

Efficient transformation of an auditory population code in a small sensory system

Jan Clemens^{a,b,1}, Olaf Kutzki^a, Bernhard Ronacher^{a,b}, Susanne Schreiber^{b,c,2}, and Sandra Wohlgenuth^{a,2}

^aBehavioral Physiology Group, Department of Biology, Humboldt-Universität zu Berlin, 10115 Berlin, Germany; ^bBernstein Center for Computational Neuroscience Berlin, 10115 Berlin, Germany; and ^cInstitute for Theoretical Biology, Department of Biology, Humboldt-Universität zu Berlin, 10115 Berlin, Germany

Edited by Terrence J. Sejnowski, Salk Institute for Biological Studies, La Jolla, CA, and approved July 12, 2011 (received for review March 23, 2011)

Optimal coding principles are implemented in many large sensory systems. They include the systematic transformation of external stimuli into a sparse and decorrelated neuronal representation, enabling a flexible readout of stimulus properties. Are these principles also applicable to size-constrained systems, which have to rely on a limited number of neurons and may only have to fulfill specific and restricted tasks? We studied this question in an insect system—the early auditory pathway of grasshoppers. Grasshoppers use genetically fixed songs to recognize mates. The first steps of neural processing of songs take place in a small three-layer feed-forward network comprising only a few dozen neurons. We analyzed the transformation of the neural code within this network. Indeed, grasshoppers create a decorrelated and sparse representation, in accordance with optimal coding theory. Whereas the neuronal input layer is best read out as a summed population, a labeled-line population code for temporal features of the song is established after only two processing steps. At this stage, information about song identity is maximal for a population decoder that preserves neuronal identity. We conclude that optimal coding principles do apply to the early auditory system of the grasshopper, despite its size constraints. The inputs, however, are not encoded in a systematic, map-like fashion as in many larger sensory systems. Already at its periphery, part of the grasshopper auditory system seems to focus on behaviorally relevant features, and is in this property more reminiscent of higher sensory areas in vertebrates.

metric | invertebrates | information theory

To increase their fitness, most animals strive to evaluate sensory signals that reveal the quality of a potential mate. What if an animal has only a few dozen neurons to preprocess this extremely important information? Optimal coding theory suggests that the creation of a sparse, decorrelated representation would be a wise investment of scarce neuronal resources (1). That is indeed what has been found in many sensory modalities and species under natural conditions (2–4).

These early sensory networks often comprise large numbers of cells and organize information in a map-like fashion, where spatial proximity of neurons reflects similarity in the selectivity for fundamental stimulus features (5, 6). These maps tend to have a complete representation of sensory space and enable subsequent processing steps to select relevant features based on attention or associative learning. Reading out such a representation by “blind” summation of responses across different neurons would be highly inefficient to recover information, because stimulus features are not only encoded by neuronal activity per se but also by neuronal identity. This type of population code is referred to as labeled-line code (7, 8). Accordingly, higher-order sensory areas need to take into account which neurons are active when producing more specific representations of behaviorally relevant stimulus aspects (9).

Do the principles derived by optimal coding theory also apply to networks with relatively few neurons and a restricted set of relevant stimuli? Here we investigate this question in the auditory periphery of grasshoppers. These insects produce genetically fixed songs to recognize and evaluate potential mates with high fidelity (10). The

involved processing stages comprise a feed-forward network of only three layers in the grasshopper’s metathoracic ganglion: 60 receptors per side faithfully encode the signal’s envelope (11, 12) and form the input stage; receptors project onto an intermediate layer of ~15 local neurons; these in turn connect to the output layer of ~20 ascending neurons (13, 14) (Fig. 1). The output of this size-constrained network is the only source of acoustic information available to the behavioral decision centers in the brain.

What transformations does the neural representation of grasshopper song undergo in this small sensory system? Are these transformations similar to those found in larger sensory systems? We find that the neural representation changes profoundly across neuronal layers: Sparseness and decorrelation of responses increase—just as in more complex systems and in accord with optimal coding theory. In the third layer, neuronal identity becomes crucial for an effective readout of the population. We show that within just two processing steps a labeled-line code is formed from a uniform representation of the stimulus at the input layer. This labeled-line code includes explicit representations of behaviorally relevant stimulus features at a surprisingly early stage of the auditory pathway and presumably does not provide the complete representation of stimulus space found in the periphery of many larger sensory systems.

Results

We recorded intracellularly from identified, single neurons at three consecutive stages within the metathoracic ganglion of migratory locusts. Auditory stimuli consisted of eight songs from representative male individuals of the species *Chorthippus biguttulus*, whose auditory system is highly homologous to that of locusts at the stages considered here (13, 15, 16). Song recognition in these animals is based on the signal’s envelope. Carrier frequency is largely irrelevant, as the frequency resolution of the auditory system is weak (17). The envelope of the grasshopper’s calling song is composed of several repetitions of a basic subunit: a syllable-pause pair. We evaluated responses to two syllable-pause pairs corresponding to the minimal signal duration necessary for song recognition by male grasshoppers (18). Despite their high homogeneity, these songs are well-discriminated at the level of auditory receptors (19) and in behavioral tests (20).

Lifetime Sparseness Increases and Reproducibility Decreases Within the Network. First, we analyzed properties of single-neuron responses. Example spike trains of different receptors and local

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

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. E-mail: clemensjan@googlemail.com.

²S.S. and S.W. contributed equally to this work.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1104506108/-DCSupplemental.

(Fig. 2*F*; receptors $61 \pm 24\%$, local neurons $76 \pm 13\%$, ascending neurons $43 \pm 15\%$). This measure of classification success was highly correlated with information ($r^2 = 0.97$; Fig. S2), but yielded significant differences only between local and ascending neurons ($P = 0.003$; all other pairs $P > 0.09$, t test).

