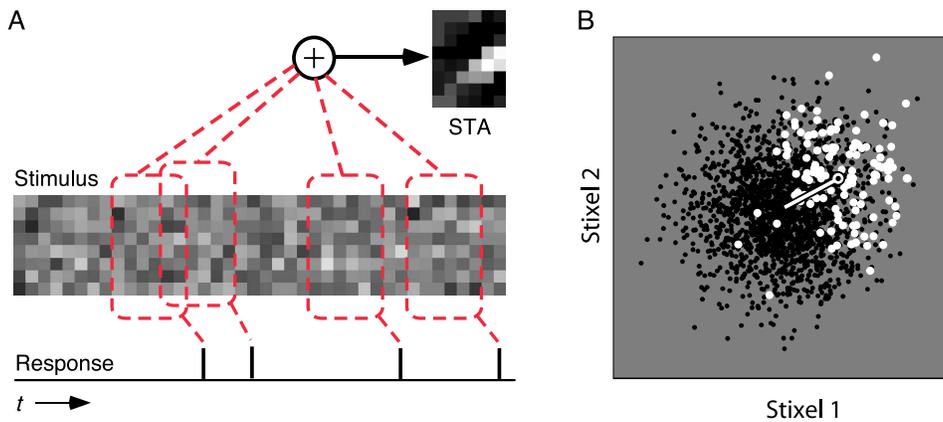


Figure 3. Two alternative illustrations of STA. (A) The STA is constructed by averaging the spike-triggered stimulus segments in red boxes (and subtracting off the average over the full set of stimulus segments). (B) Geometric (vector space) depiction of spike-triggered averaging in two dimensions. Black points indicate raw stimuli. White points indicate stimuli eliciting a spike. The STA, indicated by the line in the diagram, corresponds to the difference between the mean (center of mass) of the spike-triggered ensemble and the mean of the raw stimulus ensemble.

In the following subsections, we describe these steps in detail. In the [Experimental issues](#) section, we also stress the importance of an additional step: validating the resulting model by comparing it to neural responses from other stimuli.

Subspace (filter) estimation

In general, one can search for any deviation between the raw and spike-triggered stimulus ensembles. This can be done, for instance, using measures of information theory (Paninski, 2003; Sharpee et al., 2003, 2004). Another natural approach is to consider only changes in low-order moments between the raw and spike-triggered stimulus. Here, we focus on changes in the first and second moments, which may be computed efficiently and manipulated using a set of standard linear algebraic techniques. We also briefly discuss how the analysis relates to the Wiener/Volterra approach.

Spike-triggered average

The simplest deviation between the spike-triggered and raw stimulus distributions is a change in the mean. Assuming that the raw stimuli have zero mean, this can be estimated by computing the average of the spike-triggered ensemble (STA):

$$\hat{A} = \frac{1}{N} \sum_{n=1}^N \vec{s}(t_n), \quad (2)$$

where t_n is the time of the n th spike, $\vec{s}(t_n)$ is a vector representing the stimuli presented during the temporal window preceding that time, and N is the total number of spikes. In practice, the times t_n are binned. If there is more than one spike in a bin, then the stimulus vector for that

time bin is multiplied by the number of spikes that occurred. The STA is illustrated in [Figure 3A](#).

For an LNP model with a single linear filter, the STA provides an unbiased estimate of this filter,² provided that the input stimuli are spherically symmetric (Bussgang, 1952; Chichilnisky, 2001; Paninski, 2003), and the non-linearity of the model is such that it leads to a shift in the mean of the spike-triggered ensemble relative to the raw ensemble (see [Limitations and potential failures](#) section and [Experimental issues](#) section). This last requirement rules out, for example, a model with a symmetric nonlinearity such as full-wave rectification or squaring.

For an LNP model with multiple filters, the STA provides an estimate of a particular *linear combination* of the model filters, subject to the same restrictions on input stimuli and the form of the nonlinearity given above (Paninski, 2003; Schwartz et al., 2002). That is, the STA lies in the subspace spanned by the filters, but one cannot assume that it will exactly represent any particular filter in the model.

Spike-triggered covariance

The STA only recovers a single filter. Additional filters may be recovered seeking directions in the stimulus space in which the *variance* of the spike-triggered ensemble differs from that of the raw ensemble. Assuming that the raw stimuli have spherical covariance, this is achieved by computing the STC matrix:

$$\hat{C} = \frac{1}{N-1} \sum_{n=1}^N [\vec{s}(t_n) - \hat{A}] [\vec{s}(t_n) - \hat{A}]^T, \quad (3)$$

where the $[\cdot]^T$ indicates the transpose of the vector. Again, the t_n are binned in practice, and this means that each term should be multiplied by the number of spikes occurring in the associated time bin.

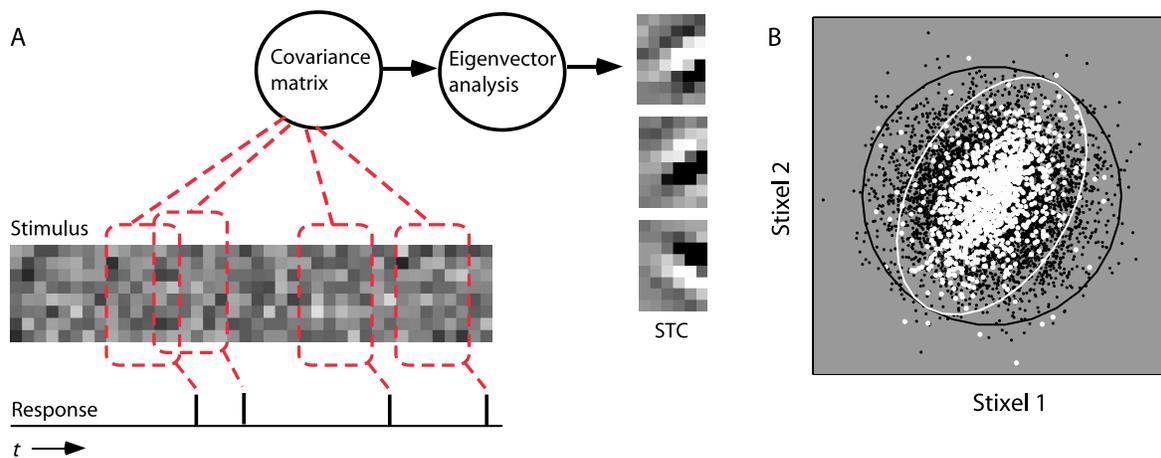


Figure 4. Two alternative illustrations of STC. (A) The STC is determined by constructing the covariance of the spike-triggered stimuli (relative to the raw stimuli), followed by an eigenvector analysis of the covariance matrix. This can result in multiple filters that represent directions in stimulus space for which the spike-triggered stimuli have lower or higher variance than the raw stimuli. (B) Geometric depiction of STC. Black points indicate raw stimuli. White points indicate stimuli eliciting a spike. Ellipses represent the covariance of each ensemble. Specifically, the distance from the origin to the ellipse along any particular direction is the standard deviation of the ensemble in that direction. Raw stimuli are distributed in a circular (Gaussian) fashion. Spike-triggered stimuli are elliptically distributed, with a reduced variance (relative to the raw stimuli) along the minor axis. The minor axis of the ellipse corresponds to a suppressive direction: Stimuli that have a large component along this direction (either positive or negative) are less likely to elicit a spike. The variance of the major axis of the ellipse matches that of the raw stimuli and, thus, corresponds to a direction in stimulus space that does not affect the neuron's firing rate.

295 The STC matrix embodies the multidimensional variance
 296 structure of the spike-triggered ensemble. Specifically, the
 297 variance of the ensemble in any direction specified by a unit
 298 vector, \hat{u} , is simply $\hat{u}^T \hat{C} \hat{u}$. The surface swept out by all
 299 such unit vectors scaled by the square root of their
 300 associated variance is a multidimensional ellipsoid. The
 301 principle axes of this ellipsoid, along with the associated
 302 variances, may be recovered as the eigenvectors and
 303 associated eigenvalues of the STC matrix. This is
 304 illustrated in Figure 4. The consistency of the STC
 305 estimate is guaranteed, provided that the input stimuli are
 306 Gaussian (Paninski, 2003) and the nonlinearity of the
 307 model is such that it leads to a change in the variance of
 308 the spike-triggered ensemble relative to the raw ensemble.
 309 Note that the Gaussianity is a more severe require-
 310 ment than the spherical symmetry required for STA
 311 analysis (see [Limitations and potential failures](#) section
 312 and [Experimental issues](#) section).

313 The STA and STC filters together form a low-dimen-
 314 sional linear subspace of the neural response. A number of
 315 groups have presented different approaches for combining
 316 the STA and STC analyses; in practice, these variants all
 317 converge to the same estimated subspace.³ Usually, the
 318 STA is subtracted prior to computing the STC filters
 319 (Brenner, Bialek & de Ruyter van Steveninck, 2000;
 320 de Ruyter van Steveninck & Bialek, 1988). It is often (but
 321 not always) the case that the STA will lie within the
 322 subspace spanned by the significant STC axes. Depending
 323 on the nonlinear properties of the response, it could

324 coincide with either high- or low-variance STC axes. To
 325 simplify visualization and interpretation of the axes, we
 326 have chosen for all of our examples to perform the STC
 327 analysis in a subspace orthogonal to the STA. Specifically,
 328 we compute STC on a set of stimuli from which the STA
 329 has been projected:

$$\vec{s} = \vec{s} - [\vec{s}^T \hat{A}] \hat{A} / |\hat{A}|^2. \quad (4)$$

330
 332

333 Comparison to Wiener/Volterra analysis

334 The STA provides an estimate of the first (linear) term in
 335 a polynomial series expansion of the system response
 336 function and, thus, is the first term of the Wiener/Volterra
 337 series. Whereas the Wiener/Volterra approach assumes that
 338 the nonlinearity is literally a polynomial, in the STA
 339 subspace approach, the nonlinearity is essentially unre-
 340 stricted. For nonlinearities such as a sigmoid, the Wiener/
 341 Volterra expansion would require many terms to capture
 342 the neural response function. An example of STA analysis
 343 for characterizing a model with a single filter and sigmoidal
 344 nonlinearity is presented in the model simulations below.

345 The second-order term in the Wiener series expansion
 346 describes the response as a weighted sum over all pairwise
 347 products of components in the stimulus vector. The weights
 348 of this sum (the second-order Wiener kernel) may be
 349 estimated from the STC matrix. However, the STC method

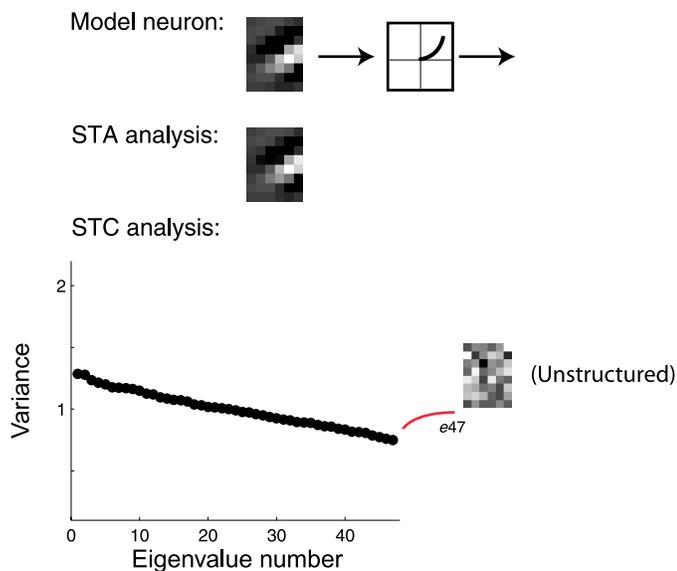


Figure 5. Eigenvalues and eigenvectors for an LNP model with a single linear filter followed by a point nonlinearity. The simulation is based on a sequence of 50,000 stimuli, with a response containing 1,891 spikes. Top: Model filter and nonlinearity. As in Figure 1, filters are 6×8 and, thus, live in a 48-dimensional space. The nonlinearity cartoon represents half squaring: Positive filter responses are squared, and negative filter responses are set to zero. Bottom: STA filter, and sorted eigenvalues of covariance matrix of stimuli eliciting spikes (STC). We plot the first 47 eigenvalues and omit the last eigenvalue, which is zero due to projecting out the STA (see Equation 4). The eigenvalues are gradually descending, and corresponding eigenvectors appear unstructured.

is *not* just a specific implementation of a second-order Wiener/Volterra model. The STC approach uses the STC matrix as a means to obtain a linear subspace, within which the nonlinearity is much less restricted. In contrast, the second-order Wiener/Volterra approach assumes a quadratic nonlinearity: This is suitable for characterizing nonlinearities such as the “energy model” (Adelson & Bergen, 1985) of complex cells in primary visual cortex (e.g., Emerson, Bergen, & Adelson, 1992; Emerson, Citron, Vaughn, & Klein, 1987; Szulborski & Palmer, 1990); however, it cannot describe response functions with nonlinearities such as divisive gain control (Albrecht & Geisler, 1991; Heeger, 1992) because these cannot be formulated as sums (or differences) of squared terms. An STA/STC approach is more flexible in capturing such nonlinearities (Rust, Schwartz, et al., 2005; Schwartz et al., 2002), as we demonstrate in the next section.

