

Nonlinear Computations Underlying Temporal and Population Sparseness in the Auditory System of the Grasshopper

Jan Clemens,^{1,2} Sandra Wohlgenuth,¹ and Bernhard Ronacher^{1,2}

¹Behavioral Physiology Group, Department of Biology, Humboldt-Universität zu Berlin, 10115 Berlin, Germany and ²Bernstein Center for Computational Neuroscience Berlin, 10115 Berlin, Germany

Sparse coding schemes are employed by many sensory systems and implement efficient coding principles. Yet, the computations yielding sparse representations are often only partly understood. The early auditory system of the grasshopper produces a temporally and population-sparse representation of natural communication signals. To reveal the computations generating such a code, we estimated 1D and 2D linear-nonlinear models. We then used these models to examine the contribution of different model components to response sparseness.

2D models were better able to reproduce the sparseness measured in the system: while 1D models only captured 55% of the population sparseness at the network's output, 2D models accounted for 88% of it. Looking at the model structure, we could identify two types of computation, which increase sparseness. First, a sensitivity to the derivative of the stimulus and, second, the combination of a fast, excitatory and a slow, suppressive feature. Both were implemented in different classes of cells and increased the specificity and diversity of responses. The two types produced more transient responses and thereby amplified temporal sparseness. Additionally, the second type of computation contributed to population sparseness by increasing the diversity of feature selectivity through a wide range of delays between an excitatory and a suppressive feature.

Both kinds of computation can be implemented through spike-frequency adaptation or slow inhibition—mechanisms found in many systems. Our results from the auditory system of the grasshopper are thus likely to reflect general principles underlying the emergence of sparse representations.

Introduction

Successful behavior is tied to the formation of specific representations of the environment to enable the discrimination of friend and foe. Typically, complex feature selectivity arises from more generic representations as one ascends a sensory pathway. This increase of specificity can lead to temporal and population sparseness (Chacron et al., 2011) (but see Willmore et al., 2011). Temporally sparse responses are characterized by well isolated firing events interleaved by periods of neuronal quiescence. Additionally, higher-order neurons are more idiosyncratic, each responding to “its own” feature. Consequently, neurons in a population are less prone to fire together, yielding population-

sparse activity. Besides being an outcome of the creation of more specific representations, sparseness also has intrinsic advantages like the efficient use of energy and neuronal bandwidth as well as the facilitation of subsequent computations involving learning and memory (Barlow, 2001; Olshausen and Field, 2004).

Here we ask how temporal and population sparseness arise in the auditory system of the grasshopper, which generates a sparse and specific representation of courtship signals in a small, three-layer feed-forward network (Fig. 1*a*) (Clemens et al., 2011).

Primary auditory receptors in the grasshopper yield a temporally dense, relatively unspecific and faithful representation of a sound's envelope (Fig. 1*b,c*; Machens et al., 2001; Gollisch et al., 2002; Rokem et al., 2006). We attempt to characterize the transformations underlying neural encoding in second- and third-order neurons—the local and ascending neurons. Local neurons create a temporally sparse representation of song from the dense inputs provided by the receptors (Fig. 1*d*). However, population sparseness is still low as different local neurons respond to very similar features. The ascending neurons in turn establish a population-sparse code with each ascending neuron responding to a more specific stimulus pattern (Fig. 1*e*).

We aimed to get insight into the mechanisms contributing to temporal and population sparseness by fitting low-dimensional models of the stimulus-response relationship to recordings of second- and third-order neurons in the grasshopper using the

Received Nov. 25, 2011; revised May 23, 2012; accepted June 3, 2012.

Author contributions: J.C. and B.R. designed research; J.C. and S.W. performed research; J.C. analyzed data; J.C. and B.R. wrote the paper.

This work was supported by grants from the Federal Ministry of Education and Research, Germany (01GQ1001A) and the Deutsche Forschungsgemeinschaft (SFB618, GK1589/1). We thank Susanne Schreiber for valuable comments on the manuscript and Henning Sprekeler for discussions. We also gratefully acknowledge the careful reading and constructive criticism of the manuscript by two anonymous reviewers.

The authors declare no competing financial interests.

Correspondence should be addressed to Jan Clemens, Behavioral Physiology Group, Department of Biology, Humboldt-Universität zu Berlin, Invalidenstrasse 43, 10115 Berlin, Germany. E-mail: clemensjan@googlemail.com.

S.W.'s present address: Freie Universität Berlin, Animal Behavior Group, Takustrasse 6, 14195 Berlin, Germany. DOI:10.1523/JNEUROSCI.5911-11.2012

Copyright © 2012 the authors 0270-6474/12/3210053-10\$15.00/0

framework of linear-nonlinear (LN) models. These models provide intuitive phenomenological depictions of the neural computations actualized by a neuron and its inputs. In their simplest form, LN models consist of a single linear filter and a static nonlinearity. As the kind of representation found at the level of ascending neurons suggests more complex transformations, we used an extension of the one-filter LN models that allowed us to describe multidimensional computations—spike-triggered covariance analysis (STC; Rust et al., 2005; Fairhall et al., 2006; Petersen et al., 2008; Fox et al., 2010).

We find two different classes of nonlinear computations contributing to sparseness in the auditory system of the grasshopper: sensitivity to the derivative of a stimulus and an AND-NOT like transformation. These abstract computations can be implemented by mechanisms ubiquitous in many neural systems and are thus likely to constitute general principles providing sparse and specific representations.

Materials and Methods

Animals, electrophysiology, and acoustic stimulation. Recordings were performed in adult locusts (*Locusta migratoria*) obtained from a local supplier and held at room temperature ($22 \pm 5^\circ\text{C}$). We recorded intracellularly from identified auditory neurons in the locust's metathoracic ganglion. Auditory neurons are organized in a three-layer feedforward network with receptors as an input layer, an intermediate layer of local neurons, and an output layer of ascending neurons. Intracellular electrophysiological recording methods are described in detail in Vogel et al. (2005). After completion of the stimulation protocol, neurons were stained with Lucifer yellow and identified by their characteristic morphology (Römer and Marquart, 1984; Stumpner and Ronacher, 1991). The dataset consists of seven types of auditory neurons from the intermediate (second-order or local neurons: TN1 $N = 10$, SN1 $N = 2$, SN3 $N = 1$, BSN1 $N = 9$) and the output layer (third-order or ascending neurons AN1 $N = 5$, AN2 $N = 1$, AN3 $N = 2$). Previous studies have shown that the local neuron BSN1 comes in two subtypes, one responding with a short burst to the onset of pulses and one firing more persistently during a pulse, most likely due to different strengths of inhibitory inputs (Stumpner, 1989). Accordingly, we refer to them as “phasic” ($N = 6$) or “tonic” ($N = 3$) subtypes of BSN1.

Natural songs of grasshoppers consist of a broadband carrier whose amplitude is modulated by a species-specific envelope. As the decisive cues for song recognition lie in this envelope, we were interested in how single neurons represent the pattern of amplitude modulation of a sound. We therefore modulated the amplitude of broadband noise (5–40 kHz) with lowpass Gaussian noise (cutoff frequency 140 Hz). The mean of this amplitude modulation was set to ≈ 10 –15 dB above each cell's threshold (thresholds ranged between 45 and 65 dB SPL). The standard deviation (SD) of the random amplitude modulations was 6 dB. We presented these noise stimuli in two variants to estimate and verify the models: one long segment lasting between 5 and 14 min for estimating the models and a shorter 6 s segment, which was repeated at least 18 times and was used for estimating the time-varying firing rate for model testing. For all further analysis we only used steady-state responses by omitting the first 400 ms of each spike train.

