

Cooperative Nonlinearities in Auditory Cortical Neurons

Craig A. Atencio, 1,2,3,4,5 Tatyana O. Sharpee, 2,4,5 and Christoph E. Schreiner 1,2,3,*

¹University of California, San Francisco/University of California, Berkeley Bioengineering Graduate Group

²W.M. Keck Foundation Center for Integrative Neuroscience

³Coleman Memorial Laboratory, Department of Otolaryngology-HNS

University of California, San Francisco, San Francisco, CA 94143, USA

⁴These authors contributed equally to this work.

⁵Present address: The Salk Institute for Biological Studies, San Diego, CA 92186, USA

*Correspondence: chris@phy.ucsf.edu

DOI 10.1016/j.neuron.2008.04.026

SUMMARY

Cortical receptive fields represent the signal preferences of sensory neurons. Receptive fields are thought to provide a representation of sensory experience from which the cerebral cortex may make interpretations. While it is essential to determine a neuron's receptive field, it remains unclear which features of the acoustic environment are specifically represented by neurons in the primary auditory cortex (AI). We characterized cat AI spectrotemporal receptive fields (STRFs) by finding both the spike-triggered average (STA) and stimulus dimensions that maximized the mutual information between response and stimulus. We derived a nonlinearity relating spiking to stimulus projection onto two maximally informative dimensions (MIDs). The STA was highly correlated with the first MID. Generally, the nonlinearity for the first MID was asymmetric and often monotonic in shape, while the second MID nonlinearity was symmetric and nonmonotonic. The joint nonlinearity for both MIDs revealed that most first and second MIDs were synergistic and thus should be considered conjointly. The difference between the nonlinearities suggests different possible roles for the MIDs in auditory processing.

INTRODUCTION

The primary auditory cortex (AI) is the main initial cortical recipient of lemniscal information and thus represents an essential station for auditory processing (Jenkins and Merzenich, 1984; Read et al., 2002). Understanding sound processing in AI has centered on characterizing the receptive fields of AI neurons. At present, however, there is no standard model of AI processing, in contrast to the primary visual cortex where the standard energy model for simple and complex cells has guided work and led to significant insights into visual processing (Adelson and Bergen, 1985; Touryan et al., 2002; Rust and Movshon, 2005).

The most widely used model to obtain the spectrotemporal receptive field (STRF) of an auditory neuron is based on the spike-

a spike (de Boer and Kuyper, 1968; Miller et al., 2002). STAs have been used to estimate many response properties of cortical cells, such as temporal and spectral modulations, stimulus selectivity, and response dynamics (deCharms et al., 1998; Theunissen et al., 2000; Depireux et al., 2001; Miller et al., 2001; Hsu et al., 2004; Woolley et al., 2005, 2006). When STAs have been used to predict responses to novel stimuli, however, they have not fully captured the processing of AI neurons (Sahani and Linden, 2003; Machens et al., 2004). This underperformance of previous STA approaches may be due to three competing issues: the STA may be biased by corre-

triggered average (STA), which is the average stimulus preceding

due to three competing issues: the STA may be biased by correlations in the stimulus ensemble, filtering was not complemented with an appropriate nonlinearity, or the STA may represent an incomplete model of AI processing.

The influence of stimulus correlations on the STA may be removed in the case of Gaussian stimuli, which are completely described by second-order correlations, though this influence may not be removed for natural stimuli, which contain higherorder correlations (Ringach et al., 2002; Paninski, 2003; Sharpee et al., 2004). In AI, neurons are driven more strongly by spectrally and temporally complex sounds rather than the pure tones generally used to define receptive fields. Thus, techniques are needed that are applicable for these more complex stimuli.

Also, while the STA may be thought of as the mean stimulus that elicits a spike, it can be considered as one feature, or stimulus dimension, that characterizes the spectrotemporal selectivity of a neuron. Other features, however, as in the case of complex cells in primary visual cortex, may be needed to adequately characterize a neuron's response (Touryan et al., 2002, 2005). Recent work has sought to overcome this limitation by combining spike triggered covariance methods with Gaussian distributed stimuli, allowing multiple stimulus features to be recovered (de Ruyter van Steveninck and Bialek, 1988; Yamada and Lewis, 1999; Rust et al., 2005; Slee et al., 2005; Fairhall et al., 2006; Maravall et al., 2007). These results, though, are well-defined only for Gaussian stimuli (Simoncelli et al., 2004; Schwartz et al., 2006).

Here, we apply a recently developed framework that avoids these constraints (Sharpee et al., 2003, 2004). By using a dynamic stimulus with short-term correlations, we first tested if the STA was an appropriate approximation to the spectrotemporal processing of Al neurons. We then determined if the model of Al processing might be extended to two features, which would represent an increase in spectrotemporal complexity.

Specifically, we modeled neurons as selective for a small number of relevant dimensions, or linear filters, out of a high-dimensional stimulus space, though within this subspace the outputs of the filters may be combined in a nonlinear manner to account for the neural responses. This linear-nonlinear model can more fully quantify the features in stimulus space that best characterize a neuron. By maximizing the mutual information between the neural responses and the projections of stimuli onto stimulus dimensions, we are able to compute these special directions in stimulus space. Since mutual information is maximized, the directions that allow us to recover the most information are the relevant dimensions. This approach removes the influence of stimulus correlations from the estimate of the relevant dimensions, and this calculation is rigorously defined even for complex, non-Gaussian sounds (Sharpee et al., 2004). In this initial application to the auditory system, we rely on the use of Gaussian stimuli to enhance comparisons to previous estimates of spectrotemporal filters. As this framework both recovers multiple spectrotemporal dimensions and provides an estimate of the nonlinearities governing the interactions between these dimensions, it provides an advance over previous work in auditory cortex in that it permits a quantitative estimate of feature selectivity and the rules by which neurons respond to stimuli.