Simulations of example model neurons

We simulate an example ideal simple cell model, for which there is only a single filter, followed by half-wave

rectification and then squaring. Specifically, the instantaneous firing rate is determined by:

$$g(\vec{s}) = r [|\vec{k} \cdot \vec{s}|^2]. \quad (5)$$

The spike-triggered analysis results are shown in Figure 5. The spike-triggered ensemble exhibits a change in the mean relative to the raw stimulus ensemble due to the asymmetric nonlinearity. We recover the STA filter by computing the change in the mean (Equation 2). Next, we consider changes in the variance between the raw and spike-triggered stimulus ensemble. For this model neuron, there is no further relationship between the stimulus space and spikes. In the limit of infinite data, the spike-triggered ensemble would be a randomly selected subset of the raw stimulus ensemble, and the variance in any direction would be identical to that of the raw stimulus set. In an experimental setting, the finiteness of the spike-triggered ensemble produces random fluctuation of the variance in different directions. As a result, there are small random increases or decreases in variance of the spike-triggered ensemble relative to the raw stimulus set. This is reflected in the eigenvalue analysis of Figure 5. Due to the random fluctuations, the sorted eigenvalues cover a range around a constant value of 1 (i.e., the variance of the raw stimulus ensemble) but are not exactly equal to this constant value.

Now, consider an example model neuron, for which there is more than a single filter. We simulate an ideal V1 complex cell model (see also simulations in Sakai & Tanaka, 2000). The model is constructed from two space-time-oriented linear receptive fields, one symmetric and the other antisymmetric (Adelson & Bergen, 1985). The linear responses of these two filters are squared and summed, and the resulting signal then determines the instantaneous firing rate:

$$g(\vec{s}) = r [(\vec{k}_1 \cdot \vec{s})^2 + (\vec{k}_2 \cdot \vec{s})^2]. \quad (6)$$

Spike-triggered analysis on the model neuron is shown in Figure 6. The STA is close to zero. This occurs because for every stimulus, there is a stimulus of opposite polarity (corresponding to a vector on opposite sides of the origin) that is equally likely to elicit a spike, and thus, the average stimulus eliciting a spike will be zero. The recovered eigenvalues indicate that two directions within this space have substantially higher variance than the others. The eigenvectors associated with these two eigenvalues correspond to the two filters in the model (formally, they span the same subspace). In contrast, eigenvectors corresponding to eigenvalues in the gradually descending region appear arbitrary in their structure.

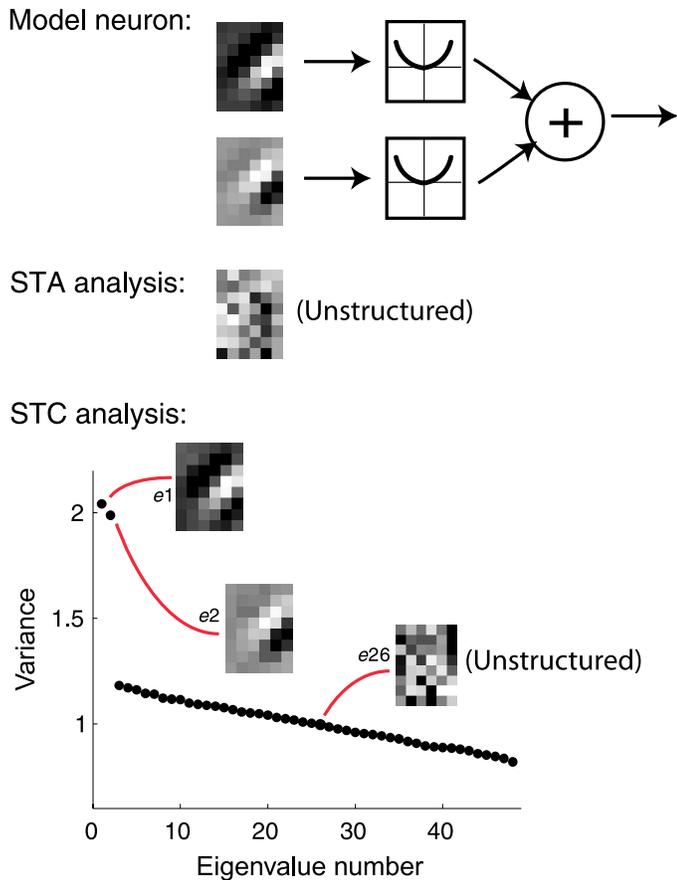


Figure 6. Eigenvalues and eigenvectors for an LNP ideal complex cell model. In this model, the Poisson spike generator is driven by the sum of squares of two oriented linear filter responses. As in Figure 1, filters are 6×8 and, thus, live in a 48-dimensional space. The simulation is based on a sequence of 50,000 raw stimuli, with a response containing 4,298 spikes. Top: Model, including two input filters, nonlinearities, and Poisson spiking. Bottom: STA filter is unstructured for the ideal complex cell. The plot also shows the eigenvalues, sorted in descending order. We plot the first 47 eigenvalues and omit the last eigenvalue which is zero due to projecting out the STA (see Equation 4). Two of the eigenvalues are substantially larger than the others and indicate the presence of two directions in the stimulus space along which the model responds. The others correspond to stimulus directions that the model ignores. Also shown are three example eigenvectors (6×8 linear filters), two of which are structured while one is unstructured.

Finally, we consider a version of a divisive gain control model (e.g., Geisler, 1992; Heeger, 1992):

$$g(\vec{s}) = r \frac{1 + [\vec{k}_1 \cdot \vec{s}]^2}{1 + (\vec{k}_2 \cdot \vec{s})^2 + .4(\vec{k}_3 \cdot \vec{s})^2}. \quad (7)$$

The analysis results are shown in Figure 7. First, we recover the STA filter, which is nonzero due to the half

squaring in the numerator. A nonsymmetrical nonlinearity of this sort is captured by changes in the mean. Next, we examine the sorted eigenvalues obtained from the STC analysis. Most of the eigenvalues descend gradually, but the last two eigenvalues lie significantly below the rest, and their associated eigenvalues span approximately the same subspace as the actual simulation filters.

Significance testing

How do we know if the recovered STA and STC filters are significant? In some cases, such as a prototypical complex cell in primary visual cortex, there is essentially no difference between the mean of the raw and spike-triggered stimuli (Rust, Schwartz, et al., 2005; Touryan et al., 2002), which leads to a weak STA. To quantify this, we test the hypothesis that the difference between the mean of the raw and spike-triggered stimulus is no different than what one would expect by chance. We specifically test whether the magnitude of the true spike-triggered stimulus STA is smaller or equal to what would be expected by chance. More specifically, we generate a distribution of random STA filters by bootstrapping: We randomly time-shift the spike train relative to the raw stimulus sequence, gather the resulting spike-triggered stimulus ensemble, and perform the STA analysis. The randomly time-shifted spike train retains all temporal structure that is present in the original spike train. We repeat this 1,000 times, each time computing the average of the stimulus subset. We can then set a significance criterion (e.g., the 95% confidence interval) within which we deem the magnitude of the true STA to be insignificant.

The issue of significance is also of importance for the STC filters. Although the low-variance eigenvalues are clearly below the gradually descending region in the illustrated example of Figure 7, the distinction is not so obvious in some experimental situations. An example in which the significance cutoff is not clear-cut is shown in Figure 8. A significance test should allow us to determine the number of eigenvector axes (filters) corresponding to significant increases or decreases in variance. That is, we would like to find changes in variance in the spike-triggered ensemble that are not just due to chance (because of the finiteness of the number of spikes) but that relate to actual neural response characteristics.

The significance testing must be done in a nested fashion because the distribution of the lowest and highest eigenvalues under the null hypothesis depends on the dimensionality of the space. We begin by assuming that none of the eigenvalues are significant. We compare the true eigenvalues to the eigenvalues of randomly selected stimuli with the same interspike interval. If the largest true eigenvalue lies outside the range of largest eigenvalues of the randomly shifted stimuli and if the smallest true eigenvalue lies outside the range of smallest eigenvalues of the randomly

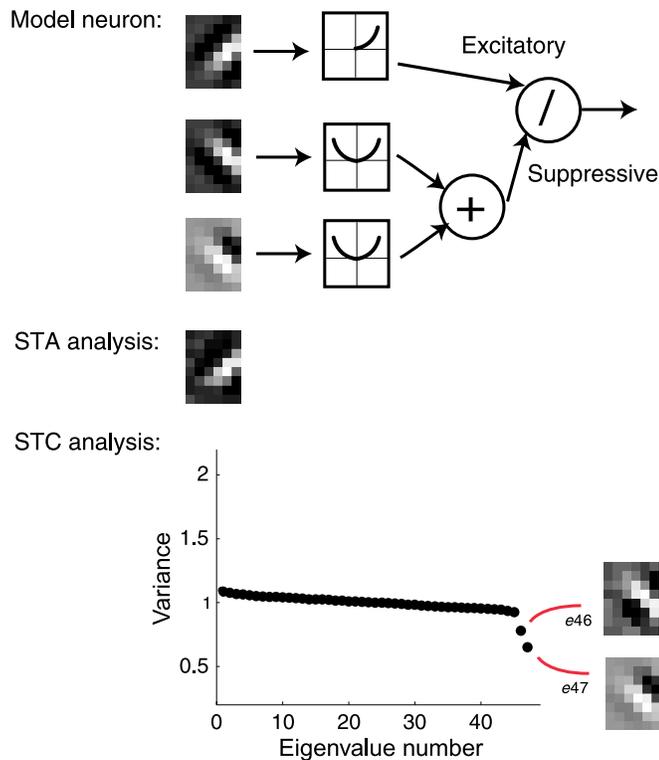


Figure 7. Eigenvalues and eigenvectors for an LNP divisive normalization model. The simulation is based on a sequence of 250,000 stimuli, with a response containing 30,444 spikes. Top: Model. Bottom: STA filter, sorted eigenvalues of covariance matrix of stimuli eliciting spikes (STC), and eigenvectors. Two of the eigenvalues are substantially lower than the others and indicate the presence of two suppressive directions in the stimulus space.

482 shifted stimuli, then we conclude that none of our axes are
 483 significant and accept the hypothesis. More specifically, to
 484 compute the randomly selected eigenvalues, we generate
 485 distributions of minimal/maximal eigenvalues by boot-

strapping: We randomly time-shift the spike train relative
 to the raw stimulus sequence, gather the resulting spike-
 triggered stimulus ensemble, perform the STA and STC
 analysis on the spike-triggered ensemble, and extract the
 minimum and maximum eigenvalues. After repeating
 1,000 times, we estimate the 95% confidence interval
 for both the largest and smallest eigenvalues. We then ask
 whether the maximal and minimal eigenvalues obtained
 from the true spike-triggered ensemble lie within this
 interval. If so, we accept the hypothesis.

Figure 8A shows that the hypothesis of no significant
 eigenvalues is unlikely to be correct for this example: The
 smallest eigenvalue lies far beyond the confidence
 interval. We therefore assume that the largest outlier
 (here, the smallest eigenvalue) has a corresponding axis
 that significantly affects the variance of the neural
 response. We thus proceed to test the hypothesis that all
 remaining axes are insignificant. To do so, we first project
 out the axis deemed significant and repeat the boot-
 strapping in the remaining subspace. Note that the
 distribution of eigenvalues (gray region in Figures 8A,
 8B, and 8C) changes as the dimensionality of the
 remaining space decreases. We continue this process in a
 nested fashion, until the largest and smallest eigenvalues
 from the true spike-triggered ensemble lie within the
 estimated confidence interval. Figure 8B shows that we
 cannot accept the hypothesis of two significant axes.
 Finally, the hypothesis of four significant axes (Figure 8B)
 is accepted and results in eigenvalues that lie within the
 confidence interval.

Filter estimation accuracy

Assuming that the recovered STA and STC filters are
 significant, we would also like to understand how accurate
 they are. The accuracy of our estimated filters depends on
 three quantities: (1) the dimensionality of the stimulus
 space, d ; (2) the number of spikes collected, N ; and (3) the

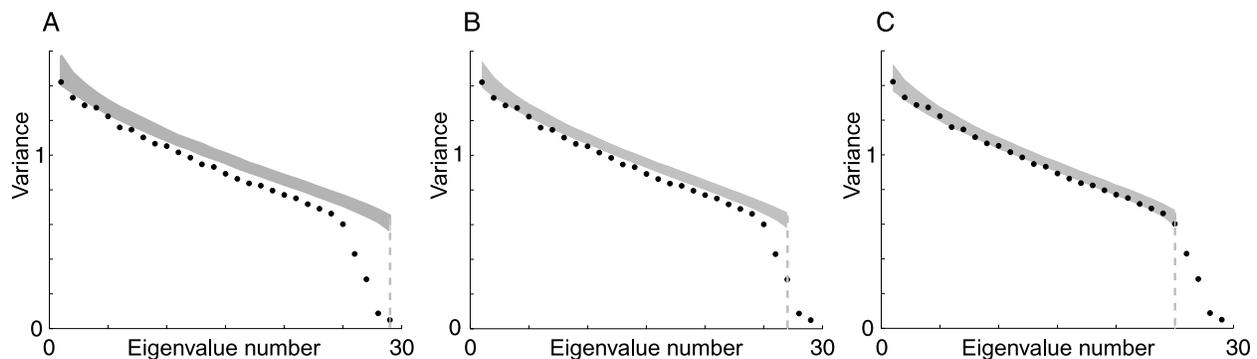


Figure 8. (A) Nested hypothesis testing. Gray solid line corresponds to 95% confidence interval, assuming no suppressive axes (B), two suppressive axes, and (C) four suppressive axes. If the hypothesis is accepted, eigenvalues should lie within the confidence interval. For the assumption of no or two suppressive axes, some eigenvalues lie below the confidence interval, indicating that the hypothesis is incorrect. In contrast, for the assumption of four suppressive axes, eigenvalues lie roughly within the confidence interval.

