Two-dimensional adaptation in the auditory forebrain

Tatyana O. Sharpee, Katherine I. Nagel and Allison J. Doupe

J Neurophysiol 106:1841-1861, 2011. First published 13 July 2011; doi:10.1152/jn.00905.2010

J Neurophysiol 106: 1841–1861, 2011.
First published July 13, 2011; doi:10.1152/jn.00905.2010.

Two-dimensional adaptation in the auditory forebrain

Tatyana O. Sharpee,[1,2] Katherine I. Nagel,[2] and Allison J. Doupe[2]

[1]*The Crick-Jacobs Center for Theoretical and Computational Biology, Computational Neurobiology Laboratory, The Salk Institute for Biological Studies, and the Center for Theoretical Biological Physics, University of California, San Diego, La Jolla; and* [2]*The Keck Center for Integrative Neuroscience, University of California, San Francisco, San Francisco, California*

Submitted 20 October 2010; accepted in final form 7 July 2011

Sharpee TO, Nagel KI, Doupe AJ. Two-dimensional adaptation in the auditory forebrain. *J Neurophysiol* 106: 1841–1861, 2011. First published July 13, 2011; doi:10.1152/jn.00905.2010.—Sensory neurons exhibit two universal properties: sensitivity to multiple stimulus dimensions, and adaptation to stimulus statistics. How adaptation affects encoding along primary dimensions is well characterized for most sensory pathways, but if and how it affects secondary dimensions is less clear. We studied these effects for neurons in the avian equivalent of primary auditory cortex, responding to temporally modulated sounds. We showed that the firing rate of single neurons in field L was affected by at least two components of the time-varying sound log-amplitude. When overall sound amplitude was low, neural responses were based on nonlinear combinations of the mean log-amplitude and its rate of change (first time differential). At high mean sound amplitude, the two relevant stimulus features became the first and second time derivatives of the sound log-amplitude. Thus a strikingly systematic relationship between dimensions was conserved across changes in stimulus intensity, whereby one of the relevant dimensions approximated the time differential of the other dimension. In contrast to stimulus mean, increases in stimulus variance did not change relevant dimensions, but selectively increased the contribution of the second dimension to neural firing, illustrating a new adaptive behavior enabled by multidimensional encoding. Finally, we demonstrated theoretically that inclusion of time differentials as additional stimulus features, as seen so prominently in the single-neuron responses studied here, is a useful strategy for encoding naturalistic stimuli, because it can lower the necessary sampling rate while maintaining the robustness of stimulus reconstruction to correlated noise.

neural coding; information theory; receptive field; spike-triggered average; spike-triggered covariance

TO ACCURATELY AND EFFICIENTLY represent events in the surrounding world, the nervous system has to take advantage of the statistical regularities present in input stimuli. These statistics, however, are not constant and routinely change over a wide range. For example, the auditory system has to tackle changes over nine orders of magnitude in sound pressure levels, from 30 dB, representing a whisper in a quiet library, through 60–70 dB, characteristic of normal conversation, to 80–90 dB of noise in city traffic. To cope with such large changes in the stimulus statistics, auditory neurons have been shown to adaptively change both how they filter incoming stimuli (Bandyopadhyay et al. 2007; Frisina et al. 1990; Krishna and Semple 2000; Kvale and Schreiner 2004; Lesica and Grothe 2008a; Lesica and Grothe 2008b; Nagel and Doupe 2006; Reiss et al. 2007; Theunissen et al. 2001; Theunissen et al. 2000; Woolley et al. 2006), and how they map the filtered stimuli onto the available range of neural firing rates (Dean et al. 2005; Nagel and Doupe 2006). With increasing sound volume, auditory neurons change their receptive fields (filters) by becoming less sensitive to the mean stimulus value, and more sensitive to deviations from the mean, either in time (Frisina et al. 1990; Krishna and Semple 2000; Lesica and Grothe 2008b; Nagel and Doupe 2006), in frequency (Bandyopadhyay et al. 2007; Lesica and Grothe 2008a; Nelken et al. 1997), or both (Lesica and Grothe 2008a; Nagel and Doupe 2008). Studies of sound onset detection by auditory neurons are also consistent with temporal summation at low sound pressure levels (Heil and Irvine 1997; 1996; Heil and Neubauer 2001), and in central auditory areas, with a shift to a derivative operation at very high sound levels (Galazyuk and Feng 2001; Heil and Irvine 1998; Phillips et al. 1984; Sullivan 1982). Finally, it should be noted that such transformations in how sounds are filtered in the auditory system offer many parallels to the adaptive filtering observed in the visual system (Barlow et al. 1957; Chander and Chichilnisky 2001; Shapley and Enroth-Cugell 1984) and are consistent with the "redundancy reduction" hypothesis (Barlow 1961).

Recent studies, however, demonstrate that neural responses in many sensory systems, including visual, auditory, and somatosensory, are affected by more than one stimulus component (Atencio et al. 2008; Chen et al. 2007; Fairhall et al. 2006; Maravall et al. 2007; Rust et al. 2005; Touryan et al. 2005; Touryan et al. 2002). The relevance of multiple stimulus dimensions to neural spiking opens up the possibility of qualitatively novel adaptive phenomena that cannot be observed in the one-dimensional (1D) model of feature selectivity. Even an unchanging multidimensional model can lead to diverse behaviors when sampled in different regimes (Hong et al. 2008), and a number of additional scenarios become possible in a truly adapting system. First, adaptation could lead to changes in the secondary dimensions that may differ from those observed in the primary dimension. Second, adaptation could also change the dependence of the firing rate on relevant stimulus components (the gain function). For instance, the gain function could rescale similarly for all relevant components, as has been recently observed in the barrel cortex (Maravall et al. 2007). Alternatively, the dependence of firing rate on stimulus could change differently for different stimulus components. This, in turn, will alter the relative contributions of each of the stimulus dimensions to neural spiking, even if the adaptation does not affect the stimulus dimensions themselves. We set out to explore these possibilities for adaptive phenomena by using the method of maximally informative dimensions (MID; Sharpee et al. 2004) to analyze multidimensional stimulus representations of neurons in the zebra finch auditory forebrain region

Address for reprint requests and other correspondence: T. O. Sharpee, The Salk Institute for Biological Studies, 10010 North Torrey Pines Rd., La Jolla, CA 92037 (e-mail: sharpee@salk.edu).

known as field L, an analog of the primary auditory cortex in mammals (Fortune and Margoliash 1992; 1995; Wild et al. 1993).

MATERIALS AND METHODS

Physiology and spike sorting. Experiments were conducted in five adult zebra finches using procedures approved by the University of California, San Francisco Institutional Animal Care and Use Committee, and in accordance with National Institutes of Health guidelines. Recordings were made using chronically implanted microdrives, as described previously (Nagel and Doupe 2006), and were all collected as part of that study, but extensively further analyzed here. Briefly, during recording, the bird was placed inside a small cage within a sound-attenuating chamber. Birds generally sat in one corner of the cage for the duration of the experiment, although their movement within the cage was not restricted. Two to three tungsten electrodes were used simultaneously. Putative single units were identified on the oscilloscope by their stable spike waveform and clear refractory period. All spikes were resorted offline based on the similarity of overlaid spike waveforms and on clustering of waveform projections in a two-dimensional (2D) principal component space. Neural recordings were considered single units if the number of violations within a 1-ms refractory period was <0.1% of the overall number of spikes. After the final recording, histological sections were examined to confirm that electrode tracks, and in some cases marker lesions, were located in field L.

Auditory stimuli. Neurons were probed with amplitude-modulated sounds whose temporal correlations approximated those characteristic of natural sounds. Each of the neurons was exposed to such sounds at low and high mean sound amplitude, as well as to low-intensity sounds scaled to have a larger variance. The stimuli were constructed from a slowly varying envelope with fixed statistical properties, as described below, and a rapidly varying carrier that could be adjusted for the frequency preference of each cell. The slowly varying envelope was generated from a log-normal distribution, such that the logarithm of the envelope was Gaussian noise with an exponential power spectrum $P(f) = \exp(-f/50\ \text{Hz})$. The Gaussian log-envelope $n(t)$ was scaled to have mean μ and variance σ. When multiplied by a noise carrier with unit standard deviation, this yielded sounds whose overall root-mean-square (RMS) amplitude (in dB) was given by $\text{RMS} = \mu + \log(10)\ \sigma^2/20$. A sample trace of log-amplitude waveform and its associated neural responses are shown in Fig. 1. We analyzed stimuli in three different distributions: *1*) "low mean/low variance" with $\mu = 30$ dB and $\sigma = 6$ dB, corresponding to RMS = 34 dB; *2*) "high mean/low variance" with $\mu = 63$ dB and $\sigma = 6$ dB, corresponding to RMS = 67 dB; and *3*) "low mean/high variance" with $\mu = 30$ dB and $\sigma = 18$ dB, corresponding to RMS = 67 dB as well. The linear voltage envelope $E(t)$ was generated from the logarithmic modulation signal $n(t)$ by exponentiation, $E(t) = 10^{-5+n(t)/20}$. A continuous stimulus (termed trial block) was generated from alternating four 5-s segments: the first segment was taken from distributions with low mean/low variance, followed by segments with low mean/high variance, another segment with low mean/low variance, and then high mean/low variance, with this sequence repeating 100 times. One-half of the 5-s segments were novel, whereas the other one-half of these segments were repeated. Responses to one such repeated segment for an example neuron are shown in Fig. 1*B*. Unique and repeated trials were randomly interleaved. Overall, a given neuron was probed with one to three trial blocks.

Finding relevant stimulus dimensions. To characterize stimulus features most relevant for eliciting spikes from field L neurons, we worked in the framework of the linear-nonlinear (LN) model (Meister and Berry 1999; Shapley and Victor 1978). According to this model, only stimulus variations along a small number of dimensions affect the spike probability. The spike probability itself can be an arbitrary nonlinear function of the relevant stimulus components. To find the relevant stimulus dimensions for each of the neurons in our data set, we used two different methods. First, we found one of the relevant stimulus dimensions using the reverse correlation method (de Boer and Kuyper 1968; Rieke et al. 1997) and correcting for stimulus correlations (Schwartz et al. 2006). To achieve this, we first computed a vector, the spike-triggered average (STA), by averaging all re-

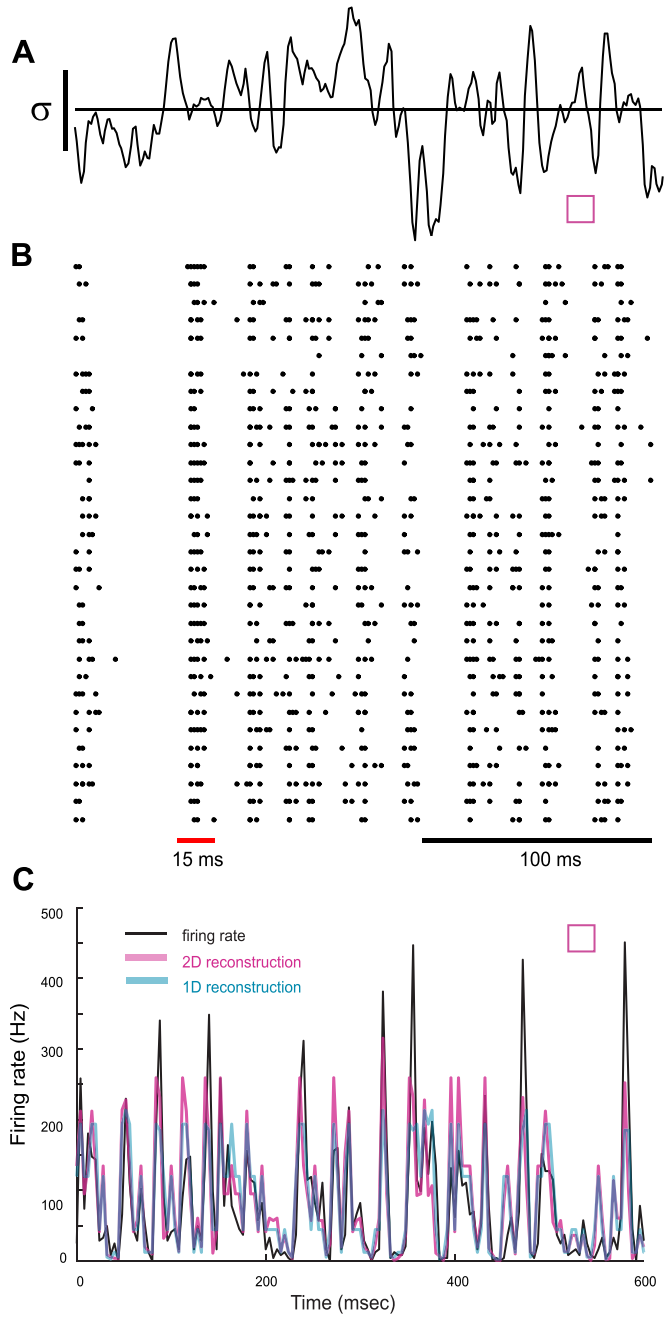

Fig. 1. Firing rate predictions using one-dimensional (1D) and two-dimensional (2D) linear/nonlinear models. *A*: a segment of the randomly varying modulation signal σ that specifies the local log-amplitude of the sound in decibels (see MATERIALS AND METHODS). *B*: responses of a single unit to 32 repetitions of the modulation signal shown in *A*. In one-half of all trials for each neuron, such repeated stimuli were presented (to use for predictions, as in *C*), while in the other one-half of trials, unrepeated versions of the modulation signal were presented and used to calculate the relevant dimensions and nonlinear gain functions of the linear-nonlinear (LN) models. *C*: real (black) and predicted firing rates for this example neuron using the 2D (pink) or 1D (blue) LN model. Neuron "eb1940".

sponses that elicited a spike (Fig. 2*A*). The STA was then multiplied by the inverse of the covariance matrix. We will refer to the resultant vector as the decorrelated STA (dSTA; Fig. 2*B*). Because multiplication by the inverse of the stimulus covariance matrix often leads to noise amplification at high temporal frequencies, we also computed the "regularized" dSTA (rdSTA; Fig. 2*C*) using a pseudoinverse (instead of the inverse) of the covariance matrix. A pseudoinverse excludes the eigenvectors of the covariance matrix that are poorly sampled (Theunissen et al. 2000). The cut-off for excluding the eigenvectors was chosen to maximize the predictive power on a test part of the data, not used in computing the STA. Overall, the features computed as rdSTA and dSTA were very similar to each other, as well as to the first features computed by maximizing information, which is described next.

In the second line of analysis, the relevant stimulus features were computed as dimensions in the stimulus space that accounted for the maximal amount of information in the neural response (Sharpee et al. 2004). The first MID was found by maximizing the following function:

$$I(\mathbf{V}) = \int dx P_V(x|\text{spike}) \log_2 \frac{P_V(x|\text{spike})}{P_V(x)} \qquad (1)$$

where $P_V(x)$ is the probability distribution of stimulus components along dimension **v**, and $P_V(x|\text{spike})$ is the analogous probability distribution computed by taking only stimulus segments that led to a spike. *Equation 1* represents the amount of information between the arrival times of single spikes and stimulus components x along a particular dimension **v** (Adelman et al. 2003; Agüera y Arcas et al. 2003; Fairhall et al. 2006; Paninski 2003; Sharpee et al. 2004). This function corresponds to the Kullback-Leibler divergence between probability distributions $P_V(x|\text{spike})$ and $P_V(x)$. Dimensions along which these two distributions differ most are thus most affected by an observation of a spike.

