

The Consequences of Response Nonlinearities for Interpretation of Spectrotemporal Receptive Fields

G. Björn Christianson,^{1,2} Maneesh Sahani,^{2*} and Jennifer F. Linden^{1,3*}

¹UCL Ear Institute, ²Gatsby Computational Neuroscience Unit, and ³Department of Anatomy and Developmental Biology, University College London, London, WC1E 6BT, United Kingdom

Neurons in the central auditory system are often described by the spectrotemporal receptive field (STRF), conventionally defined as the best linear fit between the spectrogram of a sound and the spike rate it evokes. An STRF is often assumed to provide an estimate of the receptive field of a neuron, i.e., the spectral and temporal range of stimuli that affect the response. However, when the true stimulus–response function is nonlinear, the STRF will be stimulus dependent, and changes in the stimulus properties can alter estimates of the sign and spectrotemporal extent of receptive field components. We demonstrate analytically and in simulations that, even when uncorrelated stimuli are used, interactions between simple neuronal nonlinearities and higher-order structure in the stimulus can produce STRFs that show contributions from time–frequency combinations to which the neuron is actually insensitive. Only when spectrotemporally independent stimuli are used does the STRF reliably indicate features of the underlying receptive field, and even then it provides only a conservative estimate. One consequence of these observations, illustrated using natural stimuli, is that a stimulus-induced change in an STRF could arise from a consistent but nonlinear neuronal response to stimulus ensembles with differing higher-order dependencies. Thus, although the responses of higher auditory neurons may well involve adaptation to the statistics of different stimulus ensembles, stimulus dependence of STRFs alone, or indeed of any overly constrained stimulus–response mapping, cannot demonstrate the nature or magnitude of such effects.

Key words: auditory; receptive field; auditory cortex; STRF; nonlinear; simulation

Introduction

A common goal in sensory neuroscience is to characterize a neuron in terms of a function that it computes on its input or on the input to the organism as a whole. This goal has often been pursued using a systems identification approach (de Boer, 1967; de Boer and de Jongh, 1978). In higher auditory centers, such reverse correlation analysis typically involves estimation of the spectrotemporal receptive field (STRF), which can be defined as the best-fit linear model between the spectrogram of the sound and the neuronal response it evokes (Aertsen et al., 1981; Eggermont et al., 1983; Palm and Pöpel, 1985; Eggermont, 1993; Theunissen et al., 2000).

Early research using reverse correlation methods focused on estimation of the Wiener–Volterra kernels using white noise as the driving stimulus (Marmarelis and Marmarelis, 1978). However, in higher levels of the auditory system, white noise tends to be ineffective at eliciting neuronal responses (Wang et al., 2005). This has led auditory researchers to use other stimuli, such as

dynamic random chords (DRCs) (deCharms et al., 1998; Schnupp et al., 2001; Rutkowski et al., 2002; Linden et al., 2003), natural sounds (Aertsen et al., 1981; Theunissen et al., 2000; Mächens et al., 2004), and a family of stimuli whose basic element is a ripple, a sound modulated sinusoidally in both the temporal and spectral domains (Kowalski et al., 1996a,b; Calhoun and Schreiner, 1998; Klein et al., 2000; Escabí and Schreiner, 2002; Miller et al., 2002; Qiu et al., 2003; Fritz et al., 2005). For many of these stimulus classes, the power at any two points in spectrotemporal space is uncorrelated by design, which simplifies estimation of the STRF. For the others, such as natural sounds, the impact of the correlations is removed at the analysis stage (Theunissen et al., 2000). However, in contrast to white noise, these stimuli are, with the exception of the DRC, not spectrotemporally independent; they contain non-zero third- or higher-order cross-central moments between points in spectrotemporal space.

This lack of full spectrotemporal independence is not an issue when the true underlying response function (RF) is linear. However, neuronal firing, and hence the computation performed by neurons, involves rectifying and saturating nonlinearities dictated by spiking mechanisms. Other significant nonlinearities have been directly demonstrated in the auditory cortex of bats (Suga et al., 1978), rodents (Sahani and Linden, 2003; Ahrens et al., 2008), cats (Calhoun and Schreiner, 1998), songbirds (Nagel and Doupe, 2006), and primates (Barbour and Wang, 2003) and in the inferior colliculus of barn owls (Peña and Konishi, 2001). When the RF is nonlinear, the linear fit between non-

Received April 19, 2007; revised Oct. 7, 2007; accepted Nov. 8, 2007.

This work was supported by Gatsby Charitable Foundation Grants GAT2579/GAT2623 (J.F.L.) and GAT2868 (G.B.C. and M.S.), via the Gatsby Computational Neuroscience Unit. We thank M. Ahrens and L. A. Anderson for their useful discussions regarding this work, R. Egnor for helpful comments on this manuscript, and M. Brainard, A. J. Doupe, C. Hampton, R. Egnor, and M. Hauser for providing natural stimuli.

*M.S. and J.F.L. contributed equally to this work.

Correspondence should be addressed to Jennifer F. Linden, UCL Ear Institute, University College London, London WC1X 8EE, UK. E-mail: j.linden@ucl.ac.uk.

DOI:10.1523/JNEUROSCI.1775-07.2007

Copyright © 2008 Society for Neuroscience 0270-6474/08/280446-10\$15.00/0




Figure 1. *a, b*, Examples of single frames of both the DRC stimulus (*a*) and the ripple stimulus (*b*). For STRF estimations, 75,000 frames of each stimulus were used. In the case of the DRC stimulus, each of these frames was randomly generated; in the case of the ripple stimulus, each frame had different modulation patterns along the time and frequency axes.

independent stimuli and the neuronal response can reflect statistical properties of the stimuli used in the fit rather than properties of the RF. This fact is well known in theory; in practice, its consequences for STRF analysis are not always fully appreciated. In particular, it is often assumed that STRFs always provide a reliable estimate of the receptive field of a neuron, but this is not necessarily the case.

We show here in simulation that simple, biologically plausible nonlinearities can interact with higher-order central moments in non-independent stimuli to produce STRFs with spurious receptive field elements. Moreover, we illustrate the fact that even STRFs estimated with spectrotemporally independent stimuli are dependent on the power of the stimulus. Finally, we demonstrate using natural sounds that these effects can lead to STRFs that appear to adapt to reflect stimulus structure, without any actual change in the underlying response function. Thus, the STRF of a nonlinear neuron does not necessarily reflect excitatory and inhibitory components of the underlying RF, and the structure and extent of the STRF may be stimulus dependent even when the true response function of the neuron is not.

Materials and Methods

Stimuli. All stimuli were created in “frames” of spectrograms consisting of 80 frequency bins and 30 time bins. For each type of stimulus, 75,000 spectrograms were created.

Dynamic random chord stimuli. As an example of spectrotemporally independent (and therefore also uncorrelated) stimuli, we used a DRC stimulus, one frame of which is shown in Figure 1*a*. DRC frames were generated directly in spectrotemporal space by randomly selecting 20% of the bins of the spectrogram to have zero intensity and assigning the non-zero bins to have one of five evenly distributed intensities with uniform probability. A DRC stimulus is spectrotemporally independent, in that the mean power in any given bin of the spectrogram is independent of the mean power in the other bins. In other words, knowing the power in any number of the spectrogram bins does not allow prediction of the power in any other spectrogram bin.

Ripple stimuli. As an example of uncorrelated but not independent stimuli, we used an ensemble of ripples. Each ripple spectrogram in the ensemble was assigned 1 of 128 temporal modulations (with frequencies evenly distributed from 0 through to the maximum possible), and 1 of 255 frequency modulations (again evenly sampled between 0 and the maximum possible), multiplied by a randomly assigned sign; an example of one such spectrogram is shown in Figure 1*b*. Such ensembles of ripples are spectrotemporally uncorrelated; that is, the power in any given spectrogram bin cannot be predicted from the power in any other single bin. However, these stimuli are not independent; because ripples are periodic,

the power in a given spectrogram bin can be predicted from the power in two other bins along the same line through the spectrogram.