Remarkably, information about stimulus identity retrievable from individual neurons is lowest at the output layer of the network. This is partly explained by their comparatively low firing rate (Fig. 2*A* and *E*).

Ascending Neurons Decorrelate the Neural Representation of Song.

The higher reproducibility of neuronal responses in the first two layers (receptors and local neurons) would enable the animals to discriminate individual songs better than the noisy responses of ascending neurons in the output layer (Fig. 2*D*) (12, 26). This may at first glance seem paradoxical. Obviously, the network is not enhancing the *raw* stimulus discriminability at the level of individual neurons. A likely explanation is that the population code for song undergoes a transformation at the level of ascending neurons: Whereas local neurons might constitute a relatively homogeneous population—where every neuron encodes largely the same information—ascending neurons possibly use a distributed representation, where different neurons encode different aspects of the stimulus. This might facilitate subsequent processing steps—at the expense of lower *single-cell* information about song discriminability.

A comparison of the firing-rate functions of cells within a layer suggests that this is indeed true (Fig. 1): Responses to a given stimulus were diverse across different ascending neurons, whereas responses among the groups of receptor and local neurons appeared to be more similar. We quantified the similarity of responses (same datasets as for the information estimation) by calculating Pearson's correlation coefficient between pairs of cells within each layer. The average response similarity steadily declined within the network from 0.58 ± 0.14 in receptors over 0.42 ± 0.07 in local neurons to 0.09 ± 0.16 in ascending neurons (Fig. 3*A*; $P < 4.2 \times 10^{-4}$, rank-sum test). This decorrelation was accompanied by an increase of population sparseness across the network layers (Fig. 3*B*; receptors 0.19 ± 0.04 , local neurons 0.35 ± 0.05 , ascending neurons 0.47 ± 0.03 ; $P < 5 \times 10^{-7}$, rank-sum test). Population sparseness measures to what extent only a few neurons in a population are active at any time. This is not to be confused with temporal or lifetime sparseness, which quantifies how sparsely a single neuron fires over time (Fig. 2*B*). Although the population sparseness of responses does not reach extreme values as, for example, in the olfactory system of locusts (e.g., 4, with values >0.9), ascending neurons in the auditory system can still be considered to form a decorrelated and more sparse representation of the song (compare with ref. 2).

The decorrelation of responses in the output layer—that is, the fact that different types of ascending neurons fired at different times during the stimulus—suggests also that feature selectivity may be more diverse. To show this directly, we calculated each neuron's spike-triggered average (STA) filter—the average stimulus preceding a spike—using the same set of spikes as before. Neurons at all three layers were previously well-described by simple linear–nonlinear models (e.g., 11). As the natural songs are strongly non-Gaussian (11), the STA does not represent an unbiased estimate of the neuron's filter (27). However, it allows us to assess the diversity of feature selectivity at each processing stage.

Fig. 3*D–F* shows the filters of different receptors and different types of local and ascending neurons. Evidently, the variety of stimulus features eliciting spikes increased across the three neuronal layers: All receptors and local neurons had almost identical, unimodal STAs (Fig. 3*D* and *E*, blue and red lines and Fig. 3*C*, mean correlation between pairs of filters 0.91 ± 0.08 and 0.81 ± 0.09 , respectively) and varied only little around the population average (Fig. 3*D* and *E*). In contrast, the filters of dif-

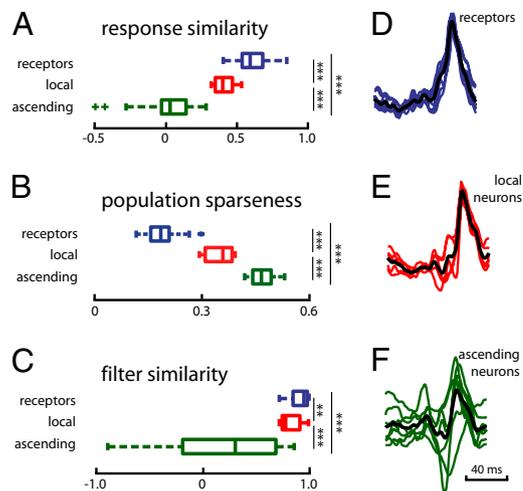


Fig. 3. Population sparseness, diversity of response pattern, and feature selectivity increase for ascending neurons. (A) Response similarity: correlation coefficients between firing-rate functions of pairs of neurons within each stage. (B) Population sparseness. (C) Filter similarity: correlation coefficients between the spike-triggered average filters of pairs of neurons within each stage. $**P > 0.01$, $***P > 0.001$. (D–F) Colored lines depict the normalized filters for individual cell types. The thick black line shows the average over all cell types in a population. Blue, receptors ($n = 10$ cells); red, local neurons ($n = 21$ cells of five types); green, ascending neurons ($n = 25$ cells of seven types).

ferent types of ascending neurons were highly dissimilar: Some resembled those of the local neurons, whereas others exhibited strong negative components (Fig. 3*F*, green lines; correlation between pairs of filters 0.22 ± 0.55). As the STAs of ascending neurons were specific to cell type (Fig. S3 and *SI Results*), we conclude that different types of ascending neurons encode different aspects of the stimulus.

Ascending Neurons Profit Most from a Multineuron Decoder. What is the consequence of the observed decorrelation at the network's output layer for the population code? In ascending neurons, information about song discriminability is likely to be distributed across different cells. As we saw in the analysis of STAs (Fig. 3*F*), individual neurons in this layer have a stronger tendency to encode different aspects of the stimulus. In this case, blind pooling (summation) across neurons is likely to destroy valuable information, whereas knowing which cell fired which spike would be highly informative for a population readout. Hence, a neuronal representation where not only the occurrence of spikes but also their neuronal identity matters is formed. Such a code corresponds to a labeled-line or distributed code (7, 28).