Constructing lower dimensional models to characterize neuronal responses. Responses to the long noise stimulus formed the basis for spike-triggered analysis (Schwartz et al., 2006). In essence, spike-triggered analysis consists of finding stimulus features influencing a neuron's spike-




Figure 1. Responses of neurons at three levels of the auditory system of grasshoppers to a subunit of song. **a**, Diagram showing the flow of information in the early auditory system of grasshoppers. **b**, Envelope of song. The black line corresponds to 40 dB, peak intensity was 78 dB. **c–e**, Normalized firing-rate functions of primary auditory receptors (**c**), secondary local neurons (**d**), and third-order ascending neurons (**e**) to the stimulus shown in **b**. Different rows correspond to different cells. While receptors respond throughout the stimulus, local neurons respond transiently but relatively uniformly at the onset of the stimulus. In contrast, ascending neurons respond more diversely with different latencies and patterns to the same stimulus.

ing by comparing the distribution of stimuli \vec{s} preceding a spike r , $p(\vec{s}|r)$, to the distribution of all stimuli, $p(\vec{s})$, and finding directions in stimulus space for which both distributions differ most. This yields LN cascade models of neural computation: a high-dimensional stimulus is reduced by linear projection to one or two feature values; then, a nonlinearity transforms the feature value(s) to the cell's firing rate.

We defined the stimulus \vec{s} as a vector corresponding to the envelope of the sound in the 64, 1 ms wide bins preceding each point in time. The 64-dimensional distribution of stimuli $p(\vec{s})$ was by construction Gaussian. $p(\vec{s}|r)$ was sampled by the spike-triggered ensemble (STE), i.e., the set of stimulus segments preceding each spike collected in response to the long noise segment.

In its simplest form, spike-triggered analysis results in calculating the difference of the mean of both distributions, yielding the spike-triggered average (STA) as a single feature: $\vec{f}_{STA} = \sum_r p(\vec{s}|r) \cdot \vec{s} - \sum_r p(\vec{s}) \cdot \vec{s}$ (the last term is the mean of all stimuli and a constant for our noise stimuli).

To characterize more complex, multidimensional feature selectivity, STC was performed. To that end, we computed the covariance matrix of the STE, $C_{\vec{s}|r}$, and subtracted the covariance matrix of all stimuli, $C_{\vec{s}}$, from it: $\Delta C = C_{\vec{s}|r} - C_{\vec{s}}$. The covariance matrix C of an arbitrary distribution $P(\vec{x})$ is given by $C_{\vec{x}} = \sum_{\vec{x}} p(\vec{x})(\vec{x} - \langle \vec{x} \rangle)(\vec{x} - \langle \vec{x} \rangle)^T$, where the angled brackets denote the average. An eigenvalue decomposition of ΔC yields stimulus directions in which the variance—and not the mean as for the STA—of the spike-triggered and the raw stimulus ensemble differ most. These directions are indicated by eigenvectors associated with non-zero eigenvalues. However, due to the finite sample size (number of spikes) most eigenvalues are non-zero. We checked the significance of the deviation of each eigenvalue from zero by computing 1 and 99% confidence intervals for the maximal/minimal eigenvalues of each recording. To that end, we generated randomized responses by shuffling the spike times and used the distribution of the largest/smallest eigenvalue from 1000 such randomized responses to derive confidence intervals. All cells in our dataset exhibited at least two significant eigenvalues at this significance level.

We performed the STC analysis in a subspace orthogonal to the STA, by projecting the STA from each stimulus vector: $\vec{s}_{\perp} = \vec{s} - (\vec{s}^T \vec{f}_{STA}) \vec{f}_{STA} / |\vec{f}_{STA}|^2$. This rendered the STC eigenvectors orthogonal to the STA and greatly facilitated the comparison of models derived from

STA and STC analysis, as a model including STC filters is a direct extension of the lower-dimensional STA-only model. As eigenvectors yielding the filters recovered by STC analysis are only defined up to an arbitrary sign, we choose the sign such that the STC filter is most similar to the negative derivative of the STA filter (Fairhall et al., 2006). Furthermore, all filters were normalized to unit-norm.

The nonlinearity is given by Bayes' rule as the ratio of the raw and the spike-triggered stimulus distribution in the stimulus subspace defined by the filter(s): $\langle r \rangle p(\vec{s}'|r)/p(\vec{s}')$. $\langle r \rangle$ is the average firing rate in the response set used for estimating the model. \vec{s}' is the stimulus projected onto a subspace defined by the STA, or the STA and the STC filter with the largest absolute non-zero eigenvalue; it can thus be either 1D or 2D. $p(\vec{s}')$ is the distribution of projection values of all stimuli and is by definition Gaussian with SD 6 dB. $p(\vec{s}'|r)$ is the distribution of projection values of the STE and was computed by kernel-density estimation.

Two kinds of model were constructed for each recording: one model consisted only of the STA and a 1D nonlinearity. We refer to it as the "STA model." The other model contained the STA filter and the filter with the largest absolute eigenvalue from the STC analysis—here called the "STC filter"—plus a 2D nonlinearity. We called it the "STC model."

Quantification of model performance. A bias-corrected version of Pearson's coefficient of correlation ρ was used to quantify how well each model predicted the neuronal response to a novel stimulus (Petersen et al., 2008). To that end, the time-varying firing rate $r(t)$ of the neuron was estimated from responses to several repetitions of a stimulus not used for model estimation by binning time at 1 ms and smoothing with a box kernel spanning two bins. A predicted response $\hat{r}(t)$ to the same stimulus was obtained from the STA and STC models.

As the neuronal response is noisy, a model of the neuron can never perform better than that noise level. Thus, the naive estimator of the correlation is downwardly biased by that noise. To correct for this bias, we estimated the noise in the response by calculating $r(t)$ from two equal-sized, exclusive subsets of the stimulus repetitions, yielding two independent estimates of the firing rate $r_1(t)$ and $r_2(t)$. The coefficient of correlation between these two estimates was then used to normalize the raw correlation: $\rho = \rho(r, \hat{r})/\rho(r_1, r_2)$.

Characterization of model structure. To characterize the shapes of the filters we used two metrics: the first was defined by the coefficient of correlation between the derivative of the STA filter and the STC filter, and the second was given by the delay between the peak of the STA filter and the peak of the STC filter with the STA filter being the reference.

We characterized to what extent each 2D nonlinearity in the STC models corresponded to a truly nonlinear combination of the STA and STC filter. This was done by fitting minimal linear and quadratic models incorporating only linear or quadratic interactions between the two filters as described by Fitzgerald et al. (2011). As the 1D STA model is a lower bound for the performance of the minimal 2D models, we quantified how well a minimal linear or quadratic model could explain the performance gain of the 2D STC model relative to the STA model: $(\rho_{\text{STC}} - \rho_{\text{STA}})/(\rho_{\text{STC}} - \rho_{\text{min}})$, where ρ_{STA} and ρ_{STC} are the bias-corrected coefficients of correlation between the empirical and the modeled firing rate and ρ_{min} is the performance of the minimal linear or quadratic model. Values close to 100% indicate that a minimal model of a given order fully explains the performance gain, values close to 0% indicate a failure to explain and hence the importance of higher-order interactions.

To better characterize the types of computation described by the 2D nonlinearity, we quantified to what extent the nonlinearities implemented one of two canonical logical operations on the output of the STA and STC filter. Logical operations were defined as follows: an (STA AND STC)-like operation corresponds to most large values of the 2D nonlinearity being concentrated in the upper right quadrant (positive values of both filters drive the cell), and an (STA AND-NOT STC)-like operation is implemented by a nonlinearity with most weight in the lower right quadrant (positive STA and negative STC outputs drive the cell). The relative weight each quadrant had in driving the cell was computed by summing over all values in a given quadrant and normalizing by the sum over all quadrants. We then summed the relative weights over the "active" quadrant(s) as defined by each canonical operator and normalized

by the number of active quadrants (see Fig. 4h, orange). This yielded values between 0 (no weight in the "active" quadrants) and 1 (all of the weight in the "active" quadrants).

