

The Contribution of Spike Threshold to Acoustic Feature Selectivity, Spike Information Content, and Information Throughput

Monty A. Escabí,^{1,2} Reza Nassiri,² Lee M. Miller,⁴ Christoph E. Schreiner,⁵ and Heather L. Read³

¹Electrical and Computer Engineering, ²Biomedical Engineering, and ³Psychology, University of Connecticut, Storrs, Connecticut 06269-1157, ⁴Center for Mind and Brain, University of California, Davis, California 95616, and ⁵W. M. Keck Center for Integrative Neuroscience, University of California, San Francisco, California 94143

Hypotheses of sensory coding range from the notion of nonlinear “feature detectors” to linear rate coding strategies. Here, we report that auditory neurons exhibit a novel trade-off in the relationship between sound selectivity and the information that can be communicated to a postsynaptic cell. Recordings from the cat inferior colliculus show that neurons with the lowest spike rates reliably signal the occurrence of stereotyped stimulus features, whereas those with high response rates exhibit lower selectivity. The highest information conveyed by individual action potentials comes from neurons with low spike rate and high selectivity. Surprisingly, spike information is inversely related to spike rates, following a trend similar to that of feature selectivity. Information per time interval, however, was proportional to measured spike rates. A neuronal model based on the spike threshold of the synaptic drive accurately accounts for this trade-off: higher thresholds enhance the spiking fidelity at the expense of limiting the total communicated information. Such a constraint on the specificity and throughput creates a continuum in the neural code with two extreme forms of information transfer that likely serve complementary roles in the representation of the auditory environment.

Key words: inferior colliculus; spike threshold; feature selectivity; mutual information; linear rate code; spectrotemporal; neural code

Introduction

Two models have been proposed to describe how information from neuronal spike patterns serves as a neural code for unique features of sound. “Feature selective” models assume that sequences of spikes from a single neuron are precisely timed to contingent constellations of stimulus parameters (features) from which the identity of a sound could be precisely determined (Barlow, 1953, 1961; Martin, 1994). Unitary action potentials from an ideal feature detector neuron would signal the occurrence of a highly stereotyped stimulus feature that is usually assumed to be of special behavioral relevance (Simmons et al., 1975; Brenowitz et al., 1997). Feature selectivity has been demonstrated for species-specific vocalizations in auditory cortex of echolocating bats (Suga and Jen, 1976; Suga et al., 1978), area X/Hv of song birds (Brenowitz et al., 1997; Doupe, 1997) and for face representation in high-level primate visual cortices (Fujita et al., 1992; Gross and Sergent, 1992). An orthogonal view of sensory feature encoding is offered by “linear rate coding” models. Linear rate

models assume that sensory features are represented by the time-varying firing rate of single neurons or neuronal populations. In principle, the response of a linear neuron is modulated by the overlap between the stimulus and its neuronal receptive field (RF). Thus, in contrast to a classical feature detector, a linear neuron could respond to range of stimulus parameters that are not precisely contingent on each other. Linear selectivities for elementary acoustic features such as spectral modulation frequency and direction have been demonstrated in peripheral auditory stations (Nelken et al., 1997; Yu and Young, 2000), the midbrain inferior colliculus (ICC) and auditory cortex (Kowalski et al., 1996; Calhoun and Schreiner, 1998; Versnel and Shamma, 1998; Schnupp et al., 2001; Escabí and Schreiner, 2002). This simpler class of acoustic selectivity is analogous to grating and directional selectivity in visual neurons (Reid et al., 1991; DeAngelis et al., 1993).

Feature-selective and linear rate-encoded responses represent two extreme response regimens that can be distinguished and quantified if appropriate stimuli are used. Distinguishing these encoding regimens requires that we differentiate the “stimulus specificity” of a neuron from its “stimulus preference.” The stimulus specificity is related to the fidelity in the association between sensory signal and evoked spikes. High specificity is achieved when a neuron responds reliably to a given stimulus or feature but is unresponsive to other stimuli. By comparison, the stimulus preference is an averaged property of the stimulus–response function. The average sound preferences of an auditory neuron,

Received May 5, 2005; revised Aug. 28, 2005; accepted Aug. 30, 2005.

This work was supported by grants from the National Institutes of Health (DC006397, DC02260), the National Science Foundation (NS0139307), and the University of Connecticut Research Foundation. We thank J.J. Chrobak for comments and discussions on this manuscript. M.A.E. directed the project, conceived the analysis and model, and wrote this manuscript. R.N. and H.L.R. contributed to the model implementation and design. H.L.R., L.M.M., and C.E.S. helped perform experiments in the cat ICC and contributed to their interpretation.

Correspondence should be addressed to Monty A. Escabí at the above address. E-mail: escabi@engr.uconn.edu.

DOI:10.1523/JNEUROSCI.1804-05.2005

Copyright © 2005 Society for Neuroscience 0270-6474/05/259524-11\$15.00/0

or its spectrotemporal receptive field (STRF), can be estimated by spike-triggered averaging all of the sound waveforms preceding a spike output (Hermes et al., 1981; deCharms et al., 1998; Keller and Takahashi, 2000; Klein et al., 2000; Theunissen et al., 2000; Escabí and Schreiner, 2002).

Calculating the information content in the sensory-related spike train is yet another way to quantify how spiking patterns contribute to the neuronal representation. Measures of spike train information content often make no assumptions about the features being encoded (de Ruyter van Steveninck et al., 1997; Borst and Theunissen, 1999) and do not assume a feature-selective or rate code model.

In the current study, we demonstrate a novel view of auditory sensory encoding by combining information theoretic approaches with measures of feature selectivity and stimulus preference. Single-cell recordings from the cat inferior colliculus demonstrate specific trade-offs between feature selectivity, spike rates, and spike information content. These trade-offs provide a continuum between the fidelity and throughput of the sensory coding and offer complementary modes of information transfer for purely linear and feature-selective strategies. A neuronal model based on sound-evoked synaptically driven input and a spike threshold accurately accounts for these trade-offs and suggests that the spike threshold is the key mechanisms controlling the specificity of the encoding regimen.

Materials and Methods

Electrophysiology. Details of the physiological recording and stimulus presentation methods for this dataset have been presented previously (Escabí and Schreiner, 2002; Escabí et al., 2003). Briefly, single-unit recordings of $n = 103$ envelope “phase-locking” neurons (Escabí and Schreiner, 2002) were obtained from the central nucleus of the inferior colliculus of five anesthetized cats. Animals were maintained in an areflexive state via continuous infusion of ketamine ($2\text{--}4 \text{ mg} \cdot \text{kg}^{-1} \cdot \text{h}^{-1}$) and diazepam ($0.4\text{--}1 \text{ mg} \cdot \text{kg}^{-1} \cdot \text{h}^{-1}$) in lactated Ringer’s solution ($1\text{--}4 \text{ mg} \cdot \text{kg}^{-1} \cdot \text{h}^{-1}$). Body temperature was maintained at $\sim 37.5^\circ\text{C}$, and heart rate, breathing rate, temperature, and peripheral reflexes were monitored throughout the experiments. Surgical methods and experimental procedures follow National Institutes of Health and United States Department of Agriculture guidelines.

Dynamic moving ripple (DMR) and/or rippled noise (RN) stimulus sequences (Escabí and Schreiner, 2002) were presented to the animal in a sound-shielded chamber (IAC, Bronx, NY). Both sounds were designed to uniformly cover spectral resolutions from 0 to 4 cycles/octaves and temporal modulation rates from 0 to 350 Hz. Independent sound sequences were presented simultaneously to each ear via a closed, binaural speaker system [electrostatic diaphragms from Stax (Saitama, Japan)]. All stimuli were presented at 30–70 dB per one-third octave for a duration of 8–20 min.

Measuring acoustic feature selectivity. We devised a second-order stimulus–response analysis to characterize the spectrotemporal selectivity of auditory neurons by analyzing the variance of the pre-event stimulus ensemble. First, we computed the STRF, which is the average stimulus pattern that evoked action potentials. The STRF of a neuron is obtained by spike-triggered averaging the stimulus waveforms:

$$\text{STRF}(\tau, X_k) = \frac{1}{N} \sum_n S(t_n - \tau, X_k), \quad (1)$$

where N is the number of action potentials, S is the stimulus spectrotemporal waveform normalized for unit variance, and the average is conditioned on the ensemble of action potentials times, $\{t_n\}$. τ and X_k correspond to the stimulus delay (relative to the spike time, t_n) and the sound frequency in octaves (for the k th frequency channel), respectively.

Next, we measured the variance of the statistically significant spike evoking stimulus patterns. To do this, the STRF was examined for a

significance level of $p < 0.002$. The subset of STRF samples that exceeded the significance criterion (outlined as white contours) (see Fig. 2), along with the overlapping pre-event sound patterns, were treated as vectors. The response-conditioned stimulus variance is given by

$$\text{var}[S\{\{t_n\}\}] = \frac{1}{L \times N} \sum_n \|S_n - \text{STRF}\|^2,$$

where $\|\cdot\|$ designates the vector norm, STRF is a vector containing L statistically significant STRF samples, and S_n is shorthand for the vector containing the pre-event waveform samples of $S(t_n - \tau, X_k)$ that overlap the significant STRF region (Fig. 2, white contours). The magnitude difference between the significant STRF and the n th pre-event spectrotemporal pattern can be expressed as follows:

$$\|S_n - \text{STRF}\|^2 = \|S_n\|^2 - 2\|S_n\| \|\text{STRF}\| \text{SI}_n + \|\text{STRF}\|^2, \quad (2)$$

where

$$\text{SI}_n = \frac{\langle \text{STRF}, S_n \rangle}{\|\text{STRF}\| \|S_n\|} \quad (3)$$

is the response-conditioned stimulus–STRF similarity index (SI) for the n th action potential, and $\langle \cdot, \cdot \rangle$ corresponds to the vector dot product. The SI is equivalent to a correlation coefficient between the pre-event stimulus pattern and the STRF and assumes values between -1 and 1 ; a value of 1 indicates that the STRF and the n th pre-event spectrotemporal pattern have identical shapes; a value of 0 indicates that they bear no resemblance; a value of -1 is unlikely to occur, because it implies that the sound pattern is 180° out of phase with the STRF.