The details of the optimization algorithm are provided in Sharpee et al. (2004) and supplemental information of Sharpee et al. (2006). Briefly, information was maximized via a sequence of 1D line optimizations along the gradient of information. During each line optimization, points that led to decreases of information were occasionally accepted with probability $\exp(-\Delta I/T)$, where ΔI is the decrease in information associated with acceptance of the new dimensions, and parameter T, "effective temperature", controls the probability of accepting decreases in information of large magnitude. Dimensions that led to an increase in information were always accepted. The optimization procedure started with the value of effective temperature $T = 1$. The effective temperature was then decreased by a factor of 0.95 after each line maximization, until it reached the value of 10^{-5}. After that, information was increased by a factor of 100, and the iteration continued. The maximum number of line maximizations was 3,000. Performance of the current dimension was evaluated on the test set after every 100 line maximizations. Upon completion of the computation, we compared dimensions that accounted for most information on the test set and the dimension that accounted for most information on the training set. These dimensions were similar; dimensions with the best performance on the test set were used as the MIDs. The search for the first MID was initialized as the STA. We have verified that optimization results starting from random initial conditions were not different (see Fig. 2, *D* and *E*). After the first MID was computed, we initialized the second dimension as a random segment of the stimulus and optimized a pair of dimensions to capture the maximal amount of information about the arrival times of the single spikes in this case. The corresponding optimization function is given by:

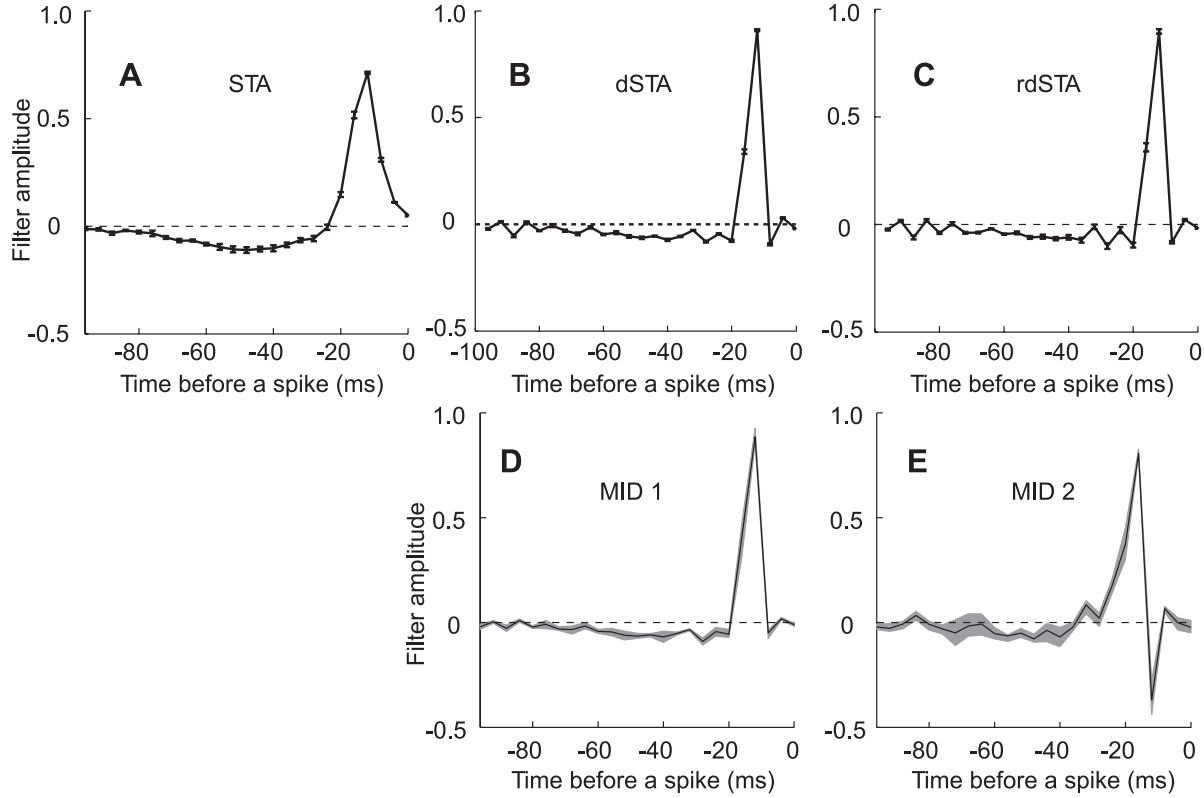

Fig. 2. Comparison of different methods for characterizing neural feature selectivity. *A*: the dimension obtained as the so-called spike-triggered average (STA). *B*: decorrelated STA (dSTA). *C*: regularized decorrelated STA (rdSTA). The first (*D*) and second (*E*) most informative dimensions (MID) (see text for method description) are shown. In *D* and *E*, the MID results are plotted for optimizations starting with the STA (black line), and starting with a random stimulus segment (gray lines denote mean 1 SD). The two cases are almost indistinguishable, which proves that the algorithm is not sensitive to initial conditions. Neuron "soba1980".

$$I(V_1, V_2) = \int dx P_{V_1,V_2}(x_1, x_2|\text{spike}) \log_2 \frac{P_{V_1,V_2}(x_1, x_2|\text{spike})}{P_{V_1,V_2}(x_1, x_2)} \quad (2)$$

where x_1 and x_2 represent stimulus components along dimensions $\mathbf{v}_1$ and $\mathbf{v}_2$, respectively. $P_{V_1,V_2}(x_1, x_2)$ represents the probability distribution of stimulus components along dimension $\mathbf{v}_1$ and $\mathbf{v}_2$, and $P_{V_1,V_2}(x_1, x_2|\text{spike})$ is the analogous probability distribution computed by taking only stimulus segments that lead to a spike. Dimensions $\mathbf{v}_1$ and $\mathbf{v}_2$ that, at the end, maximize *Eq. 2* correspond to MID1 and MID2. Typically, the MID1 profile changed little after it was optimized, together with the second dimension. One of the advantages of such optimization, which corresponds to the maximum likelihood estimation of the LN model (Kouh and Sharpee 2009), is that nonlinear aspects of the transformation from the relevant stimulus components to the neural firing rate are built into *Eqs. 1* or *2*, which eliminates the need to fit nonlinearities while searching for the relevant stimulus features.

Nonlinear gain functions. Once the relevant dimensions $\mathbf{v}_1$ and $\mathbf{v}_2$ are computed, one can determine the 2D gain function that relates the stimulus components along those dimensions to the neural firing rate. This function is given by $P_{V_1,V_2}(\text{spike}|x_1, x_2)$ and can be computed using Bayes' rule as $P(\text{spike}) P_{V_1,V_2}(x_1, x_2|\text{spike}) / P_{V_1,V_2}(x_1, x_2)$. Here, $P(\text{spike})$ is the average spike probability in a 4-ms time bin. The gain functions with respect to relevant dimensions $\mathbf{v}_1$ and $\mathbf{v}_2$ considered separately can be computed as proportional to $P_{V_1}(x_1|\text{spike})/P_{V_1}(x_1)$ and $P_{V_2}(x_2|\text{spike})/P_{V_2}(x_2)$, respectively.

Information analysis. To compute information about the stimulus carried by the arrival times of single spikes, I_{spike}, we analyzed responses to a random segment of amplitude modulations (duration 5 s) repeated between 50 and 150 times. Based on these repeated presentations, we computed the average firing rate $r(t)$ for this pseudorandom noise sequence. Information about the stimulus carried by the arrival times of single spikes can then be computed as (Brenner et al. 2000b):

$$I_{\text{spike}} = \frac{1}{T} \int dt \frac{r(t)}{\bar{r}} \log_2 \frac{r(t)}{\bar{r}} \quad (3)$$

where $\bar{r}$ is the average firing rate. [More generally, the same expression can be used to compute information carried by the arrival time of any spike pattern, in which case $r(t)$ would represent the average number of the spike patterns of interest per unit time.]

This information measure makes no assumptions about the number of relevant stimulus dimensions, nor about the shape of the nonlinear gain function describing the dependence of spike probability on the relevant stimulus components. Therefore, it can be used to quantify the performance of any model of a reduced dimensionality, such as models based on the STA or MIDs.

The values of mutual information contain a positive bias, which decreases as more data are collected (Brenner et al. 2000b; Nelken and Chechik 2007; Strong et al. 1998; Treves and Panzeri 1995). To correct for this bias, we computed information values based on different fractions of the data (80–100%). We then used linear extrapolation to find the information value projected if infinite numbers of spikes could be collected (Brenner et al. 2000b; Strong et al. 1998). This procedure was used to correct for bias in all information values (I_{spike} and information along one or two dimensions $\mathbf{v}$).

Finally, information obtained using multiple dimensions is sometimes adjusted for the addition of dimensions by subtracting the information generated by a random dimension (Fairhall et al. 2006). However, it should be noted that even a random dimension will be able to account for some fraction of the real neural response, even in the limit of infinite data (and after all of the bias corrections considered above are applied). The positive predictive power of any random dimension is not an artifact, but is due to the fact that such a dimension will always contain a nonzero component along the relevant dimensions. In the case of uncorrelated stimuli, such effects are typically small (Fairhall et al. 2006). However, the nonzero components can be substantial, up to 90%, with natural stimuli (Sharpee et al. 2004). In our data set, a single randomly selected dimension, on average, accounted for a modest amount of information (4.5 ± 1.8%), which was not subtracted, because of the above arguments.

Determining the number of relevant features based on the spike-triggered covariance method. To obtain an independent estimate of the number of relevant features in each neuron, we analyzed the spike-triggered covariance matrix (Bialek and de Ruyter van Steveninck 2005; de Ruyter van Steveninck and Bialek 1988). The number of significant eigenvalues depends on sampling (Bialek and de Ruyter van Steveninck 2005; de Ruyter van Steveninck and Bialek 1988; Rust et al. 2005; Schwartz et al. 2006). We first estimated the number of significant eigenvalues by computing the eigenvalue spectrum based on different fractions of the data, from 10 to 100%. Eigenvalues that were stable with respect to the fraction of the data used were judged to be significant. To make this analysis quantitative, we computed the eigenvalue spectrum based on spike trains shifted with respect to the stimulus by a random number with a minimum value (Bialek and de Ruyter van Steveninck 2005). The minimum value was equal to the length of the kernel: 100 ms for most cells and 160 ms for three cells. Eigenvalues computed in such a manner are limited to a band near zero (Bialek and de Ruyter van Steveninck 2005). This fact can be used to discriminate significant eigenvalues from nonsignificant eigenvalues. Those eigenvalues of the real spike-triggered covariance matrix that exceeded the range of values observed upon randomly shifting the spike train relative to the stimulus were judged to be significant. A total of 3,000 randomly shifted spike trains were analyzed for each neuron. Stimulus dimensions that were associated with the first two significant eigenvalues were similar to the STA and/or to the first two most informative features, if we accounted for the expected broadening due to stimulus correlations.

Fitting relevant dimensions with Hermite functions. To characterize the shapes of relevant dimensions, we used a basis set formed by the first three Hermite functions (Abramowitz and Stegun 1964; Victor et al. 2006). The Hermite functions are orthogonal to each other and form a complete basis, meaning that any function could, in principle, be represented in terms of a sufficient number of linear combinations of Hermite functions. For relevant dimensions of field L neurons, we found that three Hermite functions were sufficient. The first three Hermite functions are described by the following equations: $H_0(t) = \pi^{-1/4}(e^{-t^2/2})$, $H_1(t) = \sqrt{2}\pi^{-1/4} t (e^{-t^2/2})$, and $H_2(t) = \pi^{-1/4}\sqrt{2}(2t^2 - 1)e^{-t^2/2}$. We used nonlinear least squares data fitting by the Gauss-Newton method as implemented in Matlab to find the best-fitting linear coefficients with respect to $H_0[(t - t_0)/\tau]$, $H_1[(t - t_0)/\tau]$, and $H_2[(t - t_0)/\tau]$ for each of the relevant dimensions. The parameters t_0 and τ describe the centering (position of the peak in the Gaussian envelope) and scaling (the width of the Gaussian envelope). We obtained very similar fitting results, regardless of whether different relevant dimensions of one neuron were constrained to have the same pair of values for parameters t_0 and τ or not.

Estimating jitter in spike timing. We used the procedure described in Aldworth et al. (2005) to estimate spike timing jitter in our responses. In brief, this method consists of iteratively performing the following two steps. The first step was to compute the STA as an average stimulus segment preceding each spike in a spike train. Then, the timing of each spike in the spike train was adjusted to maximize correlation between the preceding stimulus segment and the current estimate of the STA. Large shifts in the timing of the spike were penalized according to a Gaussian prior. The width of the Gaussian distribution was set to the standard deviation σ of the spike time jitter distribution found during the previous iteration. On the first iteration, this value was set to 1 ms or 4 ms when working with spike trains binned at 1-ms and 4-ms resolution, respectively. The maximal time by which spikes could be shifted forward and backward in time was

set to σ of the spike time jitter distribution (which could vary during optimization) and 3 ms, respectively. Larger shifts backward in time than 3 ms were taken to violate the causality condition for eliciting a spike before a significant portion of the relevant stimulus feature occurred. Returning to step one, the new estimate of the STA was obtained based on updated latencies between stimulus segments and spikes.

We quantified the similarity between the 2D relevant spaces estimated with and without prior dejittering using a measure termed subspace projection (Rowekamp and Sharpee 2011), which ranges between 0, if subspaces do not overlap, and 1, for a perfect match. In one dimension, this measure corresponds to a dot product between two vectors normalized to length 1, whereas in two dimensions, it corresponds to computing the dot product between normals to the two estimates of the relevant plane, and taking the square root. The following expression describes this mathematically with $\vec{u}_1$ and $\vec{u}_2$ representing the first and second dimensions from one 2D subspace and $\vec{v}_1$ and $\vec{v}_2$ representing the first and second dimensions from the other 2D subspace:

$$\rho = \sqrt{\frac{\left|(\vec{u}_1\vec{v}_1)(\vec{u}_2\vec{v}_2) - (\vec{u}_1\vec{v}_2)(\vec{u}_2\vec{v}_1)\right|}{\sqrt{\left|(\vec{u}_1\vec{u}_1)(\vec{u}_2\vec{u}_2) - (\vec{u}_1\vec{u}_2)^2\right|\left|(\vec{v}_1\vec{v}_1)(\vec{v}_2\vec{v}_2) - (\vec{v}_1\vec{v}_2)^2\right|}}}$$

The resultant quantity is independent of whether individual dimensions are normalized or orthogonal.

RESULTS

2D description of feature selectivity in field L. We focused on how auditory neurons encode a single aspect of an auditory stimulus, namely its time-varying log amplitude, which has previously been shown to play a critical role in the responses of higher-order auditory neurons in this species (Gill et al. 2006; Nagel and Doupe 2006; Theunissen and Doupe 1998) and awake primates (Malone et al. 2007; Scott et al. 2011; Zhou and Wang 2010), as well as in speech recognition in humans (Shannon et al. 1995). To isolate responses to the log amplitude, we used a stimulus that segregated responses to the slowly varying amplitude modulation envelope from a rapidly varying "carrier" signal (see MATERIALS AND METHODS; Fig. 1 illustrates a sample trace of log-amplitude waveform, and the neural responses to it). The modulation envelope was designed to capture the amplitude distribution and temporal frequency of natural sounds. Rapid transitions to and from silent periods, which are common in natural sounds, lead to amplitude distributions that are strongly non-Gaussian (Escabi et al. 2003; Singh and Theunissen 2003). However, the distribution of log-amplitudes can be more closely approximated by a Gaussian (Nelken et al. 1999). Natural stimuli are also dominated by slow changes in amplitude, with power spectra that decrease as a function of temporal frequency (Lewicki 2002; Singh and Theunissen 2003; Voss and Clarke 1975). To capture both of these properties of natural sounds, the randomly varying log-amplitude was generated according to a correlated Gaussian distribution whose power spectrum decayed exponentially with the characteristic frequency of 50 Hz (Nagel and Doupe 2006). The diverse fluctuations present in this stimulus ensemble make it possible to determine how different neurons analyze incoming stimuli with minimal assumptions. Moreover, the mean and variance of this stimulus can be systematically altered to assess the dependence of coding on stimulus statistics (see MATERIALS AND METHODS).

To characterize the signal processing properties of individual neurons in the zebra finch auditory forebrain, we searched for stimulus dimensions that capture the maximum information about the recorded neural response (Sharpee et al. 2004); see Fig. 2 for a comparison with other spike-triggered methods (Schwartz et al. 2006). The method of MID allows us to ask what features, or dimensions, of a stimulus best account for the neuron's response, and to estimate what fraction of the total information carried by the neuron is captured by each of these features (Nelken and Chechik 2007). This technique extends linear methods of relating stimulus dimensions to neural responses in several ways (Christianson et al. 2008): more than one relevant stimulus dimension can be obtained; dimensions with a highly nonlinear relationship to neural firing can be extracted; finally, the MIDs are not influenced by correlations or structure in the stimulus. Furthermore, the variance in MID filters across different data subsets is the smallest possible for any unbiased method, including spike-triggered covariance and its information-theoretic generalizations (Pillow and Simoncelli 2006), because MIDs saturate the Cramer-Rao bound (Kouh and Sharpee 2009). We used the MID method sequentially to identify first the most informative dimension of our stimulus and then the second most informative dimension. The primary (maximally informative) dimension was always very similar to that derived from a STA (Atencio et al. 2008; Depireux et al. 2001; Eggermont 1993; Eggermont et al. 1983; Epping and Eggermont 1986; Hsu et al. 2004; Kim and Young 1994; Klein et al. 2000; Klein et al. 2006; Machens et al. 2004; Woolley et al. 2006; see also Fig. 2).

We found that individual neurons in field L showed significant sensitivity to (at least) two dimensions of the log amplitude of the sound envelope (for the remainder of the paper, we shall refer to this log amplitude as "the stimulus"). Although a 1D model based on just the primary stimulus component could typically replicate the major peaks in the neural firing rate, higher peaks, as well as valleys, could be predicted more accurately with a 2D linear nonlinear model (see Fig. 1*C* for an example). At the same time, the 2D model overestimated the height of some intermediate peaks. To check whether the 2D model provided a better description overall, we compared the percentage of information explained by the two models. For this example cell and stimulus condition, the 1D and 2D models accounted for 67.1 ± 0.5% and 77.5 ± 0.5% of the information carried by independent spikes, respectively. The corresponding values in terms of the percentage of variance in the firing rate accounted by the 1D and 2D models were 78.4 ± 0.3% and 87.1 ± 0.4%, respectively. All predictions were made on a novel stimulus segment not used to calculate either the relevant dimensions or the associated nonlinear gain functions. The spike-triggered covariance method, a complementary method for estimating the number of dimensions to which the firing rate is sensitive (Agüera y Arcas et al. 2003; Bialek and de Ruyter van Steveninck 2005; Brenner et al. 2000a; de Ruyter van Steveninck and Bialek 1988; Schwartz et al. 2006; Touryan et al. 2002) also indicated two significant dimensions for this cell (see MATERIALS AND METHODS), providing additional evidence for the 2D encoding realized by this cell.