Simulation of model responses to natural songs. To study the responses of the models to natural signals, we used a set of songs from eight different male grasshoppers of the species *Chorthippus biguttulus* (Clemens et al., 2011). It is well justified to use models fit to neurons recorded in one species of grasshopper, *L. migratoria*, to study the responses to signals of another species, *C. biguttulus*, as the morphological and physiological properties of neurons at the early stages of processing we are interested in are highly similar (Ronacher and Stumpner, 1988; Neuhofer et al., 2008; Creutzig et al., 2009). As the song's amplitude increases over its duration, we used the last 400 ms where the amplitude plateaued. We transformed the amplitude to a dB scale. The natural songs had a SD of 6 ± 1 dB, close to the SD of the noise stimuli used for the estimation and evaluation of the models. To cover the range of firing rates between 20 and 50 Hz observed for natural stimuli (Clemens et al., 2011), we set the average amplitude to +6 dB.

To quantify the transience of model responses, we calculated the percentage of time the firing rate was below its half-maximal value. A highly transient response will reach its maximal firing rate and then quickly return to smaller firing rates, spending little time above the half-maximal firing rate. In contrast, highly persistent responses will spend most of the time near the maximal firing rate and hence above the half-maximal firing rate. In addition, peak firing rates were estimated as the 99th percentile of the predicted firing rate to natural songs.

Quantification of temporal and population sparseness. Sparseness of the modeled responses was quantified using the measure in Willmore and Tolhurst (2001) as in the following:

$$S = 1 - \langle r(t) \rangle^2 / \langle r(t)^2 \rangle. \quad (1)$$

For temporal sparseness, the average in Equation 1 was taken over time and then S was averaged over the eight songs. For population sparseness, we constructed 500 populations for each model class by randomly combining the responses of four cells of the same model class. Then, the average in Equation 1 was taken over the four cells in a population in each individual time bin, and S was averaged over all time bins and songs. For comparison, we included empirical sparseness values of the responses of local and ascending neurons to the same set of songs (Clemens et al., 2011).

Results

We will start by describing the models of one representative cell in detail. We will then show the two classes of computation we found in our dataset and relate their properties to the generation of a sparse representation of natural communication signals.

2D models capture additional aspects of computation

To provide an intuition for the model structure obtained by STC, we will describe the model obtained for one ascending neuron AN1 in depth. The 1D STA model (Fig. 2a,b) consisted of a single STA filter, which describes the temporal feature the cell is responsive to, and a 1D nonlinearity, which transforms the output of the filter to the cell's firing rate and depicts the neuron's tuning for that feature. The cell's STA was largely unimodal, exhibiting one prominent positive lobe at 20 ms preceding the spike and a weak negative lobe between 30 and 50 ms preceding the spike (Fig. 2a). Thus, the cell was sensitive to a lowpass-filtered version of the amplitude of the stimulus. The nonlinearity was skewed toward positive filter values, indicating that the cell preferred stimuli that were similar to the STA (Fig. 2b). Firing was reduced for very large values—the cell exhibited thus a bandpass-like tuning for the STA.

The STC model (Fig. 2c,d) consisted of the STA and a second filter recovered by STC analysis. This STC filter was broader than

the STA filter, mostly positive, and led the STA (Fig. 2c). Thus, the neuron was influenced not only by the STA but also by the envelope in the 20 ms preceding the STA. In the 2D STC model the stimulus is filtered by both filters in parallel. The output values of both filters are then combined to yield the cell's firing rate. This transformation from pairs of filter values to firing rate is implemented by a 2D nonlinearity, which depicts for each pair of filter outputs the resulting firing rate of the neuron (Fig. 2d). This 2D nonlinearity shows that the cell was best driven by stimuli in the lower right quadrant, i.e., when the STA filter produced large positive output values and the STC filter yielded negative outputs. Such a combination corresponds to an AND-NOT like logical operation on the output of the filters. Thus, the STC model yielded a much richer description of neuronal feature selectivity of this cell: addition of a second filter revealed a non-linear computation performed on the stimulus which was not obvious from the STA model alone.

Generally, the stimulus transformations of auditory neurons in the grasshopper were well described by STA and STC models—model performance ranged between 0.5 and 0.8 (0.57 and 0.63 for the STA and STC models of the example presented above; Fig. 2e). The 2D STC models were able to capture additional aspects of the stimulus-response relation as they performed significantly better than the 1D STA models, increasing model performance on average by 9% ($\rho_{\text{STA}} = 0.59 \pm 0.11$, $\rho_{\text{STC}} = 0.65 \pm 0.11$, mean \pm STD, $p = 6 \cdot 10^{-6}$, sign rank). While the STC model thus explained 9% more response variance than the STA model, this gain was relatively small considering the increased complexity of the model (one additional filter and a 2D nonlinearity). However, we will show below that there existed systematic differences in the structure of the predicted responses that enabled the 2D STC models to better explain the level of sparseness observed in the auditory system of the grasshopper.

Analysis of the model structure reveals two types of computation

Looking at the model structure of all cells in our dataset, we found two principal classes of model (Fig. 3). Notably, this dichotomy was not obvious by looking at the STA filters alone: the STA filter and its nonlinearity were similar to that shown in the example (Fig. 2a,b)—the STA filter of all cells was thus mainly integrating and drove the cells for positive projection values (Figs 3, upper row, 4i). Only the incorporation of the STC filter revealed fundamental, qualitative differences between models, justifying the discrimination of two principal classes of neurons: the STC filters on the left side of the figure were all biphasic whereas those on the right were unimodal (compare Fig. 3a,b, lower row). The fact that we found only two classes of models does not exclude the existence of additional computational classes in the auditory system of the grasshopper—local or ascending neurons not recorded might implement different kinds of transformations.

Analyzing the filters and nonlinearities allowed us to interpret the computations performed by both model classes.

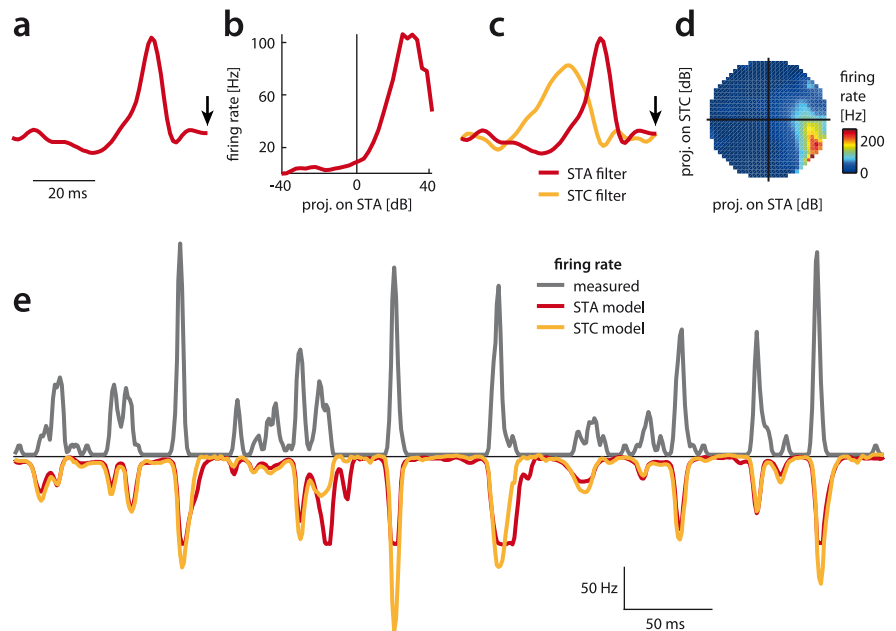


Figure 2. Structure of the STA and STC model for an ascending neuron AN1. *a*, STA filter. The arrowhead marks the time of occurrence of the spike. *b*, Nonlinearity relating the projection values of each stimulus onto the STA to the cell's firing rate. *c*, STA filter (red) and the filter obtained by STC (orange). *d*, 2D nonlinearity relating the projection values of stimuli onto the STA (*x*-axis) and STC filter (*y*-axis) to the cell's firing rate. Firing rate is color coded (see color bar). Note that this neuron responds only to a small subset of projection values. *e*, Measured response (black) and the predictions from the STA (red) and STC model (orange) were highly similar.