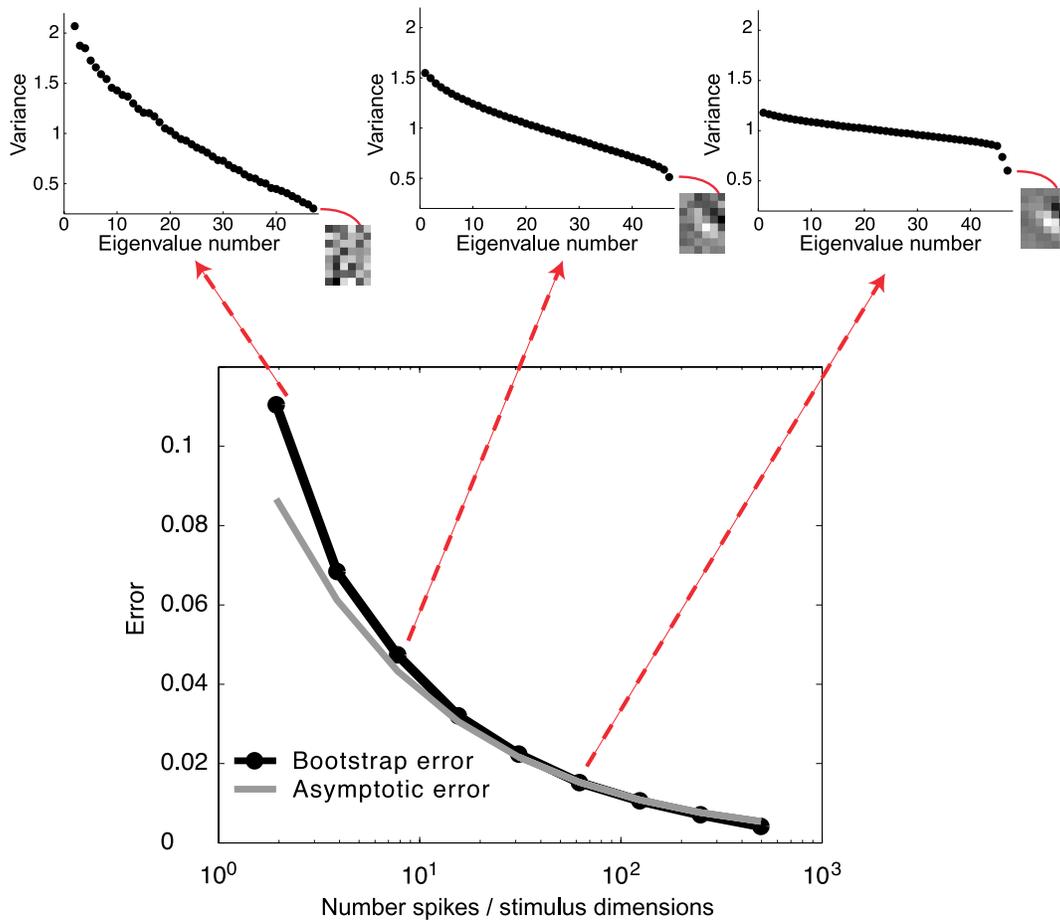


Figure 9. Accuracy in filter estimation. Simulations are shown for the divisive normalization example of Figure 7. Bottom: The error is computed as a function of the ratio of number of spikes to stimulus dimensionality. Stimulus dimensionality is held fixed for all simulations but a number of input stimuli (and thus spikes) are varied. Black line and points is the bootstrap-estimated error (mean angular error obtained from bootstrapping; see main text) of estimation of the lowest eigenvector. The gray line is the theoretical prediction of the mean angular error, computed as the square root of the stimulus dimensionality (here, 48) to number of spikes (see Equation 8 and Paninski, 2003). We multiply the theoretical prediction by a constant parameter that yields the least square error with the bootstrap-estimated error above for the last five points (because the theoretical prediction only holds for the small error regime).

523 strength of the response signal, relative to the standard
 524 deviation of the raw stimulus ensemble, σ .

525 Asymptotically, the errors decrease as (Paninski, 2003):

$$\text{MAE}(\vec{k}) = \frac{\sigma}{B(\vec{k})} \sqrt{\frac{d}{N}}, \quad (8)$$

526 where MAE indicates the mean of the angular error (the
 528 arccosine of the normalized dot product) between the
 529 estimated filter and the true filter and $B(\vec{k})$ is a
 530 proportionality factor that depends inversely on the
 531 strength of the response signal (Paninski, 2003). For
 532 example, the strength of response signal is the length of
 533 the STA vector in the limit of infinite data. An
 534 experimentalist does not have access to the strength of
 535 response signal. However, the number of spikes and
 536 number of stimulus dimensions are known, and thus, the
 537 function of Equation 8 may be used to extrapolate the

error behavior based on bootstrap estimates. To demon-
 538 strate this, we simulate an experiment on the model
 539 divisive normalization neuron.

540 We describe a bootstrap method to determine the error in
 541 filter estimation. We show that the bootstrap-estimated
 542 error is reasonably matched to the theoretical prediction of
 543 the error in Equation 8, when the ratio of number of spikes
 544 to number of stimulus dimensions is sufficiently high. We
 545 run a pilot experiment on the model divisive normal-
 546 ization neuron and collect 409,600 input samples. We
 547 consider how the ratio of stimulus dimensionality to
 548 number of spikes affects accuracy. Specifically, we hold
 549 the stimulus dimensionality fixed (which is 48 here) and
 550 vary the number of input samples (and thus spikes). For a
 551 given number of input samples, we bootstrap, drawing
 552 (with replacement) random subsets of stimuli (equal to the
 553 number of input samples). We consider the spike-triggered
 554 stimuli from this subset and compute the STA and STC.
 555

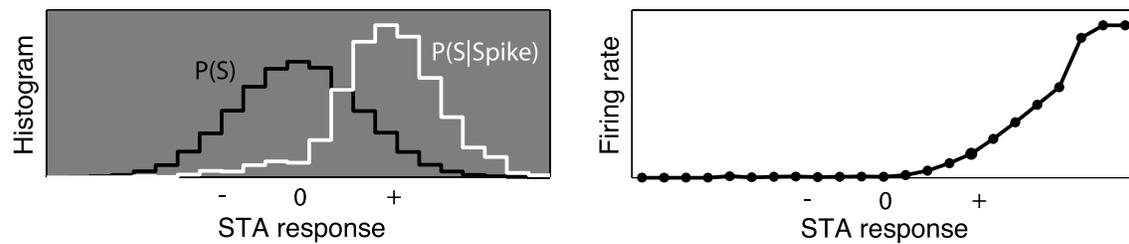


Figure 10. Nonlinearity for an LNP model with a single linear filter followed by a point nonlinearity. Left: Raw (black) and spike-triggered (white) histograms of the linear (STA) responses. Histograms have been renormalized to a maximal probability of 1. Right: The quotient of the spike-triggered and raw histograms gives an estimate of the nonlinearity that generates the firing rate.

556 We repeat this many times (here, 1,000) and derive an
 557 estimate of the mean angular error for a given STC filter.
 558 This is achieved by computing the mean of the 1,000
 559 estimated filters from the bootstrapping—we will denote
 560 this the mean estimated filter, and then, for each of the
 561 1,000 estimated filters, by computing its mean angular
 562 error with the mean estimated filter and taking an average
 563 over these computations. This analysis assumes that there
 564 are no systematic biases in the estimates (such as those
 565 shown in Figure 15).

566 In Figure 9, we plot the error estimates for the filter
 567 corresponding to the lowest eigenvalue. As the number of
 568 spikes to number of stimulus dimensions increases, the
 569 error is reduced. We also show, for three example ratios,
 570 the eigenvalues and the filter estimate corresponding to the
 571 lowest eigenvalue. For a low ratio of spike counts to
 572 stimulus dimensions, the eigenvalues descend gradually,
 573 and the smallest one is not separated from the rest; for a
 574 high ratio of spike counts to stimulus dimensions, the
 575 eigenvalues take on a pattern similar to Figure 7. Finally,
 576 we return to Equation 8: We fit this equation (and
 577 corresponding proportionality factor) to the errors derived
 578 from bootstrapping and obtain a rather good match for the
 579 low error regime. Such an analysis could be used in an
 580 experimental situation to determine data requirements for
 581 a given error level, by extrapolating the curve from values
 582 estimated from a pilot experiment. In the Experimental
 583 issues section, we elaborate on running a pilot experiment
 584 to choose a reasonable tradeoff between number of spikes
 585 and stimulus dimensionality.

587 Characterizing the nonlinearity

588 According to the LNP model, the firing rate of a neuron is
 589 given by a nonlinear transformation of the linear filter
 590 responses (Figure 2). Using the same set of stimuli and
 591 spike data as for the linear filter estimation, we seek to
 592 estimate the nonlinearity and, thus, characterize a neural
 593 model that specifies the full transformation from stimulus
 594 to neural firing rate. We therefore need to estimate the
 595 firing rate of the neuron as a function of the linear filter
 596 responses. To do so, it is important to recognize that the
 597 ratio of the frequency of occurrence of spike-triggered

stimuli to that of raw stimuli is proportional to the
 instantaneous firing rate. This can be seen using Bayes
 rule:

$$\mathcal{P}(\text{spike}|\vec{s}) = \frac{\mathcal{P}(\text{spike})\mathcal{P}(\vec{s}|\text{spike})}{\mathcal{P}(\vec{s})}, \quad (9)$$

and therefore, 602

$$\mathcal{P}(\text{spike}|\vec{s}) \propto \frac{\mathcal{P}(\vec{s}|\text{spike})}{\mathcal{P}(\vec{s})}, \quad (10)$$

where $\mathcal{P}(\text{spike}|\vec{s})$ is the instantaneous firing rate, $\mathcal{P}(\vec{s}|\text{spike})$
 is the frequency of occurrence of spike-triggered stimuli,
 and $\mathcal{P}(\vec{s})$ is the frequency of occurrence of raw stimuli.

The problem of estimating the nonlinearity can thus be
 described as one of estimating the ratio of two probability
 densities of Equation 10. The accuracy of the estimation is
 dependent on the dimensionality (number of filters) in the
 linear subspace. For one or two filters, we can use simple
 histograms to estimate the numerator and denominator of
 Equation 10. For more filters, this becomes impractical
 due to the so-called “curse of dimensionality”: The
 amount of data needed to sufficiently fill the histogram
 bins in a d -dimensional space grows exponentially with d .
 In this case, we typically need to incorporate additional
 assumptions about the form of the nonlinearity.

Consider a model LNP neuron with only a single filter
 followed by a point nonlinearity. First, we estimate the
 linear filter by computing the STA. Then, we compute the
 linear filter response for each stimulus, by taking a dot
 product of the filter with the stimulus. We do this for all
 instantiations of the spike-triggered stimuli and compute a
 histogram estimating the numerator density $\mathcal{P}(\vec{s}|\text{spike})$; we
 do this for all instantiations of the raw stimuli and
 compute a histogram estimating the denominator density
 $\mathcal{P}(\vec{s})$. The nonlinearity that determines the firing rate is
 then the ratio of these two densities or the ratio of the
 histogram values in each bin. An example is shown in
 Figure 10 (see also Chichilnisky, 2001). We plot the
 histograms of the spike-triggered and raw stimuli filter

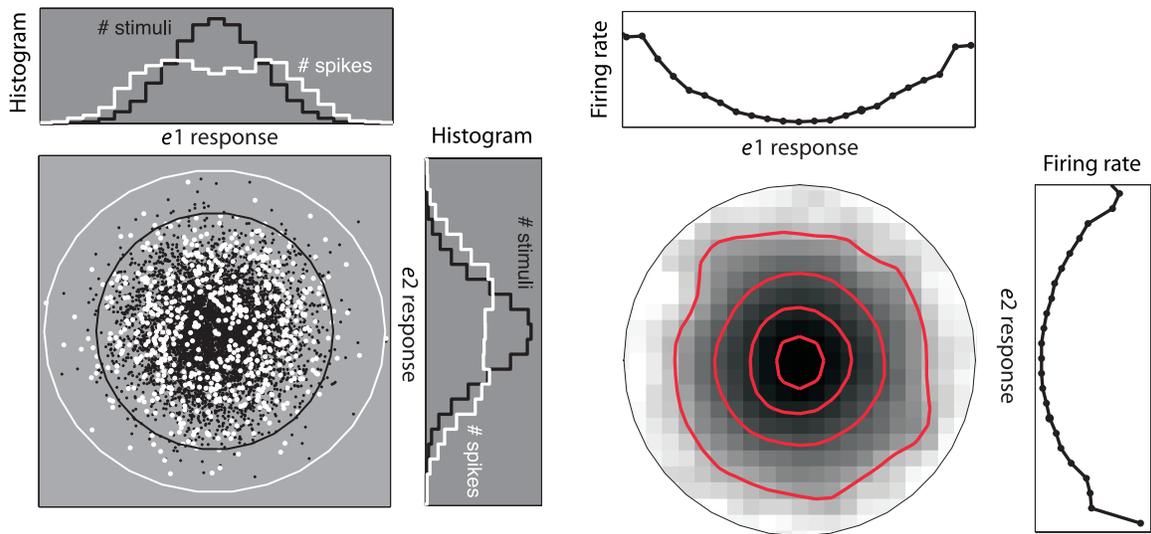


Figure 11. Nonlinearity for ideal complex cell model. This corresponds to eigenvalue and eigenvector example of Figure 6. Left: Scatter plots of stimuli projected onto estimated filters (i.e., filter responses) corresponding to first two eigenvalues (e_1 and e_2). Black points indicate the raw stimulus set. White points indicate stimuli eliciting a spike. Also shown are one-dimensional projections onto a single filter. Right: The quotient of the two-dimensional spike-triggered and raw histograms provides an estimate of the two-dimensional nonlinear firing rate function. This is shown as a circular-cropped grayscale image, where intensity is proportional to firing rate. Superimposed contours indicate four different response levels. Also shown are one-dimensional nonlinearities onto a single filter.

634 responses (Figure 10, left). We observe the nonlinearity by
 635 examining the ratio of these two histograms (Figure 10,
 636 right): The instantaneous firing rate grows monotonically
 637 and asymmetrically, that is, increases for stimuli to which
 638 the filter responds strongly and positively.