From Equation 2, two independent factors contribute to the response-conditioned stimulus variance. First, the SI determines the similarity in shape between the pre-event stimulus and the STRF. The variance is thus minimized if the SI for all stimulus patterns is near 1 . Second, the magnitude of the variance is proportional to the difference in magnitudes between the STRF and pre-event stimulus. Thus, ideally, the variance is precisely 0 if and only if the pre-event pattern of all action potentials are identical in shape to the STRF (SI is unity for all n) and their norms are identical: $\|S_n\| = \|\text{STRF}\|$ for all n . Here, we take advantage of the fact that our moving ripple stimulus has a fixed peak-to-peak amplitude (Escabí and Schreiner, 2002) and thus a constant norm. This allows us to consider how the stimulus shape (SI) contributes to the neural selectivity without the influence of stimulus norm or contrast, which we previously showed to contribute significantly to neuronal responses in the ICC (Escabí et al., 2003).

We quantified the average specificity (or selectivity) by characterizing the ensemble of SIs. Collectively, we represent these by the SI probability distribution function, $p(\text{SI})$. Our feature selectivity metric consists of measuring the bias in $p(\text{SI})$ in relationship to a hypothetical feature detector, $p_{\text{fd}}(\text{SI})$, or a randomly firing neuron, $p_r(\text{SI})$. First, we removed the bias resulting from idiosyncrasies in irregular firing by computing the SI distribution for a random neuron. This was done by generating a spike train with 12,000 Poisson distributed action potentials over the DMR stimulus duration and then using the original STRF to measure the SI distribution for the random spikes, $p_r(\text{SI})$. This distribution was typically tightly centered about an SI of 0 , although some slight skewing was observed that was dependent on STRF shape. Next, we considered the expected distribution of a hypothetical feature detector neuron, $p_{\text{fd}}(\text{SI})$. Because a hypothetical feature detector responds if and only if the shape of the pre-event stimulus precisely matches the STRF, its SI distribution consists of a single peak centered about an SI of 1 .

We next measured the bias of our real neuron relative to our random and feature detector control conditions. To do this, we first integrated the probability distribution function of our real, random, and feature detector neurons as follows:

$$P(\text{SI}) = \int_{-1}^{\text{SI}} p(\zeta) d\zeta, \quad (4)$$

and we therefore represented the behavior of each by a cumulative distribution function (CDF) (see Fig. 3, right column). The cumulative distribution of random neurons, $P_r(SI)$, contains a sharp transition at $SI = 0$. The feature detector CDF, $P_{fd}(SI)$, has a sharp transition at $SI = 1$. The CDF of real neurons $[P(SI)]$ typically had a smooth transition somewhere in between. The bias is measured as a feature selectivity index (FSI) as follows:

$$FSI = \frac{A_r - A}{A_r - A_{fd}}, \quad (5)$$

where the measurements

$$A = \int_{-1}^1 P(SI) dSI \quad (6)$$

correspond to the area underneath each cumulative distribution function for the real (A), feature detector (A_{fd}), and random (A_r) neurons. The FSI takes continuous values from 0 to 1. FSIs of 1 indicate that the neuron behaves like a hypothetical feature detector, whereas values near 0 indicate that the response specificity was low (i.e., random firing neuron).

In the case in which neurons responded binaurally, and therefore produced significant STRFs for the contralateral and ipsilateral inputs (Qiu et al., 2003), the FSI procedure was modified without loss of generality to account for the binaural activation of a neuron. To do this, SIs were measured for each ear, resulting in a contralateral $SI_{c,n}$ and an ipsilateral $SI_{i,n}$ measurement for each action potential, t_n . Because the neuronal responses are now represented by a vector of two simultaneous-SI measurements, the ensemble SI statistics can be represented as joint-binaural SI distribution, $p(SI_c, SI_i)$. The FSI could then be represented as $FSI = (V_r - V)/(V_r - V_{fd})$, where

$$V = \int_{-1}^1 \int_{-1}^1 P(SI_c, SI_i) dSI_c dSI_i$$

represents the volume measurement underneath the joint-binaural cumulative SI distribution function, $P(SI_c, SI_i)$. The validity of this binaural-FSI procedure was tested for monaural neurons, all of which produced identical measurements as the direct approach using monaural SI distribution (mean correlation coefficient, >0.99 ; $p < 0.001$).

Information calculation. Mutual information was estimated using the methods of Strong et al. (1998) as applied to 5 s ripple noise response rastergrams (125 trials) in Escabí et al. (2003). Briefly, estimates of the mutual information were obtained for $n = 42$ neurons and for simulations from our model neuron (see below, Spectrotemporal integrate-and-fire neuron model). One neuron was discarded from our sample, because it deviated substantially from the general trends and it was characterized as a significant outlier ($p < 0.006$, bootstrap estimate). The mutual information was estimated as follows:

$$I = S_{\text{Total}} - S_{\text{Noise}}, \quad (7a)$$

where

$$S_{\text{Total}} = - \sum p(w) \log_2(w) \quad (7b)$$

is the total entropy of the spike train, and

$$S_{\text{Noise}} = - \left\langle \sum p(w|t) \log_2(p(w|t)) \right\rangle \quad (7c)$$

is the entropy associated with the response variability. The distribution of response words, $p(w)$ and $p(w|t)$, were estimated from the response spike rasters (see Fig. 4A–E). The procedure was bootstrapped across multiple stimulus segments for various word lengths ($T = 5, 6, 8, 10, 15, 20, 40, 80, 160$, and 200 ms) and data fractions (100, 80, 50, 33, and 25%) and at

various temporal resolutions (1, 2, 4, 8, 16 ms). The biases associated with word length and data fraction were removed according to the procedure of Strong et al. (1998).

Spectrotemporal integrate-and-fire neuron model. We devised a simplified model neuron to examine the relationship between feature selectivity, information content, and the spiking properties of a neuron. This coupled two-compartment model consists of a linear synaptic STRF that accounts for the spectrotemporal integration of a neuron and a subsequent integrate-and-fire neuron that accounts for the cell membrane integration, thresholding, and action potential generation.

The integrate-and-fire compartment consists of a conventional leaky integrate-and-fire neuron. The membrane potential is determined from the following linear cell membrane differential equation:

$$\frac{dV(t)}{dt} + \frac{1}{\tau} V(t) = \frac{i_m(t) + n(t)}{C}, \quad (8)$$

where $V(t) = V_m(t) - E_r$ is the driven intracellular potential of the neuron [i.e., the difference between the membrane potential, $V_m(t)$, and its resting potential, E_r (−60 mV)], $i_m(t)$ is the synaptic current input, $n(t)$ is a white noise current term that is low-pass filtered by the cell membrane, and τ is the membrane time constant. Spikes are generated in the model whenever the membrane potential, V_m , exceeds a specified threshold activation potential, V_T . After activation, a 1 ms refractory period was imposed and the membrane was reset to its resting potential, at which point the integration continued.

The synaptic current activity of the neuron was computed from the synaptic STRF of the neuron, $STRF_{\text{synaptic}}$. The synaptic STRF models the synaptically driven spectrotemporal activity of the neuron without the influence of the cell membrane integration. The intracellular membrane current provided by each synaptic frequency channel is derived as a spectrotemporal convolution between the sound envelope (at a given frequency) and the synaptic STRF of the neuron as follows:

$$i_{m,k}(t) = \int STRF_{\text{synaptic}}(\sigma, X_k) S(t - \sigma, X_k) d\sigma. \quad (9a)$$

Here, $i_{m,k}$ is the membrane current for the k th frequency channel, and $S(t, X_k)$ is the sound spectrotemporal envelope. The total current that is injected into the neuron is obtained as the sum of the current contributions from each frequency channel as follows:

$$i_m(t) = G \sum_{k=1}^L i_{m,k}(t), \quad (9b)$$

where i_m is the total injected current, and G is a gain constant which adjusts the current amplitude to produce an intracellular voltage of desired variance.

To apply the spectrotemporal integrate-and-fire (STIF) model to STRFs from our neuronal data, we devised a procedure for removing the influence of the cell membrane integration from our ICC STRFs. This procedure allows us to closely approximate the synaptic STRF from our spike-derived STRFs, $STRF_{\text{spike}}$. First, although in all of our real data we did not have access to the membrane STRF directly (because we derived STRFs using extracellular spiking activity), extensive testing on our model revealed that in general $STRF_{\text{spike}} \approx \epsilon \times STRF_{\text{membrane}}$, where ϵ is an amplitude scaling constant that was strongly dependent on the normalized threshold and the internal signal-to-noise ratio (SNR). This empirical finding is consistent with Price's theorem as applied to spike-triggered averaging analysis (de Boer and Kuyper, 1968). We therefore approximated the membrane STRF by its spike-derived counterpart ($STRF_{\text{membrane}} \approx STRF_{\text{spike}}$) and then adjusted the arbitrary gain factor G in Equation 9b to produce a desired spike rate or intracellular variance.