"Derivative-like" cells

We found specimens of this first group of cells only among the second-order, local neurons: TN1, SN1, SN3, and BSN1 (the tonic subtype, termed BSN1t) (Fig. 3a). As for all cells, the STA filter was excitatory for positive projections of the stimulus (Fig. 4i, red). The STC filter of this class of models was highly similar to the negative derivative of the STA (Fig. 4a,e; correlation coefficient between the derivative of the STA and the STC filter 0.83 ± 0.11). This high correlation means that the shape of the STC filter was largely determined by that of the STA filter. Given that the STA filters were relatively uniform, this made different cells of this class respond to very similar features. In addition, both filters exhibited great overlap in time, making these cells respond to the stimulus on a short time scale of the order of the STA filter's width (Fig. 4f; delay between peaks of the STA and STC filter 3.0 ± 0.6 ms).

The nonlinearity of the STC filter was quadratic-like, rendering these cells weakly phase invariant (Fig. 4j, red). A minimal linear model of the interactions of both filters accounted for only $23 \pm 8\%$ of the performance gain of the STC model (Fig. 4g, top). The incorporation of quadratic interactions doubled this ($51 \pm 19\%$; Fig. 4g, bottom), but was not able to *fully* explain the interaction of the filters. This indicates that both filters interacted in a highly nonlinear fashion in this class of cells.

To better describe how derivative-like cells integrate the output of the STA and STC filter we fit two different canonical logical operators to the 2D nonlinearity. A cell firing only to positive projections values of both filters performs an AND operation on the two filters. A cell that responds only to positive outputs of the STA and negative outputs of the STC filter performs an AND-NOT operation. We quantified to what degree the two operations were implemented by determining the match between the empirical 2D nonlinearity (Fig. 4b,d) and a template corresponding to either logical operation (Fig. 4h, bottom; see Materials and Methods). This revealed that the nonlinearity was best explained by an

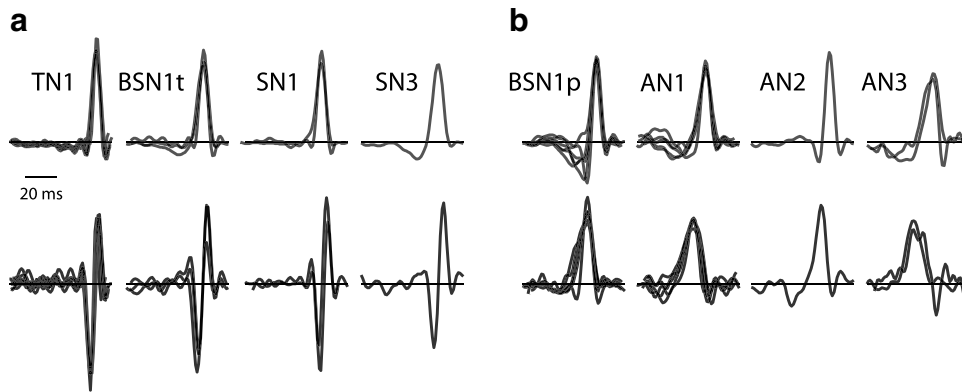


Figure 3. The cell types fall into two classes of model: STA (top row) and STC filters (bottom row) for the cell types in the dataset. Filters had a duration of 64 ms, were aligned at their peaks, and normalized for better visibility. Only the STC filters differed strongly between classes (compare **a** and **b**, lower row) while the STA filters of all cells were highly similar.




Figure 4. Filters and nonlinearities of both classes of cells. **a, c**, Filters of a derivative-like, tonic subtype of BSN1 (BSN1t) (**a**) and a leading-suppressive AN3 (**c**; STA, red; STC, orange; derivative of STA, gray). The STC filter of the cell in **a** resembled the negative derivative of the STA filter. The STC filter in **c** did not look like the negative derivative of the STA filter; however, it led the STA filter. **b, d**, Nonlinearities for the cells in **a** and **c**. Firing rate is color coded (see color bars right to each nonlinearity). **e**, Correlation coefficient between the STC filter and the negative derivative of the STA filter (derivative-like: 0.83 ± 0.11 , leading-suppressive: 0.31 ± 0.08 , $p = 2 \cdot 10^{-11}$, rank sum). Red indicates derivative-like, green leading-suppressive cells in subfigures (**e–j**). **f**, Delay between the peaks of the STA and the STC filter, taking the STA as a reference (derivative-like: 3.0 ± 0.6 ms, leading-suppressive: -7.8 ± 1.9 ms, $p = 4 \cdot 10^{-6}$, rank sum). **g**, Ability of minimal linear (top) and quadratic (bottom) models to explain the performance gain of the STC model relative to the STA model. **h**, Ability of two different canonical logical operators to explain the structure of empirical 2D nonlinearities. Template nonlinearities corresponding to an AND and AND-NOT like integration of the STA and STC filter are shown below. **i, j**, Nonlinearities for the STA and STC filter of both classes of cells (thick lines and shaded area indicate mean \pm SEM over all cells in a class).

AND-like computation on the STA and STC filter. As the STA filter primarily integrated the stimulus and as the STC filter resembled an upstroke of the stimulus (Figs. 3a, lower row, 4a), derivative-like cells thus encoded a combination of the intensity—by means of the STA filter—and the derivative—by means of the STC filter—of a sound’s envelope.

“Leading-suppressive” cells

The phasic subtype of the local neuron BSN1 (BSN1p) and the three ascending neurons, AN1 (Fig. 2), AN2, and AN3, formed the second class of models, which was thus dominated by ascending neurons

(Fig. 3b). In contrast to the derivative-like cells, where the STC filter strongly resembled the derivative of the STA filter, here, both filters were largely independent and covered a longer segment of the stimulus: the STC filter was mostly integrating and led the STA; both filters spanned between 30 and 40 ms of the stimulus (Figs. 3, 4c). This class of cells thus integrated the stimulus on a much longer time scale than the derivative-like cells. Along this line, the great range of delays between both filters (-3 to -12 ms; Fig. 4f) equipped these cells with a more diverse temporal selectivity than the comparatively uniform derivative-like cells.

The nonlinearity of the STA filter resembled that of derivative-like cells in that it also made the cells fire for positive output values (Fig. 4*i*, green). The nonlinearity of the STC filter was unimodal, with a maximum at 0 and a bias toward negative projection values (Fig. 4*j*, green). This means that positive projection values of the STC filter generally suppressed firing. While a minimal linear model of the nonlinearity only accounted for $34 \pm 10\%$ of the performance gain, a minimal quadratic model almost fully explained the performance gain of the STC models ($78 \pm 23\%$; compare Fig. 4*g*, top and bottom). This indicates that the operation of these cells is well described by a quadratic interaction of both filters. The logical operation implemented by the nonlinearity was best approximated by an AND-NOT-like logical operation of the output of both filters (Fig. 4*h*), meaning that these cells fired strongly only for positive projection values of the STA and negative projection values of the STC (Fig. 4*d*). As the peak of STC filter preceded that of the STA, we termed this model class “leading-suppressive cells.”

Contribution of model components to sparse and decorrelated coding of natural stimuli

The analysis of the encoding properties of local and ascending neurons in the grasshopper revealed two different computational classes: the local neurons in our dataset could mostly be termed derivative-like cells, while all ascending neurons were leading-suppressive cells (Figs. 3, 4). To relate the computational properties of both classes of cells to sparse coding we used the STA and STC models to predict responses to a set of natural songs and quantified the contribution of both filters to temporal and population sparseness. The range of response patterns produced by the models matched those found in actual recordings indicating that our models generalized well to natural stimuli (compare model responses in Fig. 5*a,b* with recordings in Fig. 1).