Natural stimuli. Four classes of natural sounds were used in this study: environmental sounds from the Pittsburgh database (Smith and Lewicki, 2006), speech sounds from the TIMIT (for Texas Instruments and Massachusetts Institute of Technology) speech database (Garofolo et al., 1993), a selection of tamarin vocalizations (all either contact calls or combination long calls) provided by R. Egnor and M. Hauser (Harvard University, Cambridge, MA), and Bengalese finch songs provided by C. Hampton and M. Brainard (University of California at San Francisco, San Francisco, CA). All sounds were resampled to a sampling rate of 16 kHz and passed through a filter bank consisting of 80 gamma-tone band-pass filters with center frequencies linearly distributed between 100 and 7000 Hz. The spectrogram was then given by the Hilbert envelopes of the filter-bank output, decimated to a sampling rate of 1 kHz. The stimuli were subdivided into spectrogram elements, each 80 frequency bins by 30 time bins, and then a random subset of 75,000 were chosen for use in the study.

Natural stimuli, unlike DRC stimuli and ensembles of ripple stimuli, may have (second-order) correlations. The problem of robustly compensating for the effects of these correlations on STRF estimation has been addressed in previous studies (Theunissen et al., 2000; Woolley et al., 2006). Because our primary interest was in the effects of higher-order statistics on STRF analysis, we chose to avoid the issues associated with second-order structure by numerically whitening natural stimuli before use. After this process, the off-diagonal elements of the autocorrelation matrix were all five or more orders of magnitude smaller than the diagonal elements, although any higher-order statistical structure was preserved.

Simulation of response. Spectrograms were each recast into a vector and became rows in a $75,000 \times 2400$ stimulus matrix S . This was then multiplied by one or more similarly vector-recast RF matrices \tilde{w} (2400×1), and the results were combined according to the rules below to give a $75,000 \times 1$ response vector \tilde{r} . In Figures 9 and 10, this was taken to be the response of the neuron. In other simulations, a final response $\tilde{\rho}$ ($75,000 \times 1$) was obtained by drawing 20 samples from an inhomogeneous Poisson distribution with mean parameter \tilde{r} and then averaging across the samples:

$$\tilde{\rho} = \frac{1}{20} \sum_{i=1}^{20} \tilde{\rho}_i; \quad \tilde{\rho}_i \sim \text{Poisson}(\tilde{r}).$$

For simulation of a linear RF neuron, S was multiplied by a single RF \tilde{w} (2400×1):

$$\tilde{r} = S\tilde{w}.$$

For the linear model in Figure 9, \tilde{r} was taken to represent the response, even if some entries were negative, thus preserving true linearity. In simulations with noise, the stimuli were offset so that \tilde{r} was never negative, and Poisson noise was added, as above.

Three basic nonlinear RFs were modeled. A “multiplicative RF” was modeled using the linear responses to two distinct Gaussian receptive fields in spectrotemporal space. The outputs of the linear projections were rectified and multiplied pointwise (indicated by the Schur product \odot) to give the response vector \tilde{r} :

$$\tilde{r} = |S\tilde{w}_1|^+ \odot |S\tilde{w}_2|^+.$$

A “divisive inhibition RF” was modeled using the linear response to two distinct receptive fields. Both receptive fields were Gaussian in temporal extent. The excitatory receptive field was a squared Gaussian in spectral extent, whereas the inhibitory receptive field was quadratic (see Fig. 4*a*). As with the multiplicative model, the output of each projection was rectified and then combined pointwise (with pointwise division indicated by the symbol \div):

$$\tilde{r} = |S\tilde{w}_1|^+ \div (1 + a|S\tilde{w}_2|^+).$$

The constant a was set to 7.5. Last, a “threshold RF” neuron was simulated by subtracting 0.3 times the maximum value of the linear projection and then rectifying to give the mean rate:

$$\bar{r} = |S\bar{w} - 0.3\max(S\bar{w})|^+.$$

Estimation of STRFs. All simulations used stimuli with well conditioned autocorrelation matrices, and so STRFs could be estimated by simple linear regression:

$$\hat{w} = (S^T S)^{-1} S^T \bar{r}.$$

The auto-correlation matrix $S^T S$ differed from a scaled identity matrix only because of finite-sample effects.

In plots of RFs and STRFs, black denotes the minimum response, and white denotes the maximum response.

Analytic form of the STRF for multiplicative RF. The effects of RF non-linearity are much simpler to describe analytically for the multiplicative model than for the other models used in our simulations. Therefore, we use the multiplicative RF here to illustrate quantitatively how overestimation of the true receptive field may arise from dependence of the STRF on higher-order statistics of the stimulus. Note that the multiplicative model can be viewed as a component of the second-order term in a Volterra expansion, and thus this derivation is partially relevant to any analysis of nonlinear data.

Let \bar{s} be a single $N \times 1$ stimulus frame (corresponding to a row of the stimulus matrix S defined above), which evokes a (scalar) mean response r in a multiplicative model as defined above, with $N \times 1$ RF components \bar{w}_1 and \bar{w}_2 . To avoid effects attributable to finite sampling and noise, we consider expectations with respect to both response variability and the stimulus ensemble. These expectations will be denoted by angle brackets $\langle \cdot \rangle$. Without loss of generality, we assume the stimulus set has been normalized such that $\langle \bar{s} \bar{s}^T \rangle = I$ and $\langle \bar{s} \rangle = \bar{0}$.

In the multiplicative model, the expected value of the STRF estimate is given by

$$\hat{w} = \langle \bar{s} \cdot r \rangle = \langle \bar{s} \cdot (|\bar{s}^T \bar{w}_1|^+ \cdot |\bar{s}^T \bar{w}_2|^+) \rangle. \quad (1)$$

This expectation over all stimuli can be replaced by the average over only those stimuli for which $\bar{s}^T \bar{w}_1 > 0$ and $\bar{s}^T \bar{w}_2 > 0$ (that is, those stimuli that evoke non-zero response under the multiplicative model), multiplied by a coefficient α giving the fraction of such stimuli in the overall ensemble. Writing $\langle \cdot \rangle_+$ for the restricted average and focusing on the i th element of \hat{w} , we have

$$\begin{aligned} \hat{w}_i &= \alpha \langle s_i [(\bar{s}^T \bar{w}_1) \cdot (\bar{s}^T \bar{w}_2)]_+ \rangle \\ &= \alpha \langle s_i \left(\sum_j s_j w_{1j} \right) \left(\sum_k s_k w_{2k} \right)_+ \rangle \\ &= \alpha \sum_j \sum_k w_{1j} w_{2k} \langle s_i s_j s_k \rangle_+. \end{aligned} \quad (2)$$

Thus, each element \hat{w}_i of the STRF estimate \hat{w} depends on the third-order conditional statistic $\langle s_i s_j s_k \rangle_+$, as well as on \bar{w}_1 and \bar{w}_2 .