Across the population of cells, the 2D model accounted, on average, for 66% of the information carried in the arrival of single spikes (Fig. 3). In the large majority of cells (63/74), the 2D model accounted for over one-half the total information in

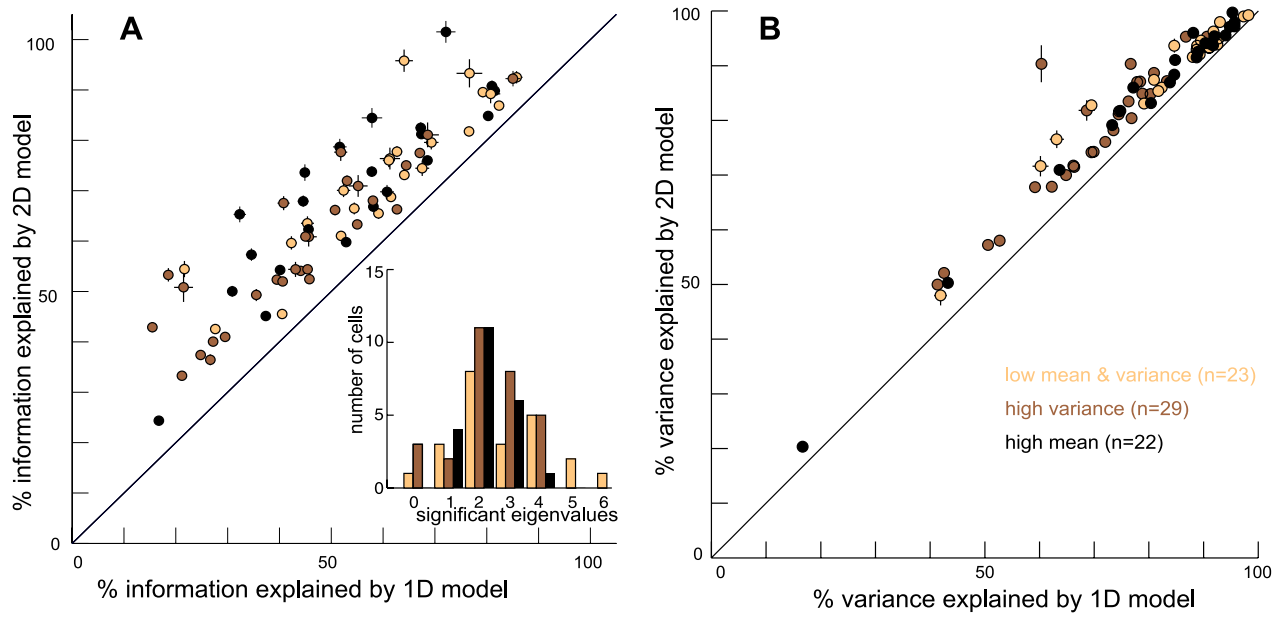

Fig. 3. A population analysis of predictive power based on 2D linear/nonlinear models. *A*: the percentage of total information between spikes and stimulus accounted for by the 2D model vs. the percentage accounted for by the 1D model based on the first MID alone. *Inset* shows a histogram of the number of spike-triggered covariance method dimensions for each cell in our data set. The distribution peaks at two dimensions for all stimulus conditions. *B*: the percentage of variance in the firing rate accounted for by the 2D model vs. the percentage accounted for by the 1D model based on the first MID alone. The percentages for information and variance were computed using a novel segment of the data. Error bars are standard errors.

the spike train, reaching values as high as 90% and 100% for some cells. Similarly, in terms of the percentage of variance in the firing rate, the 2D model accounted, on average, for 84% of this variance, again reaching values close to 100% for some cells, and greater than 90% for almost one-half of the neurons (33/74). By comparison, the 1D models accounted, on average, for 52% of the information carried in the arrival of single spikes and 72% of the variance in the firing rate. For all cells in our data set, the 2D model accounted for significantly more information ($P < 0.02$, *t*-test) and variance ($P < 0.024$, *t*-test) than the 1D models. The 2D processing across the population was also supported by the spike-triggered covariance method analysis: the distribution of the number of significant dimensions for each cell, across the population of neurons, peaked at 2 (Fig. 3*A*, *inset*). Thus auditory processing in the songbird forebrain is based on at least two dimensions, which together provide a good description of neural responses for a majority of cells in our data set.

Effects of stimulus mean on the primary and secondary dimensions. A previous analysis using a 1D linear/nonlinear model demonstrated that, for many field L neurons, the shape of the STA (equivalent to the primary dimension here) gradually changed with increasing mean stimulus amplitude from a more unimodal filter to a more biphasic one (Nagel and Doupe 2006). Here, we report that the secondary dimensions are affected just as strongly. In Fig. 4, we provide examples of the two dimensions computed under low and high mean stimulus conditions for three neurons. In each plot, zero represents the time of the spike, and the shape of the waveform preceding the spike indicates a feature of the stimulus that is strongly associated with fluctuations in the firing rate. When sounds were soft, the two dimensions describing the feature selectivity of a single neuron typically represented the low-pass filtered stimulus (the local time-average or mostly uniphasic feature) and its first time differential (biphasic feature; Fig. 4*A*, *columns 1–3*, black lines). When sounds were loud, the two dimensions represented the low-pass filtered first and second time differential of the stimulus waveform (Fig. 4*B*, columns 1–3, dark gray lines). As this figure illustrates, either the first or the second dimension could approximate the derivative of the other dimension (in each case, the calculated derivative of one of the dimensions is shown on the other dimension as a thick gray line, for comparison). For example, the second dimension of *neuron 1* was (qualitatively) the time derivative of this neuron's primary dimension under both low and high mean conditions. In example *neuron 2*, the first dimension was the time derivative of the second dimension in the low mean condition, but not in the high mean condition, where the second dimension was approximately the time derivative of the first dimension. Finally, example *neuron 3* illustrates the case where the role of the dimensions was reversed compared with *neuron 2*.

Although the simplest examples of dimensions were entirely averages or differentials, across all cells the relevant dimensions were best described by a mixture of profiles corresponding to the local time average and time differentials of the stimulus, such as the feature shown in Fig. 4*A*, *column 2*, for the first dimension. To quantitatively describe such profiles, as well as to test in general how well the primary and secondary dimensions could be approximated by linear combinations of the local time average and various differentials of the stimulus, we fitted each dimension with a linear combination of orthogonal Hermite functions of zero, first, and second order. The zeroth-order Hermite function is a Gaussian, and the first-order Hermite function is proportional to the time differential of the zeroth-order function. The second-order Hermite function is a linear combination of the Gaussian and its second time differential (see MATERIALS AND METHODS). Therefore, the three coefficients in the Hermite expansion of a given relevant stimulus dimension can capture the degree to which this dimension

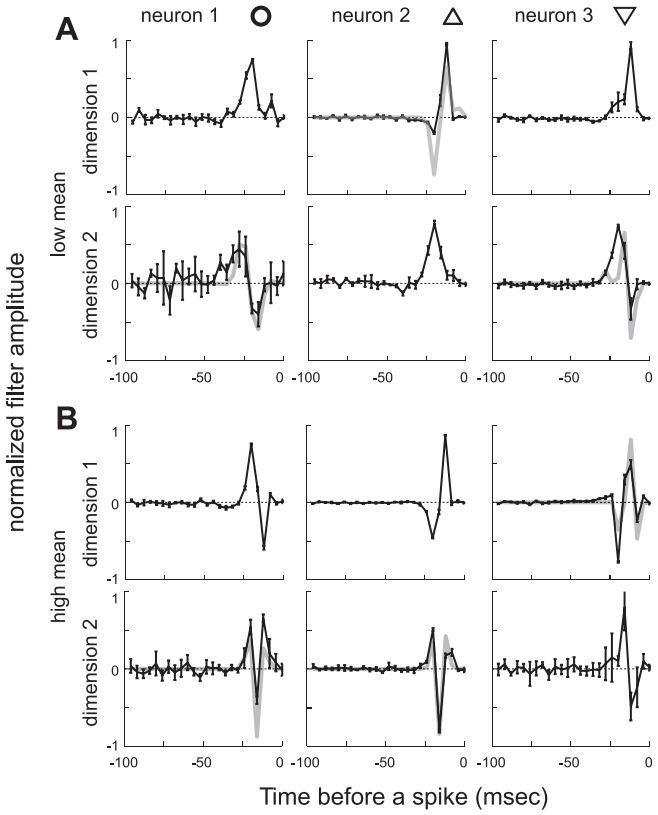

Fig. 4. Shapes of the primary and secondary stimulus dimensions are strongly affected by mean sound volume. *A*: the primary and secondary stimulus dimensions for three example neurons are shown when the mean sound volume was low. The *x*-axis represents the time before a spike (at zero), and the *y*-axis represents normalized amplitude of the filter. *B*: results for the same neurons when the mean sound amplitude was high. Dark lines show the mean and standard deviation of dimensions as derived from data. Light gray lines show first time derivatives of the other dimension under each stimulus condition. Derivative comparisons are provided for the dimensions that were better approximated by the time derivative of the other dimension in the pair. From *left* to *right*, columns neuron IDs are "ra2200", "soba1740wide", and "udon2120". See also Fig. 5 for fits using Hermite functions.

describes sensitivity to a local time average of the stimulus (H_0), its rate of change (H_1), and the second time differential or acceleration (H_2), all averaged over the duration of the Gaussian.

Overall, both the primary and secondary dimensions under all stimulus conditions could be well approximated by linear combinations of the first three Hermite functions. The fits yielded correlation coefficients of 0.93 ± 0.08 (SD) for the first stimulus dimension and 0.83 ± 0.17 (SD) for the second stimulus dimension. The fitting results (red curves) for the three example neurons from Fig. 4 are plotted together with their measured profiles in Fig. 5. Because the relevant dimensions of field L neurons could be fitted in a near-perfect manner using the first three Hermite coefficients, we could now use their relative magnitude to analyze differences between the shapes of dimensions under different stimulus conditions. A population analysis is provided in Fig. 6. In this figure, each triplet of Hermite coefficients describing a particular dimension for a neuron is plotted as a point within an equilateral triangle. Such a representation is possible because the sum of distances from a point within the triangle to each of its three sides is the same for all points within the triangle. We used a normalization such that the sum of absolute values of the Hermite coefficients for each dimension was equal to 1. Then the distance from a data point to the side of the triangle opposite to the vertex gives the magnitude of the component associated with that vertex. For example, the vertex-point labeled $|H_2|$ would correspond to the case of $H_0 = 0$, $H_1 = 0$, and $H_2 = 1$, whereas points along the side connecting vertices $|H_0|$ and $|H_1|$ would have $H_2 = 0$. Results in Fig. 6, *A–D*, provide pairwise comparisons between dimensions for all of our data. Here, data points corresponding to each cell are connected, and the dark vector shows the mean of all cells.

The population analyses demonstrate four phenomena. First, in the case of soft sounds, the primary and secondary dimensions complement each other with respect to the relative magnitude of integrator and differentiator components (Fig. 6*A*). That is, cells with a large H_0 component in the primary dimension (i.e., close to the vertex H_0) have a larger H_1 and smaller H_0 in the secondary dimension, and vice versa. This is evident both as the clustering of many individual cells across the base of the triangle, along the integration-differentiation axis, and in the orientation of the mean population vector. Second, a similar complementarity was observed in the case of loud sounds (Fig. 6*B*), but along the differentiation-acceleration axis. Here, although the magnitude of H_0 is nonzero, cells in which the primary dimension has a large H_1 component have a larger H_2 and smaller H_1 components in the secondary dimensions. Next, comparing only the primary dimension between the conditions of high and low mean, we recover our previous findings (Nagel and Doupe 2006) that the primary dimensions change from a uniphasic shape (large H_0) for soft sounds to biphasic shapes (large H_1) for loud sounds (Fig. 6*C*). The final observation is illustrated in Fig. 6*D*, where the

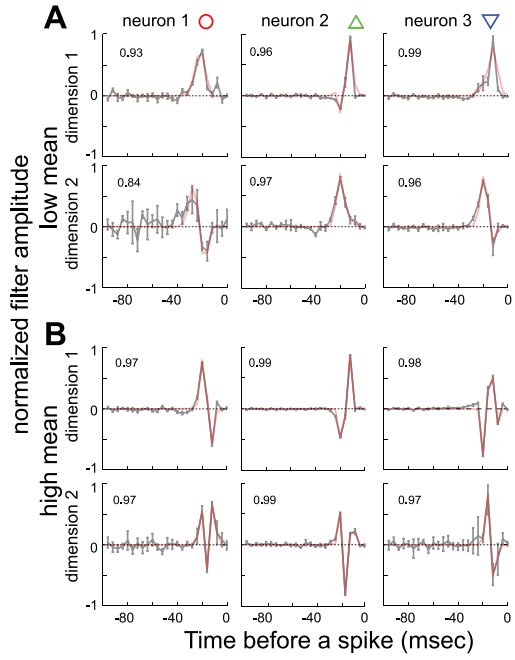

Fig. 5. The relevant dimensions at low and high sound volume fitted with Hermite functions. Notations and neurons are the same as in Fig. 4. Red lines indicate fits using the three lowest order Hermite functions. Numbers within each panel show correlation coefficients between the relevant dimensions and their fits.

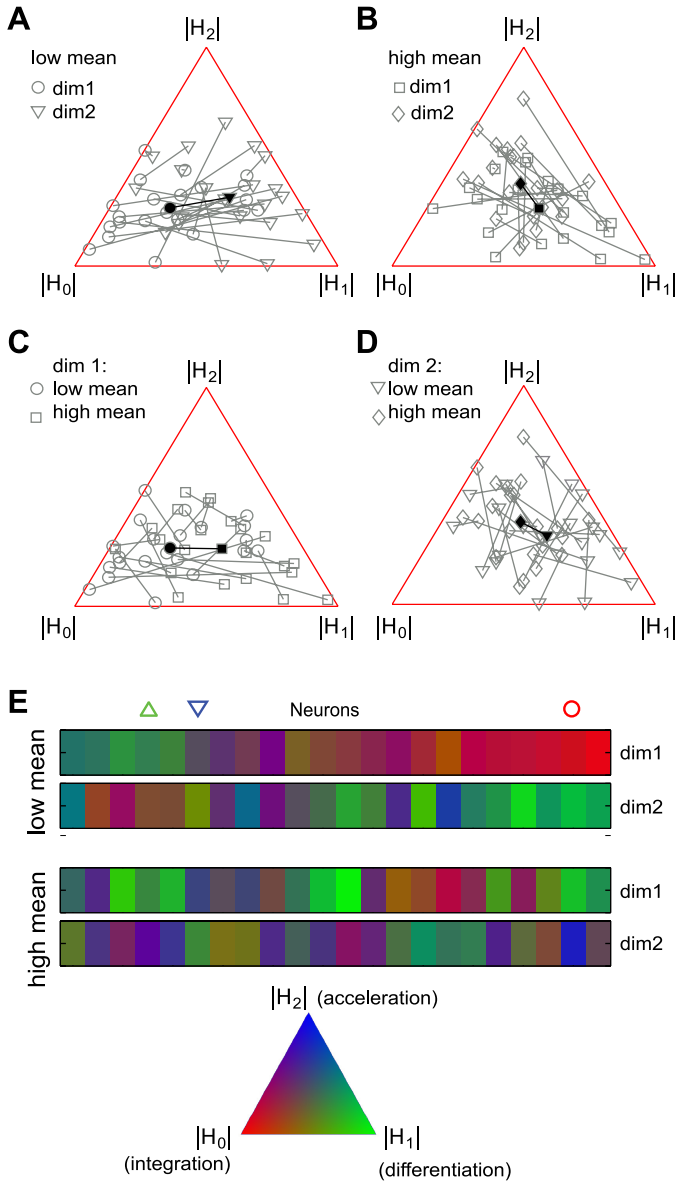

Fig. 6. A population analysis of the shapes of relevant dimensions. The triplet of Hermite coefficients is represented as a dot within an equilateral triangle. *A–D*: comparison of pairs of dimensions under either the same or different stimulus conditions. Data from each cell are linked, and the dark vector shows the population mean. *A*: comparison of two relevant dimensions under the low mean stimulus condition. *B*: same as *A*, but for the high mean stimulus condition. *C*: comparison of primary dimensions between the low and high mean stimulus conditions. *D*: same as *C*, but for the secondary dimensions. *E*: by using color to represent the location within the reference triangle (*bottom*), we could compare the primary and secondary dimensions for each neuron within and across stimulus conditions (each neuron is represented in a vertical column). Red-green colors, corresponding to integration-differentiation, dominate in the low mean stimulus condition, whereas blue-green colors (differentiation-acceleration) dominate in the high mean stimulus condition. Symbols (circle, upward and downward triangles) label points corresponding to the example neurons from Fig. 4.

secondary dimensions vary along the differentiation-acceleration axis with increasing stimulus mean.