Next, the synaptic STRF is related to the membrane STRF of the neuron, $STRF_{\text{membrane}}$, via a linear deconvolution as follows:

$$STRF_{\text{synaptic}}(\tau, X_k) = \int STRF_{\text{membrane}}(\gamma, X_k) h_{\text{membrane}}^{-1}(\tau - \gamma) d\gamma, \quad (10a)$$

which removes the membrane impulse response from the STRF of the neuron. Here, $STRF_{\text{synaptic}}$ is obtained by inverse filtering the membrane STRF with the membrane impulse response of the cell,

$h_{\text{membrane}}(\tau) = C^{-1}e^{-\tau/\tau_m}u(\tau)$, where τ is the membrane integration time constant and C is the membrane capacitance. We implemented this deconvolution in the Laplace transform domain as follows:

$$STRF_{\text{synaptic}}(s, X_k) = \frac{STRF_{\text{membrane}}(s, X_k)}{H_{\text{membrane}}(s)}, \quad (10b)$$

where s is the Laplace variable, and the membrane transfer function is $H_{\text{membrane}}(s) = C^{-1}(s + 1/\tau)^{-1}$.

To examine how the threshold and noise level of the neuron affects its spiking activity, we reduced the intracellular parameters of the neuron to only those that affected the spiking properties. Examination of the model revealed that the analysis could be simplified substantially by noticing that the model output only depends on the “relative” intrinsic drive in relation to the resting and threshold potentials. We therefore defined a dimensionless “normalized threshold” as follows:

$$N_T = \frac{V_T - E_r}{\sigma_m}, \quad (11)$$

which corresponds to the number of intracellular voltage SDs required for spike activation. The intracellular signal-to-noise ratio was defined as follows:

$$SNR = 20 \log_{10}(\sigma_m/\sigma_n) \quad (12)$$

where $\sigma_m^2 \times \sigma_n^2$ correspond to the signal and noise intracellular voltage variance, respectively (assuming no spiking).

Population model. To see whether the spiking STRF model could explain our population trends, we simulated a population of neurons by randomly choosing integrate-and-fire (IF) neuron parameters. We then optimized the model results by fitting the model to Figure 4F and tested whether it could account for our remaining population trends (see Fig. 6B–D).

To do this, we first selected a subset of six real neuron STRFs for the STIF model that closely approximated the receptive field preferences in the inferior colliculus and that are representative of our population distribution in a previous study (Qiu et al., 2003). Each receptive field was fitted with a Gabor model as described in Qiu et al. (2003). We performed this procedure to assure that receptive field size or shape in all of our simulations did not bias our results.

We next simulated a population of randomly sampled neurons by sampling STIF neuron parameters and STRFs at 240 independent conditions. At each IF neuron condition, one of the six STRF models was selected randomly, and the predicted intracellular current was then injected into the IF compartment. All IF parameters were then selected randomly from a uniform distribution (initial range: SNR = −20–10 dB; N_T = 0.5–6; τ = 5–15 ms). Optimization of the model was performed by pruning iteratively the upper and lower limit of the SNR and N_T distributions so that the linear regression slope, intercept, and SD of the model spike information versus spike rate trend (see Fig. 6A) matched the ICC population data (see Fig. 4F) as closely as possible. Pruning of the SNR and N_T upper and lower limits was performed in 5 dB and 0.5 steps, respectively. This model fitting procedure selected $n = 98$ (of 240) neurons with uniformly distributed IF parameters (optimized range: SNR = −15–0 dB; threshold, N_T = 1–4) that best matched the spike information versus spike rate data of Figure 4F. The resulting model parameters were then used to cross-validate the remaining trends (see Fig. 6B–D).




Figure 1. Hypothetical models of receptive field construction and stimulus selectivity as measured by averaging stimulus waveforms preceding action potentials. This procedure is illustrated for two idealized scenarios. In **A**, the stimulus waveforms preceding each action potential are highly stereotyped from spike to spike as expected for a neuron with high stimulus–response specificity. **B**, The same receptive field could instead be obtained from a large ensemble of stimulus–response patterns that vary from spike to spike and, therefore, contribute differently to the formation of a receptive field as for a linear integrator. The composition of the waveforms preceding action potentials is highly variable for **B**, whereas for **A**, the stimulus composition shows no variability and the neuron responds exclusively to a fixed combination of stimulus features.

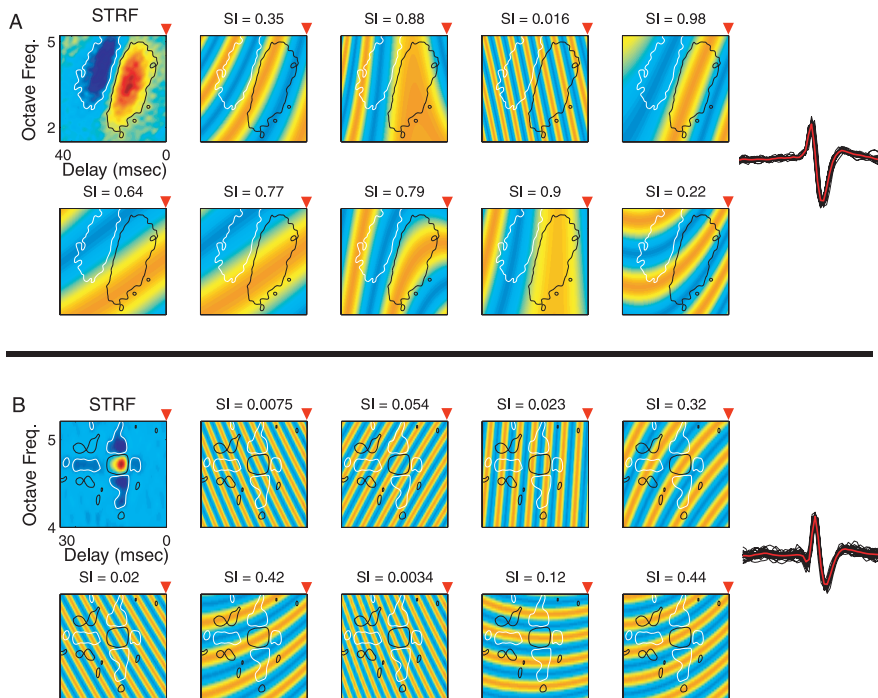


Figure 2. Representative pre-event spectrotemporal waveforms of two inferior colliculus neurons. STRFs are shown for both neurons at the far left. Significant excitatory and inhibitory receptive field domains ($p < 0.002$) are delimited by black and white contours, respectively. SI (i.e., correlation coefficients) values between the significant STRF and the overlapping pre-event waveform for each action potential are shown above each panel. **A** shows a neuron with low stimulus–response variability. Pre-event stimulus waveforms that initiate action potentials closely resemble the STRF (high SI values; near 1). Typically, excitatory and inhibitory STRF subregions overlapped ON and OFF stimulus features, respectively. For the neuron in **B**, the stimulus waveforms preceding action potentials are highly variable in their spectrotemporal composition and did not individually resemble the STRF (low SI values; near 0). Typical action potential traces are shown for reference (far right). Freq., Frequency.

Results

Acoustic feature selectivity

We examined the relationship between spiking activity and structurally rich stimulus patterns to quantify the specificity of sound encoding in the central nucleus of the inferior colliculus. We presented continuous dynamic noise sequences that probe many of the relevant spectrotemporal stimuli for the ICC. DMR and RN stimuli are broadband sounds with dynamic spectral and temporal modulations that have been shown to efficiently activate ICC neurons (Escabí and Schreiner, 2002). We next estimated the average preference of each neuron, or its STRF (Fig. 1), by spike-triggered averaging the stimulus patterns preceding each action potential. The STRF activity pattern can be interpreted in terms of intrinsic mechanisms leading to the stimulus–response activation such as excitation, inhibition, and/or suppression.

The obtained STRFs provide an initial glance at the “average” stimulus features that evoke neural responses but provide no information on how individual action potentials contribute to the average. To determine how the stimulus–response ensemble contributes to the average receptive field, we analyzed the spike-evoking stimulus patterns by considering whether these are highly conserved from one action potential to another. The possible coding schemes responsible for a particular ICC STRF pattern could extend from highly precise stimulus–response relationship to a highly variable stimulus–spike pattern. Yet, the response activity of two separate neurons could potentially lead to similar STRF patterns despite substantial differences in the specificity of the spiking activity and the stimulus patterns that initiate action potentials (Fig. 1). The relationship between the stimulus–response variance and the feature specificity of a neuron is illustrated in Figure 1 for a hypothetical feature detector and a linear integrating neuron. For a hypothetical feature detector, the pre-event stimulus patterns are highly consistent in their spectrotemporal composition (zero variance, high parameter contingency) and closely resemble the STRF of the neuron (Fig. 1A). Conversely, the STRF construction procedure could correspond to the average of a large ensemble of pre-event stimulus patterns that are highly variable (low parameter contingency) but, at least, partially overlap the average STRF. For this hypothetical linear neuron, each stimulus pattern bears little resemblance to the STRF, but together they average to the STRF pattern (Fig. 1B), as expected for a linear/energy integrating neuron (Theunissen et al., 2000; Escabí and Schreiner, 2002). In contrast, a spontaneously active random neuron would produce a zero-valued STRF in which the spectrotemporal patterns preceding each action potential have nothing in common and, consequently, are the most variable.