We are interested in determining under what conditions \hat{w} might overestimate the receptive field of a neuron. Formally, by receptive field, we mean the “dimensional support” of the model (which we will abbreviate to “support”): the stimulus dimensions that can influence the response of the model. In this context, when we say i is within the support of the model, we mean that at least one of w_{1i} and w_{2i} is non-zero, and so we want to determine the conditions under which both w_{1i} and w_{2i} are 0, but \hat{w}_i is non-zero. Accordingly, we partition each stimulus into two components:

$$s_i^I = \begin{cases} s_i & \text{if } i \text{ is in the model support} \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

$$s_i^O = \begin{cases} 0 & \text{if } i \text{ is in the model support} \\ s_i & \text{otherwise} \end{cases}$$

Substituting $s_i = s_i^I + s_i^O$ into Equation 2, we obtain the following:

$$\begin{aligned} \hat{w}_i &= \alpha \left((s_i^I + s_i^O) \left(\sum_j (s_j^I + s_j^O) w_{1j} \right) \left(\sum_k (s_k^I + s_k^O) w_{2k} \right) \right)_+ \\ &= \alpha \left((s_i^I + s_i^O) \left(\sum_j s_j^I w_{1j} \right) \left(\sum_k s_k^I w_{2k} \right) \right)_+ \\ &= \alpha \left(\sum_j \sum_k w_{1j} w_{2k} \langle s_i^I s_j^I s_k^I \rangle_+ \right) + \alpha \left(\sum_j \sum_k w_{1j} w_{2k} \langle s_i^O s_j^I s_k^I \rangle_+ \right). \end{aligned} \quad (4)$$

By construction, s_i^I and therefore the first term in the final expression of Equation 4, is zero when i is not in the dimensional support of the model. Therefore, for points outside the support, we have

$$\hat{w}_i = \alpha \sum_j \sum_k w_{1j} w_{2k} \langle s_i^O s_j^I s_k^I \rangle_+. \quad (5)$$

Thus, unless the restricted statistic $\langle s_i^O s_j^I s_k^I \rangle_+$ is zero for all i outside and j, k inside the support, estimated weights outside the support may be non-zero. If the stimulus is not independent, the statistic cannot be guaranteed to vanish.

Note that the definition of $\langle \cdot \rangle_+$ means that non-zero elements of s^I are not independent of one another within the restricted set of stimuli, even if they were in the overall ensemble. Intuitively, knowledge of one point in the support and the fact that a response was elicited puts constraints on the possible values of other points in the support. One might therefore be concerned that, even for independent stimuli, the restricted statistic might not vanish. However, as long as it is true that s_i^O for i outside the support is independent of s_j^I for j inside the support, the value of s_i^O will also be independent of the response (which depends only on s^I) and thus independent of s_j^I even after restriction to stimuli that evoke responses. Thus, for weights \hat{w}_i outside the support,

$$\begin{aligned} \hat{w}_i &= \alpha \sum_j \sum_k w_{1j} w_{2k} \langle s_i^O \rangle_+ \langle s_j^I s_k^I \rangle_+ \\ &= \alpha \sum_j \sum_k w_{1j} w_{2k} \cdot 0 \cdot \langle s_j^I s_k^I \rangle_+ \\ &= 0. \end{aligned} \quad (6)$$

Results

Multiplicative interactions

The STRF is often taken to represent the receptive field of a neuron or, in other words, is often considered to be an accurate estimate of the dimensional support (see Materials and Methods), herein referred to simply as the support, of the RF of the neuron. More precisely, it is assumed that significant weights in the STRF appear only at those time–frequency points that contribute causally to the output of the response function. For a linear response function (Fig. 2a), both spectrotemporally independent (DRC) (Fig. 2b) and non-independent but uncorrelated (ripple) (Fig. 2c) stimuli provide a consistent estimate of the support; that is, no spectrotemporal points appear in the STRFs that do not contribute to the RF in the large data limit. In the case of a nonlinear RF, for example with a multiplicative interaction (Fig. 2d), independent stimuli can still be used to obtain a conservative estimate of the RF support (Fig. 2e). This is a corollary of inde-




Figure 2. *a–c*, A linear RF formed by the addition of two Gaussian receptive fields (*a*) is correctly estimated using DRC (*b*) and ripple (*c*) stimuli. *d, e*, However, for a multiplicative RF (*d*), only the DRC estimate provides the expected STRF consisting of an approximation of the sum of the RF components (*e*). *f*, The ripple STRF includes prominent sidebands that have no basis in the true RF.

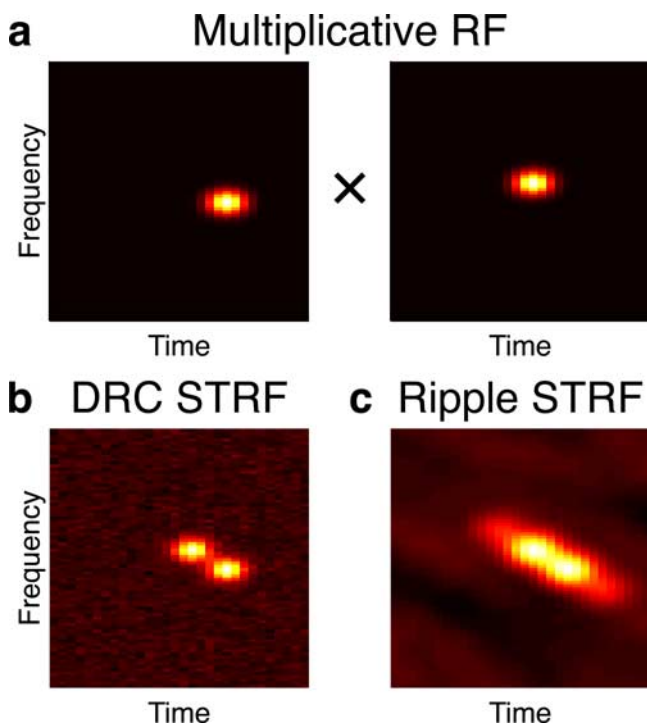


Figure 3. The effects of Figure 2 can influence STRF interpretation in a number of different ways. *a–c*, For this model sweep-sensitive neuron, with multiplicative RF components that are temporally offset and spectrally adjacent (*a*), the STRF estimated using an ensemble of ripples greatly overestimates the extent of spectrotemporal tuning (*b* vs *c*).

pendence; if one point is independent of another point, then it will also be independent of any function of that other point. Thus, regardless of the nature of the RF, points outside the support of the RF will be uncorrelated with the response of the neuron, which is, by definition, a function only of points within the support. Therefore, in the large data limit, the STRF will have zero values outside the support when estimated with a spectrotemporally independent stimulus.

However, this is not the case for ensembles of ripple stimuli, which are spectrotemporally uncorrelated but not independent. The STRF estimated using ripple stimuli shows significant sidebands of excitation (Fig. 2*f*) that do not correspond to any feature in the true RF. Although the uncorrelated nature of the ripple ensemble guarantees that points outside the support are not correlated with any single point within the support, it does not prevent points outside the support of the RF from effectively being correlated with higher-order cross moments of the stimulus within the support. Thus, if the RF depends on these higher-order moments, the stimulus outside the support may be correlated with the response of the neuron. Intuitively, in the case of this particular nonlinearity, the RF selects for stimuli that have power present simultaneously in both RF elements. Because the individual ripples are periodic in structure, the ripple ensemble possesses non-zero third-order

moments; specifying that two points are peaks, for example, means that peaks will repeat along the line drawn through the two points at integer multiples of the distance between them. The net result is that points outside the support of the RF then become correlated with combinations of points inside the support, resulting in the appearance of the sidebands in this case. This intuition is formalized in Materials and Methods, in which the analytic form of the STRF for a multiplicative model is derived, showing its dependence on a third-order statistic of the stimulus (Eq. 5).

A closer arrangement of the RF components may produce a single STRF feature with exaggerated support. A physiologically plausible example is shown in Figure 3*a*, in which two closely spaced RF components staggered in time and frequency are multiplied together. Compared with that estimated using the DRC (Fig. 3*b*), the STRF estimated using an ensemble of ripples is elongated in both the temporal and the spectral domains (Fig. 3*c*). The same principle as underlies the effect in Figure 2*f* causes this phenomenon, but the closer spacing of the RF components means that the sidebands are not visibly separated from the central region.