639 Note that the nonlinearity can be arbitrarily complicated
 640 (even discontinuous). The only constraint is that it must
 641 produce a change in the mean of the spike-triggered
 642 ensemble, as compared with the original stimulus ensemble.
 643 Thus, the interpretation of reverse correlation in the context
 644 of the LNP model is a significant departure from the Wiener/
 645 Volterra series expansion, in which even a simple sigmoidal
 646 nonlinearity would require the estimation of many terms for
 647 accurate characterization (Rieke et al., 1997).

648 Next, consider an ideal complex cell model neuron as in
 649 Equation 6. The recovered eigenvalues indicate that two
 650 directions within this space have substantially higher
 651 variance than the others (recall Figure 6). As before, we
 652 compute the raw and spike-triggered stimulus responses
 653 for each of the two filters. A two-dimensional scatter plot
 654 of these filter responses is shown in Figure 11 (left) for
 655 both the spike-triggered and raw stimuli. This is a two-
 656 dimensional depiction of samples from the numerator and
 657 denominator distributions in Equation 10. The scatter plots
 658 are similar in essence to those described in Figure 4, but
 659 the stimuli are projected onto the two filters recovered
 660 from the analysis. To estimate the two-dimensional non-
 661 linear firing rate function (Figure 11, right), we compute
 662 the two-dimensional histogram for the spike-triggered and
 663 raw stimuli responses and calculate the ratio of the
 664 histogram values in each bin. This is analogous to the

665 one-dimensional example shown in Figure 10. Similar
 666 pairs of excitatory axes and nonlinearities have been
 667 obtained from STC analysis of V1 cells in cat (Touryan
 668 et al., 2002) and monkey (Rust et al., 2004; Rust,
 669 Schwartz, et al., 2005).

670 Finally, consider the divisive normalization model in
 671 Equation 7, for which the eigenvalues and eigenvectors
 672 are shown in Figure 7. Figure 12 (left) shows a scatter plot
 673 of the STA filter response versus a suppressive filter
 674 response. The spiking stimuli lie within an ellipse, with
 675 the minor axis corresponding to the suppressive filter. This
 676 is exactly what we would expect in a suppressive system,
 677 such as that plotted in Figure 4. The two-dimensional
 678 nonlinearity is estimated by taking the quotient as before.
 679 This reveals an approximately saddle-shaped function,
 680 indicating the interaction between the excitatory and
 681 suppressive signals (Figure 12, right). Similar suppressive
 682 filters have been obtained from STC analysis of retinal
 683 ganglion cells (in both salamander and monkey; Schwartz
 684 et al., 2002) and simple and complex cells in monkey V1
 685 (Rust et al., 2004).

686 For some systems, such as H1 of the blowfly (Bialek & de
 687 Ruyter van Steveninck, 2005; Brenner, Bialek & de Ruyter
 688 van Steveninck, 2000), the dimensionality of STA and
 689 STC filters is sufficiently low (and the data set sufficiently
 690 large) to calculate the quotient of Equation 10 directly (as
 691 we have shown in the simulation examples) and thus
 692 estimate the nonlinearity. But what happens when there
 693 are more than two significant filters derived from the STA
 694 and STC analyses? There is not one single recipe; rather,
 695 there are a number of ways to try and approach this

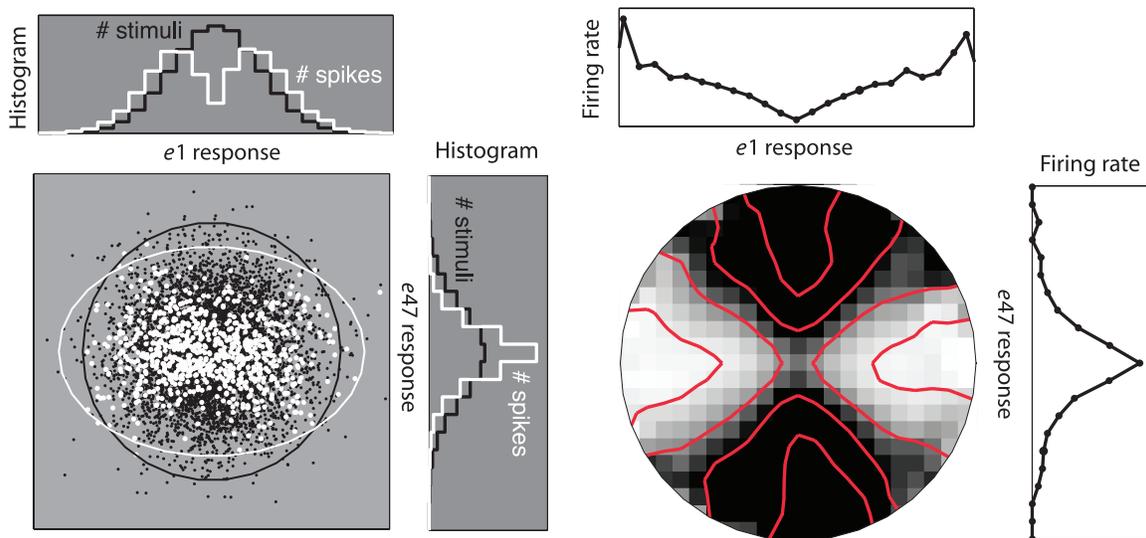


Figure 12. Nonlinearity for divisive normalization model. This corresponds to the eigenvalue and eigenvector example of Figure 7. Left: Scatter plots of stimuli projected onto estimated filters (i.e., filter responses) corresponding to STA and last suppressive eigenvector. Black points indicate the raw stimulus set. White points indicate stimuli eliciting a spike. Also shown are one-dimensional projections onto a single filter. Right: The quotient of the two-dimensional spike-triggered and raw histograms provides an estimate of the two-dimensional nonlinear firing rate function. This is shown as a circular-cropped grayscale image, where intensity is proportional to firing rate. Superimposed contours indicate four different response levels. Also shown are one-dimensional nonlinearities onto a single filter.

696 problem, and the answer depends on the particular system
697 and data at hand.

698 One approach is to consider specific classes of LNP
699 models that might be suitable for the particular neural area
700 under study. For instance, in retinal ganglion cell data, it
701 was shown that fitting a divisive normalization model to the
702 filters recovered from STA and STC provided a reasonable
703 characterization of the data (Schwartz et al., 2002). In
704 another study in area V1, the dimensionality of the filters
705 from STA and STC was too high for computing the
706 nonlinearity within the full recovered subspace (Rust,
707 Schwartz, et al., 2005). The form of nonlinearity was
708 restricted by first computing squared sums of excitatory
709 filter responses and squared sums of suppressive filter
710 responses, and only then was the nonlinearity between
711 these pooled excitatory and suppressive signals deter-
712 mined. This simplification could be made because it was
713 observed that projections of stimuli onto the recovered
714 filters within the excitatory or suppressive pools always
715 resulted in elliptical contours—suggesting sum of squares
716 operations governing the combination within each pool.
717 An alternative approach, published in this special issue,
718 assumes that the nonlinearity takes the form of a ratio of
719 Gaussians (Pillow & Simoncelli, *in press*).
721

722 Limitations and potential failures

724 The STA and STC estimates depend critically on the
725 distribution of input stimuli and on the particular non-

726 linearity of the neuron. For an LNP model with a single
727 linear filter, the consistency of the STA estimator is
728 guaranteed (e.g., irrespective of the neural nonlinearity)
729 only if the distribution of input stimuli are spherically
730 symmetric; that is, any two stimulus vectors with equal
731 vector length have an equal probability of being presented
732 (Chichilnisky, 2001). If one aims to recover a set of filters
733 using both STA and STC, then the consistency of the
734 estimator is guaranteed under the more stringent condition
735 that the stimuli be Gaussian distributed (Paninski, 2003).
736 The estimator is also guaranteed for elliptically symmetric
737 Gaussian stimuli, in which the covariance matrix is not
738 equal to the identity (see Appendix). For example, even if
739 the raw stimuli are constructed as spherical Gaussian, a
740 finite number of stimuli might, by chance, produce some
741 axes that have (slightly) higher variance than others. There
742 might also be interest in presenting to a neuron colored or
743 $1/f$ noise.

744 Note that non-Gaussian stimulus distributions can lead to
745 artifacts in the spike-triggered analysis, and the artifacts are
746 dependent on how the nonlinear response properties of the
747 neuron interact with the distribution. In the Experimental
748 issues section, we show simulated examples with non-
749 Gaussian stimuli, demonstrating how this could poten-
750 tially impact the STA and STC in a model neuron. These
751 examples do not indicate that experiments with non-
752 Gaussian stimuli and STA/STC analysis will necessarily
753 lead to artifacts, but because there is no general solution
754 for eliminating artifacts that can arise from non-Gaussian
755 stimuli, it is advisable to run experimental controls with
756 Gaussian stimuli.

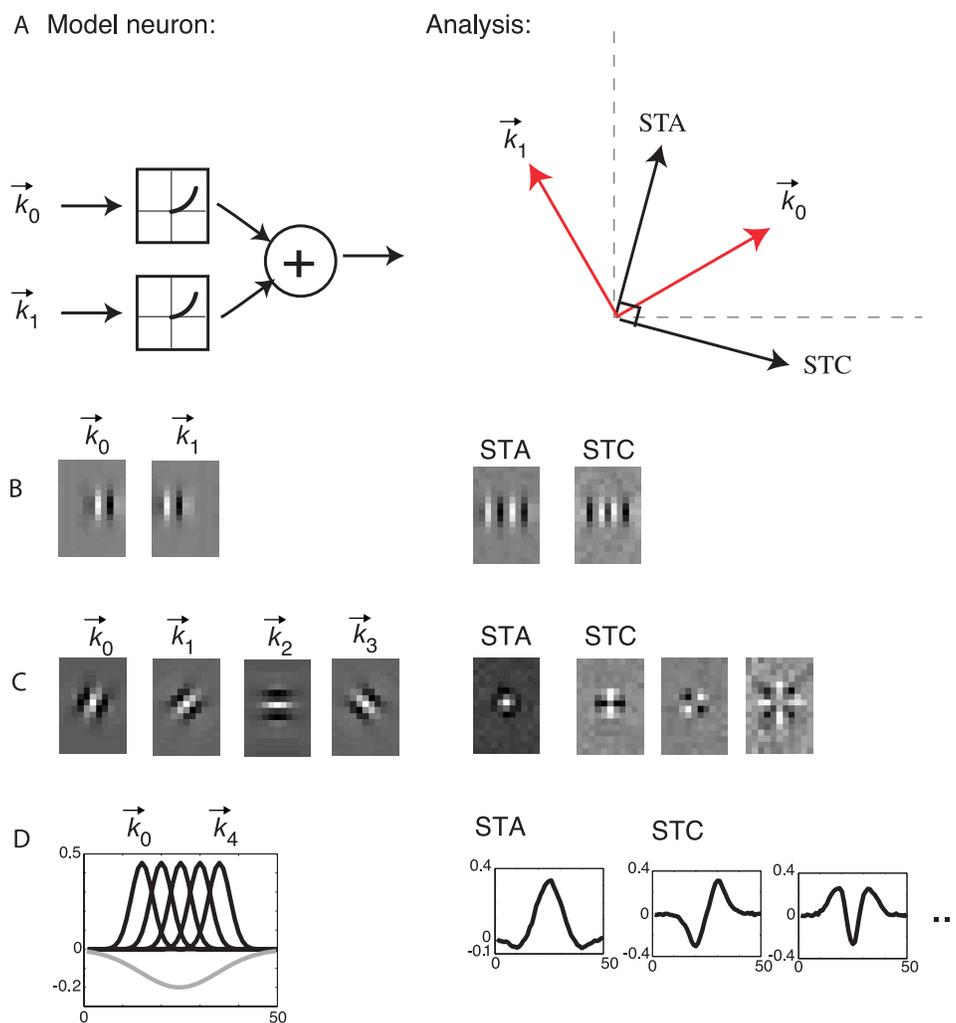


Figure 13. Interpretation issues and sum of half squares LNP model: filters. (A) Left: Model filter responses are half squared (negative values set to zero) and then added together. Note that this is different from the full squaring of the ideal complex cell. Right: Geometry of the STA and STC analysis. The STA is a vector average of the model filters. The STC is forced to be 90 deg away from the STA. Although the STA and STC filters do not equal the model filters, they do span the same subspace. (B) Example of spatially shifted model filters. Both STA and STC analysis reveal filters that are quite different from the model but span the same subspace. (C) Example of oriented filters. We extend the two-filter model to four filters that are each half squared and then added together. The STA is the average of all four filters and has a center/surround appearance rather than an oriented one. The other three STC filters are orthogonal. (D) The model neuron includes five spatially overlapping filters. The filter responses undergo a weighted sum of half squares, followed by addition of a (negative) linear surround (gray curve). The STA is a vector average of the linear filters, and the STC filters are orthogonal.

757 Even if one is careful to design an experiment and data
 758 analysis methodology that leads to accurate and artifact-
 759 free estimates, a spike-triggered analysis can still fail if the
 760 model assumptions are wrong. Two examples of failure of
 761 the LNP model are as follows: (1) there is no low-
 762 dimensional subspace in which the neural response may
 763 be described or (2) the neural response has a strong
 764 dependence on spike history (e.g., refractoriness, bursting,
 765 adaptation) that cannot be described by an inhomogeneous
 766 Poisson process. STA/STC analysis of data simulated using
 767 more realistic spike generation models, such as Hodgkin–
 768 Huxley (Agüera y Arcas & Fairhall, 2003; Agüera y Arcas,
 769 Fairhall, & Bialek, 2001, 2003; Pillow & Simoncelli,
 770 *in press*) and integrate-and-fire (Pillow & Simoncelli,

2003), produces biased estimates and artifactual filters.
 Although the STA/STC filters might in some cases still
 provide a reasonable description of a neuron's response, it
 is important to recognize that the LNP model provides
 only a crude approximation of the neural response (see
 Interpretation issues section).