In contrast, we expect neuronal identity to play only a minor role at the first two processing stages. Here, firing patterns and STAs within one stage (i.e., among receptors or among local neurons, respectively) are highly similar (Fig. 3*D* and *E*). Hence, pooling across neurons might increase information by enhancing the signal-to-noise ratio (29). We refer to codes where neuronal identity does not contribute to an optimal readout as summed-population codes, following refs. 7 and 28.

We tested how these two types of population codes perform for the decoding of song identity in the three processing layers. To this end, we implemented two decoders on the basis of the multineuron metric proposed by ref. 28: one which disregards information about neuronal identity of spikes by effectively summing up responses in a population—a “summed-population decoder”—and one which preserves this information—a “labeled-line decoder” (28) (Fig. S4 and *SI Materials and Methods*). We constructed populations of four cells, each consisting of either four

receptors, or local, or ascending neurons (*Materials and Methods*). For each of these populations, we decoded stimulus identity from responses using both decoders. First, we analyzed how much the readout of cell populations versus single cells can improve the discriminability of song. Second, we determined which of the two population decoders was able to provide more information about stimulus discriminability in each layer.

Single-cell versus population decoding. Overall, information increased significantly for populations of receptors and ascending neurons compared with single cells ($P < 7 \times 10^{-4}$, rank-sum test), but not for populations of local neurons ($P = 0.10$, rank-sum test; average information rate of populations of receptors 13.1 ± 1.2 bit/s, local neurons 13.0 ± 0.6 bit/s, ascending neurons 7.9 ± 1.6 bit/s). Evaluation of the performance of the population decoder in percent correct showed similar trends (receptors $93 \pm 5\%$, local neurons $94 \pm 2\%$, ascending neurons $72 \pm 8\%$).

Although information increased for populations in all three processing stages, a true information gain in the population readout compared with a single-cell readout can only be expected if the population decoder provides more information than the best single cell within that population. We hence calculated the information ratio between the population as decoded by the better of the two multineuron metrics (determined individually for each population) and the best individual cell in this population (in Fig. S5A and *SI Results*, we provide results for the gain with respect to the average information rate of individual neurons). Gain values significantly greater than 1.0 signal a net increase of information by reading out the population compared with the best single cell.

Gain was intermediate for receptors (Fig. 4A; 1.16 ± 0.11 ; $P < 4 \times 10^{-11}$, sign test against 1.0) and not significantly different from 1.0 for local neurons (1.05 ± 0.10 ; $P < 0.38$). Presumably, these two populations show very moderate gains because already the best single receptors or local neurons exhibited information rates close to the theoretical maximum of 15 bit/s and discriminated the stimulus almost optimally (Fig. 2D). This leaves little room for further improvement when considered as a population. Ascending neurons profited most from a population readout: Here, information in the population increased on average

1.46 ± 0.21 -fold compared with information in the best single cell ($P < 6 \times 10^{-11}$).

Note that the network converges from the receptors to the local neurons (cell numbers across layers are reduced by a factor of approximately four). Because no information is lost between these layers, this processing step could serve to compress the code.

Optimal population decoding. The ratio of the information transfers obtained from the summed-population versus the labeled-line readout indicates the disadvantages of ignoring the neuronal identity of spikes in a population (Fig. 4B and Fig. S5B): If the ratio is close to 1, then no information is gained by considering the neuronal identity of spikes; a ratio < 1 means that the labeled-line decoder yields more information, showing that it is costly to ignore information about which cell fired which spike.

The distribution for receptors and for local neurons clustered around a ratio of 1.0—that is, they were read out almost optimally using either decoder. However, there existed small yet significant trends in both cell groups: Populations of receptors were read out significantly better as a summed population (median ratio 1.04; $P = 4 \times 10^{-8}$, sign test against 1.0); populations of local neurons were decoded significantly better as a labeled line (median ratio 0.94; $P = 2 \times 10^{-6}$, sign test against 1.0). These quantitatively weak effects in receptors and local neurons are in contrast to what we found in ascending neurons: Here, the ratio was on average 0.69 and hence much smaller than 1.0 ($P < 8 \times 10^{-28}$, sign test against 1.0). Hence, a large amount of information in the population would be lost by ignoring the neuronal identity in populations of ascending neurons. The multineuron decoders confirm our hypothesis that ascending neurons implement a labeled-line code, whereas neurons in the first two layers may be efficiently read out as a summed population.

This result is also in agreement with the observed change in single-neuron information between layers (Fig. 2D): Single-neuron information increased from the first to the second layer. Considering that 60 receptors converge on 15 local neurons, noise is likely to be reduced at this stage by summation and consequently single-neuron information is enhanced. In the next step—the transition from local to ascending neurons—single-neuron information decreased. This is consistent with a labeled-line code, as information about song discriminability is distributed across the population; hence, information in *individual* neurons is lower.

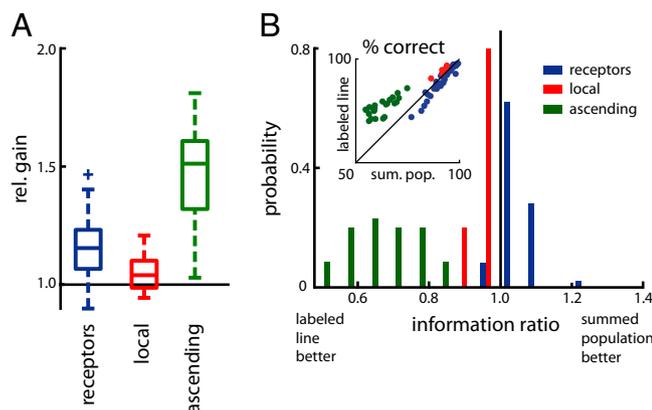


Fig. 4. Population decoding. (A) Relative gain of information through population decoding. (B) Optimal population decoder. (*Inset*) Percentage of correctly classified responses for the summed-population and for the labeled-line decoder. The histogram shows the probability distributions of the ratio of information achieved by a summed-population and a labeled-line decoder for the three processing stages (thick black line at 1.0 corresponds to equality; values < 1 indicate loss of information through summation). Blue, receptors ($n = 100$ four-cell populations); red, local neurons ($n = 160$ four-cell populations, 5 different combinations of cell types); green, ascending neurons ($n = 910$ four-cell populations, 35 different combinations of cell types); dots in the *Inset* correspond to averages over each combination of cell types.