Divisive inhibition

Similar effects can be observed using an RF model intended to simulate the inhibitory sidebands observed in two-tone mapping studies [e.g., in auditory cortex (Sutter et al., 1999; Kadia and Wang, 2003)]. Rather than explicitly modeling two-tone interaction nonlinearities, we use coextensive fields of excitation and inhibition (inspired by the results of Wehr and Zador, 2003), with the excitation dropping off more sharply than the inhibition (Fig. 4*a*). The inhibition acts divisively; similar divisive interactions are thought to play a critical role in regulating activity in visual cortex networks (Heeger, 1992; Chance and Abbott, 2000). The STRF estimated for this model using the DRC stimulus has inhibitory regions on either side of an excitatory main peak (Fig. 4*b*), which resemble the inhibitory sidebands described in two-tone mapping studies of the central auditory system (Sutter et al., 1999; Kadia and Wang, 2003). In the STRF estimated using an ensemble of ripples, however, additional banding is visible (Fig. 4*c*), extending well beyond the support of the model. Intuitively,

the lateral-inhibition-like function of the model selects for those ripples that have a frequency modulation pattern matching the spacing of the surrounds, and the resulting dominance of ripples with this frequency modulation causes additional banding to emerge in the STRF.

Thresholding

Multiplicative nonlinearities are not required to produce this sort of effect. Even a thresholding output nonlinearity (Fig. 5*a*), fundamental to the spiking response, can lead to an overestimation of support in the STRF when the stimulus is not spectrotemporally independent (Fig. 5*c*). The STRF for the rectified bimodal linear RF in Figure 5*c* is very similar to the STRF for the multiplicative nonlinearity of Figure 2*f* and for essentially the same reason; the rectification ensures that the majority of spectrograms that actually elicit responses are those that have power corresponding to both peaks in the linear RF. In addition to the side bands, another feature of the ripple-estimated STRF shared between the two cases, but more obvious in Figure 5*c*, is the increased size of the STRF elements corresponding to the true RF components. This again arises from structure in the ripple stimulus. A high threshold requires that the majority of the receptive field be stimulated to elicit a response. Because individual ripples are continuous, any stimulus that elicits a response and hence has power throughout the extent of the receptive field will also have power immediately outside the border of the receptive field. Thus, the set of all response-eliciting ripples have power extending beyond the support of the RF, and this is reflected in the STRF.

General effects of nonlinearities

In all Figures 2–5, other patterns can be seen in the STRFs, in addition to the specific features we have described. (This is true not only for the ripple-estimated STRFs but also for the DRC-estimated STRFs, which show an apparently noisy background.) Such effects are not entirely attributable to noise in the simulated responses; they remain present even when noise is excluded from the simulations. Rather, these patterns, like the specific features described previously for each model, arise from an interaction between the RF nonlinearity and the statistics of the stimulus used to estimate the STRF. (In the case of the DRC-estimated STRFs, some part of the noisy-looking background derives from non-zero moments in the stimulus that occur because the stimulus is finite in length.) These patterns are generally sensitive to minor changes in the model parameters, and their origin is difficult to describe more intuitively than with reference to the interaction between nonlinearities and stimulus statistics (e.g., see the analytic form of the STRF for the multiplicative RF in Materials and Methods).

Non-intuitive consequences of linear regression in high-dimensional spaces

The fact that nonlinearities in response functions can lead to differences in STRFs estimated using different stimuli has long been acknowledged in the literature (Marmarelis and Marmarelis, 1978; Aertsen and Johannesma, 1981; Theunissen et al., 2000; Escabí and Schreiner, 2002). Indeed, for one-dimensional regression, the point is obvious; because the linear fit is only an approx-

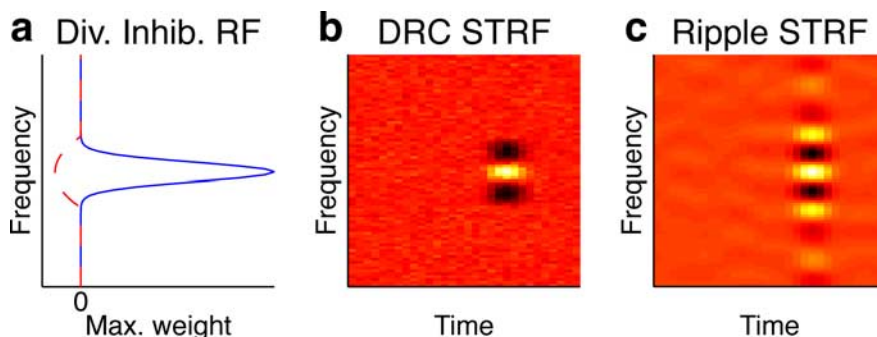


Figure 4. Divisive inhibition can cause overestimation of support. *a*, In this model RF, the maximum spectral extents of the regions of inhibition (dashed red line) and excitation (solid blue line) were identical, although the profiles had different shapes; in the temporal dimension, both extents and profile shapes were the same. *b*, Inhibitory sidebands are clearly apparent when an STRF is estimated using the DRC stimulus. *c*, In the STRF estimated with the ensemble of ripples, an alternating pattern extends beyond the simple sidebands and outside the support of the model.




Figure 5. *a, b*, A rectifying nonlinearity (*a*) applied after a simple bimodal linear RF does not have any impact on the estimation of an STRF using the independent DRC stimulus (*b*). *c*, However, with the ripple stimulus, the receptive field of the neuron is overestimated, and sidebands appear.

imation to the true nonlinear generating function, the fit will depend on the range and distribution of data to be fit (Fig. 6).

Related and equally intuitive observations apply to response prediction. Again, in one dimension, the slope of a line fit to a nonlinear function over a set of points that fall within a particular data range is generally more useful for predicting the value of the function at other points within the same range than the slope of a line fit to data in a different range would be. Likewise, in our simulations, ripple-estimated STRFs always predicted responses to novel instances of ripple stimuli better than did DRC-estimated STRFs, and DRC-estimated STRFs always predicted responses to novel instances of the DRC stimulus better than did ripple-estimated STRFs.

However, the most important implication of our simulations has no analogy in one-dimensional regression and is therefore




Figure 6. The best-fit linear regression (red line) to parabolic data (black dashed line) depends on the subset of the data used in the fit (black solid line).




Figure 7. *a, b*, The periodic structure of a ripple stimulus in spectrotemporal space (*a*) simplifies to a single point in modulation transfer function space (*b*), up to a symmetry about the origin imposed by a real-valued signal. *c–f*, The true support of the nonlinear RF (*c*) from Figure 2*c* is overestimated in spectrotemporal space when using an ensemble of ripple stimuli (*e*), but when transformed into modulation transfer function space, the estimate is conservative (*d, f*). (Note that *c* and *d* show the spectrotemporal and modulation transfer function representations of the true support of the multiplicative RF, not the DRC estimate of this support.)

less intuitive. Suppose a linear function (such as an STRF) is fit to nonlinear data (such as nonlinear neuronal responses) in a very high-dimensional space (such as spectrotemporal space, in which each time–frequency element represents a different dimension).

Suppose also that the points at which the data have been measured (e.g., stimuli used to elicit neuronal responses) are such that there are third- or higher-order statistical relationships between the values of the different input vector elements at those points (such as there are between spectrotemporal elements within the ensemble of ripple stimuli). Then the optimal linear fit may have non-zero weights even in dimensions to which the true data-generating function is insensitive; in the particular case of the nonlinear neuron,

this means that the STRF may have non-zero weights outside the support of the RF of the neuron (Figs. 2–5). This overestimation of RF support will not impair the power of the STRF to predict responses to novel stimuli, as long as those novel stimuli fall within the same region of stimulus space in which the fit was performed.