Figure 6*E* provides a further graphical comparison of dimensions across all cells and stimulus conditions, with each column corresponding to a single cell. Here, we used red, green, and blue (RGB) colors to represent position within the triangle of Hermite coefficients, as shown in the reference triangle to the bottom of Fig. 6*E*. In Fig. 6*E*, red-green colors, corresponding to integration-differentiation, dominate in the low mean stimulus condition, whereas green-blue colors (differentiation-acceleration) dominate in the high mean stimulus condition. As described above, either the integration or differentiation dimensions could play the dominant role in the low mean stimulus conditions. For example, the cell marked with a circle has a primary dimension that is integration-like, whereas the cell marked with an upward green triangle has a primary dimension that is differentiation-like. Similarly, in the case of loud sounds, the primary dimension of the cell marked with a circle is primarily differentiation-like, whereas that of the cell marked with a downward blue triangle is primarily acceleration-like.

Together, these analyses demonstrate that, when the mean stimulus was low, field L neurons encoded the temporal variations of the stimulus through measurements of its local time average and the first differential. When the mean stimulus was large, field L neurons became insensitive to the local time average of sound log-amplitude and encoded primarily its variations relative to the mean, i.e., its first time differential. However, neural encoding of the first time differential of the stimulus at high mean intensity was again 2D. This time, the two relevant dimensions were the first and second time differentials of the stimulus. If we define the waveform of interest in the high mean condition as the first time differential of the sound log-amplitude, then the first and second time differential of the sound log-amplitude are again the local time average and the first time differential of this "waveform of interest." Thus, although neurons encoded different "features" of the stimulus, depending on the mean sound amplitude, namely the time-varying log-amplitude itself for soft sounds and its first time differential for loud sounds, the encoding in both cases was based on local time-averaged measurements of the feature waveform and the first time differential of this waveform.

How fast did the relevant dimensions change shape following a change in the stimulus mean? Our previous analysis of single dimensions (Nagel and Doupe, 2006) indicated that changes in relevant dimensions were completed within 100 ms after a change in mean stimulus level. However, because of the fast time scales involved, it was not possible to monitor these changes dynamically and provide a lower bound on the adaptation time. Here we attempted to estimate the adaptation time more precisely by observing *1*) how fast the effects of adaptation to the previous stimulus condition "wear off" after switching to a new stimulus condition, and *2*) how fast the firing rate plateaus following adaptation to a new stimulus condition. In Fig. 7*A*, we compared the average neural response across the population to the identical (low mean/low variance) stimulus segment when it was preceded by either segments with high mean/low variance (black line) or by segments with low mean/high variance (gray line). We found that the population response to this identical stimulus differed, depending on the statistics of stimuli that preceded it, which is the classical definition of adaptation. The difference in the magnitude and latency of responses to the same stimulus disappeared after ~50 ms and 100 ms, respectively. Thus, in the case of adaptation to the low mean stimulus condition, one can expect to find time constants on the order of 50–100 ms. The time constants of adaptation to the high mean stimulus condition were similarly fast (Fig. 7*B*). Because of the order in

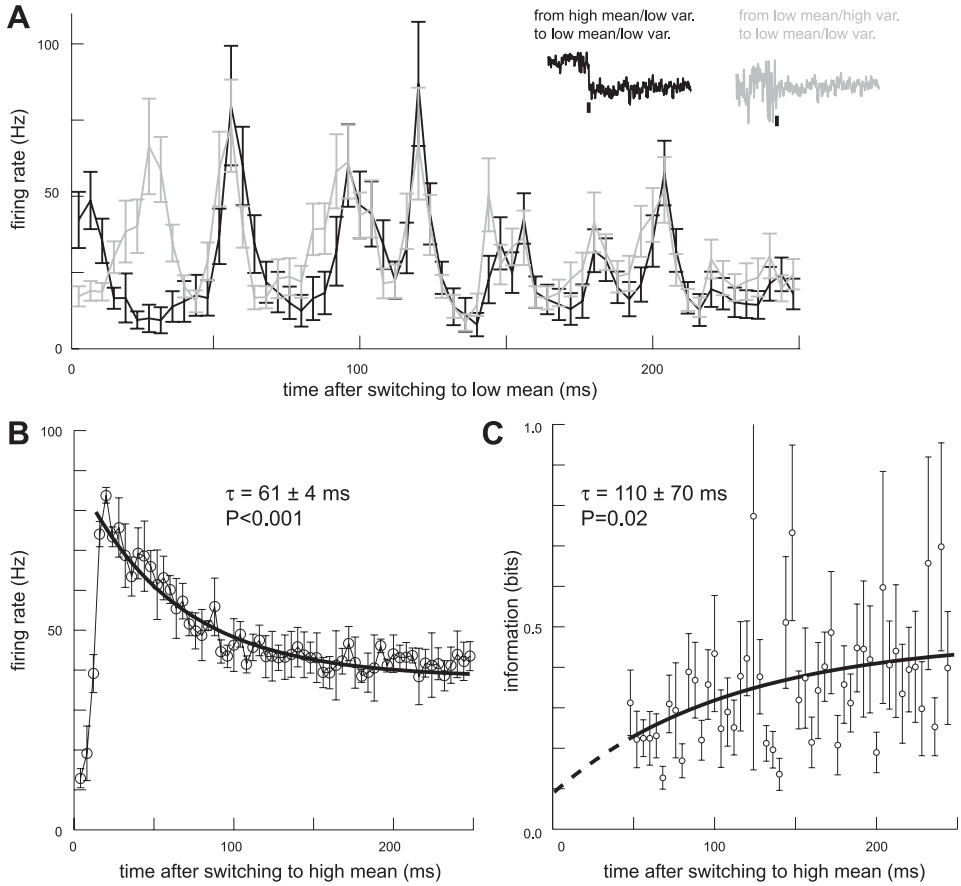

Fig. 7. The time course of adaptation. *A*: the mean population response to the same (repeated) stimulus segment in the low mean/low variance condition is different, depending on the previous state of adaptation. Neurons were previously adapted to either the high mean/low variance condition (black line) or the low mean/high variance condition (gray line). The effect of the previous adaptation condition disappeared after 100 ms. *Inset* at *top right* shows the stimulus as it transitions from two different stimulus statistics conditions to the low mean/low variance stimulus condition (transition is marked by tick marks). *B*: the mean firing rate across the population of cells recorded, after a switch from low to high mean stimulus. The mean firing rate at each moment in time was estimated for each cell from a large number of different (unrepeated) noise segments and then averaged across cells. Thus stimuli are the same, on average, at each point in time. The decrease in mean firing rate reflects the difference in average neural responses to the same average stimuli, indicating an adaptive process. The solid black line shows an exponential fit to the data (black), and the corresponding time constant is provided in the *inset*. *P* value corresponds to an *F*-test comparing an exponential fit with the null hypothesis of no time dependence. Error bars show standard errors of the mean. *C*: the first MID estimated in the stationary state provides an increasingly better description of neural responses with time following a switch to high mean stimulus condition. Each data point represents the average amount of information explained across the population of neurons to the same stimulus segment. Black line and *inset* are as in *B*. The larger variability in *C* compared with *B* stems from the fact that it is based on the population neural response to a single stimulus segment (repeated responses to the same stimulus are required for the information calculation), whereas, in *B*, averaging was done with respect to different stimulus segments and thus much more data.

which we played our stimuli, segments with high mean intensity were always preceded by segments with low mean/low variance. Therefore, we could not make a plot exactly analogous to that in Fig. 7*A*, that is, with a comparison of an identical high mean stimulus preceded by either high or low mean intensity stimuli. However, we could observe the average change in firing rate calculated from all of the different unrepeated noise segments, during the transition from low to high mean stimuli, for each cell and then across the population. Because this average rate was taken across many different noise segments, the stimuli were, on average, the same after the transition to the high mean condition. Thus the observation of the decrease in the average firing rate after the change in mean (Fig. 7*B*) also provides direct evidence of adaptation. The corresponding time constant was 61 ± 4 ms, similar to the time scale of adaptation from the high to low mean stimulus level.

When neural responses are modeled using the LN model, differences in the mean firing rate, like those seen here, can be explained by differences in the shapes either of relevant dimensions, or of nonlinear gain functions. However, in our prior work, we found that, following a switch from low- to high-stimulus mean, changes in nonlinear gain functions (that could not be explained by filter changes) took place over the time scales of seconds (unlike the fast time constants observed in gain function dynamics with changes in variance) (Nagel and Doupe 2006). This suggests that changes in the mean firing rate with changes in stimulus intensity arise primarily as a result of changes in the shapes of relevant dimensions. To assess how the relevant dimensions were changing without tackling the problem of estimating them with very fine time resolution, we instead estimated them during the last 2.5 s of each stimulus segment and then observed how well these "stationary" dimensions could predict the initial responses to the same stimulus. For this calculation, we limited filter duration to 48 ms, because most filters were zero for longer latencies, as can be seen in Fig. 4. We then computed the amount of mutual information captured by these stationary filters about neural responses soon after the switch, using 48-ms computation

windows, each time shifting the computation window by 4 ms. A time-dependent increase in the amount of information explained after a switch in the mean sound volume would indicate that the stationary form of the relevant dimensions becomes an increasingly more appropriate description of neural responses. (Mutual information already takes nonlinear relationships into account, so that this increase could not be due to a change in nonlinearities.) Indeed, we find that the predictive power increased with a time constant of 110 ± 70 ms (Fig. 7*C*), which is consistent with the measurements of firing rate adaptation (Fig. 7*B*) and of responses to the same stimulus in two different adaptation conditions (Fig. 7*A*). Thus several lines of evidence suggest that changes in the relevant dimensions in the auditory system are not instantaneous, as would be expected if they simply reflected static nonlinearities (Borst et al. 2005; Hong et al. 2008), but reflect "fast adaptation," similar to that described in the retina (Baccus and Meister 2002; Victor 1987).

Contribution of secondary dimensions to neural firing increases with stimulus variance. Prior 1D analyses of field L neurons found that, in contrast to the effects of mean amplitude, changing the stimulus variance had little effect on either the mean firing rate or the shape of the relevant dimensions (Nagel and Doupe 2006). Similarly, no changes in the shapes of multiple stimulus dimensions were found in cat dorsal cochlear nucleus (Reiss et al. 2007). Here too we found little effect of stimulus variance on the shapes of the two relevant stimulus dimensions. In all but 4 of 23 neurons for which the responses to both stimulus conditions were available, neither of the Hermite coefficients was significantly different in the different variance conditions ($P > 0.068$, see Table 1). In the remaining four neurons, at least one Hermite coefficient was affected, but no systematic trends were evident.

However, in the 1D study of field L, as well as more generally for studies of neurons in other sensory modalities (see DISCUSSION), changing the stimulus variance had a strong effect on the form of the nonlinear gain function. For multidimensional analyses, the nonlinear gain function identifies how stimulus components along the relevant dimensions determine the neural firing rate. For example, a strong output value along one component may drive the neuron to fire, whereas a strong value along another component may inhibit it (Chen et al. 2007; Maravall et al. 2007; Rust et al. 2005). A neuron may show threshold and saturation effects in its response to the stimulus components. Finally, the stimulus components can interact in nonlinear ways. For example, a coincidence-detector cell might stay quiet if only one of the components is strong, but fire robustly when both components are strong. In general, the relationship between stimulus components and neural firing may be highly nonlinear.

Table 1. *Relevant stimulus dimensions are not affected by changes in stimulus variance*

	And Hermite Coefficient 1	And Hermite Coefficient 2	And Hermite Coefficient 3
Dimension 1	$P > 0.073$	$P > 0.11$	$P > 0.091$
Dimension 2	$P > 0.08$	$P > 0.068$	$P > 0.13$

The first and second relevant dimensions were fitted using the first three Hermite functions. We used *t*-tests to compare, on the neuron-by-neuron basis, the similarity of coefficients in the Hermite expansion of relevant dimensions computed under the low mean/low variance and low mean/high variance stimulus conditions. In only 4 out of 23 neurons did one or more of the Hermite coefficients change significantly ($P < 0.05$). This table shows the smallest P values for each comparison across the remaining population of 19 neurons.

We found that, even within one stimulus condition, e.g., low mean and high variance, nonlinear gain functions had diverse relationships with respect to the two relevant stimulus components. These nonlinear relationships can be visualized by plotting the average firing rate of the neuron as a function of the normalized stimulus component along the relevant dimensions (see MATERIALS AND METHODS). Examples of such visualizations are shown in Fig. 8. The two 1D plots along the sides of each gray-scale plot show the average firing rate (*y*-axis) as a function of one of the stimulus components (*x*-axis) considered independently, while the heat map in the center shows how the firing rate (gray scale) depends on the coincidence of the two components. Figure 8, *A* and *B*, illustrate that the 2D nonlinear gain functions could take very different shapes. With respect to the primary stimulus component only, the neural firing rate functions were typically sigmoidal functions (24/29 neurons in this stimulus condition). However, in some cases, the firing rate could be significantly nonmonotonic, decreasing for large positive component values (cf. Fig. 8*B*, *b1*). The dependence of the firing rate on the secondary stimulus component was even more varied. For example, in Fig. 8*A*, *a2*, the firing rate increased for both positive and negative values of the stimulus components (although not completely symmetrically), whereas, in Fig. 8*B*, *b2*, the firing rate was suppressed by both positive and negative stimulus components. Across the data set for this stimulus condition, the firing rate was enhanced by large components (either positive or negative) along the second stimulus dimensions in 9/29 cases and suppressed in 7/29. The nonlinear gain functions of the remaining cells were of a complex shape that could not be easily classified. Thus field L neurons performed a wide variety of computations with respect to the stimulus time-average and its sequential time derivatives.

Next, we examined how these nonlinear gain functions were affected by a change in the stimulus variance. Changes in stimulus variance are known to affect the gain with respect to the primary dimension (Brenner et al. 2000a; Fairhall et al. 2001; Nagel and Doupe 2006; Reiss et al. 2007; Smirnakis et al. 1997). Recent analysis of multidimensional nonlinearities in the primary somatosensory cortex demonstrated that adaptation to variance led to a common rescaling of the gain with respect to the primary and secondary stimulus components (Maravall et al. 2007). In the case of field L neurons, we also generally found that the width of the input range over which the firing rate varied from its minimal to maximal values increased with variance. Therefore, in the following, we present and discuss nonlinear gain functions in rescaled units where the stimulus components are measured in units of their standard deviation (Fig. 8). In these rescaled coordinates, the nonlinear gain functions in different conditions should overlay each other if the gain rescales perfectly with adaptation to variance. However, the shapes of gain functions, especially with respect to secondary dimensions, in these rescaled coordinates were usually shallower in the low-variance condition and became steeper or more peaked in the high-variance condition (Fig. 8, *C* and *D*). This suggests that the contribution of secondary dimensions increases with variance. This phenomenon can be

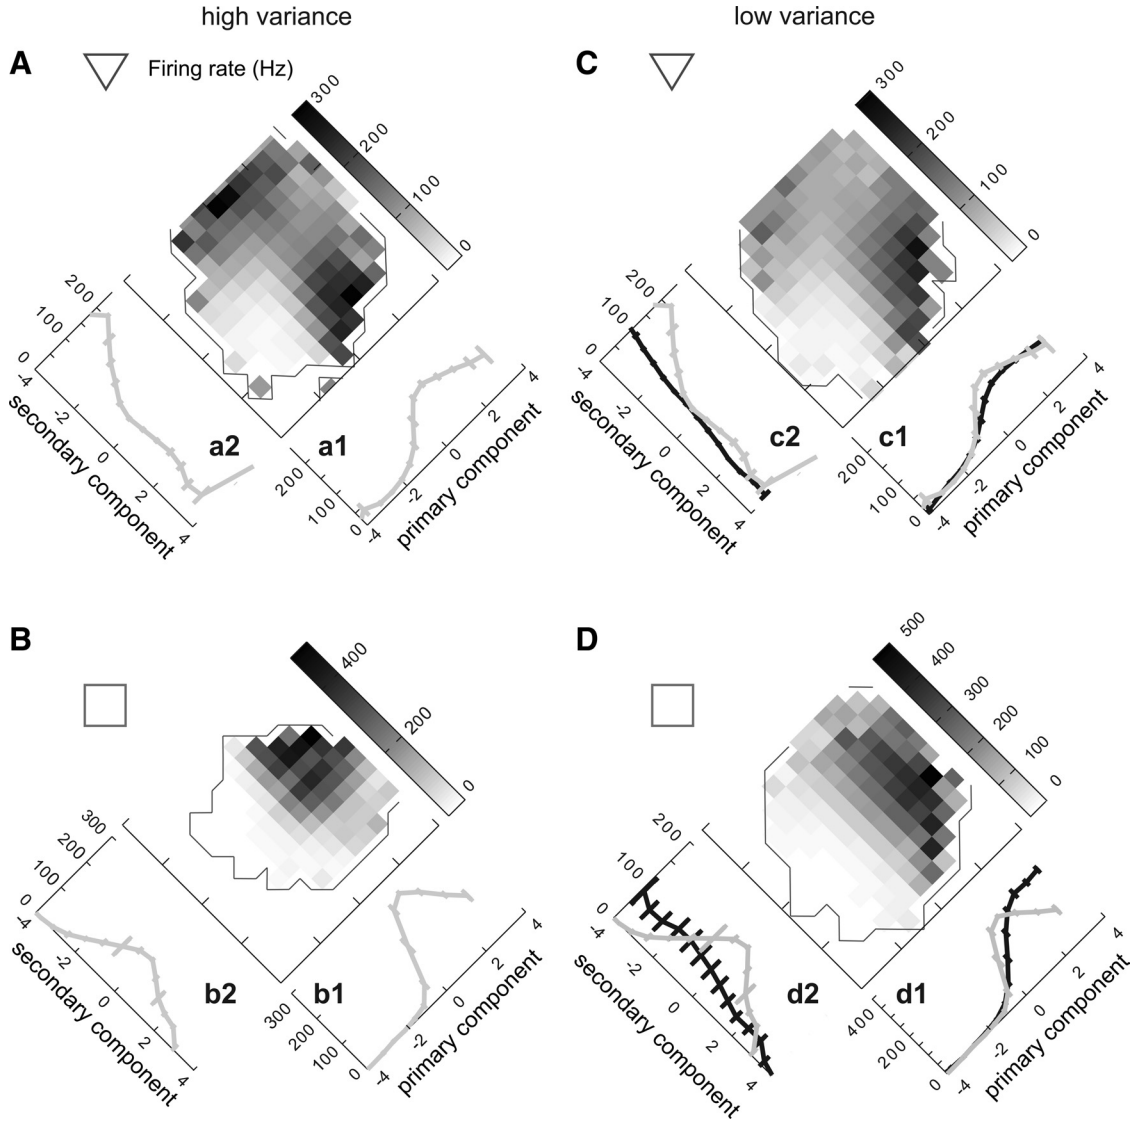

Fig. 8. Effects of stimulus variance on the 2D nonlinear gain functions for two example neurons (*A*/*B* and *C*/*D*). Gray-scale plots show the firing rate in Hz as a function of the two relevant stimulus components. Thin black lines show regions where values are >2 SE. The average firing rate as a function of individual components of the stimulus is shown in side plots (gray and black lines for the high- and low-variance conditions, respectively; the gray lines of the high variance condition are replotted in the right half of the figure for comparison). Stimulus components are normalized to have unit variance and plotted in units of standard deviation. *Top* row is one neuron ("udon2120"); *bottom* row is a different neuron ("eb1940"). Although the gain of 2D functions shows some rescaling with stimulus variance (compare columns), it is not complete for secondary dimensions, and in some cases (e.g., *bottom* row) the shape of the secondary filters can change qualitatively.