We characterized the variability of the spike-evoking stimulus patterns by using the average receptive field of each neuron as a reference template, with few assumptions about the relevant stimulus features. The variability of the stimulus–response ensemble could be ascertained by measuring the conditional variance of the stimulus patterns that initiate action potentials. Each receptive field is compared directly with the stimulus pattern preceding each action potential (Fig. 2) by computing the response-conditioned stimulus–STRF SI (equivalent to a correlation coefficient between the STRF and each pre-event stimulus pattern) (see Materials and Methods). For each action potential, the SI characterizes the degree of shape similarity between the pre-event sound pattern and the STRF (values near 1 designate a

close match; values near 0 indicate a mismatch between the STRF and the stimulus pattern).

The match in spectrotemporal shape between the pre-event DMR stimulus and the STRF for the action potentials of a neuron varied considerably across the sampled population of ICC neurons ($n = 61$). Example STRFs of two neurons and representative pre-event patterns for each are shown in Figure 2 (red arrows designate the instant of an action potential). In some instances, the stimulus patterns that evoked action potentials were highly conserved across the response ensemble (Fig. 2A). Comparing the STRF and the pre-event sound patterns overlapping the significant ($p < 0.002$) excitatory (black contour) and inhibitory (white contour) STRF subregions shows that excitatory domains are typically overlapped by ON or high-energy stimulus regions, whereas OFF or low-energy stimulus patterns typically overlap the inhibitory STRF patterns. Although the precise match in spectrotemporal shape composition between the STRF and each pattern could vary widely for this neuron (SI range, -0.02 – 1.0), it shows a strong bias toward high similarity index values (the same neuron is shown in Fig. 3B; mean SI value, 0.53). The response pattern for this neuron was accompanied by a very low spike rate (0.11 spikes/s; $n = 139$ spikes) and a highly significant STRF ($p < 0.002$; peak-to-peak STRF amplitude, 0.21 spikes/s) (Fig. 3B). The low spike rate and high SI values indicate the neuron responded selectively to a small subset of the DMR stimulus ensemble.

Other neurons exhibited only a low association between the stimulus ensemble and the measured STRF (Fig. 2B), implying lower selectivity to specific spectrotemporal stimulus parameters. Despite a highly significant STRF ($p < 0.002$; peak-to-peak STRF amplitude, 7.2 spikes/s; the same neuron is shown in Fig. 3E with the corresponding amplitude scale) and a robust firing rate (18.1 spikes/s; $n = 22,355$ spikes), the pre-event stimulus patterns for this neuron exhibited a highly variable spectrotemporal composition. Although the SI values extended over a broad range (-0.56 – 0.60) the average SI was near 0 (0.06). Accordingly, excitatory and inhibitory STRF subregions were not exclusively overlapped by ON and OFF stimulus patterns before spike initiation. This example demonstrates that the stimulus patterns that lead to the generation of action potentials do not need to precisely match the STRF of the neuron. It shows how a neuron can respond to a wide variety of stimulus patterns, each providing only a small energy contribution to the STRF structure.

This second-order analysis of the stimulus–response ensemble leads to a single similarity index measurement for each action potential. The behavior of each neuron can be characterized by the collection of SIs expressed as a distribution function (Fig. 3). For a hypothetical feature detector, the spike-evoked stimulus patterns would always precisely match the STRF, and therefore, the SI distribution consists of a single peak at $SI = 1$ (Fig. 3A–E, middle column, dashed curve). For a random neuron (see Materials and Methods), the SI distribution would be tightly centered about $SI = 0$ (Fig. 3A–E, middle column, dashed-dot curve). The SI distributions of real neurons rarely resembled those of the hypothetical feature detector (Fig. 3A–E, middle column, continuous curve) and often overlapped those of a random neuron (Fig. 3C–E, middle column).

The selectivity bias of each neuron can be obtained by converting its SI distribution to a cumulative SI distribution function (Fig. 3A–E, right panels; real neuron CDF, continuous curve; random Poisson neuron CDF, dashed-dot; feature detector CDF, dashed) and measuring the relative difference between the real neuron and the control conditions (see Materials and Methods).

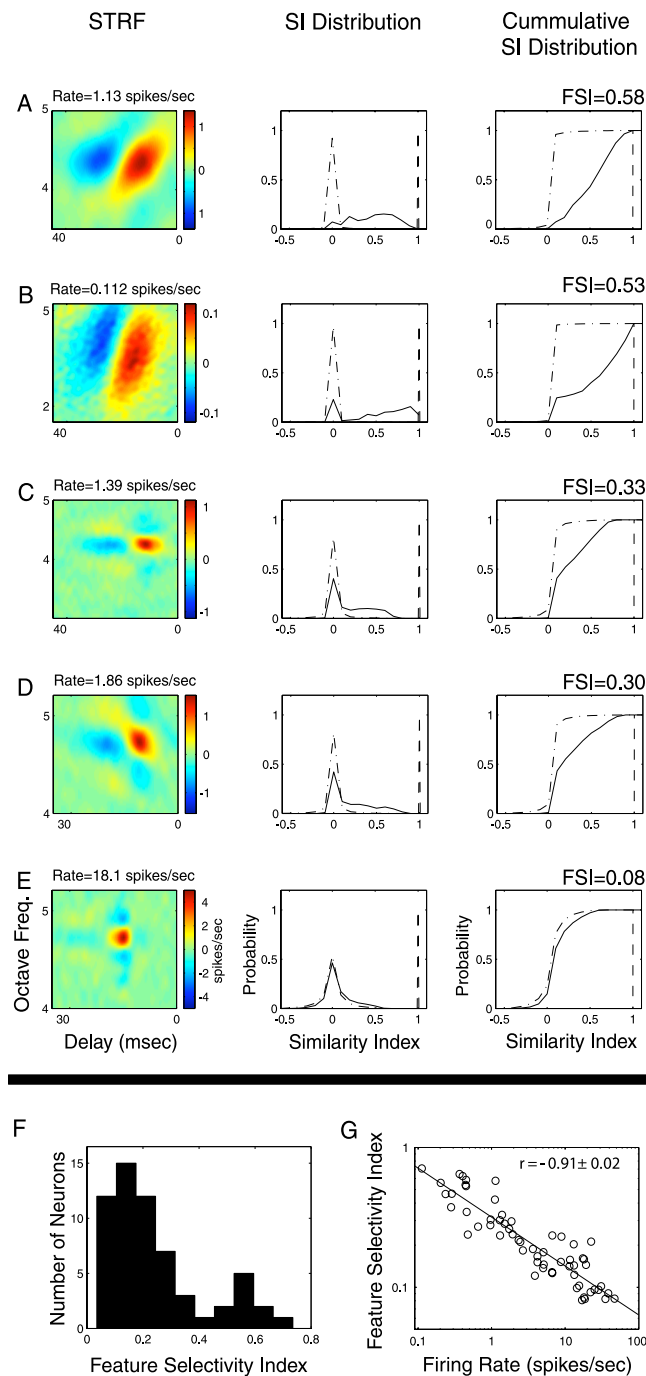


Figure 3. Trade-off in feature selectivity and spike rate throughput. Data are shown for five single neurons in the ICC (A–E). STRFs are shown for each neuron in the far left panels (color bar; units of spikes/second). Collectively, the variability of the pre-event stimulus waveforms of each neuron is represented by the probability distribution of similarity index (PDF) (middle panels) or by a cumulative SI distribution function (CDF) (right panels). Each distribution is shown for a random neuron (Poisson spike train; dashed-dotted), the real neuron (continuous), and a hypothetical feature detector (dashed). The FSI measures the bias of the real neuron CDF in relationship to the random and feature detector CDFs (see Materials and Methods). Neurons are shown according to increasing FSI (A–E). Neurons with low FSI have SI distributions biased toward the random control (high variability) and typically the highest firing rates, whereas the distributions of high FSI neurons are biased toward the feature detector control (low variance) and low spike rates. **F**, The distribution of FSI is bimodally distributed in the ICC. **G**, Firing rate and FSI shows a strong negative correlation, so that firing rates systematically decrease with increasing FSI (doubly logarithmic linear regression fit: continuous line). Freq., Frequency.