Representation dependence

Coordinate independence is a property not simply of the stimulus itself but also of its representation. An ensemble of ripple stimuli, for example, although not independent in spectrotemporal space (Fig. 7*a*), becomes only a set of discrete points in modulation transfer function space (Fig. 7*b*), in which signals are represented by the modulation of their envelopes in the temporal and spectral domains (Kowalski et al., 1996*a,b*; Calhoun and Schreiner, 1998). A linear combination of ripples with independently chosen temporal and spectral frequencies is thus independent in modulation transfer function space, so although an ensemble of ripples overestimates the true support of a nonlinear RF in the spectrotemporal domain (Fig. 7*c,e*), when the STRF is transformed into modulation transfer function space, it provides a conservative estimate of the support in the modulation transfer function domain (Fig. 7*d,f*). As many previous authors have noted, ensembles of ripples are therefore ideal stimuli for studies of neuronal response properties in modulation transfer function space (Kowalski et al., 1996*a,b*; Calhoun and Schreiner, 1998).

In contrast, unless the powers of individual elements in a DRC stimulus are chosen from a Gaussian distribution, the modulation content of the DRC stimulus at different modulation frequencies is not independent and shows periodic structure much as the ripples do in spectrotemporal space. Thus, a modulation transfer function estimate derived from a DRC is not guaranteed to provide a conservative estimate of the true modulation transfer function for a nonlinear neuron.

Power dependence

Even when the stimulus used is independent for the coordinates chosen, nonlinearities in the RF can affect the estimated STRF in a stimulus-dependent manner (Marmarelis and Marmarelis, 1978; Eggermont, 1993). This effect can confound the interpretation of apparent excitatory and inhibitory components in the STRF. Figure 8*a* illustrates a model neuron that has overlapping excitatory and inhibitory regions, in which the excitatory region is generated by a multiplicative nonlinear component and the inhibitory one is generated by a linear component. At low levels of stimulus power (Fig. 8*b*), the linear term dominates, and the STRF is inhibitory. At intermediate values, the region can essentially disappear entirely from the STRF as the two terms cancel each other out (Fig. 8*c*). Eventually, the multiplicative term dominates, and the STRF is entirely excitatory (Fig. 8*d*). In the case of

the multiplicative model, this effect arises from the fact (as seen in Eq. 4 in Materials and Methods) that, even for independent stimuli, there is a dependence of the STRF on non-negligible third-order statistics between points within the support. Although here we illustrate this effect with differing levels of stimulus power, it is also theoretically possible to see similar effects using different classes of independent stimuli that have the same power. For example, while recording responses of owl monkey auditory cortical neurons, Blake and Merzenich (2002) varied the power and spectrotemporal density of DRC stimuli and thus manipulated a combination of stimulus power and stimulus statistics. They observed emergent inhibition overlapping an excitatory region in their STRFs, which our simulations closely resemble. Thus, even when estimated using independent stimuli, STRFs are stimulus dependent, although independence does guarantee that the changes resulting from differences in the power or the nature of the stimulus used will still be constrained to a conservative estimate of the dimensional support.

Natural stimuli

It is not immediately clear that the above results hold in the case of natural stimuli, which, although not spectrotemporally independent, typically possess more complicated higher-order statistics than those considered thus far. To address this issue, we estimated STRFs for all of our model RFs using a variety of natural stimulus ensembles. The highly correlated structure of natural stimuli itself presents problems for STRF estimation (Theunissen et al., 2000), so some form of regularization is typically used during estimation of the STRFs for real neuronal responses. In simulations, however, we could avoid issues associated with regularization completely by using whitened versions of the natural stimuli that were in fact uncorrelated. This whitening allowed us to focus specifically on the effects of interactions between response nonlinearities and higher-order statistical dependencies in the stimuli. To avoid confounding the effect of signal statistics with the effect of noise, no noise was included in the model responses for the simulations shown, although inclusion of noise did not have a major influence on the results.

As expected, STRFs estimated for the linear model using whitened natural stimuli show complete recovery of the underlying RF, regardless of the stimulus ensemble used (Fig. 9*a–e*). However, for the multiplicative model, differences between the STRFs estimated with the different stimuli become apparent (Fig. 9*f–j*). In particular, the STRF estimated using ambient environmental sounds (Fig. 9*g*) displays smearing in the frequency dimension, which contrasts with the temporally elongated features of the STRFs estimated using Bengalese finch song (Fig. 9*h*) and tamarin calls (Fig. 9*i*). (These effects are consistent with the dominant features of the sounds: ambient environmental sounds include many brief spectral stacks, and Bengalese finch song and tamarin calls are composed of more elongated narrowband sounds.) Similar trends are observed in the STRFs estimated for the threshold model (Fig. 9*k–o*). For the divisive inhibition model (Fig. 9*p–t*),

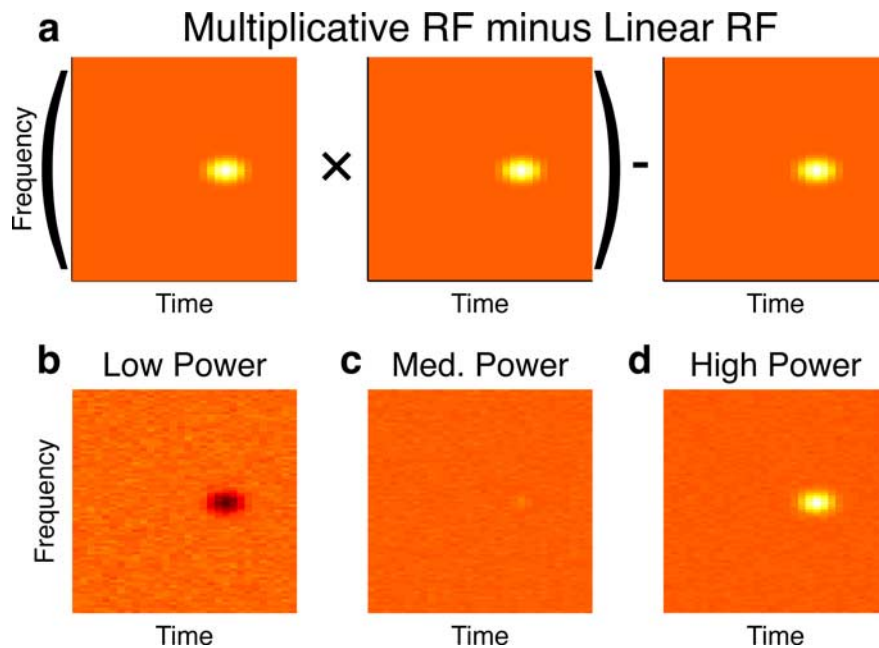


Figure 8. *a*, An RF with an inhibitory linear component and an excitatory multiplicative component, in which spectrotemporal extents of the two components are the same. *b*, At low stimulus powers, the linear component dominates, and the peak is inhibitory. *c*, *d*, As the stimulus power is increased, the influence of the multiplicative term begins to dominate; the peak diminishes (*c*) and becomes strongly excitatory (*d*). For this example, background firing was simulated by adding a constant to $\bar{\tau}$ (see Materials and Methods); orange represents this baseline response, with black being drops below this rate and white denoting increases in rate.

the inhibitory sidebands select against activation by spectrally elongated sounds, and so the frequency smearing vanishes in the case of the ambient environmental sounds (Fig. 9*q*). However, temporal elongation is still apparent near 4 kHz for the STRF estimated with finch song (Fig. 9*r*) and at lower frequencies for the STRF estimated using tamarin calls (Fig. 9*s*). In general, speech STRFs show fewer effects, although the speech STRF for the divisive inhibition model has a pronounced and temporally elongated low-frequency RF component (Fig. 9*t*).