RESULTS

Maximally Informative Dimensions versus Spike-Triggered Average

We made single-unit recordings from 247 AI neurons in the ketamine-anesthetized cat. Neurons were challenged with a continuous, dynamic, broadband moving ripple stimulus, whose spectral extent encompassed the frequency response of each neuron (Escabí and Schreiner, 2002). From the recorded neural traces single spike trains were obtained to determine the stimulus features that elicit spikes. These features were calculated with two different methods: first by computing the STA and second by computing maximally informative dimensions (MIDs). The STA model for a neuron consists of a filter followed by a nonlinearity (Figure 1A). The nonlinearity describes the response strength of the neuron as the similarity between the stimulus and the STA varies. The procedure used to obtain the MID model is shown in Figure 1B. For the case of one MID, a search is made to find the filter that maximizes the mutual information between the spiking response and the presented stimuli. Mutual information is a quantitative metric that describes the relationship between stimulus and response, and it is sensitive to the statistics of the stimulus ensemble and thus may be used with non-Gaussian stimuli. By maximizing mutual information, a filter is found, along with a nonlinearity, that provides the best model of the stimulus selectivity of the neuron. The procedure can be extended to find a second MID, which further maximizes the mutual information. This second MID was found by keeping the first MID fixed and then searching through the space of all possible dimensions to find the one that would have the most explanatory power when considered together with the first MID. During the search for two dimensions, a small percentage of computational time was spent to update the first MID. In this case the model consists of two filters, and a two-dimensional nonlinearity, which describes the stimulus selectivity of the neuron (Figure 1C). MIDs fall under the general framework of the linear-nonlinear model, which describes a set of linear filters followed by a static multidimensional nonlinearity (de Ruyter van Steveninck and Bialek, 1988; Marmarelis, 1997; Ringach, 2004; Simoncelli et al., 2004; Schwartz et al., 2006).

STAs Derived from Different Stimulus Ensembles

To determine a neuron's time-frequency receptive field, researchers have traditionally calculated the spike-triggered average (Depireux et al., 2001; Escabí and Schreiner, 2002). The STA is the average stimulus preceding a spike and is obtained by correlating stimulus segments with spike occurrences (Lee and Schetzen, 1965; de Boer and Kuyper, 1968). While the STA procedure is well defined for Gaussian stimuli, the Gaussian stimuli themselves may have different correlation structures (Schwartz et al., 2006). These structures may pose problems since stimuli with different statistics may lead to different adapted states, as shown when receptive fields were computed using natural versus white noise stimuli in the visual cortex (Sharpee et al., 2006). To determine the stability of STAs of AI neurons, we computed STAs for two types of Gaussian stimuli, dynamic moving ripple (DMR) and ripple noise (RN) (Escabí and Schreiner, 2002). DMR contains short-term correlations that approximate features of natural sounds, while RN is more noise-like and does not contain a strong correlation structure (Figures 2A and 2B; Escabí et al., 2003). The STAs for these two sets of stimuli, however, appear fundamentally similar, with only minor differences (Figures 2C1-C5 and 2D1-2D5). For the cells for which this procedure was completed the correlation between the DMR STA and the RN STA was high (Figure 2E; mean, 0.73; std, 0.079), consistent with results from the inferior colliculus (Escabí and Schreiner, 2002). Thus, in AI, STA estimates appear stable across these two stimuli and avoid potential complications that may befall other approaches (Christianson et al., 2008). We note that if receptive fields are calculated using natural stimuli, different adaptation conditions may be invoked (Sharpee et al., 2006). This may be prevalent in secondary areas which require these stimuli. Here, though, we focus solely on AI, and thus in the remaining portion of this report, we examine receptive fields calculated using the DMR.

Comparison of STAs and MIDs

Since STAs are well-defined only for Gaussian stimuli, and because they represent only one spectrotemporal feature to which a neuron is sensitive, we computed STAs and MIDs to compare the two approaches to estimating auditory receptive field models (Figure 3, each row represents one neuron). STAs display spectrotemporal diversity, with each neuron showing significantly different structure along the spectral and temporal axes, reflecting distinct spectral and temporal stimulus preferences (Figure 3, first column). The first MIDs (Figure 3, third column) represent the stimulus feature that accounts for the most information between the stimulus and the neural response. The main features in the STAs also dominate the structure of the first MID, showing that the first MID was well approximated by the



B Maximally informative dimension procedure



C Maximally informative dimension model for two filters



Figure 1. Schematic of Models of Al Receptive Fields

(A) Spike-triggered average model. The stimulus (ordinate: frequency; abscissa: time) is processed by the STA, and the probability of a spike is determined by a nonlinearity, which describes the response strength as a function of the similarity between the stimulus and the STA. The nonlinearity abscissa is in units of standard deviation. Units are calculated with respect to the expected similarity between a random stimulus and the STA. High similarity values indicate increased stimulus selectivity. Responses in the histogram are illustrative.

(B) Procedure used to calculate the MID model for a different neuron. For one MID, an iterative process is followed. A search is made for the feature, or filter, that maximizes the most mutual information between the stimulus and the spiking response. At each step of the iterative process the filter and the nonlinearity are estimated.

(C) The two MID model. The procedure in (B) can be extended to find two MIDs. In this case, after the first MID is found, it is held fixed, and an additional search is made for a second MID. Each MID linearly filters the stimulus, and the rule that governs the strength of the response will then be two-dimensional. In the non-linearity, red indicates an increase in response rate.

STA when the moving ripple stimulus was used. While differences between the first MID and STA are present, such as the slight shortening of the excitatory region for neurons that are tuned to a narrow range of frequencies (Figure 3, second and fourth rows), the STA method provides a close approximation of the first relevant stimulus dimension even when the influences of stimulus correlations have not been removed. These estimates of receptive field structure tended to be highly organized, displaying conventional excitatory-inhibitory subfields within restricted regions of time and frequency.



Figure 2. STAs Obtained Using Stimuli with Different Statistics (A) Short spectrotemporal segment of the dynamic moving ripple stimulus (DMR).

(B) Short segment of the ripple noise stimulus (RN).

(C1–C5 and D1–D5) Example STAs for five neurons. Each row represents the STAs of one neuron in response to the dynamic moving ripple (C1–C5) and to the ripple noise stimulus (D1–D5).

(E) Correlation between the dynamic moving ripple and ripple noise STAs. The STAs for the two stimuli were similar (mean correlation = 0.73; std = 0.079).

The calculation of the model for a single filter is completed when the nonlinearity has been estimated (see Experimental Procedures). The nonlinearity describes the probability of spike generation as the similarity between the filter and the stimulus varies. The nonlinearities for the STA and the first MID are shown in Figure 3, columns 2 and 4, respectively. The shape of each one-dimensional nonlinearity was characterized by an asymmetry index (ASI). The ASI is defined as (R - L)/(R + L), where R and L are the sums of all nonlinearity values corresponding to projection values greater than or less than 0, respectively. Most STA nonlinearities were asymmetric (239/247 neurons had ASIs > 0.5) and sigmoidal (monotonic) in shape. The nonlinearity for the first MID was similar in shape to the STA nonlinearity, though in some cases (<5%) it revealed a nonmonotonic function.