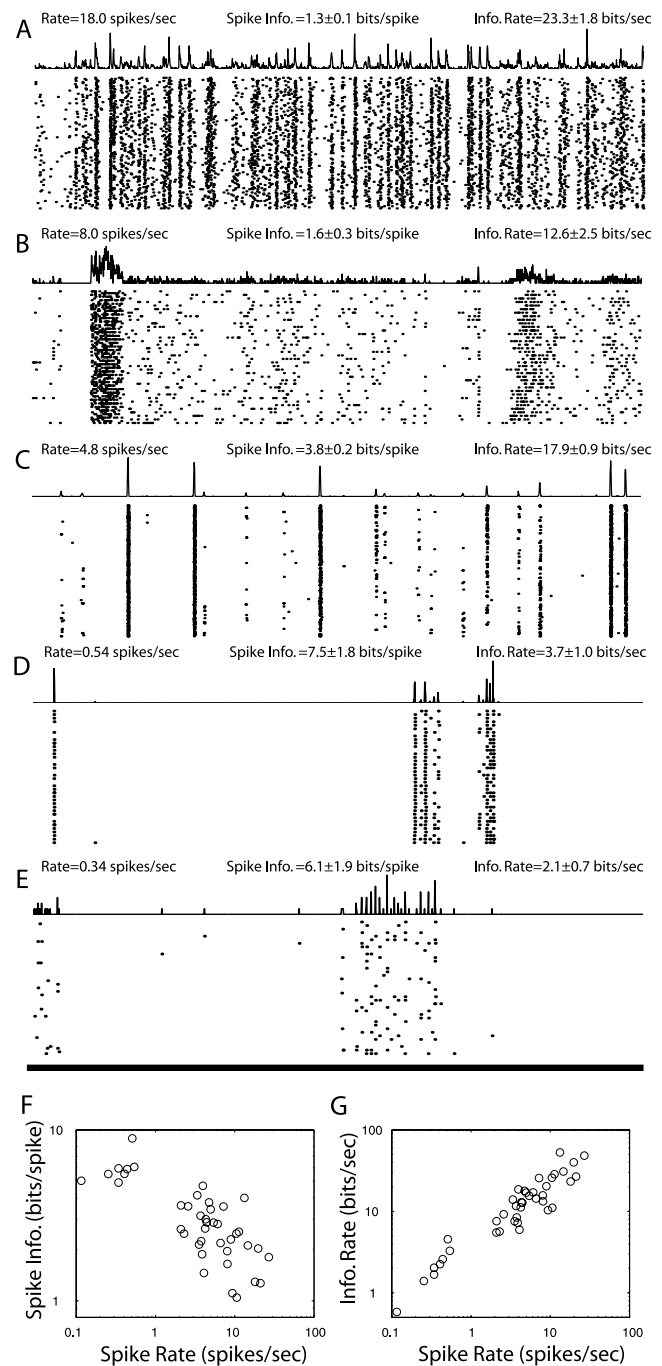


Figure 4. Trade-offs in spike information content and response throughput. **A–E**, Examples of response rastergrams (2 s segments) from five inferior colliculus neurons (sorted according to spike rate) show a precisely time-locked response signature down to millisecond resolution (shown at 1 ms bin widths). Poststimulus time histograms are superimposed above each rastergram. Single action potentials are more meaningful (high spike information) for neurons with the low spike rates. The stimulus-related information per unit time (information rate), however, appears to increase with increasing spike rate for the five neurons. As for the examples, population results show a distinct trade-off between spike information and firing rate (**F**) (mean \pm SD, $r = -0.78 \pm 0.05$; $p < 10^{-3}$). In contrast, the information rate is dominated by the firing rate of the neuron (**G**) ($r = 0.95 \pm 0.02$; $p < 10^{-3}$).

Thus, rather than forcing individual neurons to fall into either of two categories (random or feature selective), this procedure quantifies selectivity along a continuum. The FSI (see Materials and Methods) (Fig. 3A–E, far right) assumes numerical values between 0 (no feature specificity) and 1 (a neuron that behaves

like a hypothetical feature detector). The FSI of real neurons falls somewhere in between (observed range, 0.08–0.71). Because the random neuron and the feature detector neuron serve as a reference for this analysis, the FSI measures the relative feature selectivity bias of the response of the neuron.

FSIs in the ICC systematically increased as the variability of the stimulus–response ensemble decreased (Fig. 3*A–E*). Neurons with SI distributions that mostly overlapped the random neuron control typically had low FSI values. As an example, the neuron of Figure 2*B* responds to a large variety of spectrotemporal patterns (high variability), and consequently, its SI distribution consisted of values near 0 that overlapped the random control condition (Fig. 3*E*). Alternately, if the neuron responded reproducibly to a specific feature of the stimulus (low variability) (Fig. 2*A*), the SI distribution was shifted toward the feature detector control distribution (Fig. 3*B*). In such instances, the pre-event stimulus patterns were highly consistent in their spectrotemporal composition, and the STRF closely resembles the pre-event stimulus patterns (Fig. 2*A*). These examples show that the spectrotemporal feature specificity of the response of the neuron is inversely related to the variability of the stimulus–response ensemble. It also demonstrates that feature selectivity and stimulus–response variability fall along a continuum [i.e., neurons that act as pure feature detectors (at least in the case of the DMR parameter ensemble) are the exception].

Spectrotemporal feature selectivity and spike information content: trade-offs with information throughput

The spectrotemporal feature selectivity of the studied ICC neurons, estimated by their FSI to the DMR stimulus, covered a wide range of values (0.08–0.71) and was bimodally distributed (Fig. 3*F*) (verified *post hoc*, k-means cluster analysis; $p < 0.01$). Although a small number of neurons show a strong bias toward high selectivity (20%; $n = 12$; $FSI > 0.36$) the majority of ICC neurons exhibit low feature selectivity (80%; $n = 49$; $FSI < 0.36$).

High FSIs are accompanied by low spike rates, whereas neurons with low FSIs typically have high spike rates (Fig. 3*G*). Regression analysis revealed a strong negative correlation (mean \pm SD, $r = -0.91 \pm 0.02$; $p < 10^{-3}$; bootstrap) between feature selectivity and firing rates that followed a power-law relationship: $FSI \approx 0.318 \times \text{rate}^{-0.35}$ (linear regression fit on a log–log plot; predictive quality of power-law, $r = 0.79 \pm 0.05$) (Fig. 3*G*). Thus, feature selectivity places constraints on the DMR stimulus patterns that a neuron can respond to, leading to a low firing rate for highly selective neurons.

High variability in the stimulus patterns preceding action potentials limits the contribution of a single spike–stimulus pattern to the total STRF. To determine the relative contribution of individual action potentials, we computed the mutual information of $n = 42$ single neurons by analyzing 125 repeated trials of the

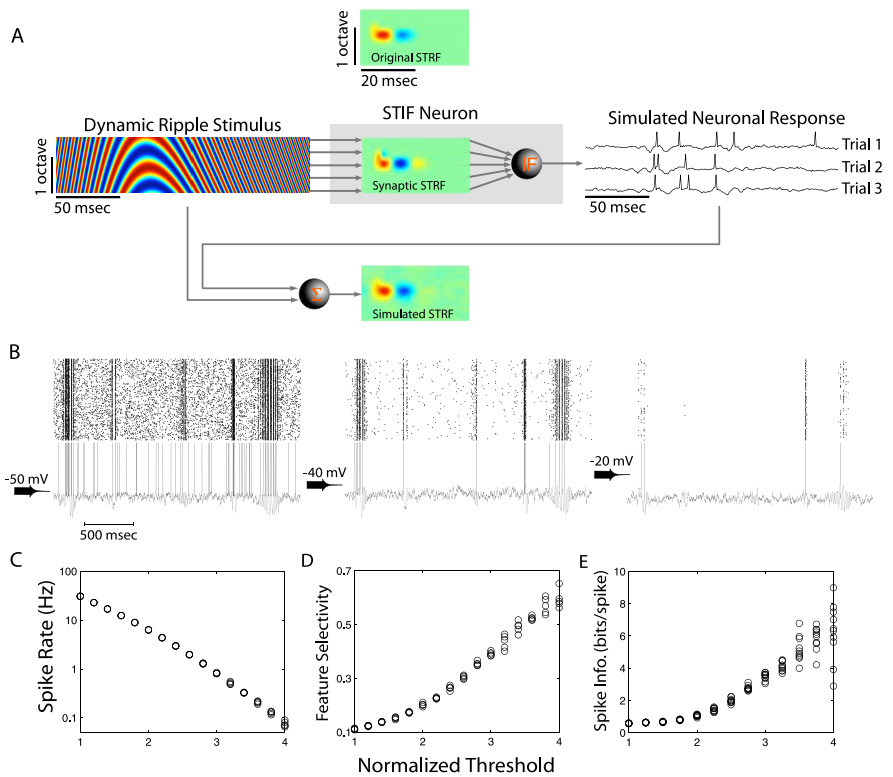


Figure 5. STIF neuron model. The STIF neuron model is shown as applied to the receptive field of an ICC neuron (original STRF; top). The dynamic moving ripple stimulus (middle left) is integrated by a synaptic STRF that accounts for the postsynaptic current input to an IF compartment. The intracellular membrane voltage is shown for three trials of the dynamic moving ripple stimulus segment. The spike-triggered average receptive field (simulated STRF; bottom) closely approximates the original receptive field as expected. Results from the STIF neuron model of **A** are shown in **B** as a function of the spike threshold. Simulated rastergrams (**B**, top) and a sample intracellular voltage waveforms (**B**, bottom) are shown for intracellular thresholds of -50 , -40 , and -20 mV (intracellular voltage variance, 10 mV with no spikes; resting potential, -60 mV). Rastergrams lose their sustained response and become increasingly phasic as the threshold level is increased. Results are shown by changing the threshold level and keeping all other model parameters fixed. Because the model results depended only on the relative synaptic activity in relation to the resting and threshold potentials, the spike threshold levels were normalized as the number of intracellular voltage SDs required for spike activation (see Materials and Methods). For **B**, these correspond to normalized thresholds of 1, 2, and 4 SDs, respectively. Spike rates systematically decrease with increasing normalized threshold (**C**), whereas feature selectivity (**D**) and spike information (**E**) increase. Results are shown for 10 simulations using different sound segments at each threshold value.

ripple noise stimulus (Strong et al., 1998). Neuronal responses in the ICC to repeated sound segments typically consisted of phasic responses of only a few milliseconds duration (Fig. 4*A–E*) that were often highly reproducible. This timing precision and response reliability was reflected in the information content per spike, which systematically increases for the examples as the spike rate is reduced (Fig. 4*A–E*, far right).

By examining the relationship between spike information content and measured spike rates for the population (Fig. 4*F*), we find that spike information is inversely related to neuronal spike rates (log rate vs log information; mean \pm SD, $r = -0.78 \pm 0.05$; $p < 10^{-3}$; bootstrap) in a manner that closely mimics the relationship for feature selectivity (Fig. 3*G*). As for feature selectivity this relationship was accurately accounted for by a power-law relationship of similar exponent (spike information $\approx 4.0 \times \text{rate}^{-0.3}$; goodness of fit, $r = 0.89 \pm 0.03$). Thus, unitary action potentials of neurons with low spike rates and high selectivity have the highest spike information content. In contrast, ICC neurons with high spike rates are characterized by lower fidelity and lower spike information, but have a large number of action potentials available to encode stimulus information.