Temporal sparseness

Temporal sparseness describes the tendency to fire in well-defined events interleaved by long stretches of relative quiescence. Generally, STC models responded more transiently to natural signals than STA models (Fig. 5*a,b*, compare red and orange traces). For derivative-like cells, this effect was often subtle, leading to a shortening and downscaling of persistent responses after onsets in the stimulus (Fig. 5*a*, black arrows). In contrast, the responses of STC models of leading-suppressive cells deviated more strongly from those of STA models. Here, the tonic responses of STA models became often purely transient in STC models (Fig. 5*b*, black arrows). To gain an intuition on how the properties of individual firing events—defined as isolated packets or bursts of spikes and evident as segregated peaks in a firing-rate profile—affect temporal sparseness, we constructed artificial firing patterns with varying degrees of response transience by either changing the duration or the magnitude of these events (Fig. 6*a*). This showed that the shortness of firing events

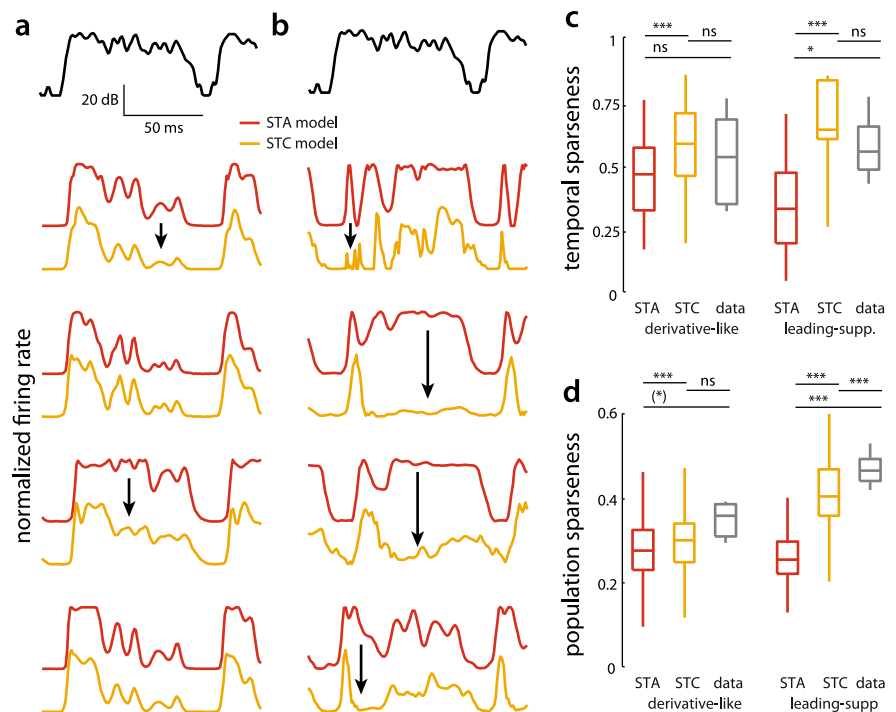


Figure 5. Sparseness of model responses to natural communication signals. *a, b*, Section of the envelope of a grasshopper song (top). Responses of STA (red) and STC models (orange) of four different specimen of derivative-like (*a*) and leading-suppressive cells (*b*). Firing rates were normalized by their maximum for plotting purposes. Black arrows indicate epochs where the response patterns of STA and STC models differed strongly. *c, d*, Temporal (*c*) and population sparseness (*d*) of STA and STC models of both classes. Sparseness values quantified from neural recordings in response to natural songs are shown in gray. *** $p < 0.001$, * $p < 0.05$, (*) $p = 0.06$, n.s. $p \gg 0.05$. Note that p values in the comparison between STA and STC values came from paired tests (sign rank), while the comparison of the STA or STC models and the data was unpaired (rank sum). Hence, the differences between STA and STC models are highly significant despite substantial overlap of the distributions.

strongly correlates with temporal sparseness (Fig. 6*a*, top). We quantified the transience of the modeled responses to natural song as the fraction of time the firing rate was below 50% of its maximum. The firing rate of a persistently firing cell will spend most of the time near the maximum, while that of a transiently firing cell will quickly fall below the half-maximal firing rate. Indeed, STC models of both classes responded much more transiently than STA models when quantified by this measure (STA vs STC: derivative-like 69 ± 25 vs $90 \pm 11\%$, leading-suppressive 60 ± 30 vs $94 \pm 10\%$; STA vs STC model: $p < 6.2 \cdot 10^{-4}$, sign rank). In addition to event shortness, the height of individual firing events also contributed to temporal sparseness (Fig. 6*a*, bottom). We quantified this response property by the 99th percentile of each cell’s firing-rate distribution and found that STC models exhibited higher peak firing rates than STA models (STA vs STC: derivative-like 238 ± 65 Hz vs 329 ± 140 Hz, leading-suppressive 102 ± 49 Hz vs 160 ± 92 Hz; STA vs STC model $p < 6.7 \cdot 10^{-3}$, sign rank).

In accord with their higher transience and peak firing rates, STC models of both classes did exhibit significantly higher temporal sparseness than the respective STA models (Fig. 5*c*). This effect was most prominent for the leading-suppressive cells. This suggests that the higher temporal sparseness of STC models was due to an amplification and shortening of firing events.

To determine how well the models could reproduce the temporal sparseness of the auditory system of grasshoppers for natural signals, we compared the values obtained from the model responses to those measured empirically in a previous study (Clemens et al., 2011). As the class of derivative-like cells was formed




Figure 6. Contribution of model features to temporal sparseness. **a**, Influence of event statistics on temporal sparseness. Firing events were defined as isolated peaks or pulses in a firing rate (see pictograms). Shown is the temporal sparseness of artificial response profiles with varying event duration (the width of a response pulse, top) and event height (the magnitude of the pulse, bottom). **b**, Role of the derivative-like STC filter in shaping temporal sparseness. The STA filter of the original model (tonic BSN1) is shown in gray. Thick red and green lines depict the modified STC filters corresponding to the response shown. Thin lines show intermediate filters. Either the positive (red, top) or the negative lobe (green, bottom) of the filter was scaled down and the temporal sparseness was determined. The response of the original model is shown in black for comparison. Gray arrowheads mark response epochs where the transience of the modified models is smaller than that of the original model; black arrowheads mark those where a secondary firing event is amplified. **c**, Influence of either the positive (red) or negative lobe (green) on the duration of the primary firing event (top, gray triangle in **b**), on the height of the secondary event (middle, black triangle in **b**), and on temporal sparseness (bottom). *x*-axis shows the log of the scaling factor. Positive values correspond to a reduction of the positive lobe, negative values to that of the negative lobe. Examples in **b** show responses for log scaling factors of -8 (top) and 8 (bottom). Temporal sparseness of the STA model shown as a black dot (bottom). **d**, Impact of the STC filter of a leading-suppressive cell (AN1) on temporal sparseness of responses to different onsets. The stimulus with successively decreasing onsets is shown on top and the output of the STA (red) and STC filter (orange) in the middle. The bottom graph depicts the output of the STC model. Temporal sparseness increased as a function of onset height (inset).

by local neurons we compared the sparseness in these models to that found in local neurons. The sparseness values of leading-suppressive models, being dominated by ascending neurons, were compared with those of ascending neurons. STA and STC models of derivative-like cells well reproduced the range of temporal sparseness values found in actual recordings of natural songs (Fig. 5c; STA 0.44 ± 0.17 , STC 0.54 ± 0.19 , data 0.52 ± 0.19). In contrast, temporal sparseness of leading-suppressive cells depended on model type. While STA models underestimated the temporal sparseness found in data, the STC model exhibited no significant difference in temporal sparseness when compared with experimental data (Fig. 5c, right; STA 0.34 ± 0.18 , STC 0.63 ± 0.23 , data 0.57 ± 0.12). Note that the STC models tended to exhibit slightly higher temporal sparseness than observed in the data.

What property of the STC model of both classes of model contributed to temporal sparseness? The STC filter of derivative-like cells employs these cells with a sensitivity to the derivative of the stimulus and presumably accentuates responses to onsets. Distorting the derivative-like STC filter by scaling down either its positive or negative lobe prolonged firing events (Fig. 6b, gray arrowheads, *c*, top) and amplified small firing events (Fig. 6b, black arrowheads, *c*, middle). Models with distorted filters thus exhibited longer firing events, less transient responses, and hence reduced temporal sparseness (Fig. 6c, bottom). In contrast, changing the delay between the STA and the STC filter had a negligible impact on temporal sparseness (Fig. 7a–c). This indicates that the differentiating shape of the STC filter in conjunction with the AND-like integration of both filters led to the increase in temporal sparseness in derivative-like cells.