Discussion

We asked whether the principles derived by optimal coding theory in the context of larger neuronal networks also apply to networks with relatively few neurons and a restricted set of relevant stimuli. Analyzing the representation of natural signals in the early auditory system of grasshoppers, we find that this small system performs transformations akin to those found in larger systems: Both temporal and population sparseness of the neuronal representation increase and responses are decorrelated. A labeled-line code for temporal features of the grasshopper's song is formed within only two processing steps. This code, however, differs fundamentally from that of larger systems, both in terms of how it is created as well as in its degree of specialization.

Grasshopper Labeled-Line Code Is Different from That of Larger Sensory Systems. Usually, labeled-line representations in larger systems are either obtained by a combination of inputs that were derived from a labeled line themselves or—at the very periphery—by neuron-specific tuning of receptor neurons, like location dependence in the retina or frequency selectivity in mammalian auditory receptor neurons. Preferred stimulus features in these labeled lines often vary systematically along an anatomical gradient following a topographic order. Auditory receptor neurons in grasshoppers, however, are relatively uniform in their selectivity for temporal features, in agreement with the finding that receptor neurons are best read out as a summed population. Hence, the

labeled-line code observed at the third neuronal layer is not explicitly derived from another labeled-line code, but has to be established *de novo* by a transformation of the stimulus representation in the first two layers. The construction of this labeled line from uniformly tuned inputs is presumably achieved by adaptation and a well-timed interplay of excitation and inhibition (30, 31).

Moreover, the grasshopper labeled-line code is less general. Whereas labeled-line codes in early sensory systems of higher animals tend to provide a complete representation of stimulus space allowing for flexibility of later readout, part of the grasshopper labeled line already explicitly encodes stimulus features of behavioral relevance.

So far, three specific examples of direct extraction of relevant features by ascending neurons have been described: Recognition of the “right” male is extremely relevant for female grasshoppers, as mating with the “wrong” one—in terms of species or quality—would severely impact a female’s reproductive success. One song parameter supporting species recognition is the pause length of the song’s subunits. The ascending neuron AN12 has been found to encode this song feature in its spike count, thereby allowing the female grasshopper to recognize the species of the male by its song (16). Another aspect of a male, which is indicative of high quality, is the ability to avoid predators. A frequent consequence of a previous encounter with predators is the loss of a hind leg. As males produce their calling song by rubbing both hind legs over their wings, a “one-legged” male will produce a distorted song with tiny gaps in the song’s syllable. The ascending neuron AN4 is strongly inhibited by songs exhibiting such gaps. A strong cotuning between the firing rate of AN4 and the behavioral response has been shown before (31), indicating that this neuron explicitly encodes a male’s quality. A third song parameter strongly influencing female choice is the onset slope of the song’s subunits. The spike count of the ascending neuron AN3 is strongly modulated by the onset slope of the pulses, rendering it a potential encoder of this feature (32). Thus, properties of the song that strongly influence behavior (pauses between syllables, gaps within syllables, and syllable onset steepness, respectively) are encoded explicitly in the spike count of ascending neurons AN12, AN4, and AN3.

Such an early specialization at the auditory periphery seems efficient, as the number of available neurons as well as the set of relevant stimuli are restricted. Usually, complete representations are found in systems comprising many neurons, where also the range of potentially relevant stimuli is large and relevance is often acquired through attention or learning (2, 4, 22, 33, 34). The auditory system of grasshoppers, however, has to discriminate only a restricted set of genetically fixed signals.

Interestingly, the structure of the auditory network seems to be older than that of many songs it processes (15). It might thus be *optimal* for the songs but it is certainly not *optimized* for them. The fact that we find behaviorally relevant information represented explicitly in the network is likely to be a consequence of evolutionary adaptation of the songs to the network and not vice versa (15, 23, 35).

Trading “What” for “When” Facilitates the Readout of Long Communication Signals. The peripheral auditory system of the grasshopper seems to fulfill two functions of optimal coding: signal compression and facilitated readout (1, 22). Compression is observed from the first to the second neuronal layer, where neuronal numbers converge by a factor of four and information rates per neuron increase. Facilitated readout is established at the neuronal output layer, as outlined in the following.

Receptors and local neurons encode temporal features of the song in the temporal structure of their spiking responses. They have to use a temporally precise and reproducible spike code to represent fine temporal features of a song (12, 36), even preserving information about when a temporal feature occurred

during a song syllable. In the recoded labeled-line representation of ascending neurons, however, the presence of a highly specific feature can be represented by the spike count across a whole syllable, alleviating the need for temporally precise responses at the cost of reduced information about a feature’s exact timing within a syllable. In that sense, the system trades the *when* of temporal features for an easily decodable spike-count representation of their *what*. The observation that ascending neurons are best read out at larger timescales than local neurons is consistent with this hypothesis (26) (Fig. S6 and *SI Materials and Methods*). Behavioral experiments suggest that the grasshopper’s brain indeed often evaluates spike count and makes little use of spike timing (37, 38). In addition, the songs exhibit a high temporal redundancy: Each song consists of 10–30 syllables, allowing for multiple “looks” at the basic subunit (39). The signal’s temporal redundancy presumably compensates for the restricted options of neuronal redundancy in this size-limited system and agrees well with a spike-count code that enables the system to increase the signal-to-noise ratio by accumulation of spikes across syllables.