Interpretation issues

There are a number of important issues that arise in
 interpreting the spike-triggered analysis. First, the number
 of filters recovered by STA and STC provides only a *lower*

771
 772
 773
 774
 775
 776
 777

778
 779

780
 781
 782

bound on the actual number of filters. The neural response may be dependent on mechanisms not identified by the STC analysis: (1) Other filters might affect the response, but the dependence is too weak and buried in the statistical error (a possibility with any experimental method—recall Figure 9); or (2) The neural response nonlinearities may not lead to a change in the mean or variance. It should be noted that although such a nonlinearity is theoretically possible, most known physiological nonlinearities *do* affect the mean, the variance, or both.

Next, the recovered filters cannot be taken literally as physiologically instantiated mechanisms. The STC filters, together with the STA, form an *orthogonal* basis for the stimulus subspace in which the responses are generated. The analysis does not yield a unique solution: A whole family of equivalent models can be constructed (by transforming to alternative sets of filters using an invertible linear transformation), which, given the same stimulus, produce the same response. Thus, even if a neuron's response is well described by an LNP model, we cannot claim to recover the actual filters that the neuron is using to compute its response. Rather, the goal is to find a set of filters that span the proper subspace; that is, with this set of filters, one can compute the same responses as with the actual set.

Figure 13 shows a simulation for an example of model neuron in which the STA and STC do not recover the actual model filters but do span the same subspace. The model neuron responds with a rate proportional to a sum of *half squares*, as opposed to the sum of squares typical of the ideal complex cell:

$$g(\vec{s}) = r[\vec{k}_1 \cdot \vec{s}]^2 + [\vec{k}_2 \cdot \vec{s}]^2.$$

The simulation results for different input filters are shown in Figure 13. Now, the STA does not result in a zero-weighted filter because the filter responses are not symmetric as in the ideal complex cell. Interestingly, the STA is not equal to either of the two excitatory filters of the model; rather, it is a vector average of the two filters. STC analysis on the stimuli perpendicular to the STA reveals an additional excitatory filter. Note that the two recovered filters together span the excitatory subspace of the original model filters. Figure 13C shows an example with four input filters of different orientations whose responses are half squared and summed; the STA takes on a more center-surround, unoriented appearance. Figure 13D shows an example of five overlapping *spatial* filters. These can be thought of as subunits, as has been proposed for retina (Hochstein & Shapley, 1976; see also Rust, Schwartz, et al., 2005 for cortical data). The nonlinear combination of these filters is followed by a subtraction of a linear surround. The resulting STA takes on the well-known spatial profile of retinal ganglion cells, and the STC filters are forced to be 90 deg apart and similar to

what is found experimentally (Pillow, Simoncelli, & Chichilnisky, 2004). The two-dimensional depiction of the nonlinearity for the above examples is interesting: The spike-triggered stimuli form a shape that resembles a portion of an annulus (Figure 14). Neurons with nonlinearities of this flavor can be seen in area V1 of the macaque (Rust, Schwartz, et al., 2005) and in retinal ganglion cells (Schwartz & Simoncelli, 2001).

Another reason why the recovered filters should not be interpreted as a physiological mechanism is that the LNP model assumes Poisson spiking. A number of authors have demonstrated that these Poisson assumptions do not accurately capture the statistics of neural spike trains (Berry & Meister, 1998; Keat, Reinagel, Reid, & Meister, 2001; Pillow, Shlens, Paninski, Chichilnisky, & Simoncelli, 2005a; Reich, Victor, & Knight, 1998). The dependence of neural responses on spike history (e.g., refractoriness, bursting, adaptation) may be captured only indirectly in the LNP model through time-delayed suppressive STC filters (Agüera y Arcas & Fairhall, 2003; Agüera y Arcas et al., 2003; Schwartz et al., 2002). For instance, during a refractory period, a neuron will not spike, and this can be captured by an LNP model with a set of suppressive STC filters in time. The suppressive filters may still provide a reasonably accurate description of the neural response but do not reveal the mechanism of refractoriness.

Finally, the labeling of whether a filter is excitatory or suppressive is crudely based on the net change in the mean or variance and may not correspond physiologically to excitation or suppression. A given filter can indeed be both excitatory and suppressive. For example, a filter might be half square rectified, yielding a positive increase in the mean, but also include a compressive squared nonlinearity (as in divisive normalization). Because the STA and STC filters are orthogonal, the analysis will extract a single filter and label it as excitatory. As before, the analysis still finds the right subspace; one can then analyze the interaction and aim to estimate a model within the subspace.

Experimental issues

We now discuss issues that arise when designing and interpreting spike-triggered experiments.

Stimulus choice

Stimulus space

The stimuli in a spike-triggered experiment need to be restricted to lie in a finite-dimensional space, and the experimentalist must choose the fundamental components (i.e., the axes) of this space. At any moment in time, the neuron is exposed to a linear combination of this set of stimulus components. In many published examples (as well

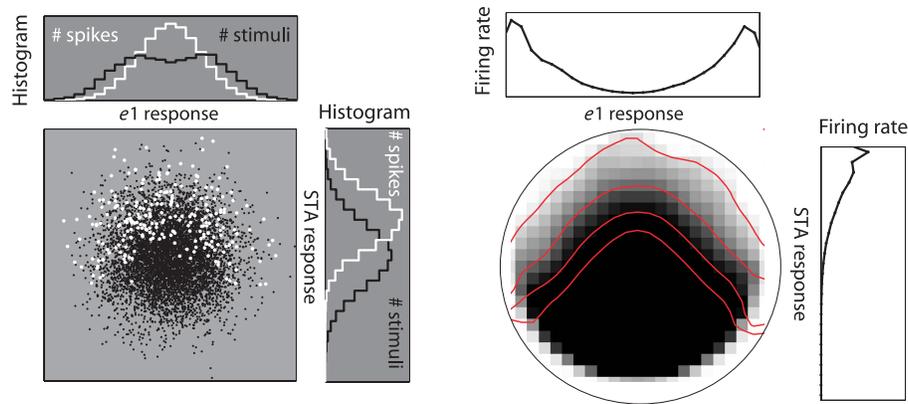


Figure 14. Interpretation issues and sum of half squares LNP model: nonlinearity. Nonlinearity is shown for model simulation of filters in Figure 13B (almost identical plots are found for Figures 13C and 13D). Left: Scatter plots of stimuli projected onto estimated filters (i.e., filter responses) corresponding to STA and first eigenvector. Black points indicate the raw stimulus set. White points indicate stimuli eliciting a spike. Also shown are one-dimensional projections onto a single filter. Right: The quotient of the two-dimensional spike-triggered and raw histograms provides an estimate of the two-dimensional nonlinear firing rate function. This is shown as a circular-cropped grayscale image, where intensity is proportional to firing rate. Superimposed contours indicate four different response levels. Also shown are one-dimensional nonlinearities onto a single filter.

890 as the examples shown in this article), the axes of the
 891 stimulus space corresponds to pixel (or stixel) intensities.
 892 However, the stimulus may be described in terms of other
 893 components, such as the amplitudes of a particular set
 894 of sinusoids (Ringach, Sapiro, & Shapley, 1997), the
 895 velocities of a randomly moving spatial pattern (Bair,
 896 Cavanaugh, & Movshon, 1997; Borghuis et al., 2003;
 897 Brenner, Bialek & de Ruyter van Steveninck 2000; de
 898 Ruyter van Steveninck & Bialek, 1988), or any other fixed
 899 set of functions. More generally, it is possible to do the
 900 analysis in a space that is a nonlinear function of the input
 901 stixels (David, Vinje, & Gallant, 2004; Nishimoto, Ishida,
 902 & Ohzawa, 2006; Theunissen et al., 2001). This is useful
 903 when one believes that the cells' response is LNP on these
 904 inputs (Rust, Simoncelli, & Movshon, 2005), although it
 905 may then be more difficult to interpret the results. The
 906 fundamental constraints on the choice of these compo-
 907 nents are that (1) the neuron should respond reasonably to
 908 stochastically presented combinations of these compo-
 909 nents and (2) the neuron's response should be well
 910 approximated by an LNP model operating in the space
 911 of these components.

912 The choice of a finite-dimensional stimulus space places
 913 a restriction on the generality of the experimental results:
 914 The response of the cell will only be characterized *within*
 915 *the subspace spanned by the stimulus components*
 916 (Ringach et al., 1997). Stated differently, without further
 917 assumptions, the model one constructs with STC can only
 918 predict stimuli responses that lie in the space defined by
 919 the experimental stimulus ensemble. For example, one
 920 cannot predict the responses to chromatic stimuli when
 921 using achromatic stimuli or to a full two-dimensional
 922 space when probing the neuron with only a single spatial
 923 dimension (as in the case of bars). Similarly, one cannot

924 use the model to predict responses to stimuli that have a
 925 finer spatial or temporal resolution than that used in the
 926 characterization.

927 To obtain a more general characterization, one needs to
 928 increase the stimulus resolution. Unfortunately, this
 929 increases the dimensionality of the stimulus space and,
 930 thus, requires more spikes to achieve the same quality of
 931 filter estimation. At the same time, the increase in
 932 resolution typically *reduces* the responsivity of the
 933 (e.g., because the effective contrast is reduced), thus
 934 making it more difficult to obtain the needed spikes.
 935 Recall that the error in filter estimation is a direct
 936 consequence of the ratio of the number of spikes to
 937 stimulus dimensionality, as in the example model neuron
 938 simulation shown in Figure 9. Therefore, it is useful to run
 939 a pilot experiment to determine the proper balance between
 940 number of spikes (e.g., duration of the experiment) to
 941 stimulus dimensionality for a particular class of neurons. In
 942 practice, it useful for a physiologist to adopt a rule of thumb
 943 for the particular system at hand: In the V1 experiments,
 944 Rust, Schwartz, et al. (2005) found that at least 100 spikes
 945 per dimension were typically needed to obtain a good
 946 characterization. Other experimental methodologies or
 947 settings (e.g., recordings from an awake behaving animal)
 948 and other classes of neurons may be more limited in the
 949 number of spikes that can be collected.
 950

951 **Stochastic stimulus distribution**

952 As stated earlier, the STC portion of the spike-triggered
 953 analysis is only guaranteed to work for Gaussian stimuli.
 954 The use of non-Gaussian white noise stimulus distributions
 955 (e.g., uniform, binary, sparse) is quite common experimen-
 956 tally, as the samples are easy to generate and the higher

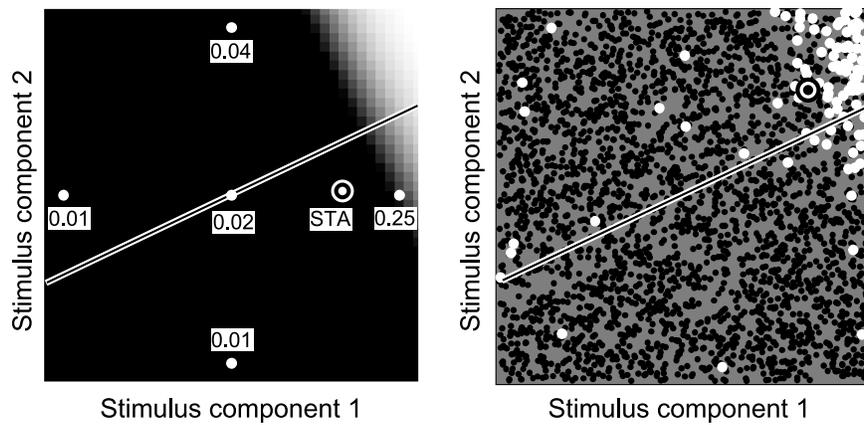


Figure 15. Simulations of an LNP model demonstrating bias in the STA for two different nonspherical stimulus distributions. The linear stage of the model neuron corresponds to an oblique axis (line in both panels), and the firing rate function is a sigmoidal nonlinearity (firing rate corresponds to intensity of the underlying grayscale image in the left panel). In both panels, the black and white “target” indicates the recovered STA. Left: Simulated response to sparse noise. The plot shows a 2-dimensional subspace of a 10-dimensional stimulus space. Each stimulus vector contains a single element with a value of ± 1 , whereas all other elements are zero. Numbers indicate the firing rate for each of the possible stimulus vectors. The STA is strongly biased toward the horizontal axis. Right: Simulated response of the same model to uniformly distributed noise. The STA is now biased toward the corner. Note that in both examples, the estimate will not converge to the correct answer, regardless of the amount of data collected.

957 contrast of the stimuli generally leads to higher average
 958 spike rates. In practice, their use is often justified by
 959 assuming that the linear filters are smooth relative to the
 960 pixel size/duration (e.g., Chichilnisky, 2001). Natural
 961 signal stimuli (such as visual scenes and auditory vocal-
 962 izations) are also non-Gaussian (Daugman, 1989; Field,
 963 1987), but their use is becoming increasingly popular
 964 (David & Gallant, 2005; David et al., 2004; Felsen & Dan,
 965 2005; Ringach, Hawken, & Shapley, 2002; Sen et al.,
 966 2000; Smyth, Willmore, Baker, Thompson, & Tolhurst,
 967 2003; Theunissen et al., 2001; for recent perspectives, see
 968 Felsen & Dan, 2005; Rust & Movshon, 2005). Natural
 969 signals can reveal response properties that occur less
 970 frequently under Gaussian white noise stimulation, such as
 971 bursting in the LGN (Lesica & Stanley, 2004), and they
 972 are often more effective in driving higher neural areas.