For the leading-suppressive cells, it was the AND-NOT like interaction of the STA and the STC filter that increased the transience of responses. The fact that the STC filter led the STA filter and suppressed firing made these cells respond strongly only to

onsets (Fig. 6d). The “inhibitory” STC filter responded tonically to stimuli with constant amplitude due to its integrating shape (Figs. 3b, bottom for the shape of STC filters, 6d, middle). This tonic suppression effectively reduced strong persistent activity, increased the height of firing events (Fig. 6d, bottom) and thereby enhanced temporal sparseness in this class of cells (Fig. 6d, inset). The delay between both filters also influenced temporal sparseness (Fig. 7g). While making the delay in the model more negative increased temporal sparseness by 10–15%, making the delay less negative tended to decrease sparseness.

Population sparseness

We calculated population sparseness by constructing four-cell populations of random combinations of models belonging to the same class. Population sparseness is high if different cells in a population do respond with different patterns to the same signal, e.g., by being selective for different stimulus features.

The model responses reveal a fundamental difference in the impact of the STC filter on response diversity (Fig. 5a,b). Derivative-like cells (Fig. 5a) exhibited relatively uniform responses with high firing rates during onsets. The STC filter did little to add to response diversity. This comes to no surprise as both filters were relatively uniform in members of this class (Figs. 3a, 4e,f). In contrast, while STA models of leading-suppressive cells also responded relatively uniformly to song, different STC models of this class exhibited a greater diversity of response patterns to the same stimulus (Fig. 5b), comparable to that found in actual recordings of ascending neurons (compare Fig. 1).

Consistent with the similarity of the STA filters, populations of STA models of both classes displayed comparable levels of population sparseness (Fig. 5d; derivative-like STA 0.28 ± 0.06 , leading-suppressive STA 0.26 ± 0.06). However, while the second filter did increase population sparseness only marginally in derivative-like cells, leading-suppressive cells profited greatly

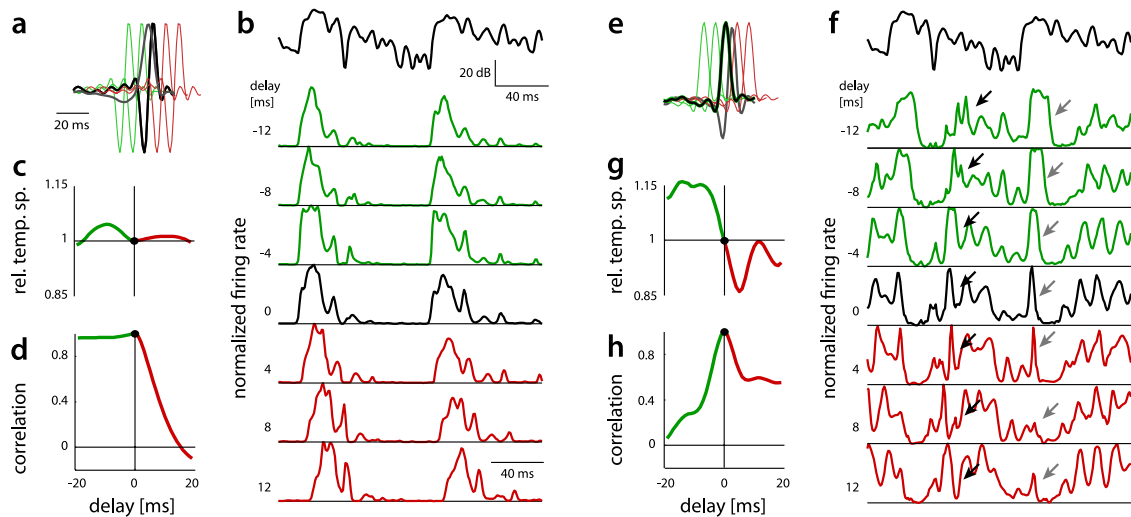


Figure 7. Impact of the delay between the STA and the STC filter on response diversity and temporal sparseness. Shown are the results for a derivative-like tonic BSN1t (*a–d*) and a leading-suppressive phasic BSN1p (*e–h*). Green marks negative, red positive temporal shifts of the STC filter. *a, e*, The original model with the STA filter (gray) and the STC filter (black) was manipulated by increasing or decreasing the delay between both filters (red and green STC filters for added delays of -8 , -4 , 4 , and 8 ms). *b, f*, Responses of the original model (0 ms delay, black) and those with positive (red) and negative (green) delays were obtained. Responses were normalized by the maximal response for plotting purposes. The stimulus is shown in black on top. Black and gray arrows in *f* indicate epochs where the delay between both filters strongly affects the response patterns. *c, g*, Relative changes in temporal sparseness as a function of the delay added. Reference is the sparseness of the unmodified STC model (green, negative delays; red, positive delays; black, original model). *d, h*, Response diversity as a function of the delay added was quantified by the correlation coefficient between the original STC model (0 ms delay) and the modified one.

from the inclusion of the second filter (STA vs STC: derivative-like 0.28 ± 0.06 vs 0.32 ± 0.06 , leading-suppressive 0.26 ± 0.06 vs 0.42 ± 0.08). While STA models of both classes and STC models of the derivative-like models exhibited population sparseness comparable to those obtained empirically for local neurons (0.35 ± 0.05), only the 2D STC models of leading-suppressive cells approached the high values reported previously for the output of the network (0.47 ± 0.03).

Why were only leading-suppressive cells but not derivative-like cells able to significantly increase population sparseness? We have shown above that leading-suppressive cells exhibited a wide range of delays between the STA and STC filter, while those of the derivative-like cells had relatively uniform delays (Fig. 4*f*). If the small range of delays limits response diversity of derivative-like cells, then artificially increasing this range should increase response diversity. Interestingly, this did not truly increase response diversity but only shifted a uniform firing event in time (Fig. 7*a,b,d*). We thus hypothesized that the derivative-like shape of the filter is the factor limiting response diversity. Indeed, distorting the shape of the STC filter by reducing its positive or negative lobe filter appeared to increase response diversity while reducing temporal sparseness (Fig. 6*b*). Hence, a derivative-like filter in combination with a small range of delays seems to be ineffective in increasing population sparseness.

In contrast to the small impact of the delay between the STA and the STC filter on response diversity in derivative-like cells, the delay had strong impact on response diversity in leading-suppressive cells. Systematically varying the delay between the STA and the STC filter altered the response patterns for these types of cells (Fig. 7*e,f*). The response to the onset of the stimulus became longer with increasingly negative delays (Fig. 7*f*, gray arrows). Positive delays fully abolished any onset response. The firing-rate modulations after the onset also changed with delay (Fig. 7*f*, black arrows). Hence, the range of delays between the STA and the STC filter decorrelated responses and increased population sparseness in leading-suppressive cells.

Discussion

Employing the framework of LN models, we found two classes of cells in the auditory system of grasshoppers. While the STA filter was similar for both classes, cells differed in their second filter: models with a derivative-like STC filter and an AND-like nonlinearity were found at the level of local neurons and models with a leading-suppressive STC filter and an AND-NOT like nonlinearity were found mainly among ascending neurons (Figs. 3, 4).

Our simulations have shown that only 2D models produce the degree of temporal and population sparseness found in the auditory system of the grasshopper (Fig. 5). While both, derivative-like and leading-suppressive cells increased temporal sparseness—though to different degrees—only the latter class of cells substantially increased population sparseness.

In the following, we discuss how the structure of both classes of models increases sparseness. In addition, we will use prior knowledge about the grasshopper to speculate on likely biophysical mechanisms, which could implement these computations.

Note that not all cells in the early auditory system of grasshoppers are included in our dataset. Other cells in network probably perform different computations, adding to the response diversity (Römer and Marquart, 1984; Stumpner and Ronacher, 1991); e.g. the ascending neuron AN4 receives fast inhibitory inputs, which leads to a suppression of responses at onsets (Fig. 7*f*). Another ascending neuron not included in the dataset is AN14, which fires only during silent parts of a stimulus.