Conclusion

Despite its limited size, the auditory system of the grasshopper shares properties of larger sensory systems: a sparse and decorrelated representation of inputs including a labeled-line population code. In contrast to larger systems, however, part of the auditory pathway seems to specialize early on for specific behaviorally relevant stimulus features. This representation is more reminiscent of higher-order areas in vertebrates. It is likely to restrict the set of stimuli that can be differentiated and hence to lower the flexibility of behavioral responses. In addition, information about precise timing of an event seems to be sacrificed for a pure detection of this event within a larger temporal window. In the context of mate selection, where signals have evolved to be sufficiently long and redundant, this may be a price worth paying and may help to invest the limited available resources specifically in the extraction of relevant information.

Materials and Methods

Recordings and Stimuli. We performed single-unit intracellular recordings from morphologically identified neurons at the three processing stages in the metathoracic ganglion of migratory locusts: receptors and local and ascending neurons. Each individual animal possesses the same set of roughly 35 unique and morphologically and physiologically identifiable types of local and ascending neurons. Cell types were identified morphologically by intracellularly injecting a fluorescent dye (Lucifer yellow). We have recorded from eight receptor neurons, five different types of local neurons (spanning the whole range of response types: TN1, BSN1, SN1, SN2, SN3), and seven different types of ascending neurons (AN1, AN2, AN3, AN4, AN11, AN12, AN14)—most of them several times (range 1–10, median frequency 3). Each recording comes from a different animal. For details on the recording procedure and stimulus presentation, refer to ref. 26.

Neurons were stimulated with eight different calling songs of male grasshoppers of the species *C. biguttulus* used previously (19). As song periodicity (duration of syllable plus pause) depends on temperature in these poikilothermic animals (10, 16), we rescaled our stimuli to a period of 100 ms and equalized the carrier spectra (19). This leaves the stimuli differing only in their envelope’s fine structure and probes the system at its limits of temporal resolution. Note that we study responses of neurons of one species of acridid grasshoppers—*Locusta migratoria*—to courtship signals of another species—*C. biguttulus*. This is well-justified, as auditory neurons in the early stages of sensory processing are morphologically and physiologically highly similar (15, 16, 31).

Estimation of Response Similarity, Reproducibility, and Sparseness. Time-varying firing-rate functions were estimated by binning spike trains at 0.05-ms resolution and smoothing them with a Gaussian filter of width $\sigma = 5$ ms. All results obtained were robust to changes of this filter’s width.

Response similarity (cell-to-cell correlations) and reproducibility (trial-to-trial correlations) were quantified as Pearson’s correlation coefficient. For response similarity, we used trial-averaged rate functions of cell pairs for each stimulus. Reproducibility was based on pairs of single-trial responses to

repeated representations of the same stimulus. Using the uncentered correlation coefficient as a measure of reproducibility gave similar results (40). Lifetime sparseness of single neurons at each processing stage was quantified from their trial-averaged firing-rate profiles using the quantity described in ref. 21. Population sparseness was quantified from the trial-averaged firing-rate profiles by calculating the measure in ref. 21 across the four cells in each population for every time bin and averaging over time. Each measure was averaged over all eight songs for plotting and statistics. So as to not bias statistics in favor of cell types we recorded more often, we additionally averaged values over all specimens of a given cell type.

Spike-Triggered Averages. We estimated STA stimuli using the same responses as those used for decoding. To correct for different firing rates, we scaled all

STAs to unit norm. Individual STAs were averaged over specimens of the same cell type. We derived estimation errors by subsampling (10 repetitions using random subsets of 80% of the spikes). Estimation errors of the normalized STAs of neurons in all three layers did not differ significantly (mean \pm SEM 9.2×10^{-3} , $P = 0.31$, one-way ANOVA).

Decoding. Song identity was decoded from neural responses using the single-neuron metric (19, 25). Populations of cells were decoded with the multi-neuron metric (28). For details on the metrics, see *SI Materials and Methods*.

ACKNOWLEDGMENTS. This work was funded by grants from the Federal Ministry of Education and Research, Germany (01GQ1001A, 01GQ0901) and the Deutsche Forschungsgemeinschaft (SFB618, GK1589/1) (to J.C., B.R., and S.S.).