After calculating the first MID, we completed a two-filter model for each AI neuron by computing the second MID (Figure 3, column 5). This second filter represents another special dimension in stimulus space that further maximizes the information between single spikes and the stimulus ensemble. In general, the spectrotemporal structure of the second MID was aligned to a similar center frequency as the first MID, though the correlation coefficient between the two filters was almost zero (mean, 0.01; std, 0.14), indicative of their orthogonality. Examination of the second MID showed that this dimension exhibited a more complex and irregular composition of excitatory and inhibitory subregions as compared to the first MID. No systematic relationship between the filters for the temporal features of the two MIDs was apparent.

The nonlinearities for the first and second MIDs reveal a striking difference (Figure 3, columns 4 and 6). In the majority of cases the nonlinearity of the first MID (and of the STA) was asymmetric and sigmoidal, while the nonlinearity of the second MID was approximately symmetric and nonsigmoidal (nonmonotonic). The symmetric nature of the second MID nonlinearity indicates an increased probability of firing when a stimulus is either correlated or anticorrelated with the filter. Because the response is invariant to the phase of the stimulus envelope, and thus to the alignment of the stimulus to the features of the second MID, the second MID does not contribute to the STA. Stimuli that are uncorrelated with the filter lead to a decreased probability of a response compared to the mean probability of a spike (represented by the dashed lines in Figure 3, columns 2, 4, and 6). The shape and structure of this nonlinearity is reminiscent of those seen for complex visual cells (Touryan et al., 2002).

In summary, the examples show that one filter accounted for the response in an asymmetric fashion that was proportional to the similarity between the filter and the stimulus. The second filter influenced firing by increasing the response probability in a manner independent of the sign of the spectrotemporal contrast in the stimulus.

Cooperative Linear-Nonlinear Model

The two MID filters and their corresponding one-dimensional nonlinearities do not constitute the complete receptive field model. While each filter may independently account for some of the mutual information encoded by a neuron, the processing of each may not be independent of the other. The rule that describes how the combined outputs of the filters lead to a spike



Figure 3. Example STAs, MIDs, and Associated Nonlinearities

Each row represents one neuron. STAs and MIDs in columns 1, 3, and 5 have frequency on the ordinate and time on the abscissa. Specific time-frequency patterns of stimulus energy lead to increased (red) or decreased (blue) responsiveness of each neuron and may be interpreted as excitatory or inhibitory regions. One-dimensional nonlinearities are shown in columns 2, 4, and 6. The ordinate represents the firing rate given the similarity between the stimulus and the filter. The dashed line represents the average firing rate over the complete stimulus presentation. Column 1: STAs have distinct excitatory and inhibitory regions. Column 2: STA nonlinearities were usually asymmetric. Column 3: first MIDs, which represent the stimulus feature that best accounts for the neuron's responses, resembled the STAs, though the temporal duration of excitation/inhibition is often decreased due to the removal of stimulus correlations. Column 4: first MID nonlinearities were similar to STA nonlinearities and were often asymmetric. Column 5: second MIDs revealed additional stimulus preferences and were not well predicted by the structure of the first MID. Column 6: second MID nonlinearities, which were usually symmetric, are structurally different from those of the first MIDs. Column 7: two-dimensional nonlinearities for the two MIDs. The ordinate/abscissa represents the similarity between the second/first MID and the stimulus. The color range (blue to red) indicates the response strength (low to high). Nonlinearities were not spatially uniform, indicating specific stimulus configurations lead to the largest response. Error bars are SEM.

can be quantified by computing a two-dimensional nonlinearity for each neuron (Figure 3, column 7). This joint nonlinearity is dependent on the simultaneous projection of the stimulus onto both the first and second MID. These examples reveal a crescent-shaped joint nonlinearity function, as expected for a neuron with a sigmoidal nonlinearity for the first filter and a symmetric nonlinearity for the second filter (Rust et al., 2005; Fairhall et al., 2006). In addition, the firing probability for the two-dimensional nonlinearity is regionally greater compared to either onedimensional nonlinearity alone.

Comparison between the STA and the First MID

The similarity between the STA and the first MID was quantified by correlating the STA and the first MID (Figure 4A). For the majority of filters, the correlations were high (mean, 0.71; std, 0.16), indicating similar spectrotemporal structure. For the nonlinearities, the correlations were similarly high (Figure 4B; mean, 0.74; std,

MID capture the same basic spectrotemporal processing patterns of AI neurons, though with the significant difference that the MID estimate is impervious to potential biases induced by non-Gaussian stimuli and stimulus correlations on the calculation of the STA. We calculated the mutual information accounted for by the STA and the first MID using a novel stimulus test set (see Experimental Procedures). Since the test set was not used to construct the STA or the first MID, it is not guaranteed that the first MID information will exceed that for the STA. These information calculations furnish a robust metric to evaluate how well each filter (plus nonlinearity) captures the actual processing of a neuron (Adelman et al., 2003). In all cases, the first MID accounted for more information than the STA (Figure 4C). We note that this result is not true by definition, since the first MID may possibly be obtained through overfitting, and thus the STA might have more information than the first MID. We found that the STA information approached

0.28). These two results demonstrate that the STA and the first



Figure 4. Comparison between STAs and First MIDs

(A) Correlation between the STA and MID1 filters. For most neurons the spectrotemporal structure of the STAs and first MIDs were highly correlated.
(B) Correlation between the STA and MID nonlinearities. The structure of the STA and MID1 nonlinearities were also similar.

(C) Comparison between the STA and MID mutual information (bits/spike). The first MID was able to account for more information than the STA.

(D) Relative comparison of STA mutual information to that of the first MID (median = 63%).

that of the first MID (Figure 4D; median, 63%), indicating that the performance of the STA was less than the ability of the first MID to predict single spikes from the stimulus ensemble.



To further establish these results, we computed the variance in the firing rates predicted by the STA and by the first MID for the test data set (Paninski, 2003; Sharpee, 2007). We formed the ratio of the first MID to STA variances and found that a higher percentage of the firing rate variance was described by the first MID compared to the STA (mean percentage, 185; std, 142; p = 0.011, t test), consistent with the higher MID mutual information compared to the STA.

One-Dimensional Nonlinearity Structure

For STAs, the majority of nonlinearities were asymmetric (ASI mean, 0.78; std, 0.14), with similar but slightly smaller values for the first MID (ASI mean, 0.57; std, 0.23; Figures 5A, 5B, and 5D). The second MID nonlinearities were on the average symmetric (Figure 5C; ASI mean, 0.01; std, 0.27). This implies that a typical AI neuron, in the extended phenomenological receptive field model, contains functional subunits that both threshold (first MID) and square (second MID) the outputs of the individual filters (Figure 5E).