quantified by the relative increase in information explained when the LN model is expanded to include the two dimensions (Fig. 9). A perfect rescaling along both dimensions, with no change in shape of the 2D gain function, would leave the relative increase in information unchanged. In contrast, either changes in the shape of the gain function, or unequal rescaling of gain with respect to the two relevant dimensions would cause this information ratio to deviate from 1. Across the population of neurons, the contribution of the second dimension was significantly larger in the high-variance compared with the low-variance condition (Fig. 9; $P = 0.016$, paired Wilcoxon test). In addition, this change reached significance in 35% of neurons considered individually (8/23). Figure 8 illustrates gain functions for the example neurons marked in Fig. 9: the difference in relative information between Fig. 8*B* and Fig. 8*D* was significant ($P < 0.01$, *t*-test), whereas, although the same trend was there in Fig. 8, *A* and *C*, it did not reach significance. In summary then, adaptation to stimulus variance can alter the degree to which relevant dimensions influence neural firing, even in cases where the relevant dimensions themselves are not affected. Thus multidimensional feature selectivity confers an extended variety of adaptive behaviors that go beyond adaptive changes available for separate encoding of stimulus components.

A possible role of spike-time jitter. Our multidimensional studies of field L suggest a striking sensitivity of neurons to sequential time derivatives of the stimulus. One mechanism that could give rise to a sensitivity of neural responses to time derivatives is jitter in spike timing (Aldworth et al. 2005; Dimitrov and Gedeon 2006; Dimitrov et al. 2009; Gollisch 2006).That is, if the neuron is sensitive to a single feature, but spikes arrive at slightly different times, it might appear that the

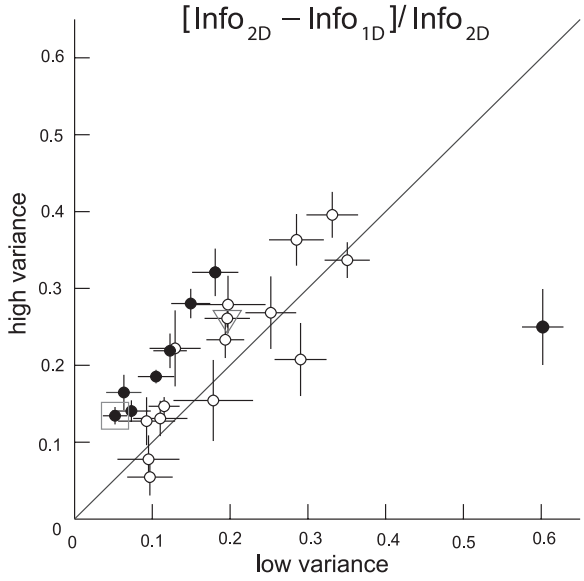

Fig. 9. Contribution of secondary dimensions to neural firing increases with stimulus variance. We measure relative information gain by computing a difference between the information accounted for by two dimensions ($Info_{2D}$) and that accounted for by the primary dimension alone ($Info_{1D}$), and then dividing by $Info_{2D}$. Such a ratio represents a way to measure the relative influence of the two dimensions on neural firing. A ratio different from one indicates a change in the shape of the nonlinear gain function relative to the stimulus probability distribution, for example, due to imperfect or uneven rescaling of firing rate gain with respect to relevant stimulus dimensions. Shown are the low mean/low variance stimulus condition (*x*-axis) vs. low mean/high variance condition (*y*-axis). Across the population of neurons, the contribution of the second dimension was significantly larger in the high variance compared with the low-variance condition ($P = 0.0156$, paired Wilcoxon test). Points that lie significantly above the line indicate that the inclusion of a second dimension increased the information more in the high variance condition than in the low; points below the line indicate more information from a second dimension in the low-variance condition. Each symbol is a neuron, black $P < 0.05$, white $P > 0.05$, *t*-test. Neurons exhibiting a significant change in the shape of the 2D gain function were characterized by a weaker contribution of the second dimension with a low-variance stimulus (7/8). The points marked with symbols are the same as in Fig. 8.

neuron is sensitive to two slightly displaced features, which could equivalently be described as a sensitivity to the feature and its first time derivative. However, as noted in Dimitrov and Gedeon (2006), time derivatives are not always due to the presence of time jitter. To check for these effects, we first estimated the amount of time jitter in these data. Two methods have been discussed for estimating jitter in spike timing in the literature. One involves separating the spike trains into "events" (Berry et al. 1997) and then measuring directly the variance in the arrival times of individual spikes within the events. This approach works well when events are clearly defined by the absence of spikes in certain segments of the repeated data. Analysis of jitter for one such event (underlined in red in Fig. 1*B*) yielded an estimate of 0.75 ms for jitter in the timing of the first spike across different trials. However, such clear separation of events is not always possible. Another method for estimating jitter in spike timing is to optimize the shape of relevant dimensions while allowing shifts in the times of the spikes with a penalty term that depends on the magnitude of the shift (Aldworth et al. 2005; Dimitrov and Gedeon 2006; Dimitrov et al. 2009; Gollisch 2006). This method has the advantage of not relying on the separation of spike trains into events. Implementing the technique described in Aldworth et al. (2005; see MATERIALS AND METHODS), we found that jitter for all cells, except one, was in the sub-millisecond range (Fig. 10*A*). The fact that the three outlier points (with spike time jitter >5 ms, see Fig. 10*A*) belonged to the same cell indicates that the optimization procedure allowed for the possibility of larger spike jitter, but also that the result that all other cells have spike time jitter of <1 ms was robust. This analysis was carried out on spike trains binned at 1-ms resolution to better estimate the amount of time jitter in the data. To check more directly for the role of time jitter, we also applied the dejittering technique to spike trains binned at the 4-ms resolution at which all other analyses of feature selectivity were carried out, and then reestimated features. Figure 10, *B* and *C*, illustrates with typical examples that the time-derivative characteristics of the shapes of relevant dimensions persist even after dejittering. Across the population, the 2D relevant spaces estimated with and without prior dejittering were highly similar. We quantified the similarity between the two spaces using a measure termed subspace projection, which ranges between 0, if subspaces do not overlap, and 1 for a perfect match (Rowekamp and Sharpee, 2011; see MATERIALS AND METHODS). Using this measure, we find that, except for one neuron in the low mean condition and three neurons in the high mean condition that had subspace projection values <0.4, the rest of the population had subspace projection values ranging from 0.85 to 0.9999 (mean 0.98), both in the low and high mean conditions. Thus the estimation of two relevant features was not affected by time jitter in almost all cases.

Possible benefits of derivative sampling. As we have seen above, when the sound level changes, both the primary and secondary dimensions undergo marked changes in their shapes. However, the systematic relationship between the dimensions, in which one of them approximates the time differential of the other dimension, persists, despite such different stimulus conditions. This experimental result raises the possibility that including time differentials is functionally significant, at the single-neuron level, for the encoding of temporally modulated sounds. One possibility that we will explore here is that the inclusion of time differentials is a useful strategy for representing continuous input signals using discrete spike-based representations, especially in the case of naturalistic stimuli.

It is well known that any continuous waveform can be represented without any information loss with a sequence of discrete measurements taken at a frequency equal to or greater than the so-called Nyquist frequency, which is twice the maximal frequency *W* in the input signal (Fig. 11*A*). However, this traditional sampling strategy is not the only lossless way to represent a continuous waveform with discrete data points. The same continuous waveform can be represented using pairs of measurements, where the values of the signal and its first time differential are recorded at one-half the Nyquist rate (Fig. 11*B*). In the absence of noise, either of these sampling strategies can be used to reconstruct the incoming continuous waveform without any bias (Shannon 1949). The two sampling strategies can be related to one another by noting that measurements of a function value and its time derivative approximate the sum and difference of two measurements of the function at nearby points in time. Therefore, one can think of the strategy of simultaneously sampling the values of the signal and its first time derivative as equivalent to merging every set of two sample points in the traditional strategy. Furthermore, one can

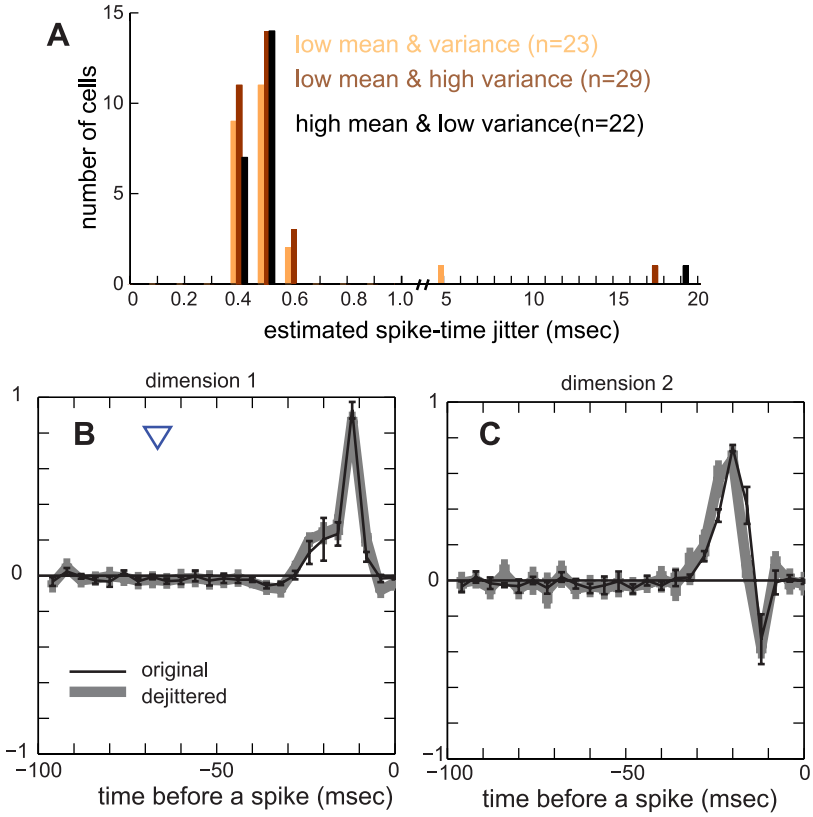

Fig. 10. The distribution of spike-time jitter. *A*: spike-time jitter was estimated following the method of Aldworth et al. (2005), whereby estimates of the relevant stimulus dimension (computed as STA) are readjusted by allowing varying delays in the arrival times of single spikes. The standard deviation of a Gaussian distribution for jitter in spike time was computed separately for each cell and stimulus condition. The three outlier points belong to the same neuron "pho1295wide" under the three different stimulus conditions. Thus nearly all cells exhibit a very low degree of spike-time jitter with submillisecond precision in spike timing. In *B* and *C*, we illustrate typical results of MID optimization obtained with (thick gray line) and without prior dejittering of spike trains binned at 4-ms resolution. The subspace projection between the two relevant spaces for this example is 0.99. Neuron "udon2120".

merge every set of three sample points in the traditional sampling strategy to arrive at a three-dimensional sampling strategy where the waveform is represented by values of its function, and both its first and second time differentials, with each set of measurements now taken at one-third of the Nyquist rate (Fig. 11*C*). More generally, and importantly for neural coding where spikes fire irregularly, time-varying stimuli do not need to be sampled at regular intervals. Instead, there simply needs to be two measurements on average within the period of the maximal signal frequency (Jerri 1977; Shannon 1949).

Although lossless stimulus reconstruction from discrete samples is possible for both the traditional sampling and for the higher-order sampling strategies that include sequential time differentials, these strategies differ in their susceptibility to noise. If high frequencies provide significant contribution to noise, then noise values change rapidly, making even nearby noise values uncorrelated (Fig. 11*D*, *top* vs. *middle* panels). This, in turn, can corrupt measurements of the time differentials, which involves subtracting one noisy point from another. This explains why the traditional sampling strategy, where only signal values are measured, is used most often in situations with high-frequency noise. The situation is different, however, when it comes to detecting sounds in a natural auditory environment. Background vocalizations, environmental sounds, and neural variability within the nervous system are all dominated by low frequencies, with a gradual decrease in the noise power spectrum as a function of frequency f, usually proportional to $\sim 1/f^2$ (Lewicki 2002; Ruderman and Bialek 1994; Singh and Theunissen 2003; Teich et al. 1997; Voss and Clarke 1975). In this situation, nearby time points tend to share the same noise offset, so subtracting them to compute a derivative may actually increase the quality of the signal waveform reconstruction (Fig. 11*D*, *bottom* panel). Field L neurons normally operate in natural settings, where signal and noise sources have a similar distribution of power across frequencies, and both are dominated by low frequencies. Therefore, the sensitivity of these neurons to both the average value of the signal amplitude and its rate of change may serve to increase the robustness of the neural reconstruction of the continuous signal. At the same time, such selectivity to multiple stimulus components could begin to explain the hypersensitivity of auditory cortical neurons to small perturbations of their acoustic input (Bar-Yosef et al. 2002), as well as the large effect that the naturalistic background noise can have on the responses of such neurons (Bar-Yosef and Nelken 2007).

To test these ideas about the benefits of multidimensional encoding quantitatively, we compared how well the stimulus waveform could be reconstructed (by calculating the average mean square error) when we used three different types of sampling [sampling either the signal values alone at the rate $1/(2W)$, the signal and first time derivative values at the rate $1/(W)$, or the signal and first and second time derivative values], and allowed for the possibility that the sampled measurements were corrupted by noise. We varied the frequency composition of the noise by varying its correlation time τ (where high-frequency noise has low τ, and vice versa), using a Gaussian distribution $C(f) = C_0\, \tau \exp(-f^2\, \tau^2/2)$. Parameter τ describes the average time over which noise correlations persist; its product with the maximal signal frequency, W, characterizes the ratio of noise correlation time to the fastest period in the signal. When the correlation time τ is short, and the product $W\, \tau$ is, therefore, small (Fig. 11*D*, *middle* panel), noise contributions at frequencies larger than W are significant.