973 However, nonspherical stimuli can produce artifacts in
 974 the STA filters, and non-Gaussian stimuli can produce
 975 artifacts in the STC filters. Figure 15 shows two simu-
 976 lations of an LNP model with a single linear filter and a
 977 simple sigmoidal nonlinearity, each demonstrating that
 978 nonspherical stimulus distributions can lead to poor
 979 estimates of the linear stage. The examples are meant to
 980 emphasize the potential for bias but do not necessarily
 981 mean that an artifact will occur in experiment. Indeed, the
 982 particular nonlinear behaviors of the neural response will
 983 determine if and how much of a bias occurs. Because we
 984 do not know the nonlinearity a priori, the safest approach
 985 is to compare the experimental linear filter estimate to a
 986 control using spherically symmetric stimuli.

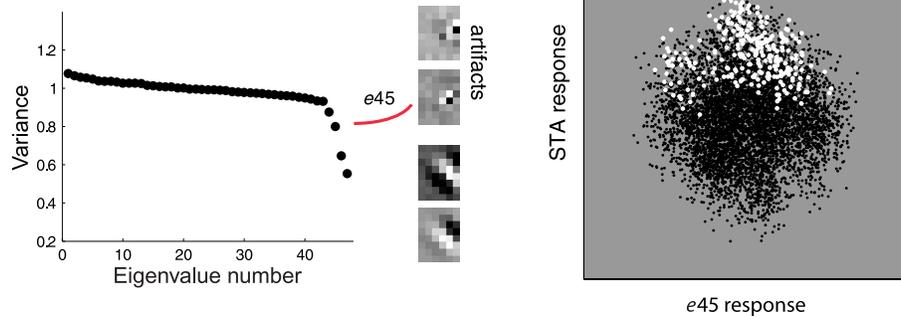
987 The first example shows a “sparse noise” experiment, in
 988 which the stimulus at each time step lies along one of the
 989 axes. As shown in the figure, the nonlinearity can result in

990 an STA that is biased toward an axis of the space. The
 991 second example uses stimuli in which each component is
 992 drawn from a uniform distribution, which produces an
 993 estimate biased toward the “corner” of the space. Note,
 994 however, that the estimate will be unbiased in the case of a
 995 purely linear neuron or of a half-wave-rectified linear
 996 neuron (Ringach et al., 1997).

997 Non-Gaussian stimuli can produce strong artifacts in the
 998 STC analysis. Figure 16A (left) shows an example
 999 simulation of the divisive normalization model with
 1000 binary stimuli. Note that in addition to the two “real”
 1001 suppressive filters of the model, the analysis also finds two
 1002 significant artifactual suppressive filters; these have a few
 1003 high-intensity stixels. Similar artifacts have been found in
 1004 experimental circumstances (Rust, Schwartz, et al., 2005).
 1005 More intuition for the artifacts can be gained by examin-
 1006 ing two-dimensional scatter plots that include an
 1007 artifactual filter response versus the STA filter response
 1008 (Figure 16A, right). The raw binary stimuli are clearly not
 1009 spherical in this two-dimensional view. Specifically, the
 1010 set tapers as one moves in the direction of the STA. This
 1011 reduction in variance of the raw stimulus happens to
 1012 coincide with the stimuli that elicit spikes (i.e., those that
 1013 have a large STA component). Thus, the spike-triggered
 1014 analysis reveals the artifactual filter as an axis of
 1015 significantly reduced variance, although it is actually not
 1016 reduced relative to the raw stimuli.

1017 There is, unfortunately, no generic recipe for reducing
 1018 artifacts. From our experience with binary stimuli, we have
 1019 found that the artifacts can be partially corrected by
 1020 adjusting the raw stimulus such that the covariance
 1021 estimated at each value of the STA is equal (conditional
 1022 whitening; Rust, Schwartz, et al., 2005). Specifically, we

A Binary stimuli simulation:



B Conditional whitening:

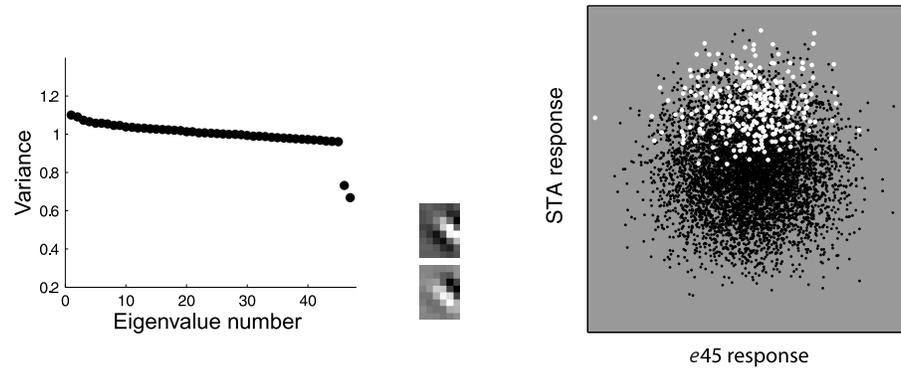


Figure 16. STC artifacts with binary stimuli. We ran the same model neuron as in Figure 7, but we replaced Gaussian stimuli with binary stimuli. (A) Left: There are four eigenvalues significantly below what one would expect by chance. Two of the corresponding eigenvectors correspond to the real model filter subspace, but two of them are artifactual. Right: Projection onto one of the artifactual filters versus the STA. The raw stimuli are nonspherical and have regions of lower variance at the top and bottom corners. The spiking stimuli appear in the upper corner because this is where the STA projection is largest. Although the variance of the raw and spike-triggered stimuli is the same when confined to this corner, the variance of the spike-triggered stimuli is significantly smaller than the variance of the entire raw ensemble, and this generates the artifactual suppressive filter (*e45*). (B) Left: After conditional whitening (see main text), there are only two significantly low eigenvalues corresponding to the model neuron subspace. Right: Projection onto the same eigenvalue as the artifactual filter above, as against the STA. The raw stimuli are now not perfectly circular but have roughly equal variance in all directions.

1023 partition the stimuli of Figure 16A (right) into horizontal
 1024 slabs according to the value of the excitatory filter
 1025 response and compute the covariance matrix for each
 1026 subset (C_n for the n th subset). The stimuli in each subset
 1027 are whitened by multiplying them by

$$E_e E_e^T + E_0 E_n D_n^{-\frac{1}{2}} E_n^T E_0, \quad (11)$$

1028 where E_e is a matrix containing the (orthogonal) excita-
 1029 tory filters (only one in this example—the STA), E_0
 1030 contains an orthogonal basis for the remainder of the
 1031 stimulus space, and E_n and D_n are the eigenvectors and
 1032 eigenvalues for the remainder of the conditional cova-
 1033 riance matrix, C_n , respectively. The first term serves to
 1034 preserve the component of the stimulus in the direction of
 1035 the STA, while the second term depicts a whitening (by
 1036 the inverse of the *raw* stimuli in that slice) in the other
 1037 dimensions.
 1038

After this conditional whitening, the stimuli are recom- 1039
 bined and STC analysis is applied on the spike-triggered set 1040
 to reestimate the filters. Figure 16B shows that following 1041
 conditional whitening, there are only two significant 1042
 suppressive eigenvalues corresponding to the real model 1043
 filter subspace. 1044

We have described an example of binary stimulus 1045
 artifacts and partially correcting for those artifacts. There 1046
 is generally no known fix for artifacts, but there are several 1047
 things that can be done to check for artifacts: 1048

1. It is helpful to examine the projection of the raw and 1049
 spike-triggered stimuli onto pairs of filters recovered 1050
 from the analysis; if these are not spherical, then the 1051
 filters can include artifacts. However, it is important to 1052
 remember that the stimulus space is huge, and 1053
 projection onto two dimensions might appear spheri- 1054
 cally symmetric but does not guarantee spherical 1055
 symmetry in the full space. 1056

1057 2. It is sometimes useful to run a model neuron
 1058 simulation with the given stimuli and see if artifactual
 1059 filters emerge. The simplest simulation one can run is
 1060 an LNP model with a single linear filter: If a
 1061 significant STC filter is found, this is indicative of an
 1062 artifactual axis in simulation. Here, we have demon-
 1063 strated a slightly more involved example of a divisive
 1064 normalization simulation. However, it is important to
 1065 realize that we have control only over the stimulus
 1066 ensemble; we have no control over the nonlinear
 1067 behaviors of the neural response, and the artifacts
 1068 depend on these nonlinearities. We can explore in
 1069 simulation nonlinearities that have been attributed to
 1070 neurons, and this has proved helpful in some cases.

1071 3. It is recommended to compare experimentally the
 1072 filter subspace recovered with a given stimulus
 1073 ensemble with that recovered with Gaussian
 1074 stimuli (recording from the same neuron); differ-
 1075 ences in the outcome between the two stimulus
 1076 types could indicate estimation biases or failures
 1077 of the model.

1078 Touryan, Felsen, and Dan (2005) compared STC
 1079 analysis in area V1 for white noise and natural images.
 1080 To partially correct for the natural image stimuli, they first
 1081 whitened the stimuli in the ensemble. Although this
 1082 cannot correct for the nonspherical nature of the stimuli,
 1083 they showed that the first two eigenvectors (representing
 1084 complex cells in their data) were similar for white noise
 1085 and natural images. The natural images required far fewer
 1086 raw stimuli to achieve the same result, probably because
 1087 they are more effective at eliciting spikes. They also found
 1088 additional significant (and artifactual) filters that were
 1089 compared with artifactual filters arising in a simulation
 1090 with natural images.

1091 Other techniques have been designed to cope directly
 1092 with non-Gaussian input, such as images, and thus bypass
 1093 this limitation of the STC approach. The basic idea is quite
 1094 simple: Instead of relying on a particular statistical moment
 1095 (e.g., mean or variance) for comparison of the spike-
 1096 triggered and raw stimulus distributions, one can use a more
 1097 general comparison function that can identify virtually any
 1098 difference between the two distributions. A natural choice
 1099 for such a function is information-theoretic: One can
 1100 compare the *mutual information* between a set of filter
 1101 responses and the probability of a spike occurring
 1102 (Paninski, 2003; Sharpee et al., 2003, 2004). This
 1103 approach is promising because it places essentially no
 1104 restriction on the stimulus ensemble. A drawback is that
 1105 the estimation problem is significantly more complicated;
 1106 it is more expensive to compute and may get trapped in
 1107 local optima. However, it has been successfully applied to
 1108 estimate one- or two-dimensional subspace models in
 1109 simulation and from physiological data in response to
 1110 natural images (Paninski, 2003; Sharpee et al., 2003, 2004,
 1111 2006). Other techniques, based on artificial neural net-
 1112 works (Lau, Stanley, & Dan, 2002; Lehky, Sejnowski, &

Desimone, 1992), have also been developed and applied to
 natural images (Prenger, Wu, David, & Gallant, 2004).

Validation

Validation is useful to evaluate the degree to which the
 recovered model is an accurate description of the neural
 response. At the very least, it is worthwhile verifying that the
 model, when fit to one run of white noise stimulation, can
 then predict responses to another run. Because the model is a
 rate model, this is most directly done by measuring
 responses to repeated stimuli and comparing their average
 (the PSTH) against that predicted from the model. Another
 possibility is to “play back” as stimuli the eigenvectors that
 were found in the spike-triggered analysis to verify that they
 affect the neuron’s response as expected (Rust, Schwartz,
 et al., 2005; Touryan et al., 2002). This requires that one
 perform the analysis and stimulus generation online during
 the experiment. Playing back the eigenvectors is also
 helpful for determining the importance of the individual
 model components that are recovered from the analysis; for
 example, the weakest components might have only a minor
 impact on the neural response.

It is also of interest to test how well the model generalizes
 to other stimuli: If one characterizes the model with a set of
 bars, how well does the model predict the response to a
 single bar? If one characterizes the model with high
 contrast stimuli, how well does it predict the response to
 low contrast stimuli? Ultimately, we would like a model
 that predicts the response to any arbitrary stimulus.
 Validating the model on different stimuli can help assess
 the robustness of the model and when it breaks, and, in turn,
 can identify the need for further improving spike-triggered
 analysis techniques.

Discussion

We have described a set of spike-triggered techniques for
 characterizing the functional response properties of neu-
 rons using stochastic stimuli. In general, there is a tradeoff
 between restricting the subspace dimensionality (as in the
 STA and STC approaches) versus restricting the non-
 linearity (as in the Wiener/Volterra approaches). Here, we
 have focused specifically on STA and STC analyses. These
 methods rely on an assumption that the response of the
 neuron is governed by an initial linear stage that serves to
 reduce the dimensionality of the stimulus space. The linear
 stage is followed by a nonlinearity upon which we place
 fairly minimal constraints. Having worked with these
 methods in both retina and V1, we have found that many
 experimental and analysis issues are quite tricky. We have
 presented examples with model neuron simulations, high-
 lighting similarities with experiments where possible.

1167 Estimation of the linear subspace can be corrupted by
 1168 three distinct sources of error, which we have discussed in
 1169 this article. First, there are errors due to the finiteness of the
 1170 data. The rate at which these decrease with increasing data
 1171 is given in Equation 8 and illustrated in Figure 9. Second,
 1172 there are biases that can arise from the interaction of the
 1173 neural nonlinearities and use of non-Gaussian stimuli.
 1174 Examples are shown in Figure 15. Finally, there are errors
 1175 due to model failure.