Two-Dimensional Nonlinearity Structure

The rule that describes how the combined outputs of the MID filters lead to a spike was quantified by computing a two-dimensional nonlinearity for each neuron (see Figure 3, column 7, for examples). This joint nonlinearity is dependent on the simultaneous projection of the stimulus onto both the first and second MID. The two-dimensional nonlinearities in Figure 3 reveal a crescent-shaped function, as expected for a neuron with a sigmoidal nonlinearity for the first filter and a symmetric nonlinearity for the second filter (Rust et al., 2005; Fairhall et al., 2006). In addition, the firing rate for the two-dimensional nonlinearity is regionally greater compared to either one-dimensional nonlinearity alone

Figure 5. Population Analysis of 1D Nonlinearity Structure for STAs and MIDs

Asymmetry of the 1D nonlinearities was determined by comparing the values corresponding to positive similarities to those for negative similarities. Asymmetry Index values (ASIs) near +1 or -1 indicate highly asymmetric nonlinearities. ASIs near 0 indicate a symmetry.

(A–C) Frequency histogram of the asymmetry of the nonlinearities. The STA and MID_1 nonlinearities are highly asymmetric, while the MID_2 nonlinearities were symmetric.

(D) Comparison of STA and first MID nonlinearity structure. Most AI neurons that had an STA with an asymmetric nonlinearity also had an MID₁ with a similarly structured nonlinearity. (E) Comparison of first and second MID nonlinearity structure. Auditory neurons usually had an MID₁ nonlinearity that was asymmetric paired with an MID₂ nonlinearity that was symmetric.



Figure 6. Population Analysis of the Structure of 2D Nonlinearities

(A) Peak response rates in the two-dimensional nonlinearities are plotted against the sum of the peak rates in each of the MID₁ and MID₂ onedimensional nonlinearities. For most neurons the joint nonlinearity contained a higher firing rate, indicating that special stimulus configurations may excite AI neurons in a manner that cannot be predicted by the independent processing of the MIDs. (B) Structural analysis of two-dimensional nonlinearities. The frequency histogram of the inseparability of the two-dimensional nonlinearities across all AI neurons is shown. Inseparability index values near 0 indicate that the 2D nonlinearity can be approximated by a product of two

1D nonlinearities. For most AI neurons this approximation was inappropriate, as indicated by the nonzero mode of the distribution, implying that the joint stimulus processing of the two MIDs contains more information than the independent processing of each MID.

(Figure 3). This increased response probability is not simply a product of the different and independent nonlinearities but reflects a cooperativity, or synergy, when both filters act simultaneously on the stimulus. Indeed, for most neurons the peak response rate in the two-dimensional nonlinearity is greater than the sum of the peak rates in each of the two one-dimensional MID nonlinearities (Figure 6A). Thus, the joint probability distribution governing the response contains more information than a simple product of each filter's one-dimensional nonlinearity.

The degree of cooperativity in the nonlinearity can be quantified by determining its separability into its marginals. Twodimensional nonlinearities that are highly separable imply, if stimuli are uncorrelated, that the processing of the neuron is well characterized by two independent filters. Highly inseparable nonlinearities may imply that the probability distribution describing the neural response cannot be captured by a simple product of two one-dimensional functions. Therefore, the joint distribution likely includes more information than a product of its marginals. An inseparability index quantifies the degree of separability for the population of AI neurons (Figure 6B). Inseparability indices near 0 indicate a nonlinearity that can be described as a product of two independent distributions. Values near 1 indicate the inadequacy of this description. Inseparability indices in AI ranged from nearly separable to moderately inseparable (mean, 0.27; std, 0.10), though all neurons showed some degree of inseparability, indicating that it is more appropriate to characterize the two filters by a two-dimensional rule and not simply by the product of the two one-dimensional nonlinearities. Another way to check for the degree of synergy between the two relevant dimensions is to compare the information captured by these two features together or separately. We consider this next.

Model Comparison

Given a more complete two-dimensional linear-nonlinear receptive field model, the performance of the model can be quantified and tested against other possible AI receptive field models. A natural measure of performance is the mutual information between the stimulus projections onto the MIDs and the neural responses (Adelman et al., 2003; Sharpee et al., 2003; Fairhall et al., 2006), since mutual information is sensitive to subtle differences in the probability distributions which govern the likelihood of a neural response. In the results that follow, information values are directly related to the predictive power of each model, since for each model the values were calculated using a novel test stimulus set that was not used to estimate the MIDs (see Experimental Procedures). The results of this analysis show that the information captured by the two-filter model (Figure 7A) in every case exceeds the one-filter model. The one-filter model usually accounted for 62% (population median) of the information in the two-filter model (Figure 7B). This result provides insight into the number of features needed to adequately represent the response properties of AI neurons. For some neurons the STA, or one-dimensional model, is adequate, though a clear advantage is conferred on the predictions based on a model with at least two maximally informative dimensions.

For every neuron, the first dimension accounted for more information than the second, in agreement with the definition (Figure 7C). The information contributed by the second filter was approximately 25% (median) of the first filter's information (Figure 7D).

The joint two-filter model provides more information than the sum of two independent filters. The ratio of the joint information to the sum of the independently calculated information describes the amount of synergy in the processing of the two filters (Figure 7F). The synergy also quantifies the earlier observation that the two-dimensional nonlinearities of AI neurons are partially inseparable (Figure 6B). For most neurons, the receptive field model with the joint two-dimensional nonlinearity accounts for more information than a model of independently calculated filters and nonlinearities (Figures 7E and 7F). If the joint and independent MID receptive field models account for the same amount of information, the synergy is 100%. The median synergy was 128%, which was significantly greater than 100% (p < 0.01, Wilcoxon signed-rank test; Bain and Engelhardt, 1992). Thus, in most cases there is synergy between the maximally informative dimensions, as demonstrated by the points above the solid diagonal line that represents unity slope (Figure 7E). For 37% of neurons, the synergy exceeded 150% (Figure 7F). Thus, the responses of these neurons are most accurately modeled from a nonlinear rule that is a function of the joint stimulus processing of the two filters.

The relative contribution of each MID to the jointly conveyed information predicts the degree of synergy (Figure 8). Comparing



Figure 7. Population Analysis of Information Processing of Maximally Informative Dimensions

(A) Comparison of mutual information (bits/spike) for the two MID model versus the one MID model. The two MID model (ordinate) always accounted for more information than the single MID model (abscissa).

(B) Frequency histogram of the relative contribution of the first MID to the two MID model information. The MID₁ accounted for approximately 62% of the information in the combined model.