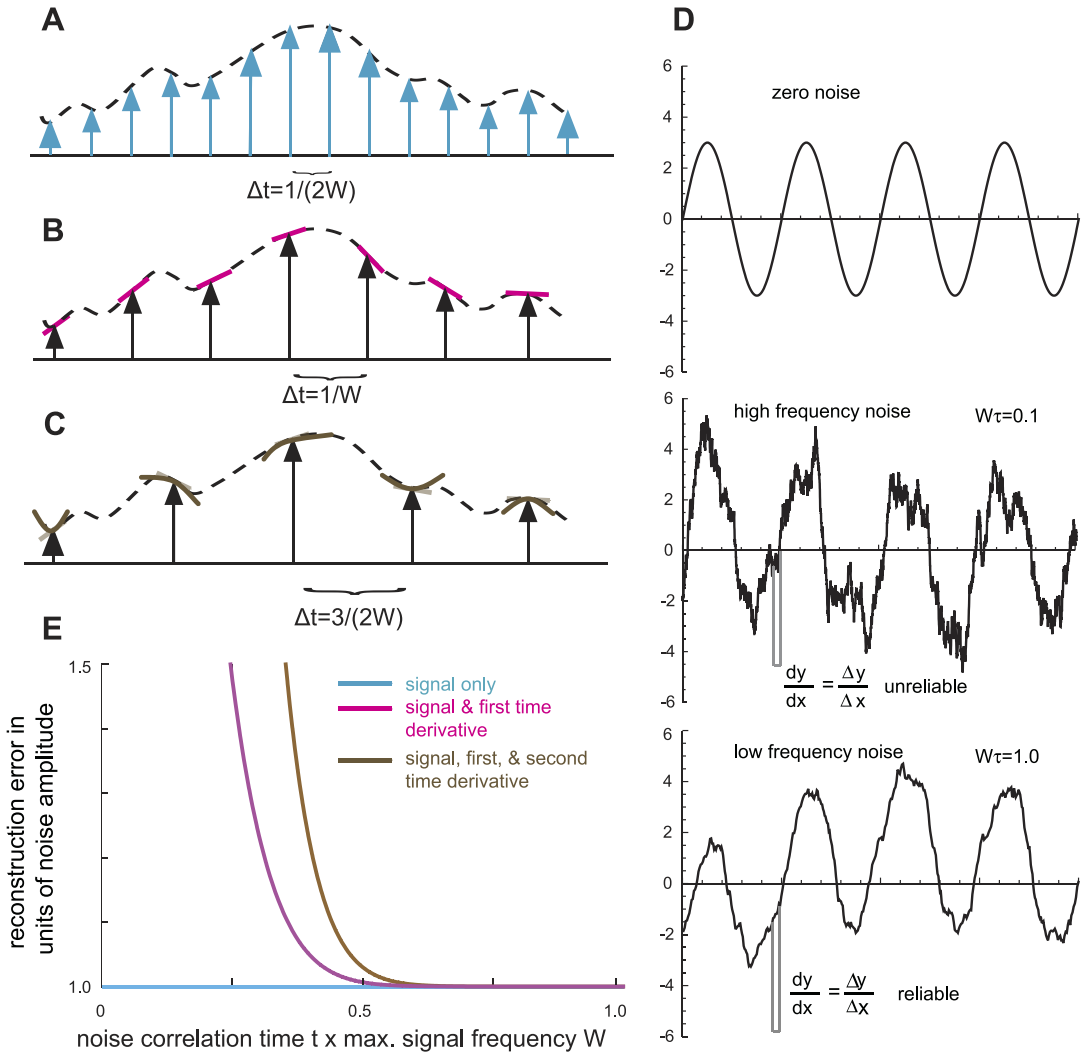

Fig. 11. Different sampling strategies and their susceptibility to noise. *A*: in traditional sampling, function values are measured at an average rate of $2W$, where W is the signal bandwidth. *B*: when the signal and its time differential are sampled simultaneously, measurements can be done at the reduced rate of W. *C*: when the signal and its first and second time differentials are measured together, the sampling rate can be further reduced to $2W/3$. All three strategies are equivalent in the absence of noise. *D*: illustration of the effects of noise on the computation of derivatives. *Top*: signal with no noise. *Middle*: in the presence of high-frequency noise, nearby time points have very different noise values, making the computation of time-derivatives noisy ($W\tau = 0.1$). *Bottom*: in the presence of low-frequency noise, nearby time points generally have similar noise values, making the derivative computation more robust ($W\tau = 1.0$). *E*: mean reconstruction error in a linear model where sample values are corrupted by noise. The results are plotted as a function of noise correlation time τ. Blue line, traditional sampling of signal values; *left* magenta line, sampling of signal and first time derivatives; *right* brown line, sampling of signal, first and second time derivatives.

In contrast, when τ is long and the product $W\tau$ of order 1 and larger (Fig. 11*D*, *bottom* panel), noise contributions are concentrated at frequencies below W, potentially allowing for a reliable computation of the derivative.

The traditional sampling strategy based on sampling the signal values alone has an important advantage whereby its reconstruction error is always equal to the noise variance, independent of how noise power is distributed across frequencies (Jerri 1977; Shannon 1949). Thus the reconstruction error in this case is independent of the noise correlation time. Because the mean square error always increases proportionally to the noise variance, we report results for this and other sampling strategies in units of noise standard deviation in Fig. 11*E*. In these units, the reconstruction error of the traditional sampling strategy is indicated by a line at $y = 1$ in Fig. 11*E* and will be used as the benchmark for other sampling strategies.

In contrast to the sampling of signal values alone, reconstruction errors for sampling strategies that rely on measurement of derivatives can depend strongly on the noise correlation time τ, just as suggested by the intuition in Fig. 11*D* (see APPENDIX A for details of the derivation). When the noise correlation time τ is much shorter than $\sim 1/(2W)$, which is the minimum time interval between samples taken at the Nyquist limit for a function with a maximum frequency of W, the sampling strategies that involve derivative measurements perform much more poorly than the traditional sampling (see *left* portion of Fig. 11*E*). This agrees with idea illustrated in Fig. 11*D* that high-frequency (low τ) noise can corrupt the estimate of time derivatives. However, when the noise correlation time increases beyond the value of $1/(2W)$ (i.e., beyond $W\tau = 0.5$, cf. *center* and *right* portions of Fig. 11*E* when noise becomes lower frequency), the mean square error associated with mea-

suring time derivatives improves dramatically, reaching the benchmark of one as $W\tau$ increases further (magenta and brown lines, Fig. 11*E*). Importantly for this idea, in many natural sensory environments, $W\tau$ is expected to be on the order of 1. This occurs because, in such settings, not only are stimuli dominated by lower frequencies, but the main sources of "noise" are also from other natural stimuli. When signals and noise are of the same origin, correlations between noise values persist over similar time scales as those between the signal values, making W τ close to 1. Therefore, for naturalistic situations, multidimensional sampling strategies can be just as viable as the traditional sampling strategy. We note that the relative comparisons of the mean square error remain valid, even if the reconstruction takes steps to optimally suppress noise in ways that extend Wiener filtering arguments for continuous signals to the case of discrete samples (see APPENDIX B).

An important advantage of combined sampling of function values and their time differentials is that samples can now be taken at reduced frequencies (Fig. 11). The duration of the neural filter T or of its spike waveform, which correlates positively with T (Nagel and Doupe 2008), sets the limit on the sampling rate $\Delta t \sim 1/T$ that can be sustained by the neuron. Therefore, a neuron with a broad relevant dimension may not be able to sustain the same sampling rate as a neuron with a narrower relevant stimulus dimension. To encode the same fast-changing stimulus, the neuron with a broad relevant stimulus dimension can instead rely on a multidimensional sampling strategy that is based on combinations of the local time-average of the signal and its time derivatives (Fig. 11, *B* and *C*). Although higher-order derivatives are progressively more susceptible to noise, such strategies become acceptable (Fig. 11*E*) when noise is strongly correlated in time, with a strong contribution from low temporal frequencies, as in natural stimuli.

How would these effects play out in the context of stimulus encoding with a population of neurons? Here, in addition to lowering the required sampling rate, multidimensional encoding may also reduce the overall effective noise level. This is because the additive noise that we consider has two components. The first component arises because the stimulus values are not known precisely. The second component arises because the sampling times are not known precisely. It has been shown previously (Jerri 1977) that the second component can be mapped onto the first one, i.e., one can consider the reconstruction where the sampling times are known exactly, but with increased uncertainty in the stimulus values. This is the simple calculation that we did. However, in a model with populations of neurons, each of which is sampling the stimulus at many times, the noise in sampling time, or jitter, becomes more increasingly important. In such a model, the noise standard deviation relative to which the reconstruction error is measured may be further reduced by multidimensional encoding, because the reduced number of samples in such sampling reduces the second component of the noise. On the other hand, more neurons might be needed to encode the many 2D stimulus combinations of the mean and its time derivatives, compared with encoding of the 1D mean stimulus value. We, therefore, envision that the optimal dimensionality of stimulus encoding might be determined by the trade-off between the number of neuronal responses necessary to reconstruct a multidimensional stimulus and the decrease in sampling and accompanying jitter noise possible with multidimensional encoding. In light of our experimental observation of 2D encoding, however, the goal of the simple analysis here of the mean square error is to provide some intuition about possible benefits of two- and higher-dimensional encoding strategies.

DISCUSSION

Adaptation lies at the heart of sensory signaling. Although a great deal has been learned about how adaptation affects primary feature selectivity in sensory neurons, we are only beginning to determine how it reshapes complementary stimulus features that, for example, increase selectivity and contribute to sparseness of neural responses.

In this paper, we demonstrated that the firing rate of single neurons in the field L region of the auditory forebrain was affected by at least two features of the sound log-amplitude waveforms. A strikingly systematic relationship between the relevant dimensions was observed, in which one of them approximated the time differential of the other dimension, although either dimension could play the dominant role. This systematic relationship was preserved even when both dimensions underwent marked changes in their shape following a change in the stimulus mean. Unlike adaptation to stimulus mean, adaptation to stimulus variance did not change the relevant dimensions, but altered their relative contributions to neural firing. Finally, the systematic relationship observed between the primary and secondary dimensions can be seen as a strategy for minimizing the sampling rate while preserving the accuracy of reconstruction of stimuli correlated over long times (like natural sounds).

Auditory neurons have been known to adjust their filtering properties in response to changes in mean sound amplitude and auditory contrast (Frisina et al. 1990; Krishna and Semple 2000; Kvale and Schreiner 2004; Lesica and Grothe 2008a; Nagel and Doupe 2006; Nagel and Doupe 2008; Rees and Moller 1987) or with increasing ambient noise (Lesica and Grothe 2008b) or higher moments of the stimulus distribution, such as kurtosis (Kvale and Schreiner 2004). At low sound volumes, neurons predominantly show averaging or integrating characteristics with respect to both sinusoidally (Frisina et al. 1990; Krishna and Semple 2000; Rees and Moller 1987) and randomly amplitude-modulated sounds (Lesica and Grothe 2008a; Nagel and Doupe 2006; Nagel and Doupe 2008). This allows them to detect faint signals embedded in noise at the expense of fine temporal resolution (Lesica and Grothe 2008b). With increasing sound volume, neurons increase their "differentiating" characteristics, enhancing their temporal resolution for important sound features such as onsets and offsets. Such intensity-dependent changes in neural filtering occur both in single neurons when the mean of injected currents is increased (Mainen and Sejnowski 1995) and in the visual system, especially the retina (Atick and Redlich 1992; Baccus and Meister 2002; Chander and Chichilnisky 2001; Enroth-Cugell and Lennie 1975; Kuffler et al. 1957; Srinivasan et al. 1982), where neurons lose their differentiating surrounds in dim light (low signal to noise). Our results demonstrate that, even at the soft sound pressure levels (30 dB), where integrating characteristics dominate in most neurons, both the integrating and the differentiating components already make nonlinear contributions to neural firing of single neurons, regardless of whether these

neurons are primarily driven by the integrating or differentiating components. At high stimulus mean, the two relevant quantities switch to the first and second time derivatives of stimuli. These two findings could not have been predicted based on previous analyses of single relevant dimensions (Nagel and Doupe 2006).

A sensitivity to multiple stimulus components has been observed in several sensory systems, including the retina (Fairhall et al. 2006) and olfactory sensory neurons (Kim et al. 2011), the primary visual (Rust et al. 2005; Touryan et al. 2002), auditory (Atencio et al. 2008; Zhou and Wang 2010), and somatosensory (Maravall et al. 2007) cortices. In the auditory system, multidimensional description of neural coding may provide a complementary way to account for dynamic effects in studies of neural sound detection in the quiet (Heil et al. 2008). In that work, much of the variance in the first-spike latency can be explained by a 1D model, where a fixed threshold is applied to the low-pass filtered sound amplitude (Heil and Neubauer 2001; 2003). At the same time, full description of the data from auditory nerve fibers and auditory cortical neurons also requires the introduction of time dependence, where the threshold value increases sublinearly with latency (Heil and Neubauer 2001; 2003). The 2D model with a fixed threshold that is a function of both the low-pass filtered amplitude and its first time derivative may provide another way of capturing this phenomenon. This observation raises the possibility that the selectivity for multiple features that we describe could first arise early in the ascending auditory system, but be augmented by processing at many stages. On the other hand, recent data demonstrate that the importance of the second dimension increases greatly between the input and higher layers of cat primary auditory cortex, suggesting that multiple-feature selectivity is a property of higher auditory processing stages (Atencio et al. 2009).

In the context of auditory and other sensory data, the finding that neural responses in the auditory forebrain of songbirds are affected by multiple stimulus components provides further evidence for the common principles of sensory encoding across multiple species and sensory modalities. Furthermore, the information gain from adding the second dimension relative to the information of the 1D model was very similar in our data set (median 25%, mean 36%) to values previously reported in the primary mammalian auditory cortex (median 25%) (Atencio et al. 2008) and in the retina (Fairhall et al. 2006). In the case of the retina, the relative information gain could be more than 100% for some cells, and this was also true in our data set, which contained five such cases. Thus, the 2D encoding we observe in the field L is just as significant as in other sensory systems. Furthermore, even smaller increases, on the order of 10% in predictive power, are considered "noteworthy" (Gill et al. 2008, see also Prenger et al. 2004). Therefore, our analysis of how two filters per neuron are affected by adaptation to the stimulus mean and variance should be of relevance to a number of other sensory systems.

We found that, in field L, the relevant stimulus dimensions had specific profiles, representing consecutive time differentials of the stimulus waveform. In the case of soft sounds, this feature was the log-amplitude itself, whereas, with loud sounds, the feature of interest was the first time differential of the log-amplitude. Similar types of derivative-based encoding have also been observed in dynamic models of single-neuron spiking (Hong et al. 2007) and in slices from the brain stem nucleus magnocellularis (Slee et al. 2005). Combined with our theoretical analysis of the mean square reconstruction error, the observation of derivative-based encoding again speaks to the common problems and solutions employed in different neural systems. Moreover, because single-neuron dynamics can only explain derivatives as secondary, not as primary, dimensions, an additional finding here is that time derivatives may play either the primary or secondary role. This suggests that relevant dimensions in field L neurons must also be shaped by network mechanisms, as well as neuron dynamics.

In sensory coding, sensitivity to the velocity and acceleration in whisker deflections has been observed in the somatosensory cortex (Maravall et al. 2007). That study examined adaptation to stimulus variance, but not to stimulus mean, and did not report changes in the stimulus dimensions. Our finding that relevant dimensions in the auditory forebrain also do not change following adaptation to variance provides a third line of evidence for the similarities of adapting strategies across the sensory systems. At the same time, our results on filter changes in response to increases in stimulus mean further suggest that, in the somatosensory cortex, adaptation to stimulus mean might also trigger changes in the multiple relevant dimensions.

Analysis of neural responses to stimuli of different variance revealed the possibility of diverse and, to our knowledge, novel effects of adaptation on the shape of multidimensional nonlinear gain functions. Previously, gain control mechanisms have been primarily studied with respect to the primary stimulus dimensions. A recent 2D study in the somatosensory cortex found that adaptation to stimulus variance did not change the first and second stimulus dimensions, but rescaled the firing rate gain with respect to the dimensions by a common factor proportional to stimulus variance (Maravall et al. 2007). Note that, in that system, the shapes of nonlinear gain functions were similar across both neurons and stimulus conditions, with firing rate increasing as a square of stimulus components along the relevant dimensions. In contrast, in field L, we observed a great variety of nonlinear gain functions, even for a given stimulus condition. Moreover, following a switch to a higher stimulus variance, the contribution of the second component increased on average across the population, an effect that would not be expected for a common rescaling of firing rate gain along both the primary and secondary stimulus dimensions. For neurons with an especially strong increase in the contribution of the second component with variance, the shape of the 2D firing rate function was changed altogether. Thus, even though changes in stimulus variance do not affect the shape of relevant dimensions, a number of adaptive phenomena are possible that take advantage of the multidimensional character of neural feature selectivity.

With respect to the mechanism through which the 2D feature selectivity might arise, one possibility is that sensitivity to both a feature and its first time derivative arises as a result of threshold crossing (Agüera y Arcas and Fairhall 2003; Agüera y Arcas et al. 2003; Deweese 1996; Fairhall et al. 2006; Hong et al. 2007; Slee et al. 2005). However, in these situations, the feature itself is always dominant. In contrast, we found that either the average or the derivative of the feature could play the dominant role in neural coding. In addition, we found that neurons perform diverse computations within this 2D input space that are usually not sensitive to the particular sign of the

derivative components (Fig. 8, *A*, *a2*, and *B*, *b2*), as would be the case if the selectivity to the derivative feature arose as a result of threshold crossing. Another mechanism that could give rise to a sensitivity of neural responses to time derivatives is jitter in spike timing (Aldworth et al. 2005; Dimitrov and Gedeon 2006; Dimitrov et al. 2009; Gollisch 2006). However, spike timing jitter in the responses of these neurons was quite small ($<$ 1 msec for nearly all cells; see red underlined event in Fig. 1b for an example, as well as Fig. 10 for population data), just as in mammalian auditory cortex (DeWeese et al. 2003) and other sensory systems (Bair 1999; Bialek et al. 1991; Bryant and Segundo 1976; Mainen and Sejnowski 1995; Petersen 2007; Tiesinga et al. 2008). The fact that the precision of spike timing is much smaller than the temporal scale of the relevant dimensions rules out the possibility that the sensitivity to a time derivative is due to jitter in spike times. Finally, neither the threshold crossing mechanism nor time jitter effects are likely to explain a change in the filter shapes with increasing mean sound amplitude.