Which of these two possible extremes conveys the most stim-

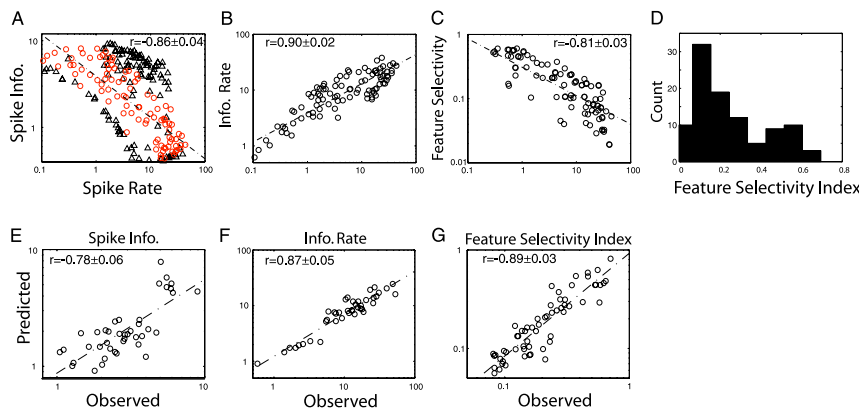


Figure 6. Population simulation of the STIF model accounts for observed trade-offs in the ICC. The STIF model was simulated at 240 randomly sampled threshold and signal-to-noise conditions. The optimal threshold and SNR parameter range were determined by matching the model data from **A** against the ICC spike information versus spike rate trend (Fig. 4*F*) (see Materials and Methods). Optimal (red circles) and nonoptimal (black triangles) model parameters are shown for **A**; only the optimal values are shown for the subsequent panels (black circles). Trade-offs in response throughput (spike rate and information rate) and spiking fidelity (spike information and feature selectivity) are accurately predicted by the STIF model. The model accurately replicates the spike information versus spike rate trend (**A**; compare with Fig. 4*F*), spike rate versus information rate (**B**; compare with Fig. 4*G*), spike rate versus feature selectivity (**C**; compare with Fig. 3*G*), and the bimodal nature of feature selectivity (**D**; compare with Fig. 3*F*). Dashed-dot curves in **A–C** represent the optimal power-law fits to the model data. The predicted and observed spike information (**E**), information rate (**F**), and feature selectivity (**G**) are closely matched for the experimental data and model. Predictions for each metric were obtained by substituting the spike rate of each neuron into the power-law fits from the STIF model data in **A–C**.

ulus information to the neuronal encoding? The overall information conveyed by a neuron can be expressed as the information rate, that is, the product of information per spike and the spike rate of the neuron: $I_{\text{rate}} = R \times I_{\text{spike}}$ (units in bits/second). The relationship between neuronal spike rates and information rates in the ICC shows a strong positive correlation ($r = 0.95 \pm 0.02$; $p < 10^{-3}$) (Fig. 4*G*) and is well accounted for by a power-law fit: information rate $\approx 4.0 \times \text{rate}^{0.7}$ ($r = 0.87 \pm 0.05$). Thus, despite a significant trade-off for spike information and spike rate (Fig. 4*F*), the mean spike rate of the neuron (and not the information per spike) dominates the total communicated information.

Intracellular basis for stimulus selectivity and its trade-off with spike rate and spike information

We hypothesize that the observed trade-offs in the fidelity of the neuronal encoding (feature selectivity, spike information) and its throughput (spike rate, information rate) could be explained by thresholding in the action potential generation mechanism. Conceptually, neurons with higher spike thresholds should produce lower spike rates and would exhibit higher specificity in the stimulus–response relationship and therefore greater fidelity. It is not clear, however, whether such a simple mechanism could account for the bimodal nature of the feature selectivity in the ICC population and the paradoxical result that low-fidelity neurons convey the most information by virtue of their higher spike rates. We devised a simplified STIF neuronal model to test whether it could account for the observed trends.

The neuronal model consists of a synaptic STRF that takes into account the spectrotemporal integration of the synaptic afferents (see Materials and Methods) (Fig. 5*A*). The synaptic current produced by this compartment is used to drive an integrate-and-fire neuron model that accounted for the cell membrane integration and spike generation (see Materials and Methods). Example traces of the model output for a segment of the DMR stimulus show that the spike rates decrease as the relative spike activation threshold is increased (Fig. 5*B,C*); the response spike

specificity, however, improves for higher thresholds as can be seen in the trends for spike information and feature selectivity (Fig. 5*D,E*). The model results resemble data from the real ICC neurons in that the response rasters lose their sustained response and become increasingly phasic, and thus more precise, as the spike rate is decreased (Fig. 5*B*).

In a large-scale simulation, we randomly sampled neuronal parameters for 240 STIF model neurons to mimic the random sampling of the neuronal data for the ICC population. Each neuron consisted of a representative ICC STRF (Qiu et al., 2003) (see Materials and Methods), which was tested by computing its FSI, spike information, information rate, and spike rate at a randomly chosen threshold condition, membrane time constant, and SNRs. The range of model parameters used (SNR and threshold level) was iteratively adjusted so that the relationship between spike rates and spike information for the model population (Fig. 6*A*) covered a comparable range as in the ICC data (Fig. 4*F*). Although the model parameters were strictly adjusted for the spike

information and spike rate relationship, the general features of the remaining trends in the ICC (information rate and feature selectivity) emerged naturally from the model simulations (Figs. 6*B–D*).

The relationship between spike rate and feature selectivity obtained from the STIF model (Fig. 6*C*) closely resembles the ICC population data (Fig. 3*G*). Stimulus patterns that lead to activation of the neuron model are significantly more precise on a spike-to-spike basis (FSI vs spike rate: $r = -0.81 \pm 0.03$; $p < 10^{-3}$) and convey more stimulus related information (Fig. 6*A*) (spike information vs spike rate: $r = -0.86 \pm 0.04$; $p < 10^{-3}$) for neurons in which the spike rate is low. Despite this negative correlation between spike information and spike rate, information rates are proportional to the mean spike rates (Fig. 6*B*) ($r = 0.90 \pm 0.02$; $p < 10^{-3}$) as observed in the ICC data (Fig. 4*G*). Therefore, the average driven activity of each neuron dominates their information-carrying capacity. Surprisingly, the model accurately replicates the bimodal nature of the feature selectivity for the ICC (compare Figs. 6*D*, 3*F*) despite the fact that the original distribution of model parameters was continuously defined (see Materials and Methods).

To quantify how well the STIF model accounts for the mutual information and feature selectivity trends, we fitted each of the model scatter plots to a power-law function (Fig. 6*A–C*, dashed-dot curves) and used these simulated curves to predict the original population data. The power-law relationships from the STIF model accurately predicted the feature selectivity of individual neurons from their spike rate measurements (regression slope = 1.0; $r = 0.89 \pm 0.03$) (Fig. 6*G*). Furthermore, the STIF model accounts for the spike information (regression slope = 0.80; $r = 0.78 \pm 0.06$) (Fig. 6*E*) and information rate trends (regression slope = 0.76; $r = 0.87 \pm 0.05$) (Fig. 6*F*) and replicates the strong negative correlation between these two variables.

Which neuronal factors contribute to the observed trade-offs? We addressed this question by examining the relationships between the intracellular parameters of each model neuron and

each of the measured response metrics (Fig. 7). Although the intracellular SNR shows a significant but subtle correlation with feature selectivity (Fig. 7B) ($r = 0.44 \pm 0.07$; $p < 10^{-3}$) and spike information (Fig. 7F) ($r = 0.28 \pm 0.09$; $p < 10^{-3}$), detailed examination revealed that this correlation does not account for the dominant trade-off in the spiking fidelity because SNR would have to be negatively correlated with spike rate (data not shown) ($r = 0.04 \pm 0.09$; $p > 0.9$). Similarly, the membrane time constant shows no significant correlation with feature selectivity (Fig. 7A) ($r = 0.0 \pm 0.1$; $p > 0.9$) and only a weak correlation with spike information (Fig. 7E) ($r = -0.2 \pm 0.1$; $p < 0.05$) that does not account for the observed trends in neuronal fidelity. The most pronounced trend resulted from the spike activation threshold of the model. Spike rates decrease systematically with increasing threshold (Fig. 7D) ($r = -0.964 \pm 0.008$; $p < 10^{-3}$), whereas the feature selectivity and spike information systematically increase (Fig. 7C) (FSI vs threshold: $r = 0.83 \pm 0.03$; $p < 10^{-3}$) (Fig. 7G) (spike information vs threshold: $r = 0.86 \pm 0.02$; $p < 10^{-3}$). The close agreement of the response trends between the neuronal data and model thus suggests that the observed trade-offs between spiking fidelity and response throughput are controlled by the spike threshold of the neuron.

Discussion

In the ICC, as in the model, neuronal fidelity (spike information, feature selectivity) and throughput (spike rate, information rate) were systematically traded off. Neurons with high spike rates can convey large amounts of stimulus-related information despite unreliable spiking. In contrast, neurons exhibiting sparse responses can signal the occurrence of stereotyped parameter constellations or features with action potentials that convey high information value. The precise nature of this trade-off was accurately reproduced in a population of model neurons by simply considering a distribution of spike threshold levels. Threshold has the effect of enhancing the spiking fidelity at the expense of limiting the communicated information in a manner that closely matches the physiological data. It is plausible that variations in effective spiking threshold in ICC could arise from variation of ascending inhibition, voltage-gated conductance, cell surface areas, input resistance differences, or anesthesia level (see details below). Because these are common mechanisms of altering threshold level in all spiking neurons, the results outline a general encoding principle.