(C) Comparison between the first MID and second MID information. The MID_1 information was greater than the MID_2 information.

(D) Frequency histogram of the relative comparison between the first and second MID information. The MID_2 information was approximately 30% of the MID_1 information.

(E) Comparison between the information of the two MID model when the filters process stimuli jointly versus independently. The combined, joint processing model accounted for more information than a model of two independently processing MIDs.

(F) Frequency histogram of the cooperative processing, or synergy, of the two MIDs. The median synergy between the MIDs was 128%; for 37% of neurons it exceeded 150%.

the degree of single-dimension contribution to the synergy between maximally informative dimensions reveals a strong negative correlation (r = -0.96, p < 0.01, t test). In addition, a tendency exists that the more balanced the contribution of the two linear filters the higher is the synergy between the processing of each filter (r = 0.26, p < 0.0003, t test). These findings suggest that as the complexity of the spectrotemporal processing of Al neurons increases, the less independent is the processing by each



Figure 8. Comparison between Synergy and First MID Contribution Ordinate: synergy. Abscissa: ratio of the first MID information to the two MID information. The correlation between the data points was significant (r = -0.96, p < 0.01, t test).

filter. As a consequence, the amount of information conveyed by the interaction between the various filters is considerably greater than the independent processing of each filter by itself.

DISCUSSION

The purpose of this study was to apply a linear-nonlinear framework to the analysis of AI spectrotemporal receptive fields that combines multiple linear filters with nonlinear processing stages. The traditional approach to analyzing the spectrotemporal processing of AI neurons has been to calculate the spike triggered average, or linear spectrotemporal receptive field. The STA represents the mean stimulus that elicits a spike, though it may be biased by stimulus correlations and statistics, and it describes only one spectrotemporal stimulus feature that influences the neuron's response. By applying the linear-nonlinear model to describe AI receptive fields, previous limitations of spike-triggered techniques were largely avoided. The linear receptive field portion was determined by maximizing the mutual information between the estimated receptive field and the neural response (Sharpee et al., 2003, 2004). An inherent feature of the methodology is that it corrects for stimulus correlations.

This framework revealed that STAs are in reasonable agreement with the first maximally informative dimension that emerged in the linear-nonlinear model (Figure 4). Further, across every tested neuron for which a significant number of spikes were obtained, there were no cases in which a first informative dimension was found but an STA was not (for six cells significant spiking was observed, though no STAs or MIDs emerged). This stands in contrast to the situation in the visual cortex, where complex cells cannot be analyzed using spike-triggered averaging but may be quantified using more nonlinear techniques (Touryan et al., 2002, 2005; Sharpee et al., 2004; Felsen et al., 2005; Rust et al., 2005). The accompanying nonlinearities of the STA and the first MID were also highly correlated. Thus, at the level of AI the STA is a good approximation to the first maximally informative dimension, and in the case of ripple stimuli commonly used in AI studies this provides validation for previous receptive field estimates (Klein et al., 2000; Miller et al., 2001, 2002; Fritz et al., 2005).

While the first maximally informative dimension accounts for the main feature in stimulus space to which a neuron responds, other contributing stimulus features may exist. Maximizing information allowed the derivation of a second relevant dimension in the linear-nonlinear model. This second MID influences a neuron's firing, though in a manner substantially different from the first relevant dimension. The second MID processes stimuli in an envelope-phase invariant manner by increasing the probability of a spike regardless of the contrast polarity of the stimulus. The only requirement is that the stimulus be moderately correlated or anticorrelated with this second dimension (Figure 3). This style of processing is reminiscent of the nonlinearities associated with complex cells or of those in current extended models of simple cell processing in the visual cortex, where the nonlinearities are estimated as even ordered, or approximately squaring, functions (Rust et al., 2005; Chen et al., 2007).

The shape of STAs may depend on the stimulus set used to obtain the filter estimates, especially when using non-Gaussian statistics (Sharpee et al., 2006). STAs obtained for different Gaussian stimuli showed no clear differences (Figure 2; Escabí and Schreiner, 2002). MIDs are likely to be more stable than STAs as shown by a study in the visual cortex even for non-Gaussian stimuli due to the inherent properties of the analysis framework (Sharpee et al., 2006). We therefore expect that in the auditory cortex, any potential complications that may befall other approaches, due to nonlinear interactions between filter components (Christianson et al., 2008), will be minimized by the MID methodology and likely be of minor consequence for the results and the interpretations presented here.

Identifying two relevant stimulus dimensions permits the computation of the joint two-dimensional nonlinearity which governs the probability of spiking and describes a rule by which neural responsiveness can be quantified. The joint nonlinearity was not simply a product of each relevant dimension's individual nonlinearity but reflected specific, nonpredictable interactions that resulted in a cooperative or synergistic response behavior. Significant synergy, in excess of 125% for 53% of the neurons (Figure 7), imposes a requirement for accurate modeling of AI spiking responses. Models that do not account for this previously unknown computational richness may not perform well simply because they lack the ability to quantify how the different dimensions interact to influence neural responsiveness, and thus the predictive power of these models will be diminished.

Analysis of the contribution of each filter to the stimulus information in the spike train showed that the relative contributions fell along a continuum, perhaps in analogy to the classification of simple and complex cells in the visual cortex, which may also fall along a continuum (Skottun et al., 1991; Mechler and Ringach, 2002; Priebe et al., 2004). An implication for Al is that for some neurons the first dimension, or even the STA, is a sufficient approximation to the overall processing. For many other neurons, however, a greater diversity of computation is seen. For these neurons, one dimension is not an adequate description, and the standard spike-triggered average model is not appropriate. Coupled with the insight that the multiple stimulus dimensions of many Al neurons operate synergistically, it becomes clear that the extended linear-nonlinear model provides a much more complete understanding of neuronal coding strategies (Figure 8). The application of the linear-nonlinear model to Al using information maximization represents a significant advance over previous approaches, since it permits quantification of several stimulus features that influence a neuron's firing, how feature selective a neuron is, and how this feature selectivity interacts synergistically to influence a neuron's response. This behavior and its quantification can be construed as a generalization of the combination-sensitivity principle that has been demonstrated for cortical and subcortical neurons in acoustically specialized animals such as bats and birds (Suga et al., 1978; Margoliash and Fortune, 1992; Portfors and Wenstrup, 2002).