- Barlow H (2001) Redundancy reduction revisited. *Network* 12:241–253.
- Vinje WE, Gallant JL (2000) Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* 287:1273–1276.
- DeWeese MR, Wehr M, Zador AM (2003) Binary spiking in auditory cortex. *J Neurosci* 23:7940–7949.
- Laurent G (2002) Olfactory network dynamics and the coding of multidimensional signals. *Nat Rev Neurosci* 3:884–895.
- Yaksi E, von Saint Paul F, Niessing J, Bunschuh ST, Friedrich RW (2009) Transformation of odor representations in target areas of the olfactory bulb. *Nat Neurosci* 12:474–482.
- Kaschube M, et al. (2010) Universality in the evolution of orientation columns in the visual cortex. *Science* 330:1113–1116.
- Aronov D, Reich DS, Mechler F, Victor JD (2003) Neural coding of spatial phase in V1 of the macaque monkey. *J Neurophysiol* 89:3304–3327.
- Jia H, Rochefort NL, Chen X, Konnerth A (2010) Dendritic organization of sensory input to cortical neurons in vivo. *Nature* 464:1307–1312.
- Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435:1102–1107.
- von Helversen D (1972) Gesang des Männchens und Lautschema des Weibchens bei der Feldheuschrecke *Chorthippus biguttulus* (Orthoptera, Acrididae). *J Comp Physiol A Neuroethol Sens Neural Behav Physiol* 81:381–422.
- Machens CK, et al. (2001) Representation of acoustic communication signals by insect auditory receptor neurons. *J Neurosci* 21:3215–3227.
- Rokem A, et al. (2006) Spike-timing precision underlies the coding efficiency of auditory receptor neurons. *J Neurophysiol* 95:2541–2552.
- Stumpner A, Ronacher B (1994) Neurophysiological aspects of song pattern recognition and sound localization in grasshoppers. *Am Zool* 34:696–705.
- Kumar A, Rotter S, Aertsen A (2010) Spiking activity propagation in neuronal networks: Reconciling different perspectives on neural coding. *Nat Rev Neurosci* 11:615–627.
- Neuhofer D, Wohlgemuth S, Stumpner A, Ronacher B (2008) Evolutionarily conserved coding properties of auditory neurons across grasshopper species. *Proc Biol Sci* 275:1965–1974.
- Creutzig F, et al. (2009) Timescale-invariant representation of acoustic communication signals by a bursting neuron. *J Neurosci* 29:2575–2580.
- Hennig RM, Franz A, Stumpner A (2004) Processing of auditory information in insects. *Microsc Res Tech* 63:351–374.
- Ronacher B, Hennig RM (2004) Neuronal adaptation improves the recognition of temporal patterns in a grasshopper. *J Comp Physiol A Neuroethol Sens Neural Behav Physiol* 190:311–319.
- Machens CK, et al. (2003) Single auditory neurons rapidly discriminate conspecific communication signals. *Nat Neurosci* 6:341–342.
- Einhäupl A, Stange N, Hennig RM, Ronacher B (2011) Attractiveness of acoustic communication signals correlates with their robustness against noise. *Behav Ecol* 22:791–799.
- Willmore B, Tolhurst DJ (2001) Characterizing the sparseness of neural codes. *Network* 12:255–270.
- Olshausen BA, Field DJ (2004) Sparse coding of sensory inputs. *Curr Opin Neurobiol* 14:481–487.
- Smith EC, Lewicki MS (2006) Efficient auditory coding. *Nature* 439:978–982.
- Quiroga RQ, Panzeri S (2009) Extracting information from neuronal populations: Information theory and decoding approaches. *Nat Rev Neurosci* 10:173–185.
- van Rossum MC (2001) A novel spike distance. *Neural Comput* 13:751–763.
- Wohlgemuth S, Ronacher B (2007) Auditory discrimination of amplitude modulations based on metric distances of spike trains. *J Neurophysiol* 97:3082–3092.
- Sharpee TO, et al. (2006) Adaptive filtering enhances information transmission in visual cortex. *Nature* 439:936–942.
- Houghton CJ, Sen K (2008) A new multineuron spike train metric. *Neural Comput* 20:1495–1511.
- Schneider DM, Woolley SM (2010) Discrimination of communication vocalizations by single neurons and groups of neurons in the auditory midbrain. *J Neurophysiol* 103:3248–3265.
- 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.
- Ronacher B, Stumpner A (1988) Filtering of behaviourally relevant temporal parameters of a grasshopper's song by an auditory interneuron. *J Comp Physiol A Neuroethol Sens Neural Behav Physiol* 163:517–523.
- Krahe R, Budinger E, Ronacher B (2002) Coding of a sexually dimorphic song feature by auditory interneurons of grasshoppers: The role of leading inhibition. *J Comp Physiol A Neuroethol Sens Neural Behav Physiol* 187:977–985.
- Leonardo A (2005) Degenerate coding in neural systems. *J Comp Physiol A Neuroethol Sens Neural Behav Physiol* 191:995–1010.
- Zhaoqing L (1988) A saliency map in primary visual cortex. *Trends Cogn Sci* 6:9–16.
- Clemens J, Weschke G, Vogel A, Ronacher B (2010) Intensity invariance properties of auditory neurons compared to the statistics of relevant natural signals in grasshoppers. *J Comp Physiol A Neuroethol Sens Neural Behav Physiol* 196:285–297.
- 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.
- Creutzig F, et al. (2010) Timescale-invariant pattern recognition by feedforward inhibition and parallel signal processing. *Neural Comput* 22:1493–1510.
- von Helversen D, von Helversen O (1998) Acoustic pattern recognition in a grasshopper: Processing in the time or frequency domain? *Biol Cybern* 79:467–476.
- Viemeister NF, Wakefield GH (1991) Temporal integration and multiple looks. *J Acoust Soc Am* 90:858–865.
- Schreiber S, Fellous J-M, Whitmer D, Tiesinga PH, Sejnowski TJ (2003) A new correlation-based measure of spike timing reliability. *Neurocomputing* 52:54:925–931.

Supporting Information

Clemens et al. 10.1073/pnas.1104506108

SI Results

Specificity of Spike-Triggered Averages. In Fig. 3 in the main text, we showed that spike-triggered averages (STAs)—and hence the feature selectivity—become more diverse at the level of ascending neurons. The decoding approach (Fig. 4 in the main text) led us to conclude that ascending neurons are optimally read out as a labeled line, indicating that each ascending neuron signals a specific aspect of the stimulus.

The second and third processing layers each consist of several cell types. To further support our findings, we also looked at the similarity of each cell type's STA across different individuals/recordings by computing Pearson's correlation coefficient between the average STA of a cell type and the STA of each individual cell of that type. In order for the STA (feature) to be specific for the cell type, this "intratype" similarity should be larger than the similarity across different cell types of that layer ("intertype" similarity).

As we considered each receptor a different type, inter- and intratype similarities for receptors are identical (Fig. S3, blue box plots). For local neurons (Fig. S3, red box plots), both the STAs of the same cell type and of different cell types are highly similar (intratype similarity 0.95 ± 0.04 , mean \pm SD; intertype similarity 0.81 ± 0.10). Whereas individual STAs of the same type are significantly more similar than those of different types ($P = 0.019$, rank-sum test), overall similarity at the level of local neurons is high. In ascending neurons, however, STAs of the same cell type are much more similar than those of different types (intratype similarity 0.85 ± 0.17 , intertype similarity 0.22 ± 0.55 ; $P = 0.002$, rank-sum test). This cell specificity of STA filters further supports our hypothesis that each type of ascending neuron encodes a specific aspect of the stimulus.