Our theoretical analysis demonstrates that multidimensional encoding based on sequences of time derivatives can be just as robust to distortions from noise in natural environments as encoding based on just one primary stimulus feature per neuron. Multidimensional encoding can be realized either at the level of individual neurons, as we observed here for single field L neurons, or at the level of neural populations, where each neuron would be sensitive to just one stimulus feature that would take different shapes for different neurons. However, one of the key advantages of multidimensional sampling is that it reduces the necessary sampling rate. This reduction can allow even neurons with broad spikes and/or filters to participate in encoding of fast-changing stimuli by measuring both the average and the rate of change of the stimulus. Such sampling rate limitations are less likely to be relevant for encoding in which stimulus features are distributed across a population of different neurons. Our result suggests, however, that multidimensional encoding might be a necessary feature of stimulus encoding by single neurons and could be more pronounced among pyramidal neurons compared with fast-spiking interneurons with narrower spikes.

The theoretical arguments here also raise the possibility that the dimensionality of neural representations can be affected by adaptation to the statistics of the noise, even when the mean and variance of the sound log-amplitude are held constant. The prediction is that the contribution of the secondary stimulus dimensions would increase with increasing contributions of low temporal frequencies in either the stimulus or noise distributions. Thus, as neurons are probed with naturalistic noise of increasing correlation time, the second dimension might play an increasingly larger role relative to the first. Recent work (Garcia-Lazaro et al. 2006) indirectly supports this hypothesis. There, the authors demonstrated that the responses of neurons in the primary auditory cortex are more reproducible when probed with stimuli with $\sim f^{-2}$ power spectra, compared with sounds with faster or slower decreases of stimulus power with frequency. This suggests that the sampling strategy neurons use is optimized for naturalistic, as opposed to other kinds of signals.

In summary, adaptation to mean sound amplitude affects both the primary and secondary stimulus dimensions. For both soft and loud sounds, the relevant dimensions represent consecutive time differentials, although the order of the differentials changes from zero (integration) and first, in the case of soft sounds, to first and second, in the case of loud sounds. These data suggest that the selectivity to combinations of time differentials that is observed in a number of sensory systems, from single neuron dynamics to somatosensory cortex, and now in auditory encoding, is a consequence of having to encode continuous natural stimuli using discrete spikes.

APPENDIX A: DERIVATION OF RECONSTRUCTION ERRORS FOR DIFFERENT SAMPLING STRATEGIES

In this section, we provide the analytic expressions for the dependence of the reconstruction error on the noise correlation function, for different sampling strategies. In all cases, we will assume that a time-varying signal $s(t)$ is band-limited to frequencies $|f| < W$ Hz. The process of reconstruction of the continuous signal from discrete samples is based on convolving the sequence of these samples with certain functions, which we refer to as convolution kernels, to estimate the value of the function in between the sampling points. The explicit solutions for these kernels that would allow us to perfectly reconstruct the signal in the absence of noise were obtained by Shannon (1949) for the traditional 1D sampling strategy, and by Jerri (1977) for multidimensional sampling strategies. The corresponding expressions in the time and frequency domains are as follows (Bracewell 1986).

For the 1D sampling strategy, there is just one kernel $u^{(1D)}(t)$ to be convolved with a sequence of function values taken at intervals of $\Delta t = 1/(2W)$:

$$u^{(1D)}(t) = \frac{\sin \pi t/\Delta t}{\pi t/\Delta t}, \quad U^{(1D)}(f) = 1/(2W) \tag{A1}$$

where the two expressions describe this kernel in the time and Fourier domains, respectively.

For the 2D sampling strategy where the values of the function and its first time derivative are measured at intervals of $\Delta t = 1/W$, there are two convolution kernels. The kernel to be convolved with the sequence of function values is given by:

$$v_0^{(2D)}(t) = \left(\frac{\sin \pi t/\Delta t}{\pi t/\Delta t}\right)^2, \quad V_0^{(2D)}(f) = \frac{W - |f|}{W^2} \tag{A2}$$

in the time and Fourier domains, respectively. The corresponding expressions for the second kernel to be convolved with the stream of derivative measurements can be written as:

$$V_1^{(2D)}(f) = \frac{|f|}{2\pi i W^2 f}, \quad v_1^{(2D)}(t) = t v_0(t) \tag{A3}$$

Both $V_0(f) = V_1(f) = 0$ for frequencies $|f| > W$.

Finally, in the case of the 3D sampling strategy where the values of the function, together with its first two time derivatives are recorded at intervals of $\Delta t = 3/(2W)$, there are three different convolution kernels. The three kernels are to be convolved with three different streams of samples: $w_0^{(3D)}(t)$ with function values, $w_1^{(3D)}(t)$ with values of the first derivative, and $w_2^{(3D)}(t)$ with values of the second derivative. The corresponding expressions in the time domain are:

$$w_0^{(3D)}(t) = \left(\frac{\sin \pi t/\Delta t)}{\pi t/\Delta t}\right)^3 \left(1 + \frac{\pi^2 t^2}{2\Delta t^2}\right),$$
$$w_1^{(3D)}(t) = t\left(\frac{\sin \pi t/\Delta t}{\pi t/\Delta t}\right)^3, w_2^{(3D)}(t) = \frac{t^2}{2}\left(\frac{\sin \pi t/\Delta t}{\pi t/\Delta t}\right)^3 \tag{A4}$$

Each of the three reconstruction strategies, either 1D described by *Eq. A1*, 2D described by *Eqs. A2* and *A3*, or 3D described by *Eq. A4*

would yield the zero reconstruction error between the incoming signal $s(t)$ and the reconstructed waveform $\hat{s}(t)$ in the absence of noise (i.e., they are unbiased reconstruction strategies). In the presence of noise, the mean reconstruction error $\langle \delta s^2(t) \rangle$ is proportional to variance C_0 and may also depend on the noise correlation function $C(t - t') = <n(t)n(t')>$.

In the case of the traditional 1D sampling strategies, Shannon (1949) showed that the mean reconstruction error depends only on the noise variance and does not depend on other noise statistics, including the time dependence of the noise correlation function:

$$\langle \delta s^2(t) \rangle = C_0 \tag{A5}$$

This happens because the overlap integral between the convolution kernels *Eq. A1* centered at different sampling points is zero. This is no longer true for higher dimensional sampling strategies. For the 2D sampling of function values and their first time derivatives, we find that the mean reconstruction error is given by:

$$\langle \delta s^2(t) \rangle^{(2D)} = \frac{2}{3} C_0 - \frac{\Delta t^2}{2\pi^2} \ddot{C}(0) + \frac{2}{\pi^2} \sum_{m=1}^{\infty} \frac{C(m\Delta t)}{m^2} - \frac{2\Delta t}{\pi^2} \sum_{m=1}^{\infty} \frac{\dot{C}(m\Delta t)}{m} \tag{A6}$$

where $\dot{C}$ and $\ddot{C}$ are the first and second time derivatives, respectively of the noise correlation function $C(t)$. For the 3D sampling of function values and their first and second time derivatives, we find that the mean reconstruction error is:

$$\begin{aligned} \langle \delta s^2(t) \rangle^{(3D)} = {} & \frac{143}{160} C(0) + \frac{3\Delta t^2}{16\pi^2} \ddot{C}(0) + \frac{3\Delta t^4}{32\pi^4} C^{(4)}(0) \\ & - \frac{3}{4} \sum_{k=1}^{\infty} \frac{(-1)^k}{\pi^2 k^2} C(k\Delta t)\left(1 + \frac{6}{\pi^2 k^2}\right) \\ & + \frac{3}{2} \sum_{k=1}^{\infty} \frac{(-1)^k}{\pi^3 k^3} \frac{\Delta t}{\pi} \dot{C}(k\Delta t)\left(3 + \frac{\pi^2}{2} k^2\right) - \frac{9}{4} \frac{\Delta t^2}{\pi^2} \sum_{k=1}^{\infty} \ddot{C}(k\Delta t) \\ & + \frac{3\Delta t^3}{\pi^3} \sum_{k=1}^{\infty} \frac{(-1)^k}{4\pi k} C^{(3)}(k\Delta t) \end{aligned} \tag{A7}$$

where $C^{(3)}$ and $C^{(4)}$ are the third and fourth time derivatives, respectively of the noise correlation function.

The above expressions (*Eqs. A2–A7*) were obtained in the general case, without specifying the function $C(t - t')$. In Fig. 11*E*, we analyze the dependence of the mean square error as a function of noise correlation time τ for a Gaussian noise model with $C(t - t') = C_0 \exp[-(t - t')^2/2\tau^2]$. In the frequency domain, this function corresponds to the noise power spectrum $C(f) = C_0\tau \exp(-f^2\tau^2/2)$. Because the mean reconstruction error $\langle \delta s^2(t) \rangle$ is always proportional to variance C_0, we plot results relative to this value. With this normalization, the mean square error for 1D sampling (*Eq. A5*) becomes a straight line at the value of 1, cf. blue line. In comparison, the mean square error for the 2D sampling strategy (*Eq. A6*) diverges at small noise correlation times, but then approaches the unity line. Thus, with increasing noise correlation time, the 2D sampling strategy becomes indistinguishable from the 1D sampling strategy. This behavior is also observed for the 3D sampling strategy (brown line computed according to *Eq. A7*). However, the mean reconstruction error for the 3D sampling strategy is always greater than that of the 2D sampling strategy and thus requires slightly longer noise correlation times to reach the unity line.

APPENDIX B: OPTIMAL WIENER RECONSTRUCTION WITH DISCRETE SAMPLING

The above derivations compared the mean reconstruction error using convolution kernels that yield zero reconstruction error in the absence of noise. However, if the noise correlation structure is known, then it might be beneficial to use somewhat different kernels that might not give zero reconstruction error in the absence of noise, but would minimize the error in the presence of noise with this particular structure. This idea was explored by Wiener (1964) for encoding continuous signals (without taking effects due to sampling). Here, we show that the optimal convolution kernels for different sampling strategies represent a product of Wiener kernels that are designed to optimally suppress noise for continuous signals and the Shannon sampling kernels that interpolate signal values between the sampling points. In other words, the optimal reconstruction of the signal waveform can be thought of as a two-stage process: prefiltering of signals according to Wiener, followed by interpolation according to Shannon equations. Therefore, the relative comparisons of the mean reconstruction error considered above, cf. *Eqs. A5–A7*, remain valid, even if the reconstruction includes the optimal suppression of noise that is matched to its structure.

We begin by reviewing the Wiener arguments for how to best filter the measured signal $c(t) = s(t) + n(t)$ that is a mixture of both the "true" signal $s(t)$ and noise $n(t)$. One searches for such a filter that would minimize the mean square error between the reconstructed and true signal. The derivation is easiest to carry out in the Fourier domain, where the mean square error between the true signal power spectrum $S(f)$ and its reconstruction $\hat{S}(f)$ is given by

$$\int df \left| \hat{S}(f) - S(f) \right|^2 \tag{B1}$$

where the Fourier transform of the reconstructed signal is $\hat{S}(f) = R(f)[S(f) + N(f)]$, where $N(f)$ is the Fourier transform of the noise, and $R(f)$ is the Fourier transform of the optimal filter. The smallest error is achieved when

$$R_W(f) = \frac{\langle |S(f)|^2 \rangle}{\langle |S(f)|^2 \rangle + \langle |N(f)|^2 \rangle} \tag{B2}$$

where $\langle |S(f)|^2 \rangle$ is the power spectrum of the signal, and $\langle |N(f)|^2 \rangle$ is the power spectrum of the noise. One notices that, in the case of auditory signal processing in a natural environment where the average power spectrum of the noise and signal are similar to each other across frequencies, the Wiener filter reduces to an almost identity transformation, and thus is not expected to reduce the reconstruction error. In addition, in what follows, we show that Wiener filtering can be linearly combined, in an identical manner, with reconstructions for different sampling strategies.

In the traditional sampling strategy, we have to find such a convolution filter $u(t)$ that, when applied to the stream of function values (corrupted by noise) taken at intervals of $\Delta t = 1/(2W)$, would yield the reconstructed signal $\hat{s}(t)$ with the smallest mean square distance from the true signal $s(t)$. The reconstructed signal is given by

$$\hat{s}(t) = \sum_{k=-\infty}^{\infty} [s(k\Delta t) + n(k\Delta t)] u(t - k\Delta t) \tag{B3}$$

In the Fourier representation, this convolution can be written as $\hat{S}(f) = U(f)2W[S(f) + N(f)]$, where we took into account that signals have nonzero power only for $|f| < W$. In the absence of noise, the Fourier transforms of the reconstructed signal $\hat{S}(f)$ and the true signal $S(f)$ would match if $U(f) = 1/(2W)$, which is another way to derive the Shannon result we quoted above in *Eq. A5* (Shannon 1949). In the presence of noise, the mean square error is given by

$$\left\langle \int df \left| \hat{S}(f) - S(f) \right|^2 \right\rangle = \left\langle \int df \left| U(f)2W[S(f) + N(f)] - S(f) \right|^2 \right\rangle \tag{B4}$$

which is minimized when the Fourier transform of the convolution kernel is

$$U(f) = \frac{1}{2W} R_W(f) \quad (B5)$$

Here, we took into account the expression for the optimal Wiener kernel in the continuous case (*Eq. B2*). Thus the optimal reconstruction from discrete samples of the function values can be achieved by two sequentially applied filtering operations (in any order): one to optimally filter the continuous signal (either before sampling or after the reconstruction) according to Wiener, and second to interpolate between the sampling points according to Shannon formula, *Eq. A5*.

The same turns out to be true in the case of other, higher dimensional sampling strategies. For example, minimizing the mean square error between the Fourier transform of the "true" signal $S(f)$ and its reconstruction $\hat{S}(f)$ in the 2D sampling strategies, we find that the smallest reconstruction error is obtained when the stream of function values is convolved with the filter

$$V_0(f) = \frac{W - |f|}{W^2} R_W(f) \quad (B6)$$

whereas the stream of first-order derivatives is convolved with

$$V_1(f) = \frac{|f|}{2\pi i W^2 f} R_W(f) \quad (B7)$$

In both *Eq. B6* and *Eq. B7*, $R_W(f)$ is the same Wiener filter as in *Eq. B2* and *Eq. B5*. Thus we again see that the optimal reconstruction of the underlying continuous signal can be done as a combination of two operations: *1*) interpolation between samples with the convolution kernel appropriate for a given sampling strategy (compare *Eqs. B6*, *B7* with *Eqs. A2*, *A3* that do not specifically take advantage of noise structure); and *2*) optimal filtering with the Wiener kernel $R_W(f)$ from *Eq. B2*.

In summary, the optimal sampling and reconstruction of the continuous signal based on discrete sample points can be done in two steps. First, the sampled values need to be convolved with kernels that give zero reconstruction error in the absence of noise. These kernels differ, depending on the sampling strategy, for example, when function values are measured at equally spaced intervals, or when a pair of function values are measured at nearby points in time, yielding measurements of the function values and their time derivatives. The resultant continuous waveform can then be further filtered according to the Wiener optimal filtering procedure. The second step is the same, regardless of the sampling strategy, and will have the same effect on the accuracy of the signal representation. Thus the relative comparisons between different reconstruction strategies remain valid, even if the reconstruction includes optimal noise suppression following Wiener arguments.

ACKNOWLEDGMENTS

We thank William Bialek, Brian Wright, and Michael Stryker for many useful discussions.

GRANTS

This research was supported by National Institutes of Health (NIH) Grants MH055987 and MH077970 (A. J. Doupe), a Howard Hughes Medical Institute predoctoral fellowship (K. I. Nagel); NIH grants MH068904 and EY019493, National Science Foundation (NSF) Grant IIS-0712852, Alfred P. Sloan Fellowship, the Searle Scholar Award, the McKnight Scholar Award, the Ray Thomas Edwards Career Award in Biomedical Sciences, the W. M. Keck Research Excellence Award, and the Center for Theoretical Biological Physics (NSF PHY-0822283) (T. O. Sharpee). Computing resources were provided by the NSF under the following NSF programs: Partnerships for Advanced Computational Infrastructure at the San Diego SuperComputer Center, Distributed Terascale Facility, and Terascale extensions.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the author(s).

REFERENCES

Abramowitz M, Stegun IA. *Handbook of Mathematical Functions*. Washington, DC: US Government Printing Office, 1964.

Adelman TL, Bialek W, Olberg RM. The information content of receptive fields. *Neuron* 40: 823–833, 2003.

Agüera y Arcas B, Fairhall AL. What causes a neuron to spike? *Neural Comput* 15: 1789–1807, 2003.

Agüera y Arcas B, Fairhall AL, Bialek W. Computation in a single neuron: Hodgkin and Huxley revisited. *Neural Comput* 15: 1715–1749, 2003.

Aldworth ZN, Miller JP, Gedeon T, Cummins GI, Dimitrov AG. Dejittered spike-conditioned stimulus waveforms yield improved estimates of neuronal feature selectivity and spike-timing precision of sensory interneurons. *J Neurosci* 25: 5323–5332, 2005.

Atencio CA, Sharpee TO, Schreiner CE. Cooperative nonlinearities in auditory cortical neurons. *Neuron* 58: 956–966, 2008.

Atencio CA, Sharpee TO, Schreiner CE. Hierarchical computation in the canonical auditory cortical circuit. *Proc Natl Acad Sci U S A* 106: 21894–21899, 2009.

Atick JJ, Redlich AN. What does the retina know about natural scenes? *Neural Comput* 4: 196–210, 1992.