A functional interpretation of the components of the two-filter model remains speculative, especially in light of the strong interaction between the two nonlinearities. However, based on insights from the simple and complex cells in the visual system (Felsen et al., 2005), it is not unreasonable to postulate that the first filter in combination with the asymmetric nonlinearity may be a feature detector tuned to a narrow range of stimulus constellations and with high sensitivity to the envelope phase. The closer the match between stimulus and filter, the stronger is the response. By contrast, the second filter, with the associated symmetric nonlinearity, may correspond more to an envelope phase-insensitive detector tuned to a broader range of stimulus parameter constellations and acting to improve the saliency of auditory features (Felsen et al., 2005).

EXPERIMENTAL PROCEDURES

Electrophysiology

Electrophysiological methods and stimulus design are similar to previous reports (Miller and Schreiner, 2000; Miller et al., 2002). Young adult cats were given an initial dose of ketamine (22 mg/kg) and acepromazine (0.11 mg/kg), and anesthesia was maintained with pentobarbital sodium (Nembutal, 15–30 mg/kg) during the surgical procedure. The animal's temperature was maintained with a thermostatic heating pad. Bupivicaine was applied to incisions and pressure points. Surgery consisted of a tracheotomy, reflection of the soft tissues of the scalp, craniotomy over AI, and durotomy. After surgery, the animal was maintained in an areflexive state with a continuous infusion of ketamine/diazepam (2–10 mg/kg/hr ketamine, 0.05–0.2 mg/kg/hr diazepam in lactated Ringer solution). All procedures were in strict accordance with the University of California, San Francisco Committee for Animal Research.

All recordings were made with the animal in a sound-shielded anechoic chamber (IAC, Bronx, NY), with stimuli delivered via a closed speaker system (diaphragms from Stax, Japan). Simultaneous extracellular recordings were made using multichannel silicon recording probes (kindly provided by the University of Michigan Center for Neural Communication Technology). The probes contained 16 linearly spaced recording channels, with each channel separated by 150 μ m. The impedance of each channel was 2–3 M Ω . Probes were carefully positioned orthogonally to the cortical surface and lowered to depths between 2300 and 2400 μ m using a microdrive (David Kopf Instruments, Tujunga, CA).

Neural traces were band-pass filtered between 600 and 6000 Hz and were recorded to disk with a Neuralynx Cheetah A/D system at sampling rates between 18 and 27 kHz. The traces were sorted off-line with a Bayesian spike-sorting algorithm (Lewicki, 1994). Each probe penetration yielded 8–16 active channels, with ~1–2 single units per channel. Stimulus-driven neural activity was recorded for ~75 min at each location.

Stimulus

For any recording position neurons were probed with pure tones, then two presentations of a 15 or 20 min dynamic moving ripple stimulus, followed by approximately 20 min of complete silence, during which time spontaneous activity was recorded. Each pure tone was presented five times. The level and frequency of each pure tone was chosen randomly from 15 different levels

(5 dB SPL spacing) and 45 different frequencies. The dynamic ripple stimulus, which has a Gaussian amplitude distribution, was a temporally varying broadband sound (500-20,000 or 40,000 Hz) composed of approximately 50 sinusoidal carriers per octave, each with randomized phase (Escabí and Schreiner, 2002). The magnitude of a carrier at any time was modulated by the spectrotemporal envelope, consisting of sinusoidal amplitude peaks ("ripples") on a logarithmic frequency axis that changed over time. Two parameters defined the envelope: a spectral and a temporal modulation parameter. Spectral modulation rate was defined by the number of spectral peaks per octave, or ripple density. Temporal modulations were defined by the speed and direction of the peaks' change. Both the spectral and temporal modulation parameters were varied randomly and independently during the nonrepeating stimulus. Spectral modulation rate varied slowly (max. rate of change 1 Hz) between 0 and 4 cycles per octave; the temporal modulation rate varied between -40 Hz (upward sweep) and 40 Hz (downward sweep), with a maximum 3 Hz rate of change. Both parameters were statistically independent and unbiased within those ranges. Maximum modulation depth of the spectrotemporal envelope was 40 dB. Mean intensity was set at 70 or 80 dB SPL. Ripple noise (RN) stimuli were constructed from the sum of 16 independent DMR stimuli resulting in a stimulus of similar spectral and temporal content as the DMR stimuli but with reduced local correlations (Escabí and Schreiner, 2002).

Analysis

Data analysis was carried out in MATLAB (Mathworks, Natick, MA). Before receptive field analysis the ripple stimulus was downsampled to have a resolution of 12 carriers per octave spectrally and 5 ms temporally. We first used the reverse correlation method to derive the spectrotemporal receptive field (STRF), which is the average spectrotemporal stimulus envelope immediately preceding a spike (STA) (Aertsen and Johannesma, 1980; deCharms et al., 1998; Klein et al., 2000; Theunissen et al., 2000; Escabí and Schreiner, 2002). Positive (red) regions of the STA indicate that stimulus energy at that frequency and time tended to increase the neuron's firing rate, and negative (blue) regions indicate where the stimulus envelope induced a decrease in firing rate. For each STA we computed the nonlinear function that related the stimulus to the probability of a spike. Each stimulus segment s that elicited a spike was projected onto the STA via the inner product $z = s \cdot STA$. The projection values were then binned to obtain the probability distribution P(z|spike). We then projected all stimuli onto the STA without regard to a spike occurrence and formed the prior stimulus distribution, P(z). The projection values that comprised P(z|spike) and P(z) were transformed to units of standard deviation by normalizing relative to the mean, μ , and standard deviation, σ , of P(z), using $x = (z - \mu)/\sigma$. The nonlinearity for the STA was then computed as $P(spike|x) = P(spike) \frac{P(x|spike)}{P(x)}$, where P(spike) is the average firing rate of the neuron (Aguera y Arcas et al., 2003).

To obtain the MIDs, we followed previously reported methodologies (Sharpee et al., 2004, 2006). The first MID is the direction in stimulus space that accounts for the most mutual information between stimulus and response. The first MID was obtained through an iterative procedure, where the relevance of any "candidate" dimension v_1 was quantified by computing the mutual information between the occurrence of single spikes and projections of the stimulus, s, onto v_1 . We searched through different directions in the stimulus space until convergence. This estimation procedure automatically corrects for stimulus correlations. The second MID was then found as the dimension in the stimulus space that, together with the first MID, further maximized the information. Further MIDs were not calculated due to data limitations and computational considerations. One-dimensional nonlinearities for the first and second MIDs were computed in the same manner as the nonlinearity for the STA.