1176 There are a number of interesting directions for future
 1177 research. First, the LNP model can be extended to
 1178 incorporate some spike history dependence, by recursively
 1179 feeding back the spiking output into the linear input stage.
 1180 This “recursive LNP” model (also referred to as a general
 1181 linear model [GLM]) has appeared in recent literature
 1182 (Pillow, Paninski, Uzzell, Simoncelli, & Chichilnisky,
 1183 2005; Truccolo, Eden, Fellows, Donogue, & Brown,
 1184 2005) and may allow the introduction of some adaptation
 1185 effects, as well as shorter timescale effects such as
 1186 refractoriness, bursting, or rapid gain adjustments. This
 1187 model can no longer be directly fit to data with STA and
 1188 STC and requires more complex fitting procedures. In
 1189 addition, the techniques described here can be adjusted for
 1190 the analysis of multineuronal interactions (e.g., Nykamp,
 1191 2003; Okatan, Wilson, & Brown, 2005; Pillow, Shlens,
 1192 Paninski, Chichilnisky, & Simoncelli, 2005b). Such
 1193 methods have been applied, for example, in visual cortex
 1194 (Tsodyks, Kenet, Grinvald, & Arieli, 1999), motor cortex
 1195 (Paninski, Fellows, Shoham, Hatsopoulos, & Donoghue,
 1196 2004), and hippocampus (Harris, Csicsvari, Hirase,
 1197 Dragoi, & Buzsáki, 2003). Also, neurons adapt to stimuli
 1198 over multiple timescales (Brenner, Bialek & de Ruyter
 1199 van Steveninck, 2000; Fairhall, Lewen, Bialek, & de
 1200 Ruyter van Steveninck, 2001), and it would be interesting
 1201 to extend current approaches to incorporate adaptation.
 1202 Finally, it would be desirable to develop techniques that
 1203 can be applied to a cascaded series of LNP stages. This
 1204 will be essential for modeling responses in higher order
 1205 sensory areas, which are presumably constructed from
 1206 more peripheral responses. Specifically, if the afferent
 1207 responses that arrive in a particular neural area are
 1208 reasonably understood, then one may be able to arrange
 1209 to perform the spike-triggered analysis in the space of the
 1210 afferents (Rust, Simoncelli, et al., 2005).

1211 Appendix

1213 We describe how to compute STA and STC for
 1214 elliptically symmetric Gaussian stimuli. If the distribution
 1215 of stimuli is elliptically symmetric, then a modified
 1216 STA can be computed as follows (e.g., Theunissen et al.,
 1217 2001):

$$\hat{A}' = C^{-1}\hat{A}, \quad (\text{A1})$$

where

$$C = \sum_n \vec{s}(t_n)\vec{s}^T(t_n) \quad (\text{A2})$$

1220 is the covariance matrix of the raw stimuli such that $C =$
 1221 $V^T D V$ (we assume that the mean stimulus is zero). Note
 1222 that this solution is a regression estimate for a linear
 1223 mapping from stimuli to spikes. The surprising result is
 1224 that one can use linear regression on a one-dimensional
 1225 LN model if the input vectors are elliptically distributed.
 1226

1227 As in the case of STA, STC can be generalized to the case
 1228 of an elliptically symmetric stimulus distribution. Here, the
 1229 natural choice is to solve for stimulus dimensions in which
 1230 the *ratio* of variances of the spike-triggered and raw
 1231 stimulus ensembles is either large or small. Mathemati-
 1232 cally, we write this ratio in a direction specified by unit
 1233 vector \hat{u} as:

$$r(\hat{u}) = \frac{\hat{u}^T \hat{C} \hat{u}}{\hat{u}^T C \hat{u}}. \quad (\text{A3})$$

1234 The solution to this problem can be computed directly
 1235 using a generalized eigenvector analysis. Specifically, we
 1236 first solve for the whitening transform in the denominator,
 1237 computing the eigenvalues D and eigenvectors V of the
 1238 covariance matrix of the raw stimuli. We set
 1239 $X = V(\sqrt{D})^{-1}$ and $\hat{u} = X\hat{v}$, obtaining:
 1240

$$r(\hat{u}) = \frac{\hat{v}^T X^T \hat{C} V \hat{v}}{\hat{v}^T \hat{v}}. \quad (\text{A4})$$

1241 This is now equivalent to solving a standard eigenvector
 1242 problem, calculating the eigenvalues and eigenvectors of
 1243 $X^T \hat{C} X$.
 1244

Acknowledgments

1245 We are very grateful to E. J. Chichilnisky, Tony
 1246 Movshon, and Liam Paninski for helpful discussions.

1252 Commercial relationships: none.
 1253 Corresponding author: Odelia Schwartz.
 1254 Email: odelia@salk.edu.
 1255 Address: The Salk Institute, 10010 North Torrey Pines
 1256 Road, La Jolla, CA 92037, USA.

Footnotes

1257 ¹It should be noted that a Wiener/Volterra approach has
 1258 also been applied within a subspace, but under the
 1259 assumption of a low—order polynomial nonlinearity (e.g.,
 1260 Emerson et al., 1987, 1992; Szulborski & Palmer, 1990).

1263 ²Note that the STA estimate is unbiased but it does *not*, in
 1264 general, correspond to a maximum likelihood estimate
 1265 (Dayan & Abbott, 2001).
 1266 ³Note that recent work (Pillow & Simoncelli, 2006)
 1267 suggests an information—theoretic objective that com-
 1268 bines the STA and STC optimally.

References

- 1271 Adelson, E. H., & Bergen, J. R. (1985). Spatiotemporal
 1272 energy models for the perception of motion. *Journal*
 1273 *of the Optical Society of America A, Optics and*
 1274 *Image Science*, 2, 284–299. [PubMed]
- 1275 Agüera y Arcas, B., & Fairhall, A. L. (2003). What causes a
 1276 neuron to spike? *Neural Computation*, 15, 1789–1807.
 1277 [PubMed]
- A006 Agüera y Arcas, B., Fairhall, A. L., & Bialek, W. (2001).
 1279 What can a single neuron compute? In *Adv. Neural*
 1280 *Information Processing Systems (NIPS*00)* (Vol. 13,
 1281 pp. 75–81).
- 1282 Agüera y Arcas, B., Fairhall, A. L., & Bialek, W. (2003).
 1283 Computation in a single neuron: Hodgkin and Huxley
 1284 revisited. *Neural Computation*, 15, 1715–1749. [PubMed]
- 1285 Albrecht, D. G., & Geisler, W. S. (1991). Motion
 1286 sensitivity and the contrast-response function of
 1287 simple cells in the visual cortex. *Visual Neuroscience*,
 1288 7, 531–546. [PubMed]
- 1289 Anzai, A., Ohzawa, I., & Freeman, R. D. (1999). Neural
 1290 mechanisms for processing binocular information:
 1291 I. Simple cells. *Journal of Neurophysiology*, 82,
 1292 891–908. [PubMed] [Article]
- A007 Bair, W., Cavanaugh, J. R., & Movshon, J. A. (1997).
 1294 Reconstructing stimulus velocity from neuronal
 1295 responses in area MT. In M. C. Mozer, M. I. Jordan,
 1296 & T. Petsche (Eds.), *Adv. Neural Information Process-*
 1297 *ing Systems (NIPS*96)* (Vol. 9, pp. 34–40). MIT Press.
- 1298 Berry, M. J. II, & Meister, M. (1998). Refractoriness
 1299 and neural precision. *The Journal of Neuroscience*,
 1300 18, 2200–2211. [PubMed] [Article]
- 1301 Bialek, W., & de Ruyter van Steveninck, R. (2005).
 1302 Features and dimensions: Motion estimation in fly
 1303 vision. <http://arxiv.org/abs/q-bio/0505003>
- 1304 Borghuis, B. G., Perge, J. A., Vajda, I., van Wezel, R. J.,
 1305 van de Grind, W. A., & Lankheet, M. J. (2003). The
 1306 motion reverse correlation (MRC) method: A linear
 1307 systems approach in the motion domain. *Journal of*
 1308 *Neuroscience Methods*, 123, 153–166. [PubMed]
- 1309 Brenner, N., Bialek, W., & de Ruyter van Steveninck R.
 1310 (2000). Adaptive rescaling maximizes information
 1311 transmission. *Neuron*, 26, 695–702. [PubMed]
 1312 [Article]
- Bussgang, J. J. (1952). Cross-correlation functions of
 amplitude-distorted gaussian signals. *RLE Technical*
 Reports, 216. 1313 1315
- Chichilnisky, E. J. (2001). A simple white noise analysis
 of neuronal light responses. *Network*, 12, 199–213.
 [PubMed] 1317 1318
- Daugman, J. G. (1989). Entropy reduction and decorrela-
 tion in visual coding by oriented neural receptive
 fields. *IEEE Transactions on Bio-medical Engineer-*
 ing, 36, 107–114. [PubMed] 1319 1320 1321 1322
- David, S. V., & Gallant, J. L. (2005). Predicting neuronal
 responses during natural vision. *Network*, 16, 239–260.
 [PubMed] 1323 1324 1325
- David, S. V., Vinje, W. E., & Gallant, J. L. (2004).
 Natural stimulus statistics alter the receptive field
 structure of v1 neurons. *The Journal of Neuroscience*,
 24, 6991–7006. [PubMed] [Article] 1326 1327 1328 1329
- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuro-*
 science: *Computational and mathematical modeling*
 of neural systems. MIT Press. 1330 1331 1332
- DeAngelis, G. C., Ohzawa, I., & Freeman, R. D. (1993).
 Spatiotemporal organization of simple cell receptive
 fields in the cat's striate cortex: II. Linearity of
 temporal and spatial summation. *Journal of Neuro-*
 physiology, 69, 1118–1135. [PubMed] 1333 1334 1335 1336 1337
- deBoer, E., & Kuyper, P. (1968). Triggered correlation.
 In *IEEE Transactions on Bio-medical Engineering*
 (Vol. 15, pp. 169–179). 1338 1339 1340
- de Ruyter van Steveninck, R., & Bialek, W. (1988).
 Coding and information transfer in short spike
 sequences. In *Proceedings of the Royal Society*
 of London B, *Biological Sciences* (Vol. 234,
 pp. 379–414). 1341 1342 1343 1344 1345
- Eggermont, J. J., Johannesma, P. M., & Aertsen, A. M.
 (1983). Reverse-correlation methods in auditory
 research. *Quarterly Reviews of Biophysics*, 16,
 341–414. [PubMed] 1346 1347 1348 1349
- Emerson, R. C., Bergen, J. R., & Adelson, E. H. (1992).
 Directionally selective complex cells and the compu-
 tation of motion energy in cat visual cortex. *Vision*
 Research, 32, 203–218. [PubMed] 1350 1351 1352 1353
- Emerson, R. C., Citron, M. C., Vaughn, W. J., & Klein,
 S. A. (1987). Nonlinear directionally selective sub-
 units in complex cells of cat striate cortex. *Journal of*
Neurophysiology, 58, 33–65. [PubMed] 1354 1355 1356 1357
- Fairhall, A. L., Lewen, G. D., Bialek, W., & de Ruyter
 Van Steveninck, R. R. (2001). Efficiency and
 ambiguity in an adaptive neural code. *Nature*, 412,
 787–792. [PubMed] 1358 1359 1360 1361
- Felsen, G., & Dan, Y. (2005). A natural approach to
 studying vision. *Nature Neuroscience*, 8, 1643–1646.
 [PubMed] 1362 1363 1364