The effects illustrated in Figure 9 are less dramatic than those shown previously using ripple stimuli. It must be reiterated, however, that even these more subtle effects are attributable entirely to higher-order stimulus statistics and are not a result of either noise or regularization. Moreover, the effects depend on an interaction between the model and higher-order stimulus statistics. Changes in the RF model parameters can lead to dramatic results (Fig. 10). Results with ripple stimuli provide an intuitive explanation for this strong dependence on the RF nonlinearity. In ripple STRFs, structure outside the RF support is far more pronounced for the divisive inhibition model (Fig. 4*c*) than for the multiplicative model (Fig. 2*f*). This is because the inhibitory sidebands of the divisive inhibition RF are strongly selective for a particular category of ripples (those with frequency modulations fitting the off–on–off spacing of the model) and hence for those higher-order statistics. The RF models used in this study were chosen for their simplicity and were not designed to select for any features present in the natural stimuli, so it is not surprising that interactions between these RF nonlinearities and higher-order statistics of natural stimuli are relatively subtle (Fig. 9). However, recordings in the auditory cortex of marmosets (Wang et al., 2005) suggest that neurons in the central auditory system can be highly selective for complex stimuli. In fact, the observation that independent stimuli can drive auditory neurons poorly is one of

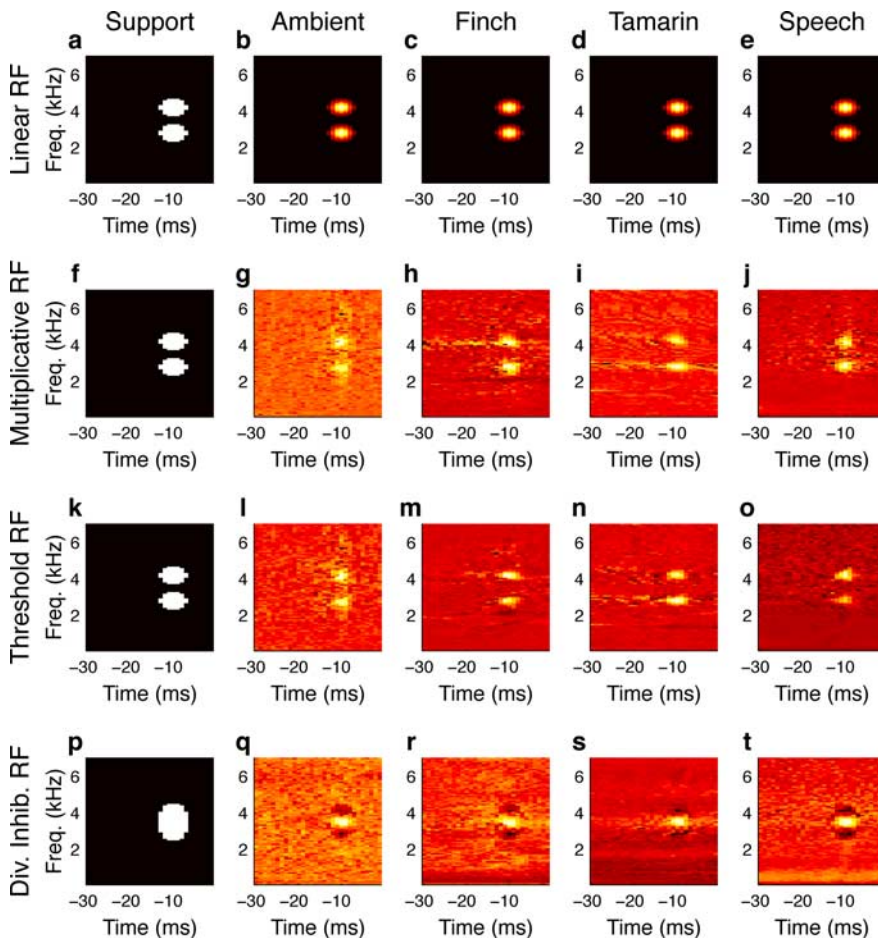


Figure 9. *a–t*, The STRF of a model neuron varies with the stimulus used to estimate it. Although the STRF estimated for a linear model (*a–e*) is invariant with the stimulus used, STRF estimates are stimulus dependent for the multiplicative RF (*f–j*), the threshold RF (*k–o*), and the divisive inhibition RF (*p–t*). *a, f, k, p*, Support of the models. All time–frequency combinations that can contribute to the response are shown in white. *b, g, l, q*, STRFs estimated using ambient environmental sounds. *c, h, m, r*, STRFs estimated using Bengalese finch song. *d, i, n, s*, STRFs estimated using cotton-top tamarin vocalizations. *e, j, o, t*, STRFs estimated using speech.

the primary motivations for the use of non-independent stimuli in reverse correlation studies (Klein et al., 2000). Thus, it is a real concern that the effects observed in STRFs estimated using natural stimuli would become much more pronounced in the presence of neuronal nonlinearities that selected strongly for natural stimulus statistics.

Discussion

We have demonstrated through simulation that even simple, biologically plausible nonlinearities can have powerful effects on STRF structure. For a neuron with a linear RF, the STRF will be an accurate estimate of the support of the RF (i.e., the receptive field of the neuron). However, when nonlinearities are present in the RF, decorrelation of the stimulus will no longer ensure accurate estimation of the support. Higher-order statistics of an uncorrelated but not independent (or neither uncorrelated nor independent) stimulus can interact with the RF nonlinearities to produce features in the STRF that do not correspond to any actual component of the RF. This effect, in which the support of the STRF can exceed that of the true RF, does not impair the performance of the STRF on prediction of responses to stimuli similar to those used for STRF estimation (although it may impair prediction of responses to stimuli with distinctly different characteristics).

Rather, it impairs our ability to interpret the STRF to derive parameters relating to the true RF support. Thus, when estimated with stimuli that are not independent, STRFs cannot reliably be used to determine the spectrotemporal extent of receptive fields for nonlinear neurons.

STRFs estimated with spectrotemporally independent stimuli will not have features that lie outside the support of the RF; however, the estimate of RF support will be conservative. Therefore, it is possible for the extent of the receptive field to be underestimated (Fig. 8). Moreover, for STRFs estimated with either independent or non-independent stimuli, the concept of the “sign” of STRF components is problematic. In a nonlinear RF, the excitatory or inhibitory action of stimulus power at any time–frequency point can depend on the total power in the stimulus; a linear approximation to the RF (i.e., the STRF) will therefore be stimulus dependent even when estimated using independent stimuli, and interpretation of inhibition and excitation will be limited to the specific stimulus used for the STRF estimation. Thus, for independent as well as non-independent stimuli, stimulus dependence of STRFs may arise through interactions between constant RF nonlinearities and the properties of the stimulus ensemble. This fact is implicit in the previous theoretical work on STRFs (Eggermont, 1993); here we have presented explicit examples of such interactions using biologically plausible nonlinearities. We also made the previously unacknowledged point that, for STRFs estimated with non-independent stimuli, these interactions can create features in the STRF that lie outside the true receptive field of a nonlinear neuron.