The mutual information between projections onto individual filters and single spikes was computed according to $I(v) = \int dx P(x|spike) \log_2 \left[\frac{P(x|spike)}{P(x)}\right]$. The filter *v* was either the STA, the first MID, or the second MID. The mutual information between single spikes and both MIDs was calculated as $I(MID_1, MID_2) = \iint dx_1 dx_2 P(x_1, x_2|spike) \log_2 \left[\frac{P(x_1, x_2|spike)}{P(x_1, x_2)}\right]$, where x_1 and x_2 represent the projections of the stimulus onto the first and second MIDs, respectively. The two-dimensional nonlinearity was calculated via $P(spike|x_1, x_2) = P(spike) \frac{P(x_1, x_2|spike)}{P(x_1, x_2)}$.

All estimates of relevant stimulus dimensions (STA, first and second MID) were computed as an average of four jackknife estimates (Efron and Tibshirani, 1994). Each jackknife estimate was computed by using a different $\frac{3}{4}$ of the data (the training data set), thus leaving a different $\frac{1}{4}$ of the data as a test data set. The test data set was used for estimating the information values accounted for by the jackknife estimate of the relevant dimensions. Information values were calculated using different fractions of the test data set for each neuron. To accomplish this, the information values were calculated over the first 50%, 60%, 70%, 80%, 90%, 92.5%, 95%, 97.5%, and 100% of the test data set. The information calculated from these data fractions was plotted against the inverse of the data fraction percentage (1/50, 1/60, 1/70, etc.). We extrapolated the information values to infinite data set size by fitting a line to the plot and taking the ordinate intersect as the information value for unlimited data size.

The synergy between the two MIDs was defined as $100\frac{I(MD_2,MD_2)}{I(MD_1)+I(MD_2)}$, where each mutual information value was obtained via the extrapolation procedure. The shape of each one-dimensional nonlinearity was characterized by an

asymmetry index (ASI). The ASI is defined as (R - L)/(R + L), where R and L are the sums of all nonlinearity values that correspond to projection values greater than or less than 0, respectively. The index ranges from -1 to 1, with 0 representing a nonlinearity that is completely symmetric for positive and negative projection values, implying that correlated and anticorrelated stimuli equally influence the probability of a neural response. ASIs near 1 or -1 indicate neurons that have an increased probability of spiking when the stimulus is either positively or negatively correlated with the filter, respectively.

The inseparability of the two-dimensional nonlinearity was determined by performing singular value decomposition on $P(spike|x_1, x_2)$ (Depireux et al., 2001). The inseparability index was defined as $1 - \sigma_1^2 / \sum_i \sigma_i^2$ where σ_1 is the largest singular value. The inseparability index, which ranges between 0 and 1, describes how well $P(spike|x_1, x_2)$ may be described by a product of two one-dimensional nonlinearities, with values near 0 corresponding to a nonlinearity for which this description is appropriate.

ACKNOWLEDGMENTS

We thank Andrew Tan, Marc Heiser, Kazuo Imaizumi, and Benedicte Philibert for experimental assistance; Mark Kvale for the use of his SpikeSort 1.3 Bayesian spike-sorting software; and Rob Froemke and Robert Liu for earlier comments on the manuscript. This work was supported by National Institutes of Health grants DC02260 and MH077970 and by Hearing Research Inc. (San Francisco, CA).

Received: October 11, 2007 Revised: January 31, 2008 Accepted: April 22, 2008 Published: June 25, 2008

REFERENCES

Adelman, T.L., Bialek, W., and Olberg, R.M. (2003). The information content of receptive fields. Neuron 40, 823–833.

Adelson, E.H., and Bergen, J.R. (1985). Spatiotemporal energy models for the perception of motion. J. Opt. Soc. Am. A 2, 284–299.

Aertsen, A., and Johannesma, P.I. (1980). Spectro-temporal receptive fields of auditory neurons in the grassfrog. I. Characterization of tonal and natural stimuli. Biol. Cybern. 38, 223–234.

Aguera y Arcas, B., Fairhall, A.L., and Bialek, W. (2003). Computation in a single neuron: Hodgkin and Huxley revisited. Neural Comput. *15*, 1715–1749.

Bain, L.J., and Engelhardt, M. (1992). Introduction to Probability and Mathematical Statistics, Second Edition (Boston: PWS-KENT).

Chen, X., Han, F., Poo, M.M., and Dan, Y. (2007). Excitatory and suppressive receptive field subunits in awake monkey primary visual cortex (V1). Proc. Natl. Acad. Sci. USA *104*, 19120–19125.

Christianson, G.B., Sahani, M., and Linden, J.F. (2008). The consequences of response nonlinearities for interpretation of spectrotemporal receptive fields. J. Neurosci. 28, 446–455.

de Boer, R., and Kuyper, P. (1968). Triggered correlation. IEEE Trans. Biomed. Eng. *15*, 169–179.

deCharms, R.C., Blake, D.T., and Merzenich, M.M. (1998). Optimizing sound features for cortical neurons. Science 280, 1439–1443.

Depireux, D.A., Simon, J.Z., Klein, D.J., and Shamma, S.A. (2001). Spectrotemporal response field characterization with dynamic ripples in ferret primary auditory cortex. J. Neurophysiol. *85*, 1220–1234.

Efron, B., and Tibshirani, R.J. (1994). An Introduction to the Bootstrap (Boca Raton: Chapman & Hall/CRC).

Escabí, M.A., and Schreiner, C.E. (2002). Nonlinear spectrotemporal sound analysis by neurons in the auditory midbrain. J. Neurosci. 22, 4114–4131.

Escabí, M.A., Miller, L.M., Read, H.L., and Schreiner, C.E. (2003). Naturalistic auditory contrast improves spectrotemporal coding in the cat inferior colliculus. J. Neurosci. *23*, 11489–11504.

Fairhall, A.L., Burlingame, C.A., Narasimhan, R., Harris, R.A., Puchalla, J.L., and Berry, M.J., 2nd. (2006). Selectivity for multiple stimulus features in retinal ganglion cells. J. Neurophysiol. *96*, 2724–2738.

Felsen, G., Touryan, J., Han, F., and Dan, Y. (2005). Cortical sensitivity to visual features in natural scenes. PLoS Biol. 3, e342. 10.1371/journal.pbio.0030342.

Fritz, J.B., Elhilali, M., and Shamma, S.A. (2005). Differential dynamic plasticity of A1 receptive fields during multiple spectral tasks. J. Neurosci. *25*, 7623–7635.

Hsu, A., Woolley, S.M., Fremouw, T.E., and Theunissen, F.E. (2004). Modulation power and phase spectrum of natural sounds enhance neural encoding performed by single auditory neurons. J. Neurosci. 24, 9201–9211.

Jenkins, W.M., and Merzenich, M.M. (1984). Role of cat primary auditory cortex for sound-localization behavior. J. Neurophysiol. *52*, 819–847.