- 1365 Field, D. J. (1987). Relations between the statistics of
1366 natural images and the response properties of cortical
1367 cells. *Journal of the Optical Society of America A,*
1368 *Optics and Image Science, 4,* 2379–2394. [[PubMed](#)]
- 1369 Geisler, D. (1992). Two-tone suppression by a saturating
1370 feedback model of the cochlear partition. *Hearing*
1371 *Research, 63,* 203–211.
- 1372 Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G., &
1373 Buzsáki, G. (2003). Organization of cell assemblies in
1374 the hippocampus. *Nature, 424,* 552–556. [[PubMed](#)]
- 1375 Heeger, D. J. (1992). Normalization of cell responses in
1376 cat striate cortex. *Visual Neuroscience, 9,* 181–197.
1377 [[PubMed](#)]
- 1378 Hochstein, S., & Shapley, R. M. (1976). Linear and
1379 nonlinear spatial subunits in Y cat retinal ganglion
1380 cells. *The Journal of Physiology, 262,* 265–284.
1381 [[PubMed](#)] [[Article](#)]
- 1382 Horwitz, G. D., Chichilnisky, E. J., & Albright, T. D.
1383 (2005). Blue–yellow signals are enhanced by spatio-
1384 temporal luminance contrast in macaque V1. *Journal*
1385 *of Neurophysiology, 93,* 2263–2278. [[PubMed](#)]
1386 [[Article](#)]
- 1387 Jones, J. P., & Palmer, L. A. (1987). The two-dimensional
1388 spatial structure of simple receptive fields in cat striate
1389 cortex. *Journal of Neurophysiology, 58,* 1187–1211.
1390 [[PubMed](#)]
- 1391 Keat, J., Reinagel, P., Reid, R. C., & Meister, M. (2001).
1392 Predicting every spike: A model for the responses of
1393 visual neurons. *Neuron, 30,* 803–817. [[PubMed](#)]
1394 [[Article](#)]
- 1395 Korenberg, M. J., Sakai, H. M., & Naka, K. (1989).
1396 Dissection of the neuron network in the catfish inner
1397 retina: III. Interpretation of spike kernels. *Journal of*
1398 *Neurophysiology, 61,* 1110–1120. [[PubMed](#)]
- 1399 Lau, B., Stanley, G. B., & Dan, Y. (2002). Computational
1400 subunits of visual cortical neurons revealed by
1401 artificial neural networks. *Proceedings of the*
1402 *National Academy of Sciences of the United States*
1403 *of America, 99,* 8974–8979. [[PubMed](#)] [[Article](#)]
- 1404 Lehky, S. R., Sejnowski, T. J., & Desimone, R. (1992).
1405 Predicting responses of nonlinear neurons in monkey
1406 striate cortex to complex patterns. *Journal of Neuro-*
1407 *science, 12,* 3568–3581. [[PubMed](#)] [[Article](#)]
- 1408 Lesica, N. A., & Stanley, G. B. (2004). Encoding of
1409 natural scene movies by tonic and burst spikes in the
1410 lateral geniculate nucleus. *The Journal of Neuro-*
1411 *science, 24,* 10731–10740. [[PubMed](#)] [[Article](#)]
- 1412 Marmarelis, P. Z., & Marmarelis, V. Z. (1978). *Analysis*
1413 *of physiological systems: The white noise approach.*
1414 London: Plenum Press.
- 1415 McLean, J., & Palmer, L. A. (1989). Contribution of
1416 linear spatiotemporal receptive field structure to
velocity selectivity of simple cells in area 17 of cat. 1417
Vision Research, 29, 675–679. [[PubMed](#)] 1418
- Meister, M., Pine, J., & Baylor, D. A. (1994). Multi-neuronal 1419
signals from the retina: Acquisition and analysis. 1420
Journal of Neuroscience Methods, 51, 95–106. 1421
[[PubMed](#)] 1422
- Nishimoto, S., Ishida, T., & Ohzawa, I. (2006). Receptive 1423
field properties of neurons in the early visual cortex 1424
revealed by local spectral reverse correlation. *The* 1425
Journal of Neuroscience, 26, 3269–3280. [[PubMed](#)] 1426
- Nykamp, D. (2003). White noise analysis of coupled 1427
linear-nonlinear systems. *SIAM Journal on Applied* 1428
Mathematics, 63, 1208–1230. 1429
- Nykamp, D. Q., & Ringach, D. L. (2002). Full identi- 1430
fication of a linear-nonlinear system via cross- 1431
correlation analysis. *Journal of Vision, 2*(1), 1–11, 1432
<http://journalofvision.org/2/1/1/>, doi:10.1167/2.1.1. 1433
[[PubMed](#)] [[Article](#)] 1434
- Okatan, M., Wilson, M. A., & Brown, E. N. (2005). 1435
Analyzing functional connectivity using a network 1436
likelihood model of ensemble neural spiking activity. 1437
Neural Computation, 17, 1927–1961. [[PubMed](#)] 1438
- Paninski, L. (2003). Convergence properties of some 1439
spike-triggered analysis techniques. *Network, 14,* 1440
437–464. [[PubMed](#)] 1441
- Paninski, L., Fellows, M., Shoham, S., Hatsopoulos, N., & 1442
Donoghue, J. (2004). Nonlinear population models 1443
for the encoding of dynamic hand position signals in 1444
MI. *Neurocomputing*. Paper presented at Computa- 1445
tional Neuroscience, Alicante, Spain, July 2003. 1446
- Pillow, J. W., Paninski, L., Uzzell, V. J., Simoncelli, 1447
E. P., & Chichilnisky, E. J. (2005). Prediction and 1448
decoding of retinal ganglion cell responses with a 1449
probabilistic spiking model. *The Journal of Neuro-* 1450
science, 25, 11003–11013. [[PubMed](#)] [[Article](#)] 1451
- Pillow, J. W., Shlens, J., Paninski, L., Chichilnisky, E. J., & 1452
Simoncelli, E. P. (2005a). Modeling multineuronal 1453
responses in primate retinal ganglion cells. In *Compu-* 1454
tational and Systems Neuroscience (CoSyNe), Salt 1455
Lake City, UT. 1456
- Pillow, J. W., Shlens, J., Paninski, L., Chichilnisky, E. J., & 1457
Simoncelli, E. P. (2005b). Modeling the correlated 1458
spike responses of a cluster of primate retinal ganglion 1459
cells. In *Annual Meeting, Neurosciences*, Washington, 1460
DC. 1461
- Pillow, J. W., & Simoncelli, E. P. (2003). Biases in white 1462
noise analysis due to non-Poisson spike generation. In 1463
Neurocomputing (Vol. 52–54, pp. 109–115). Elsevier. 1464
Paper presented at Computational NeuroScience 1465
(CNS*02), Chicago, 21–25 July 2002. 1466
- Pillow, J. W., & Simoncelli, E. P. (2006). Dimensionality 1467
reduction in neural models: An information-theoretic 1468

- 1469 generalization of spike-triggered average and covari- 1521
 1470 ance analysis. *Journal of Vision* (in press). 1522
- 1471 Pillow, J. W., Simoncelli, E. P., & Chichilnisky, E. J. 1523
 1472 (2003). Characterization of nonlinear spatiotemporal 1524
 1473 properties of macaque retinal ganglion cells using 1525
 1474 spike-triggered covariance. In *Annual Meeting, Neu-* 1526
 1475 *rosciences*, New Orleans.
- 1476 Pillow, J. W., Simoncelli, E. P., & Chichilnisky, E. J. 1527
 1477 (2004). Characterization of macaque retinal ganglion 1528
 1478 cell responses using spike-triggered covariance. In 1529
 1479 *Computational and Systems Neuroscience (CoSyNe)*, 1530
 1480 Cold Spring Harbor, NY. 1531
- 1481 Prenger, R., Wu, M. C., David, S. V., & Gallant, J. L. 1532
 1482 (2004). Nonlinear V1 responses to natural scenes 1533
 1483 revealed by neural network analysis. *Neural Networks,* 1534
 1484 *17*, 663–679. [PubMed] 1535
- 1485 Reich, D. S., Victor, J. D., & Knight, B. W. (1998). The 1536
 1486 power ratio and the interval map: Spiking models and 1537
 1487 extracellular recordings. *The Journal of Neuro-* 1538
 1488 *science*, *18*, 10090–10104. [PubMed] [Article] 1539
- 1489 Reid, R. C., & Alonso, J. M. (1995). Specificity of 1540
 1490 monosynaptic connections from thalamus to visual 1541
 1491 cortex. *Nature*, *378*, 281–284. [PubMed] 1542
- 1492 Rieke, F., Warland, D., de Ruyter van Steveninck, R. R., & 1543
 1493 Bialek, W. (1997). *Spikes: Exploring the neural code.* 1544
 1494 Cambridge, MA: MIT Press. 1545
- 1495 Ringach, D. L. (2004). Mapping receptive fields in 1546
 1496 primary visual cortex. *The Journal of Physiology,* 1547
 1497 *558*, 717–728. [PubMed] [Article] 1548
- 1498 Ringach, D. L., Hawken, M. J., & Shapley, R. (2002). 1549
 1499 Receptive field structure of neurons in monkey primary 1550
 1500 visual cortex revealed by stimulation with natural 1551
 1501 image sequences. *Journal of Vision*, *2*(1), 12–24, 1552
 1502 <http://journalofvision.org/2/1/2/>, doi:10.1167/2.1.2. 1553
 1503 [PubMed] [Article] 1554
- 1504 Ringach, D. L., Sapiro, G., & Shapley, R. (1997). A subspace 1555
 1505 reverse-correlation technique for the study of visual 1556
 1506 neurons. *Vision Research*, *37*, 2455–2464. [PubMed] 1557
- 1507 Rust, N. C., & Movshon, J. A. (2005). In praise of artifice. 1558
 1508 *Nature Neuroscience*, *8*, 1647–1650. [PubMed] 1559
- 1509 Rust, N. C., Schwartz, O., Movshon, J. A., & Simoncelli, 1560
 1510 E. P. (2004). Spike-triggered characterization of 1561
 1511 excitatory and suppressive stimulus dimensions in 1562
 1512 monkey V1. In *Neurocomputing*. Elsevier. Paper 1563
 1513 presented at Computational NeuroScience 1564
 1514 (CNS*03), Alicante Spain, July 2003. 1565
- 1515 Rust, N. C., Schwartz, O., Movshon, J. A., & Simoncelli, E. P. 1566
 1516 (2005). Spatiotemporal elements of macaque V1 recep- 1567
 1517 tive fields. *Neuron*, *46*, 945–956. [PubMed] 1568
- 1518 Rust, N. C., Simoncelli, E. P., & Movshon, J. A. (2005). 1569
 1519 How macaque MT cells compute pattern motion. In 1570
 1520 *Annual Meeting, Neurosciences*, Washington, DC. 1571
- Sakai, H. M., & Naka, K. (1987). Signal transmission in 1572
 the catfish retina: V. Sensitivity and circuit. *Journal*
of Neurophysiology, *58*, 1329–1350. [PubMed] 1573
- Sakai, K., & Tanaka, S. (2000). Spatial pooling in the 1574
 second-order spatial structure of cortical complex 1575
 cells. *Vision Research*, *40*, 855–871. [PubMed] 1576
- Schwartz, O., Chichilnisky, E. J., & Simoncelli, E. P. 1577
 (2002). Characterizing neural gain control using 1578
 spike-triggered covariance. In T. G. Dietterich, S. 1579
 Becker, & Z. Ghahramani (Eds.), *Adv. Neural* 1580
*Information Processing Systems (NIPS*01)* (Vol. 14, 1581
 pp. 269–276). Cambridge, MA: MIT Press. 1582
- Schwartz, O., & Simoncelli, E. P. (2001). A spike- 1583
 triggered covariance method for characterizing divi- 1584
 sive normalization models. In *First Annual Meeting,* 1585
Vision Sciences Society, Sarasota. 1586
- Sen, K., Wright, B. D., Doupe, A. J., & Bialek, W. (2000). 1587
 Discovering features of natural stimuli relevant to 1588
 forebrain auditory neurons. In *Society for Neuro-* 1589
science Abstracts. 1590
- Sharpee, T., Rust, N. C., & Bialek, W. (2003). Max- 1591
 imizing informative dimensions: Analyzing neural 1592
 responses to natural signals. In S. Becker, S. Thrun, & 1593
 K. Obermayer (Eds.), *Adv. Neural Information* 1594
*Processing Systems (NIPS*02)* (Vol. 15). Cambridge, 1595
 MA: MIT Press. 1596
- Sharpee, T., Rust, N. C., & Bialek, W. (2004). Analyzing 1597
 neural responses to natural signals: Maximally informa- 1598
 tive dimensions. *Neural Computation*, *16*, 223–250. 1599
 [PubMed] 1600
- Sharpee, T. O., Sugihara, H., Kurgansky, A. V., Rebrik, 1601
 S. P., Stryker, M. P., & Miller, K. D. (2006). 1602
 Adaptive filtering enhances information transmission 1603
 in visual cortex. *Nature*, *439*, 936–942. [PubMed] 1604
- Simoncelli, E. P., Pillow, J., Paninski, L., & Schwartz, O. 1605
 (2004). Characterization of neural responses with 1606
 stochastic stimuli. In M. Gazzaniga (Ed.), *The* 1607
cognitive neurosciences (3rd ed.). MIT Press. 1608
- Smyth, D., Willmore, B., Baker, G. E., Thompson, I. D., & 1609
 Tolhurst, D. J. (2003). The receptive-field organization 1610
 of simple cells in primary visual cortex of ferrets under 1611
 natural scene stimulation. *The Journal of Neuro-* 1612
science, *23*, 4746–4759. [PubMed] [Article] 1613
- Szulborski, R. G., & Palmer, L. A. (1990). The two- 1614
 dimensional spatial structure of nonlinear subunits in 1615
 the receptive fields of complex cells. *Vision Research,* 1616
30, 249–254. [PubMed] 1617
- Theunissen, F. E., David, S. V., Singh, N. C., Hsu, A., 1618
 Vinje, W. E., & Gallant, J. L. (2001). Estimating 1619
 spatio-temporal receptive fields of auditory and visual 1620
 neurons from their responses to natural stimuli. 1621
Network, *12*, 289–316. [PubMed] 1622

- 1573 Touryan, J., Felsen, G., & Dan, Y. (2005). Spatial
1574 structure of complex cell receptive fields measured
1575 with natural images. *Neuron*, 45, 781–791. [[PubMed](#)]
1576 [[Article](#)]
- 1577 Touryan, J., Lau, B., & Dan, Y. (2002). Isolation of
1578 relevant visual features from random stimuli for
1579 cortical complex cells. *The Journal of Neuroscience*,
1580 22, 10811–10818. [[PubMed](#)] [[Article](#)]
- 1581 Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue,
1582 J. P., & Brown, E. N. (2005). A point process
1583 framework for relating neural spiking activity to
1584 spiking history, neural ensemble, and extrinsic
1595 covariate effects. *Journal of Neurophysiology*, 93, 1585
1074–1089. [[PubMed](#)] [[Article](#)] 1586
- Tsodyks, M., Kenet, T., Grinvald, A., & Arieli, A. (1999). 1587
Linking spontaneous activity of single cortical neu- 1588
rons and the underlying functional architecture. 1589
Science, 286, 1943–1946. [[PubMed](#)] 1590
- Volterra, V. (1959). *Theory of functionals and of integro* 1591
and integro-differential equations. New York: Dover. 1592
- Wiener, N. (1958). *Nonlinear problems in random theory*. 1593
New York: Wiley. 1594

AUTHOR QUERIES

AUTHOR PLEASE ANSWER ALL QUERIES

- AQ1: Please provide city and state for this affiliation, web address (URL), if available.
- AQ2: Please provide web address (URL) and e-mail address, if available.
- AQ3: Please provide city and state for this affiliation, web address (URL) and e-mail address if available.
- AQ4: Please check if these keywords are appropriate (none were provided).
- AQ5: Please check if the web address (URL) is correct.
- AQ6: Please provide editor name(s) as well as publisher name and location.
- AQ7: Please provide publisher location.
- AQ8: Please provide publisher name and location.
- AQ9: Chichilnisky (2001a) and Chichilnisky (2001b) are the same; hence, only one was captured (and was thus changed to Chichilnisky, 2001).
- AQ10: Please provide editor name(s) and publisher location.
- AQ11: Please provide updated publication status (e.g., volume number, page range).
- AQ12: Please provide editor name(s), publisher location, and page range.
- AQ13: Please provide editor name(s), page range, and publisher name and location.
- AQ14: Please provide page range.
- AQ15: Please provide page range and publisher location.

END OF AUTHOR QUERIES