Higher feature selectivity limits the subset of stimulus features that a neuron can potentially respond to as evident in our analysis. By design, the features in the DMR are highly variable, and consequently, average firing rates of neurons with high selectivity are low. Presumably, if the DMR were biased to include many more epochs of the preferred feature, neurons with high FSIs could generate higher mean spike rates. Neurons with low FSIs and highly significant STRFs generally had exceptionally high spike rates. An ideal linear integrating neuron would respond to a large collection of stimulus patterns, provided that stimulus energy is presented within its spectrotemporal filter. Such an idealized neuron should phase-lock to a variety of structured and

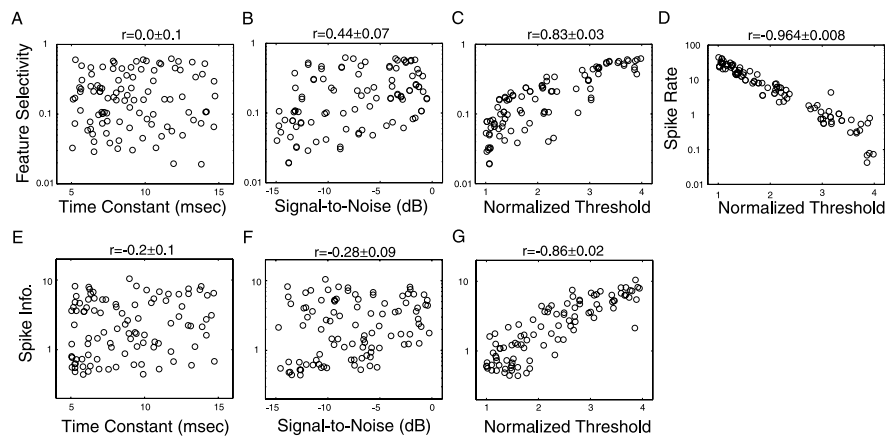


Figure 7. Neuronal factors contributing to the observed trade-offs in feature selectivity, information throughput, and spike information content. The membrane time constant and SNR of the STIF neuron showed only a weak correlation with feature selectivity (A and B) and spike information (E and F) that did not account for the observed trade-offs. Raising the requirements for spike activation by increasing the normalized threshold level for the model population reduced spike rates (D) and increased the selectivity (C) and information per spike (G). This simultaneous reduction of the population throughput and increase in neuronal fidelity that results from the spike threshold explains the observed trade-offs.

unstructured inputs with equal efficacy (Escabí and Schreiner, 2002), assuming that sufficient stimulus energy is provided. Neurons with high selectivity may also require highly correlated inputs (so that the stimulus pattern closely resembles the average STRF) to initiate temporally phase-locked action potentials (Escabí and Schreiner, 2002; Hsu et al., 2004). Such an additional requirement would force the neuron to respond exclusively to stimulus patterns that resemble the average.

The relationship between neuronal fidelity and response throughput across the sampled population likely represents a continuum in the neuronal code for the ICC. A distributed code could have a significant impact on sensory information transfer from ICC to auditory thalamus. Single action potentials from high spike rate ICC neurons are not per se very meaningful, but can convey large amounts of information if spikes are pooled together, as required for a “linear” rate code. In contrast, neurons with low spike rates and high feature selectivity would generate sparse responses and low information rates. However, single action potentials from high-FSI neurons are very informative, highly reliable, and are much more likely to convey meaningful information in the precise timing of single neuronal events (Figs. 4 and 5). Whereas the information per spike and feature selectivity of ICC neurons (Figs. 3G, 4F) and our model (Fig. 6A, C) were inversely related to driven spike rates, overall information rates were dominated by the average firing rate (Figs. 4G, 6B).

An intriguing result predicted by the neuronal analysis and model is that neurons with very similar receptive field preferences could employ vastly different neuronal encoding strategies. Conceptually, two neurons with identical average spectrotemporal preferences could exhibit different levels of selectivity. A neuron with a highly precise stimulus–response relationship would respond reproducibly to a specific stimulus pattern or a specific stimulus combination as for an ideal feature detector neuron (Fig. 1A) or a representative ICC cell (Fig. 2A). The low variance (high FSI) for this example ICC cell indicates that it acts as a frequency modulation up-sweep detector and distinguishes this cell from a neuron with similar-looking STRF that responds to any component within this sweep at different latencies (low FSI) (as for Figs. 1B, 2B). Thus, the concept of selectivity is closely related to the variance of the spike-evoking stimulus ensemble

and not to the mean stimulus as previously suggested (deCharms et al., 1998). As a consequence, a meaningful interpretation of STRFs requires knowledge of the response variance expressed in terms of quantitative measures such as FSI or spike information. Neurons with identical STRFs may reflect either the averaged response to a large stimulus ensemble or an exclusive response pattern to a highly stereotyped stimulus feature. This could occur, for instance, if two neurons share similar intracellular activity, but have different spike thresholds or resting potentials as observed for populations of ICC neurons (Kuwada et al., 1997; Sivaramakrishnan and Oliver, 2001).

Numerous nonlinearities have been implicated as serving an essential role for neuronal coding of sensory information. Because linear receptive field models are often not very predictive on their own, their usefulness for modeling single neuron activities has been recently brought into question (Bar-Yosef et al., 2002; Sahani and Linden, 2003; Machens et al., 2004). Nonlinear ionic conductances in the spiking mechanisms are perhaps the most pronounced nonlinearities in single neurons, because they are responsible for converting the continuous intracellular activity to a binary spiking pattern. Higher thresholds would effectively increase the order of the system nonlinearity, thus increasing the response selectivity. Recent extensions of a linear temporal receptive field model that incorporates a simple spike generating threshold has been shown to accurately predict the spiking activity of visual neurons (Keat et al., 2001). Our data further reveal that a linear RF model with an appropriate spike generating nonlinearity accurately reflects trade-offs in the fidelity and throughput as observed in the ICC population.

The threshold nonlinearity in the spike generation mechanism is a key attribute of all spiking neurons and the resulting trade-offs in spiking fidelity and response throughput could represent a general property of the neuronal code. The spike threshold “level” on its own is insufficient for generating the observed population trade-offs, because the resting potential and strength of the synaptic drive also contribute significantly to the observed finding. Conceptually, a higher spike threshold voltage has an equivalent functional outcome for the model as a lower resting potential or a smaller synaptic current drive. Our concept of “normalized” threshold was thus introduced to account for these three factors that together influence the firing rate and spiking fidelity of the model. Although the spike threshold nonlinearity shapes some aspects of neuronal responses in the visual and auditory system (Casseday et al., 1994; Kempter et al., 1998; Bringuier et al., 1999; Priebe et al., 2004), its role in shaping response selectivity and information transmission has not been reported previously. These findings therefore address issues of neuronal information and stimulus coding beyond modality-specific details that likely apply to spiking neurons in general.

Although in the neuronal model, spiking fidelity and throughput are collectively controlled by the combined influence of the spike threshold, resting potential, and the relative size of the synaptic drive, an equivalent functional outcome could be implemented in real neurons with a variety of mechanisms. A reduction in the resting potential via shunting conductance could extend the intracellular voltage range required to generate action potentials leading to an “effective” higher spike threshold level and selectivity. Similarly, a sustained inhibitory input could substantially hyperpolarize the neuron, thus extending the requirements for spike initiation (Casseday et al., 1994; Kuwada et al., 1997). Alternately, the input resistance of a cell could also play a role, because neurons with lower input resistance would require a strong and highly concerted synaptic current drive to reach spike

activation (Sivaramakrishnan and Oliver, 2001). Additional enhancement can be provided by nonlinear synaptic conductances (Reyes, 2001; Sivaramakrishnan and Oliver, 2001; Svirkis et al., 2002), which together with the spike threshold could enhance the preference for a particular stimulus component. Collectively, such subthreshold mechanisms combined with the highly nonlinear spike generation threshold would have a common effect of shaping the suprathreshold response rate and refining the stimulus requirements necessary for spike initiation, thereby enhancing (or reducing) the stimulus–response selectivity. Although the level of anesthesia has been shown to alter neuronal responses in the ICC (Kuwada et al., 1989; Ramachandran et al., 1999), it is unlikely that it alone accounts for the observed trade-offs, because a reduced level of excitability (as expected for ketamine; NMDA antagonist) would primarily increase the effective threshold across our population of neurons. This could potentially bias our results toward higher selectivity; however, it would not account for the negative correlation between firing rate and spiking fidelity. A nonlinear mechanism such as the spike threshold is required to produce such an effect.

The possible coding schemes that can be achieved with a simple change in effective threshold level extend from a temporally imprecise “rate code” at low thresholds and high spike rates to a precise “timing code” at higher thresholds with lower spike rates and highly informative action potentials. Lower thresholds are more likely to generate a greater number of unreliable action potentials, because the synaptic activity from a larger number of inputs would surpass the threshold level required for spike activation. Such unreliable spiking activity could, however, be integrated and averaged at the postsynaptic cell to convey a meaningful message. In contrast, higher thresholds could potentially enhance the spiking efficacy in the transmission of information at a postsynaptic neuron, because higher thresholds would presumably favor the intracellular activity of the strongest synaptic link (Swadlow and Gusev, 2002). All or none binary spiking selectivity observed in cortical auditory neurons may in fact be related to a similar threshold mechanism (DeWeese et al., 2003).