Klein, D.J., Depireux, D.A., Simon, J.Z., and Shamma, S.A. (2000). Robust spectrotemporal reverse correlation for the auditory system: optimizing stimulus design. J. Comput. Neurosci. *9*, 85–111.

Lee, Y.W., and Schetzen, M. (1965). Measurement of the Wiener kernels of a non-linear system by cross-correlation. Int. J. Control *2*, 237–254.

Lewicki, M.S. (1994). Bayesian modeling and classification of neural signals. Neural Comput. *6*, 1005–1030.

Machens, C.K., Wehr, M.S., and Zador, A.M. (2004). Linearity of cortical receptive fields measured with natural sounds. J. Neurosci. 24, 1089–1100.

Maravall, M., Petersen, R.S., Fairhall, A.L., Arabzadeh, E., and Diamond, M.E. (2007). Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. PLoS Biol. *5*, e19. 10.1371/journal. pbio.0050019.

Margoliash, D., and Fortune, E.S. (1992). Temporal and harmonic combination-sensitive neurons in the zebra finch's HVc. J. Neurosci. *12*, 4309–4326.

Marmarelis, V.Z. (1997). Modeling methodology for nonlinear physiological systems. Ann. Biomed. Eng. 25, 239–251.

Mechler, F., and Ringach, D.L. (2002). On the classification of simple and complex cells. Vision Res. 42, 1017–1033.

Miller, L.M., and Schreiner, C.E. (2000). Stimulus-based state control in the thalamocortical system. J. Neurosci. 20, 7011–7016.

Miller, L.M., Escabí, M.A., and Schreiner, C.E. (2001). Feature selectivity and interneuronal cooperation in the thalamocortical system. J. Neurosci. *21*, 8136–8144.

Miller, L.M., Escabí, M.A., Read, H.L., and Schreiner, C.E. (2002). Spectrotemporal receptive fields in the lemniscal auditory thalamus and cortex. J. Neurophysiol. 87, 516–527.

Paninski, L. (2003). Convergence properties of three spike-triggered analysis techniques. Network *14*, 437–464.

Portfors, C.V., and Wenstrup, J.J. (2002). Excitatory and facilitatory frequency response areas in the inferior colliculus of the mustached bat. Hear. Res. *168*, 131–138.

Priebe, N.J., Mechler, F., Carandini, M., and Ferster, D. (2004). The contribution of spike threshold to the dichotomy of cortical simple and complex cells. Nat. Neurosci. 7, 1113–1122.

Read, H.L., Winer, J.A., and Schreiner, C.E. (2002). Functional architecture of auditory cortex. Curr. Opin. Neurobiol. *12*, 433–440.

Ringach, D.L. (2004). Mapping receptive fields in primary visual cortex. J. Physiol. 558, 717–728.

Ringach, D.L., Hawken, M.J., and Shapley, R. (2002). Receptive field structure of neurons in monkey primary visual cortex revealed by stimulation with natural image sequences. J. Vis. 2, 12–24.

Rust, N.C., and Movshon, J.A. (2005). In praise of artifice. Nat. Neurosci. 8, 1647–1650.

Rust, N.C., Schwartz, O., Movshon, J.A., and Simoncelli, E.P. (2005). Spatiotemporal elements of macaque v1 receptive fields. Neuron *46*, 945–956.

Sahani, M., and Linden, J.F. (2003). How linear are auditory cortical responses? In Advances in Neural Information Processing Systems, S. Becker, S. Thrun, and K. Obermayer, eds. (Cambridge, MA: MIT Press), pp. 109–116.

Schwartz, O., Pillow, J.W., Rust, N.C., and Simoncelli, E.P. (2006). Spike-triggered neural characterization. J. Vis. 6, 484–507.

Sharpee, T.O. (2007). Comparison of information and variance maximization strategies for characterizing neural feature selectivity. Stat. Med. *26*, 4009–4031.

Sharpee, T., Rust, N.C., and Bialek, W. (2003). Maximally informative dimensions: analyzing neural responses to natural signals. In Advances in Neural Information Processing 15, S. Becker, S. Thrun, and K. Obermayer, eds. (Cambridge, MA: MIT Press).

Sharpee, T., Rust, N.C., and Bialek, W. (2004). Analyzing neural responses to natural signals: maximally informative dimensions. Neural Comput. *16*, 223–250.

Sharpee, T.O., Sugihara, H., Kurgansky, A.V., Rebrik, S.P., Stryker, M.P., and Miller, K.D. (2006). Adaptive filtering enhances information transmission in visual cortex. Nature 439, 936–942.

Simoncelli, E.P., Paninski, L., Pillow, J., and Schwartz, O. (2004). Characterization of neural responses with stochastic stimuli. In The New Cognitive Neurosciences, Third Edition, M. Gazzaniga, ed. (Cambridge, MA: MIT Press).

Skottun, B.C., De Valois, R.L., Grosof, D.H., Movshon, J.A., Albrecht, D.G., and Bonds, A.B. (1991). Classifying simple and complex cells on the basis of response modulation. Vision Res. *31*, 1079–1086.

Slee, S.J., Higgs, M.H., Fairhall, A.L., and Spain, W.J. (2005). Two-dimensional time coding in the auditory brainstem. J. Neurosci. *25*, 9978–9988.

Suga, N., O'Neill, W.E., and Manabe, T. (1978). Cortical neurons sensitive to combinations of information-bearing elements of biosonar signals in the mustache bat. Science 200, 778–781.

Theunissen, F.E., Sen, K., and Doupe, A.J. (2000). Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. J. Neurosci. 20, 2315–2331.

Touryan, J., Lau, B., and Dan, Y. (2002). Isolation of relevant visual features from random stimuli for cortical complex cells. J. Neurosci. 22, 10811–10818.

Touryan, J., Felsen, G., and Dan, Y. (2005). Spatial structure of complex cell receptive fields measured with natural images. Neuron 45, 781–791.

Woolley, S.M., Fremouw, T.E., Hsu, A., and Theunissen, F.E. (2005). Tuning for spectro-temporal modulations as a mechanism for auditory discrimination of natural sounds. Nat. Neurosci. *8*, 1371–1379.

Woolley, S.M., Gill, P.R., and Theunissen, F.E. (2006). Stimulus-dependent auditory tuning results in synchronous population coding of vocalizations in the songbird midbrain. J. Neurosci. *26*, 2499–2512.

Yamada, W.M., and Lewis, E.R. (1999). Predicting the temporal responses of non-phase-locking bullfrog auditory units to complex acoustic waveforms. Hear. Res. *130*, 155–170.