Temporal sparseness

Temporal sparseness increases if transient firing is accentuated and persistent firing is attenuated, leading to responses with short firing events interleaved by long silent epochs (Figs. 5*a,b*, 6*a*). The two model classes achieve this transformation by two different computations. The derivative-like STC filter of most local neurons leads to a differentiation of the stimulus (Figs. 6*b*, 7*b*). The leading-suppressive STC filter of the ascending neurons quenches prolonged responses (Fig. 6*d*).

Both operations can be subsumed under the phenomenon of spike-frequency adaptation, the decrease of neuronal firing in response to prolonged stimulation (Benda et al., 2001; Lundstrom et al., 2008; Tripp and Eliasmith, 2010). The relation between spike-frequency adaptation and temporal sparseness has been reported previously (Farkhooi et al., 2009; Houghton, 2009; Nawrot, 2012). Adaptation has been described in terms of differentiation (Lundstrom et al., 2008) or as a highpass filter (Benda et al., 2001), transformations that reduce temporal correlations and thereby increase temporal sparseness (Wang et al., 2003; Tripp and Eliasmith, 2010). The STC filter in derivative-like cells thus likely implements adaptation by determining the derivative of the stimulus (Figs. 3*a*, 6*b*). Derivative-like, 2D models as described in our study have been found in many sensory systems (Brenner et al., 2000; Slee et al., 2005; Fairhall et al., 2006; Atencio et al., 2008; Fox et al., 2010; Kim et al., 2011; Sharpee et al., 2011), suggesting that this model structure instantiates—in addition to its contribution to temporal sparseness—beneficial properties, like adaptation to stimulus statistics and robust and efficient encoding of time-varying stimuli (Fairhall et al., 2001; Sharpee et al., 2011).

Adaptation can be implemented by cell-intrinsic mechanisms via adaptation currents (Wang et al., 2003) or in a network via synaptic depression, feedback inhibition (Papadopoulou et al., 2011), and slow feedforward inhibition (Assisi et al., 2007; Creutzig et al., 2009). In our system, the derivative-like STC filter is likely to be implemented by cell-intrinsic adaptation currents, especially for those cells (TN1, SN1), which receive only excitatory inputs from receptors (Römer and Marquart, 1984). In two cells (BSN1t and SN3), the derivative-like filter could be shaped by additional inhibitory inputs.

Note that a derivative-like STC filter can also be the result of imprecise spike-timing (Dimitrov et al., 2006). However, this is more likely if such jitter is greater than the time scale of the filter (Fairhall et al., 2006; Sharpee et al., 2011). We calculated spike-time jitter as the SD of the timing of individual firing events across trials as in Desbordes et al. (2008). The jitter was always much smaller than the filter width, for derivative-like 10 times and for leading-suppressive cells 5 times smaller (jitter: derivative-like 0.6 ± 0.4 ms, leading-suppressive 1.3 ± 0.5 ms; filter width: derivative-like 5.3 ± 0.9 ms, leading-suppressive 6.8 ± 1.7 ms) (Figs. 2*e*, 3). This renders jitter an unlikely source of the derivative-like filter.

Although the filters of derivative-like cells are relatively similar across cells, the temporal sparseness values cover a relatively broad range (Fig. 5*c*). This reflects a great diversity in the fine structure of the nonlinearity, which enables different degrees of selectivity and hence temporal sparseness. While all nonlinearities implemented an AND-like integration of the STA and STC filter (Fig. 4*h*), the thresholds of individual cells of this class ranged between 45 and 65 dB. This great diversity of thresholds increases the dynamic range of the network and justifies the existence of many different types of local neurons with highly similar filters (Fig. 3*a*).

In the leading-suppressive cells, temporal sparseness is increased by shutting off persistent responses via slow suppression. The two properties of the leading-suppressive models contributing to this transformation are the delay between the STA and the STC filter (Figs. 4*e*, 7*e–f*) and the AND-NOT like nonlinearity (Fig. 4*h*), which are possibly implemented in a network: the STA filter corresponds to excitatory inputs and drives the cell; as the STC filter leads the STA and is suppressive, the cell will only fire strongly if the stimulus preceding the STA is relatively soft. Such an implementation is highly likely for the phasic BSN1, AN1, and

AN3, for which strong, slow inhibitory inputs have been shown in dendritic recordings (Römer and Marquart, 1984; Hildebrandt et al., 2009). For another leading-suppressive cell type in our dataset—AN2—a strong afterhyperpolarization has been shown to underlie adaptation (Hildebrandt et al., 2009); yet, this cell also receives contralateral inhibition that can be slower than the excitation. The STC filter of this cell type is thus likely to be a combination of both cell-intrinsic adaptive currents and inhibitory inputs.

That the most likely mechanisms contributing to temporal sparseness are adaptation currents in derivative-like cells and inhibitory and excitatory inputs in leading-suppressive cells suggests a change in the factors dominating response properties of neurons in subsequent stages. The derivative-like cells were local neurons, which primarily pool the responses of receptors and probably gain their sensitivity to the derivative via a cell-intrinsic mechanism. In contrast, most leading-suppressive cells were ascending neurons; their computations are most likely shaped by the connectivity with local neurons and thus by network properties.

Population sparseness

Our results show that differentiation of the stimulus and slow inhibition increase temporal sparseness by reducing persistent firing (Figs. 5, 6). This property in itself does not necessarily lead to population sparseness. For population sparseness to be high, cells in a population need to exhibit little tendency to fire together by being selective for different features of a stimulus.

The ability of derivative-like filters to increase population sparseness was relatively small (Fig. 5*d*). As the STA filters were similar and the STC filter was heavily constrained by the shape of the STA filter—on average 83% of the STC filter's shape of each cell was explained by the STA filter—this second filter added little response diversity across cells (Figs. 3*a*, 4*e*). Accordingly, the derivative-like cells exhibited very similar feature selectivity and responded uniformly to a stimulus.

In contrast, the STC filter of leading-suppressive cells strongly increased population sparseness—up to the values observed in the auditory system of the grasshopper (Fig. 5*d*). A highly diverse feature selectivity in these cells is established through a large range of delays between the excitatory STA filter and the suppressive STC filter (Fig. 4*f*). As argued above, the model structure of leading-suppressive cells is probably generated by slow feedforward inhibition (Luo et al., 2010). The role of excitation and inhibition in shaping temporal filters and in decorrelating responses between cells in a population has been appreciated previously (Schmuker and Schneider, 2007; Wiechert et al., 2010; George et al., 2011). In addition to the filters of leading-suppressive cells being diverse, the AND-NOT like joint nonlinearity equips these models with a highly nonlinear operation to select a small set of stimuli suitable for firing (Fig. 4*d,h*). This narrows the tuning of leading-suppressive cells and reduces the overlap between responses of different cells of this type. The AND-NOT like computation also leads to a “delayed anti-coincidence detection”—the cells fire strongly only if the stimulus at different delays is not loud (Borst et al., 2005). This can yield a combinatorial and synergistic code (Osborne et al., 2008; Schneiderman et al., 2011).

Conclusion

Not diversity and changes in the linear STA filter but nonlinear integration of two filters governs the transformation of the code from a dense and uniform one to a temporally and population-sparse one in the grasshopper (Pitkow and Meister, 2012). The

shape of the second filter equips neurons with transformations that decorrelate responses in time and across cells in a population. Additionally, a 2D nonlinearity allows neurons to specifically select a small subset of the feature space spanned by the STA and the STC filter. Mechanisms implementing these abstract computations are ubiquitous in many nervous systems; the transformations found in the grasshopper are thus likely to constitute general principles underlying the transformation of neural representations.