Information Gain Relative to the Population Average. In the main text, we quantified the information gain by relating the information of a four-cell population (the larger value of those obtained with the summed-population and the labeled-line decoder) with that of the best cell in that population. Alternatively, we also considered the gain with respect to the average information of all four cells comprising that population (Fig. S5A). Clearly, this measure of information gain yields higher values: Receptors exhibit an average gain of 1.99, local neurons 1.49, and ascending neurons 2.49. Thus, the gain relative to the average information in the population is 1.4- to 1.7-fold greater than the gain relative to the best cell. This is due to an upward bias in this alternative measure: The more cells one includes in a population, the more likely it is to "hit" a highly informative one. The receptors with their high spread of single-neuron information values (Fig. 2D in the main text) are especially susceptible to this bias. We hence decided to quantify information gain relative to the best cell in each population as a more conservative and less biased measure.

SI Materials and Methods

Decoding. We quantified information in neural responses using a decoding approach (1). Although we thereby underestimate the full information in the statistical sense, we probably come closer to what a concrete, biologically plausible system can read out from the spike trains we study here.

Single-neuron metric. The spike-train dissimilarity of single neurons was quantified using the van Rossum metric (2). Spike trains were binned with a resolution of 0.05 ms and filtered with an α function: $\alpha(t) = \Theta(t) t \exp(-t/\tau)$, where $\Theta(t)$ is Heaviside's function.

The parameter τ governs the temporal resolution of the metric. The Euclidean distance between all pairs of responses (eight repetitions of eight song segments of different males, duration 200 ms each) yields a distance matrix that forms the base for the classification algorithm outlined below.

Multineuron metric. Population data were combined from single-cell recordings of four individual cells. This was justified, as neural activity in the network is entirely stimulus-driven. Hence, neurons are conditionally independent: There are no "noise" correlations between neurons, only signal correlations (3). Because we were interested in how the population code changed between processing stages, we created three different classes of four-cell populations, combining different types of either receptors or local or ascending neurons. Thus, each population was characterized by a unique combination of four different cell types of a single layer. So as not to overrepresent those populations that consist of cell types we have recorded more often, we averaged information rates and gains for each kind of population (i.e., combination of cell types) for plotting and statistics.

For a formal derivation of the multineuron metric, see ref. 4. Application of this metric amounts to filtering the spike trains with an α function, embedding the spike trains from multiple cells into a vector space, and then taking the Euclidean distance between different spike trains. The resulting distance matrix for each population is then used to quantify stimulus discriminability through the classification algorithm. Thus, the only difference from the single-cell metric is that the spike trains of the cells comprising a population are embedded in a vector space.

The multineuron metric allows for different kinds of embedding, which is controlled by the "independence" parameter θ —the "angle between cells." This parameter allows interpolating between two versions of a population code: a summed-population code and a labeled-line code. At $\theta = 0^\circ$, the metric corresponds to a summed-population code, where responses of different cells are embedded colinearly. Information about which cell fired which spike is lost. This is optimal only if differences in the firing pattern between cells in a population are not relevant for the decoding tasks or if cells in a population are similarly tuned—this applies in our case to receptors and local neurons. In contrast, information about each spike's origin is fully retained in a labeled-line code, which is implemented at $\theta = 90^\circ$ (orthogonal embedding). This is desirable, if cells are tuned differently and represent different aspects of a stimulus, like the ascending neurons.

To illustrate that the labeled-line decoder incorporates information about which neuron fired which spikes—the neuronal identity of spikes—we provide a simplified example of how three different stimuli can be distinguished with the summed-population and the label-line decoder, respectively, based on surrogate responses from two neurons. Fig. S4A shows the surrogate spike trains of both cells in response to the three different stimuli. To simplify the argument and without loss of generality, we reduce these spike trains to spike counts, which corresponds to applying a filter with a large time constant τ . In response to stimulus 1, cell A (green) and cell B (blue) fire three spikes each. Stimulus 2 evokes one spike in cell A and five spikes in cell B. The response pattern for stimulus 3 is inverted: Now cell A fires five spikes and cell B only one.

The summed-population decoder sums these spike counts before computing pairwise distances between all stimuli. As the sum of spikes in cell A and cell B is the same, the population response to all three stimuli is represented by a 6; they cannot be distinguished. In contrast, the labeled-line decoder does not pool

the response of the two cells in the population. Here each response is represented by an ordered pair of spike counts, which is different for each stimulus. This is also reflected in the resulting distance matrices (Fig. S4B). As the summed-population spike counts are the same for all three stimuli, the distance matrix has all zero entries and the summed-population decoder cannot discriminate between the three stimuli (information 0 bit). The labeled-line decoder, however, discriminates all three stimuli, as all off-diagonal entries in the distance matrix exhibit nonzero entries (information $\log_2 3 = 1.6$ bit). The labeled-line decoder can distinguish stimulus 1 from both stimuli 2 and 3. In particular, it can also disambiguate stimuli 2 and 3, which differ only in neuronal identity of responses (both stimuli evoke one spike in one cell and five in the other, but in a different order). This ordering is the major difference for the summed-population decoder, and reflects the role of neuronal identity for the labeled-line decoder.

Classifier. Responses were classified using a nearest-neighbor clustering algorithm as in ref. 5. Nearness was given by the single or multineuron metrics. We randomly selected one template spike train from each of the eight songs. The remaining spike trains were then classified as being evoked by the song to which the nearest template belonged. This was repeated many times, always with a new, randomly selected set of templates. We organized the classification results in a confusion matrix $H(s, s')$, which shows the frequency with which a spike train being evoked by song s was classified as being evoked by song s' . The average of this matrix's main diagonal denotes the fraction of correctly decoded spike trains.