Many recent studies have demonstrated stimulus dependence in STRFs estimated using different stimuli (Theunissen et al., 2000; Blake and Merzenich, 2002; Escabí and Schreiner, 2002; Valentine and Eggermont, 2004; Woolley et al., 2006) or stimulus dependence in other measures of neuronal responses to different complex sounds (Bar-Yosef et al., 2002; Nagel and Doupe, 2006). There is therefore little doubt that central auditory neurons are in fact nonlinear (but see below), but the nature of that nonlinearity remains a topic of some debate. One possibility is that stimulus-dependent changes in STRFs may reflect some underlying process of adaptation; that is, a change in the stimulus causes an otherwise predominantly linear mapping to adapt, nonlinearly, to the altered context (e.g., as suggested by Woolley et al., 2006). Another (or additional) possibility is that stimulus-dependent changes in STRFs may arise not from changes in RF parameters but instead from a stationary nonlinearity in the response function (e.g., as suggested by Theunissen et al., 2000; Escabí and Schreiner, 2002; Valentine and Eggermont, 2004). [A similar point has been raised recently in studies of adaptation of motion detection in the fly visual system (Borst et al., 2005).] Adaptation is ubiquitous in the brain and surely plays an important role in

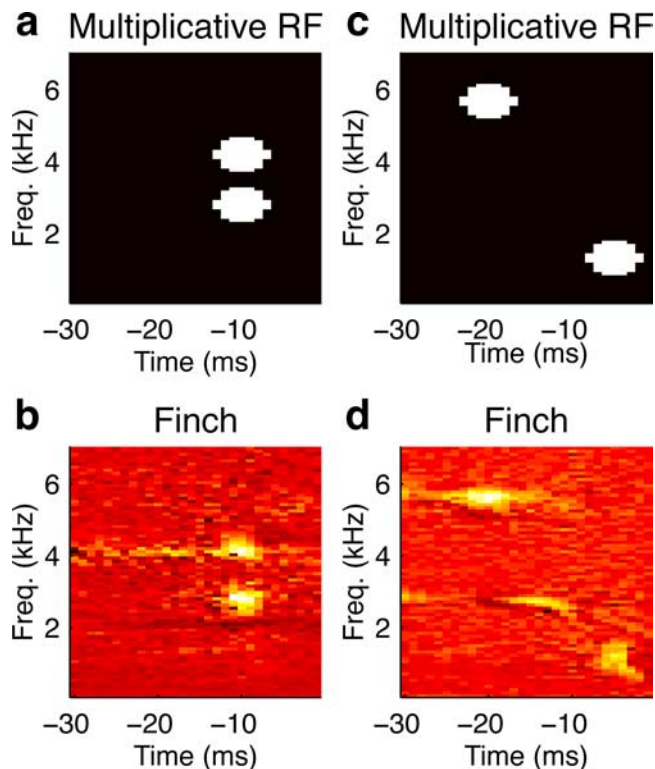


Figure 10. Minor changes in the model parameters can influence interactions with stimulus statistics. *a, b*, The multiplicative RF used in this paper (*a*; only model support shown, as in Fig. 9) leads to spectral elongation in the STRF estimated using Bengalese finch song (*b*; same as in Fig. 9). *c, d*, However, changing the locations of the RF components (*c*) leads to both more pronounced spectral elongation and the appearance of a clear third peak in the STRF (*d*).

neural processing, but the effects of stationary response function nonlinearities described here complicate quantification of the extent and form of adaptation. This does not mean that STRF-like analysis is inappropriate for use in adaptation studies, however. To demonstrate adaptation, it is sufficient to show that the parameters of the response function change as a function of time, more gradually than the stimulus itself (Shechter and Depireux, 2006). Alternatively, it is possible to look at changes in STRFs estimated using the same stimuli embedded within an adapting context of other stimuli or task demands (Fritz et al., 2005), thus avoiding the issue of stimulus-dependent estimation entirely. A critical point not addressed by previous work, however, is that the differences between STRFs resulting from static nonlinearities may only indirectly reflect the nature of the underlying response function. Thus, if changing the stimulus results in new areas of the STRF displaying power, it does not necessarily mean that additional portions of the receptive field have been “uncovered.” Such variation might instead reflect only the interaction between a constant nonlinear response and the statistics of the stimulus.

We should note here that there is still some debate in the literature about whether STRFs tend to be stimulus dependent and therefore whether central auditory neurons have nonlinear response functions. Previously, for example, Klein et al. (2006) examined the stability of STRFs estimated for auditory cortical neurons using a variety of stimulus classes and concluded that stimulus dependence of the STRFs was minimal (and therefore that the neuronal response functions could reasonably be described as linear). However, the stimuli they used were all derived from ripple stimuli and may possess similar higher-order statistics. More importantly, their analysis focused on gross-scale

structure of the STRFs, whereas our results suggest that nonlinearities may express themselves in a more subtle manner. The DRC- and ripple-estimated STRFs of Figure 2, *e* and *f*, for instance, have a correlation coefficient of 0.83; this contrasts with the mean correlation coefficient of 0.64 for STRFs in the study by Klein et al. (2006). Although the STRFs we estimate with DRC and ripple stimuli are well correlated, the differences between them would greatly complicate analysis of receptive field structure.

The issues we described here are not exclusive to STRF analysis. They are in fact general consequences of the fundamental problem of model mismatch. Single- and two-tone frequency response areas (FRAs) illustrate the same concept. The single- and two-tone FRAs of a neuron are generally different, and the fact that the single-tone FRA cannot predict two-tone interactions does not mean that the response function of the neuron changes with the number of tones presented, but rather that the response function cannot be completely characterized in terms of linear sums of the responses to single tones alone. STRF analysis makes the explicit assumption that the true RF is linear. This is, however, a very difficult assumption to verify in practice, and thus great caution must be taken in interpreting STRFs. Fundamentally, the most useful interpretation of the STRF is predictive: that is, it is the linear model that best predicts the response to the chosen stimuli. Although it may be useful to compare STRFs estimated with the same sound stimulus, it is difficult to derive meaningful conclusions from the comparison of STRFs derived using different stimulus sets. Similarly, extracting the true RF from comparison of any collection of STRFs is a difficult problem, in which even the analysis of such basic details as the spectrotemporal extent of the RF support must be approached with caution. Just as the single-tone frequency response function is not the whole RF, the (perhaps misleadingly named) STRF is only an accurate description of the receptive field in the case of a neuron with a linear RF.

In light of the issues presented here, it may seem a reasonable question to ask, why use STRF analysis at all? Just as single-tone FRAs have not been invalidated as a tool for auditory neuroscience by the biological implausibility of the assumption that auditory neurons can be characterized by their responses to single tones alone, STRF analysis is not invalidated by the overly constraining assumption that auditory neuronal response functions are spectrotemporally linear. Indeed, in many circumstances, the simplicity of the STRF, like the simplicity of the single-tone FRA, can be an advantage for analysis. A generalized nonlinear model, such as the Wiener–Volterra series, may seem preferable to STRFs, but fitting increasing orders of complexity in such models requires an exponentially increasing amount of data. More sophisticated approaches to modeling linear–nonlinear–Poisson (LNP) neural responses (that is, response functions in which one or more linear components are combined in a static nonlinearity) have been proposed. These include “conditional whitening” in the context of the spike-triggered covariance method (Rust et al., 2005) and information maximization (Sharpee et al., 2004) (for review, see Simoncelli et al., 2004). These methods might go some way toward alleviating the problems discussed here, although they cannot resolve the issues entirely (for a discussion, see Schwartz et al., 2006). More significantly, the issue of model mismatch remains for these methods, and it is not clear whether the LNP assumptions are any easier to validate than those of STRF models. In our view, the STRF is a good starting point for modeling auditory neuronal responses, because the analysis makes a simple assumption and requires a minimum amount of data. By

examining how the linear approximation to the RF changes with modifications in the stimulus, it may be possible to explore the nature of any underlying RF nonlinearities (Kvale and Schreiner, 1997; Blake and Merzenich, 2002). Moreover, by incorporating specific nonlinearities into models that extend STRF analysis and comparing the predictive power of those models with that of STRFs, it is possible to improve our understanding of how auditory neurons process complex stimuli (Ahrens et al., 2008). Thus, STRF analysis is an important tool for characterizing auditory neuronal responses, when used with full awareness of the possible consequences of linear approximation to likely nonlinear neuronal response functions.