One hypothesis regarding the functional segregation of high and low feature selectivity neurons in the population data are that these may correspond to anatomically distinct cell classes, such as disk and stellate cells in the ICC (Oliver and Morest, 1984), or distinct populations of ascending input (Ramachandran et al., 1999). Collicular neurons can differ in many aspects, including intrinsic membrane properties and synaptic receptor characteristics resulting in distinct physiological response types (Sivaramakrishnan and Oliver, 2001). Although it is tempting to speculate that the differences in feature selectivity reflect biophysical properties or morphological differences between cells that affect the intracellular requirements for spike activation, more detailed studies are necessary to establish such relationships. Our model, however, does not require a bimodal distribution of biophysical properties, because the selectivity distribution was accounted for with a continuously defined spike threshold.

Our data and model suggest that thresholding not only affects the average driven activity of a neuron, but it also constrains the rate and specificity of the communicated information in a manner that allows for complementary neuronal codes. Neurons with low spike rates, high thresholds, and high selectivity convey the most information with single spikes that can be used for detecting specific instances of the sensory signal. In contrast, neurons with low thresholds and high spike rates convey the most information if spikes are pooled together and are therefore well suited for

encoding sensory information as a rate code. Given that thresholding is a common mechanism of all spiking neurons, such trade-offs in the fidelity and information throughput of the encoded message may represent a general feature of the neuronal code.

References

- Barlow H (1953) Summation and inhibition in the frog retina. *J Physiol (Lond)* 119:69–78.
- Barlow H (1961) The coding of the sensory message. In: *Current problems in animal behavior* (Thorpe WH, Zangwill OL, eds), pp 217–234. Cambridge, MA: MIT.
- Bar-Yosef O, Rotman Y, Nelken I (2002) Responses of neurons in cat primary auditory cortex to bird chirps: effects of temporal and spectral context. *J Neurosci* 22:8619–8632.
- Borst A, Theunissen FE (1999) Information theory and neural coding. *Nat Neurosci* 2:947–957.
- Brenowitz EA, Margoliash D, Nordeen KW (1997) An introduction to bird-song and the avian song system. *J Neurobiol* 33:495–500.
- Bringuier V, Chavane F, Glaeser L, Fregnac Y (1999) Horizontal propagation of visual activity in the synaptic integration field of area 17 neurons. *Science* 283:695–699.
- Calhoun BM, Schreiner CE (1998) Spectral envelope coding in cat primary auditory cortex: linear and non-linear effects of stimulus characteristics. *Eur J Neurosci* 10:926–940.
- Casseday JH, Ehrlich D, Covey E (1994) Neural tuning for sound duration—role of inhibitory mechanisms in the inferior colliculus. *Science* 264:847–850.
- DeAngelis GC, Ohzawa I, Freeman RD (1993) Spatiotemporal organization of simple-cell receptive fields in the cat's striate cortex. II. Linearity of temporal and spatial summation. *J Neurophysiol* 69:1118–1135.
- de Boer R, Kuypers P (1968) Triggered correlation. *IEEE Trans Biomed Eng* 15:169–179.
- deCharms RC, Blake DT, Merzenich MM (1998) Optimizing sound features for cortical neurons. *Science* 280:1439–1443.
- de Ruyter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W (1997) Reproducibility and variability in neural spike trains. *Science* 275:1805–1808.
- DeWeese MR, Wehr M, Zador AM (2003) Binary spiking in auditory cortex. *J Neurosci* 23:7940–7949.
- Doupe AJ (1997) Song- and order-selective neurons in the songbird anterior forebrain and their emergence during vocal development. *J Neurosci* 17:1147–1167.
- Escabí MA, Schreiner CE (2002) Nonlinear spectrotemporal sound analysis by neurons in the auditory midbrain. *J Neurosci* 22:4114–4131.
- Escabí MA, Miller LM, Read HL, Schreiner CE (2003) Naturalistic auditory contrast improves spectrotemporal coding in the cat inferior colliculus. *J Neurosci* 23:11489–11504.
- Fujita I, Tanaka K, Ito M, Cheng K (1992) Columns for visual features of objects in monkey inferotemporal cortex. *Nature* 360:343–346.
- Gross CG, Sergent J (1992) Face recognition. *Curr Opin Neurobiol* 2:156–161.
- Hermes DJ, Aertsen AM, Johannesma PI, Eggermont JJ (1981) Spectrotemporal characteristics of single units in the auditory midbrain of the lightly anaesthetized grass frog (*Rana temporaria* L) investigated with noise stimuli. *Hear Res* 5:147–178.
- Hsu A, Woolley SM, Fremouw TE, Theunissen FE (2004) Modulation power and phase spectrum of natural sounds enhance neural encoding performed by single auditory neurons. *J Neurosci* 24:9201–9211.
- Keat J, Reinagel P, Reid RC, Meister M (2001) Predicting every spike: a model for the responses of visual neurons. *Neuron* 30:803–817.
- Keller CH, Takahashi TT (2000) Representation of temporal features of complex sounds by the discharge patterns of neurons in the owl's inferior colliculus. *J Neurophysiol* 84:2638–2650.
- Kempler R, Gerstner W, van Hemmen JL (1998) How the threshold of a neuron determines its capacity for coincidence detection. *Biosystems* 48:105–112.
- Klein DJ, Depireux DA, Simon JZ, Shamma SA (2000) Robust spectrotemporal reverse correlation for the auditory system: optimizing stimulus design. *J Comput Neurosci* 9:85–111.
- Kowalski N, Depireux DA, Shamma SA (1996) Analysis of dynamic spectra in ferret primary auditory cortex. II. Prediction of unit responses to arbitrary dynamic spectra. *J Neurophysiol* 76:3524–3534.
- Kuwada S, Batra R, Stanford TR (1989) Monaural and binaural response properties of neurons in the inferior colliculus of the rabbit: effects of sodium pentobarbital. *J Neurophysiol* 61:269–282.
- Kuwada S, Batra R, Yin TC, Oliver DL, Haberly LB, Stanford TR (1997) Intracellular recordings in response to monaural and binaural stimulation of neurons in the inferior colliculus of the cat. *J Neurosci* 17:7565–7581.
- Machens CK, Wehr MS, Zador AM (2004) Linearity of cortical receptive fields measured with natural sounds. *J Neurosci* 24:1089–1100.
- Martin KA (1994) A brief history of the “feature detector.” *Cereb Cortex* 4:1–7.
- Nelken I, Kim PJ, Young ED (1997) 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.
- Oliver DL, Morest DK (1984) The central nucleus of the inferior colliculus in the cat. *J Comp Neurol* 222:237–264.
- Priebe NJ, Mechler F, Carandini M, Ferster D (2004) The contribution of spike threshold to the dichotomy of cortical simple and complex cells. *Nat Neurosci* 7:1113–1122.
- Qiu A, Schreiner CE, Escabí MA (2003) Gabor analysis of auditory mid-brain receptive fields: spectro-temporal and binaural composition. *J Neurophysiol* 90:456–476.
- Ramachandran R, Davis KA, May BJ (1999) Single-unit responses in the inferior colliculus of decerebrate cats. I. Classification based on frequency response maps. *J Neurophysiol* 82:152–163.
- Reid RC, Soodak RE, Shapley RM (1991) Directional selectivity and spatio-temporal structure of receptive fields of simple cells in cat striate cortex. *J Neurophysiol* 66:505–529.
- Reyes A (2001) Influence of dendritic conductances on the input-output properties of neurons. *Annu Rev Neurosci* 24:653–675.
- Sahani M, Linden JF (2003) How linear are auditory cortical responses? In: *Advances in neural information processing systems 15* (Becker S, Thrun S, Obermayer K, eds), pp 109–116. Cambridge, MA: MIT.
- Schnupp JW, Morsic-Flogel TD, King AJ (2001) Linear processing of spatial cues in primary auditory cortex. *Nature* 414:200–204.
- Simmons JA, Howell DJ, Suga N (1975) Information content of bat sonar echoes. *Am Sci* 63:204–215.
- Sivaramakrishnan S, Oliver DL (2001) Distinct K currents result in physiologically distinct cell types in the inferior colliculus of the rat. *J Neurosci* 21:2861–2877.
- Strong SP, de Ruyter van Steveninck RR, Bialek W, Koberle R (1998) On the application of information theory to neural spike trains. *Pac Symp Biocomput* 621–632.
- Suga N, Jen PH (1976) Disproportionate tonotopic representation for processing CF-FM sonar signals in the mustache bat auditory cortex. *Science* 194:542–544.
- Suga N, O'Neill WE, Manabe T (1978) Cortical neurons sensitive to combinations of information-bearing elements of biosonar signals in the mustache bat. *Science* 200:778–781.
- Svirskis G, Kotak V, Sanes DH, Rinzel J (2002) Enhancement of signal-to-noise ratio and phase locking for small inputs by a low-threshold outward current in auditory neurons. *J Neurosci* 22:11019–11025.
- Swadlow HA, Gusev AG (2002) Receptive-field construction in cortical inhibitory interneurons. *Nat Neurosci* 5:403–404.
- Theunissen FE, Sen K, Doupe AJ (2000) Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. *J Neurosci* 20:2315–2331.
- Versnel H, Shamma SA (1998) Spectral-ripple representation of steady-state vowels in primary auditory cortex. *J Acoust Soc Am* 103:2502–2514.
- Yu JJ, Young ED (2000) Linear and nonlinear pathways of spectral information transmission in the cochlear nucleus. *Proc Natl Acad Sci USA* 97:11780–11786.