References

- Assisi C, Stopfer M, Laurent G, Bazhenov M (2007) Adaptive regulation of sparseness by feedforward inhibition. *Nat Neurosci* 10:1176–1184.
- Atencio CA, Sharpee TO, Schreiner CE (2008) Cooperative nonlinearities in auditory cortical neurons. *Neuron* 58:956–966.
- Barlow H (2001) Redundancy reduction revisited. *Network* 12:241–253.
- Benda J, Bethge M, Hennig RM, Pawelzik KR, Herz AV (2001) Spike-frequency adaptation: phenomenological model and experimental tests. *Neurocomputing* 38:105–110.
- Borst A, Flanagan VL, Sompolinsky H (2005) Adaptation without parameter change: dynamic gain control in motion detection. *Proc Natl Acad Sci U S A* 102:6172–6176.
- Brenner N, Bialek W, de Ruyter van Steveninck R (2000) Adaptive rescaling maximizes information transmission. *Neuron* 26:695–702.
- Chacron MJ, Longtin A, Maler L (2011) Efficient computation via sparse coding in electrosensory neural networks. *Curr Opin Neurobiol* 21:752–760.
- Clemens J, Kutzki O, Ronacher B, Schreiber S, Wohlgenuth S (2011) Efficient transformation of an auditory population code in a small sensory system. *Proc Natl Acad Sci U S A* 108:13812–13817.
- Creutzig F, Wohlgenuth S, Stumpner A, Benda J, Ronacher B, Herz AV (2009) Timescale-invariant representation of acoustic communication signals by a bursting neuron. *J Neurosci* 29:2575–2580.
- Desbordes G, Jin J, Weng C, Lesica NA, Stanley GB, Alonso JM (2008) Timing precision in population coding of natural scenes in the early visual system. *PLoS Biol* 6:e324+.
- Dimitrov AG, Gedeon T (2006) Effects of stimulus transformations on estimates of sensory neuron selectivity. *J Comput Neurosci* 20:265–283.
- Fairhall AL, Lewen GD, Bialek W, de Ruyter Van Steveninck RR (2001) Efficiency and ambiguity in an adaptive neural code. *Nature* 412:787–792.
- Fairhall AL, Burlingame CA, Narasimhan R, Harris RA, Puchalla JL, Berry MJ 2nd (2006) Selectivity for multiple stimulus features in retinal ganglion cells. *J Neurophysiol* 96:2724–2738.
- Farkhooi F, Muller E, Nawrot MP (2009) Sequential sparsing by successive adapting neural populations. *BMC Neurosci* 10 [Suppl 1]:O10.
- Fitzgerald JD, Sincich LC, Sharpee TO (2011) Minimal models of multidimensional computations. *PLoS Comput Biol* 7:e1001111.
- Fox JL, Fairhall AL, Daniel TL (2010) Encoding properties of haltere neurons enable motion feature detection in a biological gyroscope. *Proc Natl Acad Sci U S A* 107:3840–3845.
- George AA, Lyons-Warren AM, Ma X, Carlson BA (2011) A diversity of synaptic filters are created by temporal summation of excitation and inhibition. *J Neurosci* 31:14721–14734.
- Gollisch T, Schütze H, Benda J, Herz AV (2002) Energy integration describes sound-intensity coding in an insect auditory system. *J Neurosci* 22:10434–10448.
- Hildebrandt KJ, Benda J, Hennig RM (2009) The origin of adaptation in the auditory pathway of locusts is specific to cell type and function. *J Neurosci* 29:2626–2636.
- Houghton C (2009) Studying spike trains using a van Rossum metric with a synapse-like filter. *J Comput Neurosci* 26:149–155.
- Kim AJ, Lazar AA, Slutskiy YB (2011) System identification of *Drosophila* olfactory sensory neurons. *J Comput Neurosci* 30:143–161.
- Lundstrom BN, Higgs MH, Spain WJ, Fairhall AL (2008) Fractional differentiation by neocortical pyramidal neurons. *Nat Neurosci* 11:1335–1342.
- Luo SX, Axel R, Abbott LF (2010) Generating sparse and selective third-order responses in the olfactory system of the fly. *Proc Natl Acad Sci U S A* 107:10713–10718.
- Machens CK, Stemmler MB, Prinz P, Krahe R, Ronacher B, Herz AV (2001) Representation of acoustic communication signals by insect auditory receptor neurons. *J Neurosci* 21:3215–3227.
- Nawrot MP (2012) Dynamics of sensory processing in the dual olfactory pathway of the honeybee. *Apidologie* 43:1–23.
- Neuhöfer D, Wohlgenuth S, Stumpner A, Ronacher B (2008) Evolutionarily conserved coding properties of auditory neurons across grasshopper species. *Proc R Soc B* 275:1965–1974.
- Olshausen BA, Field DJ (2004) Sparse coding of sensory inputs. *Curr Opin Neurobiol* 14:481–487.
- Osborne LC, Palmer SE, Lisberger SG, Bialek W (2008) The neural basis for combinatorial coding in a cortical population response. *J Neurosci* 28:13522–13531.
- Papadopoulou M, Cassenaer S, Nowotny T, Laurent G (2011) Normalization for sparse encoding of odors by a wide-field interneuron. *Science* 332:721–725.
- Petersen RS, Brambilla M, Bale MR, Alenda A, Panzeri S, Montemurro MA, Maravall M (2008) Diverse and temporally precise kinetic feature selectivity in the VPm thalamic nucleus. *Neuron* 60:890–903.
- Pitkow X, Meister M (2012) Decorrelation and efficient coding by retinal ganglion cells. *Nat Neurosci* 15:628–635.
- Rokem A, Watzl S, Gollisch T, Stemmler M, Herz AV, Samengo I (2006) Spike-timing precision underlies the coding efficiency of auditory receptor neurons. *J Neurophysiol* 95:2541–2552.
- Römer H, Marquart V (1984) Morphology and physiology of auditory interneurons in the metathoracic ganglion of the locust. *J Comp Phys A* 155:249–262.
- Ronacher B, Stumpner A (1988) Filtering of behaviourally relevant temporal parameters of a grasshopper's song by an auditory interneuron. *J Comp Phys A* 163:517–523.
- Rust NC, Schwartz O, Movshon JA, Simoncelli EP (2005) Spatiotemporal elements of macaque v1 receptive fields. *Neuron* 46:945–956.
- Schmuker M, Schneider G (2007) Processing and classification of chemical data inspired by insect olfaction. *Proc Natl Acad Sci U S A* 104:20285–20289.
- Schneidman E, Puchalla JL, Segev R, Harris RA, Bialek W, Berry MJ 2nd (2011) Synergy from silence in a combinatorial neural code. *J Neurosci* 31:15732–15741.
- Schwartz O, Pillow JW, Rust NC, Simoncelli EP (2006) Spike-triggered neural characterization. *J Vis* 6:484–507.
- Sharpee TO, Nagel KI, Doupe AJ (2011) Two-dimensional adaptation in the auditory forebrain. *J Neurophysiol* 106:1841–1861.
- Slee SJ, Higgs MH, Fairhall AL, Spain WJ (2005) Two-dimensional time coding in the auditory brainstem. *J Neurosci* 25:9978–9988.
- Stumpner A (1989) Physiological variability of auditory neurons in a grasshopper. *Naturwissenschaften* 76:427–429.
- Stumpner A, Ronacher B (1991) Auditory interneurons in the metathoracic ganglion of the grasshopper *Chorthippus biguttulus*: I. Morphological and physiological characterization. *J Exp Biol* 158:391–410.
- Tripp BP, Eliasmith C (2010) Population models of temporal differentiation. *Neural Comput* 22:621–659.
- Vogel A, Hennig RM, Ronacher B (2005) Increase of neuronal response variability at higher processing levels as revealed by simultaneous recordings. *J Neurophysiol* 93:3548–3559.
- Wang XJ, Liu Y, Sanchez-Vives MV, McCormick DA (2003) Adaptation and temporal decorrelation by single neurons in the primary visual cortex. *J Neurophysiol* 89:3279–3293.
- Wiechert MT, Judkewitz B, Riecke H, Friedrich RW (2010) Mechanisms of pattern decorrelation by recurrent neuronal circuits. *Nat Neurosci* 13:1003–1010.
- Willmore BD, Tolhurst DJ (2001) Characterizing the sparseness of neural codes. *Network* 12:255–270.
- Willmore BD, Mazer JA, Gallant JL (2011) Sparse coding in striate and extra-striate visual cortex. *J Neurophysiol* 105:2907–2919.