References

- Aertsen AM, Johannesma PI (1981) A comparison of the spectro-temporal sensitivity of auditory neurons to tonal and natural stimuli. *Biol Cybern* 42:145–156.
- Aertsen AM, Olders JH, Johannesma PI (1981) Spectro-temporal receptive fields of auditory neurons in the grassfrog. III. Analysis of the stimulus-event relation for natural stimuli. *Biol Cybern* 39:195–209.
- Ahrens M, Linden JF, Sahani M (2008) Nonlinearities and contextual influences in auditory cortical responses modeled with multilinear spectro-temporal methods. *J Neurosci*, in press.
- Bar-Yosef O, Rotman Y, Nelken I (2002) Responses of neurons in cat primary auditory cortex to bird chirps: effects of temporal and spectral context. *J Neurosci* 22:8619–8632.
- Barbour DL, Wang X (2003) Auditory cortical responses elicited in awake primates by random spectrum stimuli. *J Neurosci* 23:7194–7206.
- Blake DT, Merzenich MM (2002) Changes of AI receptive fields with sound density. *J Neurophysiol* 88:3409–3420.
- Borst A, Flanagan VL, Sompolinsky H (2005) Adaptation without parameter change: dynamic gain control in motion detection. *Proc Natl Acad Sci USA* 102:6172–6176.
- Calhoun BM, Schreiner CE (1998) Spectral envelope coding in cat primary auditory cortex: linear and nonlinear effects of stimulus characteristics. *Eur J Neurosci* 10:926–940.
- Chance FS, Abbott LF (2000) Divisive inhibition in recurrent networks. *Network* 11:119–129.
- de Boer E (1967) Correlation studies applied to the frequency resolution of the cochlea. *J Aud Res* 7:209–217.
- de Boer E, de Jongh HR (1978) On cochlear encoding: potentialities and limitations of the reverse-correlation technique. *J Acoust Soc Am* 63:115–135.
- deCharms RC, Blake DT, Merzenich MM (1998) Optimizing sound features for cortical neurons. *Science* 280:1439–1444.
- Eggermont J (1993) Wiener and Volterra analyses applied to the auditory system. *Hearing Res* 10:191–202.
- Eggermont JJ, Aertsen AM, Johannesma PI (1983) Quantitative characterization procedure for auditory neurons based on the spectro-temporal receptive field. *Hearing Res* 10:167–190.
- Escabi MA, Schreiner CE (2002) Nonlinear spectrotemporal sound analysis by neurons in the auditory midbrain. *J Neurosci* 22:4114–4131.
- Fritz J, Elhilali M, Shamma S (2005) Active listening: task-dependent plasticity of spectrotemporal receptive fields in primary auditory cortex. *Hearing Res* 206:159–176.
- Garofolo JS, Lamel LF, Fisher WM, Fiscus JG, Pallett DS, Dahlgren NL, Zue V (1993) TIMIT acoustic-phonetic continuous speech corpus. Philadelphia: Linguistic Data Consortium.
- Heeger DJ (1992) Normalization of cell responses in cat striate cortex. *Vis Neurosci* 9:181–197.
- Kadia SC, Wang X (2003) Spectral integration in A1 of awake primates: neurons with single- and multiplexed tuning characteristics. *J Neurophysiol* 89:1603–1622.
- Klein DJ, Depireux DA, Simon JZ, Shamma SA (2000) Robust spectrotemporal reverse correlation for the auditory system: optimizing stimulus design. *J Comput Neurosci* 9:85–111.
- Klein DJ, Simon JZ, Depireux DA, Shamma SA (2006) Stimulus-invariant processing and spectrotemporal reverse correlation in primary auditory cortex. *J Comput Neurosci* 20:111–136.
- Kowalski N, Depireux D, Shamma S (1996a) Analysis of dynamic spectra in ferret primary auditory cortex. I. Characteristics of single-unit responses to moving ripple spectra. *J Neurophysiol* 76:3503–3523.
- Kowalski N, Depireux D, Shamma S (1996b) Analysis of dynamic spectra in ferret primary auditory cortex. II. Prediction of unit responses to arbitrary dynamic spectra. *J Neurophysiol* 76:3524–3534.
- Kvale M, Schreiner CE (1997) Perturbative M-sequences for auditory systems identification. *Acta Acoust* 83:653–658.
- Linden JF, Liu RC, Sahani M, Schreiner CE, Merzenich MM (2003) Spectrotemporal structure of receptive fields in areas AI and AAF of mouse auditory cortex. *J Neurophysiol* 90:2660–2675.
- Machens CK, Wehr MS, Zador AM (2004) Linearity of cortical receptive fields measured with natural sounds. *J Neurosci* 24:1089–1100.
- Marmarelis P, Marmarelis V (1978) Analysis of physiological systems: the white noise approach. New York: Plenum.
- Miller LM, Escabi MA, Read HL, Schreiner CE (2002) Spectrotemporal receptive fields in the lemniscal auditory thalamus and cortex. *J Neurophysiol* 87:516–527.
- Nagel KI, Doupe AJ (2006) Temporal processing and adaptation in the songbird auditory forebrain. *Neuron* 51:845–859.
- Palm G, Pöpel B (1985) Volterra representation and Wiener-like identification of nonlinear systems: scope and limitations. *Q Rev Biophys* 18:135–164.
- Peña JL, Konishi M (2001) Auditory spatial receptive fields created by multiplication. *Science* 292:249–252.
- Qiu A, Schreiner CE, Escabi MA (2003) Gabor analysis of auditory mid-brain receptive fields: spectro-temporal and binaural composition. *J Neurophysiol* 90:456–476.
- Rust NC, Schwartz O, Movshon JA, Simoncelli EP (2005) Spatiotemporal elements of macaque V1 receptive fields. *Neuron* 46:945–956.
- Rutkowski RG, Shackleton TM, Schnupp JW, Wallace MN, Palmer AR (2002) Spectrotemporal receptive field properties of single units in the primary, dorsocaudal and ventrostral auditory cortex of the guinea pig. *Audiol Neurootol* 7:214–217.
- Sahani M, Linden JF (2003) How linear are auditory cortical responses? In: *Advances in neural information processing systems*, Vol 15 (Becker S, Thrun S, Obermayer K, eds), pp 109–116. Cambridge, MA: MIT.
- Schnupp JW, Morsic-Flogel TD, King AJ (2001) Linear processing of spatial cues in primary auditory cortex. *Nature* 414:200–204.
- Schwartz O, Pillow JW, Rust NC, Simoncelli EP (2006) Spike-triggered neural characterization. *J Vis* 6:484–507.
- Sharpee T, Rust NC, Bialek W (2004) Analyzing neural responses to natural signals: maximally informative dimensions. *Neural Comput* 16:223–250.
- Shechter B, Depireux DA (2006) Response adaptation to broadband sounds in primary auditory cortex of the awake ferret. *Hear Res* 221:91–103.
- Simoncelli EP, Pillow J, Paninski L, Schwartz O (2004) Characterization of neural responses with stochastic stimuli. In: *The cognitive neurosciences* (Gazzaniga M, ed), pp 327–338. Cambridge, MA: MIT.
- Smith E, Lewicki M (2006) Efficient auditory coding. *Nature* 439:800–805.
- Suga N, O'Neill WE, Manabe T (1978) Cortical neurons sensitive to combinations of information-bearing elements of biosonar signals in the mustache bat. *Science* 200:778–781.
- Sutter ML, Schreiner CE, McLean M, O'Connor KN, Loftus WC (1999) Organization of inhibitory frequency receptive fields in cat primary auditory cortex. *J Neurophysiol* 82:2358–2371.
- Theunissen FE, Sen K, Doupe AJ (2000) Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. *J Neurosci* 20:2315–2331.
- Valentine PA, Eggermont JJ (2004) Stimulus dependence of spectrotemporal receptive fields in cat primary auditory cortex. *Hear Res* 196:119–133.
- Wang X, Lu T, Snider RK, Liang L (2005) Sustained firing in auditory cortex evoked by preferred stimuli. *Nature* 435:341–346.
- Wehr M, Zador AM (2003) Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* 426:442–446.
- Woolley SMN, Gill PR, Theunissen FE (2006) Stimulus-dependent auditory tuning results in synchronous population coding of vocalizations in the songbird midbrain. *J Neurosci* 26:2499–2512.