Baccus SA, Meister M. Fast and slow contrast adaptation in retinal circuitry. *Neuron* 36: 909–919, 2002.

Bair W. Spike timing in the mammalian visual system. *Curr Opin Neurobiol* 9: 447–453, 1999.

Bandyopadhyay S, Reiss LA, Young ED. Receptive field for dorsal cochlear nucleus neurons at multiple sound levels. *J Neurophysiol* 98: 3505–3515, 2007.

Bar-Yosef O, Nelken I. The effects of background noise on the neural responses to natural sounds in cat primary auditory cortex. *Front Comput Neurosci* 1: 3, 2007.

Bar-Yosef O, Rotman Y, Nelken I. Responses of neurons in cat primary auditory cortex to bird chirps: effects of temporal and spectral context. *J Neurosci* 22: 8619–8632, 2002.

Barlow HB. Possible principles underlying the transformation of sensory messages. In: *Sensory Communication*, edited by Rosenblith WA. Cambridge, MA: MIT, 1961, p. 217–234.

Barlow HB, Fitzhugh R, Kuffler SW. Change of organization in the receptive fields of the cat's retina during dark adaptation. *J Physiol* 137: 338–354, 1957.

Berry MJ, Warland DK, Meister M. The structure and precision of retinal spike trains. *Proc Natl Acad Sci U S A* 94: 5411–5416, 1997.

Bialek W, de Ruyter van Steveninck RR. Features and dimensions: motion estimation in fly vision (Online). http://arxiv.org/abs/q-bio/0505003.

Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D. Reading a neural code. *Science* 252: 1854–1857, 1991.

Borst A, Flanagin VL, Sompolinsky H. Adaptation without parameter change: Dynamic gain control in motion detection. *Proc Natl Acad Sci U S A* 102: 6172–6176, 2005.

Bracewell RN. *The Fourier Transform and Its Applications*. New York: McGraw-Hill, 1986.

Brenner N, Bialek W, de Ruyter van Steveninck R. Adaptive rescaling maximizes information transmission. *Neuron* 26: 695–702, 2000a.

Brenner N, Strong SP, Koberle R, Bialek W, de Ruyter van Steveninck RR. Synergy in a neural code. *Neural Comput* 12: 1531–1552, 2000b.

Bryant HL, Segundo JP. Spike initiation by transmembrane current: a white-noise analysis. *J Physiol* 260: 279–314, 1976.

Chander D, Chichilnisky EJ. Adaptation to temporal contrast in primate and salamander retina. *J Neurosci* 21: 9904–9916, 2001.

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

Christianson GB, Sahani M, Linden JF. The consequences of response nonlinearities for interpretation of spectrotemporal receptive fields. *J Neurosci* 28: 446–455, 2008.

de Boer E, Kuyper P. Triggered correlation. *IEEE Trans Biomed Eng* 15: 169–179, 1968.

de Ruyter van Steveninck RR, Bialek W. Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc R Soc Lond B Biol Sci* 234: 379–414, 1988.

Dean I, Harper NS, McAlpine D. Neural population coding of sound level adapts to stimulus statistics. *Nat Neurosci* 8: 1684–1689, 2005.

Depireux DA, Simon JZ, Klein DJ, Shamma SA. Spectro-temporal response field characterization with dynamic ripples in ferret primary auditory cortex. *J Neurophysiol* 85: 1220–1234, 2001.

Deweese M. Optimization principles for the neural code. *Network* 7: 325–331, 1996.

DeWeese MR, Wehr M, Zador AM. Binary spiking in auditory cortex. *J Neurosci* 23: 7940–7949, 2003.

Dimitrov AG, Gedeon T. Effects of stimulus transformations on estimates of sensory neuron selectivity. *J Comput Neurosci* 20: 265–283, 2006.

Dimitrov AG, Sheiko MA, Baker J, Yen SC. Spatial and temporal jitter distort estimated functional properties of visual sensory neurons. *J Comput Neurosci* 27: 309–319, 2009.

Eggermont JJ. Wiener and Volterra analyses applied to the auditory system. *Hear Res* 66: 177–201, 1993.

Eggermont JJ, Epping WJ, Aertsen AM. Stimulus dependent neural correlations in the auditory midbrain of the grassfrog (Rana temporaria L.). *Biol Cybern* 47: 103–117, 1983.

Enroth-Cugell C, Lennie P. The control of retinal ganglion cell discharge by receptive field surrounds. *J Physiol* 247: 551–578, 1975.

Epping WJ, Eggermont JJ. Sensitivity of neurons in the auditory midbrain of the grassfrog to temporal characteristics of sound. II. Stimulation with amplitude modulated sound. *Hear Res* 24: 55–72, 1986.

Escabi MA, Miller LM, Read HL, Schreiner CE. Naturalistic auditory contrast improves spectrotemporal coding in the cat inferior colliculus. *J Neurosci* 23: 11489–11504, 2003.

Fairhall AL, Burlingame CA, Narasimhan R, Harris RA, Puchalla JL, Berry MJ 2nd. Selectivity for multiple stimulus features in retinal ganglion cells. *J Neurophysiol* 96: 2724–2738, 2006.

Fairhall AL, Lewen GD, Bialek W, de Ruyter Van Steveninck RR. Efficiency and ambiguity in an adaptive neural code. *Nature* 412: 787–792, 2001.

Fortune ES, Margoliash D. Cytoarchitectonic organization and morphology of cells of the field L complex in male zebra finches (Taenopygia guttata). *J Comp Neurol* 325: 388–404, 1992.

Fortune ES, Margoliash D. Parallel pathways and convergence onto HVc and adjacent neostriatum of adult zebra finches (Taeniopygia guttata). *J Comp Neurol* 360: 413–441, 1995.

Frisina RD, Smith RL, Chamberlain SC. Encoding of amplitude modulation in the gerbil cochlear nucleus. I. A hierarchy of enhancement. *Hear Res* 44: 99–122, 1990.

Galazyuk AV, Feng AS. Oscillation may play a role in time domain central auditory processing. *J Neurosci* 21: RC147, 2001.

Garcia-Lazaro JA, Ahmed B, Schnupp JW. Tuning to natural stimulus dynamics in primary auditory cortex. *Curr Biol* 16: 264–271, 2006.

Gill P, Woolley SM, Fremouw T, Theunissen FE. What's that sound? Auditory area CLM encodes stimulus surprise, not intensity or intensity changes. *J Neurophysiol* 99: 2809–2820, 2008.

Gill P, Zhang J, Woolley SM, Fremouw T, Theunissen FE. Sound representation methods for spectro-temporal receptive field estimation. *J Comput Neurosci* 21: 5–20, 2006.

Gollisch T. Estimating receptive fields in the presence of spike-time jitter. *Network* 17: 103–129, 2006.

Heil P, Irvine DR. First-spike timing of auditory-nerve fibers and comparison with auditory cortex. *J Neurophysiol* 78: 2438–2454, 1997.

Heil P, Irvine DR. On determinants of first-spike latency in auditory cortex. *Neuroreport* 7: 3073–3076, 1996.

Heil P, Irvine DR. The posterior field P of cat auditory cortex: coding of envelope transients. *Cereb Cortex* 8: 125–141, 1998.

Heil P, Neubauer H. Temporal integration of sound pressure determines thresholds of auditory-nerve fibers. *J Neurosci* 21: 7404–7415, 2001.

Heil P, Neubauer H. A unifying basis of auditory thresholds based on temporal summation. *Proc Natl Acad Sci U S A* 100: 6151–6156, 2003.

Heil P, Neubauer H, Brown M, Irvine DR. Towards a unifying basis of auditory thresholds: distributions of the first-spike latencies of auditory-nerve fibers. *Hear Res* 238: 25–38, 2008.

Hong S, Aguera y Arcas B, Fairhall AL. Single neuron computation: from dynamical system to feature detector. *Neural Comput* 19: 3133–3172, 2007.

Hong S, Lundstrom BN, Fairhall AL. Intrinsic gain modulation and adaptive neural coding. *PLoS Comput Biol* 4: e1000119, 2008.

Hsu A, Woolley SM, Fremouw TE, Theunissen FE. Modulation power and phase spectrum of natural sounds enhance neural encoding performed by single auditory neurons. *J Neurosci* 24: 9201–9211, 2004.

Jerri AJ. The Shannon sampling theorem–its various extensions and applications: a tutorial review. *Proc IEEE Inst Electr Electron Eng* 65: 1565–1596, 1977.

Kim AJ, Lazar AA, Slutskiy YB. System identification of Drosophila olfactory sensory neurons. *J Comput Neurosci* 30: 143–161, 2011.

Kim PJ, Young ED. Comparative analysis of spectro-temporal receptive fields, reverse correlation functions, and frequency tuning curves of auditory-nerve fibers. *J Acoust Soc Am* 95: 410–422, 1994.

Klein DJ, Depireux DA, Simon JZ, Shamma SA. Robust spectrotemporal reverse correlation for the auditory system: optimizing stimulus design. *J Comput Neurosci* 9: 85–111, 2000.

Klein DJ, Simon JZ, Depireux DA, Shamma SA. Stimulus-invariant processing and spectrotemporal reverse correlation in primary auditory cortex. *J Comput Neurosci* 20: 111–136, 2006.

Kouh M, Sharpee TO. Estimating linear-nonlinear models using Renyi divergences. *Network* 20: 49–68, 2009.

Krishna BS, Semple MN. Auditory temporal processing: responses to sinusoidally amplitude-modulated tones in the inferior colliculus. *J Neurophysiol* 84: 255–273, 2000.

Kuffler SW, Fitzhugh R, Barlow HB. Maintained activity in the cat's retina in light and darkness. *J Gen Physiol* 40: 683–702, 1957.

Kvale MN, Schreiner CE. Short-term adaptation of auditory receptive fields to dynamic stimuli. *J Neurophysiol* 91: 604–612, 2004.

Lesica NA, Grothe B. Dynamic spectrotemporal feature selectivity in the auditory midbrain. *J Neurosci* 28: 5412–5421, 2008a.

Lesica NA, Grothe B. Efficient temporal processing of naturalistic sounds. *PLos One* 3: e1655, 2008b.

Lewicki MS. Efficient coding of natural sounds. *Nat Neurosci* 5: 356–363, 2002.

Machens CK, Wehr MS, Zador AM. Linearity of cortical receptive fields measured with natural sounds. *J Neurosci* 24: 1089–1100, 2004.

Mainen ZF, Sejnowski TJ. Reliability of spike timing in neocortical neurons. *Science* 268: 1503–1506, 1995.

Malone BJ, Scott BH, Semple MN. Dynamic amplitude coding in the auditory cortex of awake rhesus macaques. *J Neurophysiol* 98: 1451–1474, 2007.

Maravall M, Petersen RS, Fairhall AL, Arabzadeh E, Diamond ME. Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS Biol* 5: e19, 2007.

Meister M, Berry MJ. The neural code of the retina. *Neuron* 22: 435–450, 1999.

Nagel K, Doupe AJ. Temporal processing and adaptation in the songbird auditory forebrain. *Neuron* 51: 845–859, 2006.

Nagel KI, Doupe AJ. Organizing principles of spectro-temporal encoding in the avian primary auditory area field L. *Neuron* 58: 938–955, 2008.

Nelken I, Chechik G. Information theory in auditory research. *Hear Res* 229: 94–105, 2007.

Nelken I, Kim PJ, Young ED. Linear and nonlinear spectral integration in type IV neurons of the dorsal cochlear nucleus. II. Predicting responses with the use of nonlinear models. *J Neurophysiol* 78: 800–811, 1997.

Nelken I, Rotman Y, Bar Yosef O. Responses of auditory-cortex neurons to structural features of natural sounds. *Nature* 397: 154–157, 1999.

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

Petersen CC. The functional organization of the barrel cortex. *Neuron* 56: 339–355, 2007.

Phillips CG, Zeki S, Barlow HB. Localization of function in the cerebral cortex. Past, present and future. *Brain* 107: 327–361, 1984.

Pillow JW, Simoncelli EP. Dimensionality reduction in neural models: an information-theoretic generalization of spike-triggered average and covariance analysis. *J Vis* 6: 414–428, 2006.

Prenger R, Wu MC, David SV, Gallant JL. Nonlinear V1 responses to natural scenes revealed by neural network analysis. *Neural Netw* 17: 663–679, 2004.

Rees A, Moller AR. Stimulus properties influencing the responses of inferior colliculus neurons to amplitude-modulated sounds. *Hear Res* 27: 129–143, 1987.

Reiss LA, Bandyopadhyay S, Young ED. Effects of stimulus spectral contrast on receptive fields of dorsal cochlear nucleus neurons. *J Neurophysiol* 98: 2133–2143, 2007.

Rieke F, Warland D, de Ruyter van Steveninck R, Bialek WB. *Spikes: Exploring the Neural Code*. Cambridge, MA: MIT, 1997.

Rowekamp RJ, Sharpee TO. Analyzing multicomponent receptive fields from neural responses to natural stimuli. *Network*. doi:10.3109/0954898x.2011.566303.

Ruderman DL, Bialek W. Statistics of natural images: scaling in the woods. *Phys Rev Lett* 73: 814–817, 1994.

Rust NC, Schwartz O, Movshon JA, Simoncelli EP. Spatiotemporal elements of macaque v1 receptive fields. *Neuron* 46: 945–956, 2005.

Schwartz O, Pillow JW, Rust NC, Simoncelli EP. Spike-triggered neural characterization. *J Vis* 6: 484–507, 2006.

Scott BH, Malone BJ, Semple MN. Transformation of temporal processing across auditory cortex of awake macaques. *J Neurophysiol* 105: 712–730, 2011.

Shannon CE. Communication in the presence of noise. *Proceedings of the IRE* 37: 10–21, 1949.

Shannon RV, Zeng FG, Kamath V, Wygonski J, Ekelid M. Speech recognition with primarily temporal cues. *Science* 270: 303–304, 1995.

Shapley RM, Enroth-Cugell C. Visual adaptation and retinal gain controls. *Prog Retin Eye Res* 3: 263–346, 1984.

Shapley RM, Victor JD. The effect of contrast on the transfer properties of cat retinal ganglion cells. *J Physiol* 285: 275–298, 1978.

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

Sharpee TO, Sugihara H, Kurgansky AV, Rebrik SP, Stryker MP, Miller KD. Adaptive filtering enhances information transmission in visual cortex. *Nature* 439: 936–942, 2006.

Singh NC, Theunissen FE. Modulation spectra of natural sounds and ethological theories of auditory processing. *J Acoust Soc Am* 114: 3394–3411, 2003.

Slee SJ, Higgs MH, Fairhall AL, Spain WJ. Two-dimensional time coding in the auditory brainstem. *J Neurosci* 25: 9978–9988, 2005.

Smirnakis SM, Berry MJ, Warland DK, Bialek W, Meister M. Adaptation of retinal processing to image contrast and spatial scale. *Nature* 386: 69–73, 1997.

Srinivasan MV, Laughlin SB, Dubs A. Predictive coding: a fresh view of inhibition in the retina. *Proc R Soc Lond B Biol Sci* B216: 427–459, 1982.

Strong SP, Koberle R, de Ruyter van Steveninck RR, Bialek W. Entropy and information in neural spike trains. *Phys Rev Lett* 80: 197–200, 1998.

Sullivan WE 3rd. Possible neural mechanisms of target distance coding in auditory system of the echolocating bat Myotis lucifugus. *J Neurophysiol* 48: 1033–1047, 1982.

Teich MC, Heneghan C, Lowen SB, Ozaki T, Kaplan E. Fractal character of the neural spike train in the visual system of the cat. *J Opt Soc Am A Opt Image Sci Vis* 14: 529–546, 1997.

Theunissen FE, David SV, Singh NC, Hsu A, Vinje WE, Gallant JL. Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network* 12: 289–316, 2001.

Theunissen FE, Doupe AJ. Temporal and spectral sensitivity of complex auditory neurons in the nucleus HVc of male zebra finches. *J Neurosci* 18: 3786–3802, 1998.

Theunissen FE, Sen K, Doupe AJ. Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. *J Neurosci* 20: 2315–2331, 2000.

Tiesinga P, Fellous JM, Sejnowski TJ. Regulation of spike timing in visual cortical circuits. *Nat Rev Neurosci* 9: 97–107, 2008.

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

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

Treves A, Panzeri S. The upward bias in measures of information derived from limited data samples. *Neural Comput* 7: 399–407, 1995.

Victor JD. The dynamics of the cat retinal x-cell centre. *J Physiol* 386: 219–246, 1987.

Victor JD, Mechler F, Repucci MA, Purpura KP, Sharpee T. Responses of V1 neurons to two-dimensional hermite functions. *J Neurophysiol* 95: 379–400, 2006.

Voss RF, Clarke J. "1/f noise" in music and speech. *Nature* 258: 317–318, 1975.

Wiener N. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications*. Cambridge, MA: MIT, 1964.

Wild JM, Karten HJ, Frost BJ. Connections of the auditory forebrain in the pigeon (Columba livia). *J Comp Neurol* 337: 32–62, 1993.

Woolley SM, Gill PR, Theunissen FE. Stimulus-dependent auditory tuning results in synchronous population coding of vocalizations in the songbird midbrain. *J Neurosci* 26: 2499–2512, 2006.

Zhou Y, Wang X. Cortical processing of dynamic sound envelope transitions. *J Neurosci* 30: 16741–16754, 2010.