Estimation of information. The mutual information of this confusion matrix $I(s, s')$ was used as a proxy for the information content of the neural responses $I(s, r)$ (1). Information is given by $I(s, r)$

$$\propto I(s, s') = \sum_{s, s'} p(s, s') \log_2 \frac{p(s, s')}{p(s)p(s')},$$

where $p(s)$ is the entry in the confusion matrix, $p(s) = \sum_{s'} p(s, s') = 1/8$ is the prior stimulus probability, and $p(s') = \sum_s p(s, s')$ is the marginal over the decoded stimuli (6). Mutual information is 0 bit when the confusion matrix is uniformly distributed, that is, when each entry has the value $1/64$. It is maximal [for eight stimuli $\log_2(8) = 3$ bit] when there is a one-to-one relationship between spike trains and classes, for example, when all entries are concentrated at the matrix's diagonal. As this measure is upwardly biased, we calcu-

lated the same quantity 10 times for random assignments between responses and stimulus classes and subtracted this bias from the naive estimator $I(s, s')$ (7). We expressed information either as a rate in bit/s by dividing the information by the stimulus length (maximal information rate being thereby $8/0.2 \text{ s} = 15 \text{ bit/s}$) or as information per spike (bit per spike) by normalizing the information rate by the cell's firing rate. Firing rate was quantified as the spike count divided by the length of the spike train segment (200 ms).

Optimization of the metric's parameters. Classification performance is a function of the metric's temporal resolution τ . We optimized information with a grid search for τ ranging from 0.25 to 64 ms (nine values, spaced linearly on a logarithmic scale). The τ used for decoding are shown in Fig. S6. Receptors exhibited an intermediate range of τ between 4 and 8 ms with two outliers at 16 and 32 ms. The τ of local neurons were significantly smaller ($P = 0.01$), spanning a range of 3–4.2 ms. Ascending neurons had the highest τ between 6.7 and 42 ms, being significantly greater than those of local neurons ($P = 0.003$). For population decoding with the multineuron metric, we used a single optimal τ for all cells in a population.

In the main text, we consider only the information rates obtained for two “extreme-value decoders” at $\theta = 0^\circ$ (summed-population) and at $\theta = 90^\circ$ (labeled-line) for each population. We have also determined information at the optimal θ for each population by a grid search in the interval $[0^\circ, 90^\circ]$. As either of the two decoders at 0° or 90° yielded near-optimal performance for any population (median information loss 2%), we decided to consider only those two for all analyses.

Statistics. All plots and statistics were based on average values for each cell type or type of population, that is, over all recordings of a cell type for the analysis of single cells and over unique, unordered 4-tuples for populations of cells. Tests—if not stated otherwise—were either parametric (t test) or nonparametric (two-sided Wilcoxon's rank-sum test), depending on the outcome of a Jarque-Bera test for normality with a significance level $\alpha = 0.05$. No correction for multiple comparisons was performed to avoid false negatives, as we were interested in the outcome of each individual pairwise comparison, not in the general detection of statistical differences between groups.

All analysis was done in MATLAB (The MathWorks).

1. Quian Quiroga R, Panzeri S (2009) Extracting information from neuronal populations: Information theory and decoding approaches. *Nat Rev Neurosci* 10:173–185.
2. van Rossum MC (2001) A novel spike distance. *Neural Comput* 13:751–763.
3. Brody CD (1999) Disambiguating different covariation types. *Neural Comput* 11:1527–1535.
4. Houghton CJ, Sen K (2008) A new multineuron spike train metric. *Neural Comput* 20:1495–1511.

5. Machens CK, et al. (2003) Single auditory neurons rapidly discriminate conspecific communication signals. *Nat Neurosci* 6:341–342.
6. Victor JD, Purpura K (1997) Metric-space analysis of spike trains: Theory, algorithms and application. *Network* 8:127–164.
7. Aronov D, Reich DS, Mechler F, Victor JD (2003) Neural coding of spatial phase in V1 of the macaque monkey. *J Neurophysiol* 89:3304–3327.

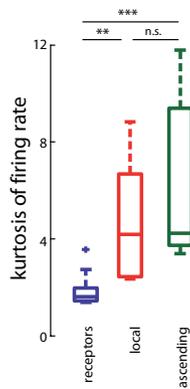


Fig. S1. Kurtosis of the firing-rate distribution as an alternative measure of lifetime sparseness. Receptors 1.9 ± 0.7 , local neurons 4.8 ± 2.7 , ascending neurons 6.2 ± 3.6 ; receptors versus local neurons $P = 0.008$, receptors versus ascending neurons $P = 2.1 \times 10^{-4}$, local neurons versus ascending neurons $P = 0.5$. n.s., nonsignificant. $P > 0.05$, $**P > 0.01$, $***P < 0.001$, rank-sum test.

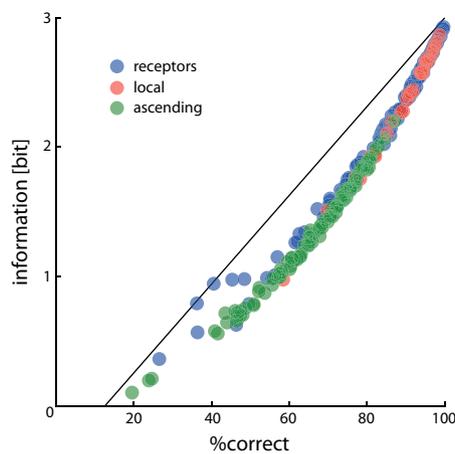


Fig. S2. Two measures of decoder performance are highly correlated. In our dataset, both the mutual information and the percentage of correct classification yield highly similar results for decoding of single cells as well as of populations ($r^2 = 0.97$).

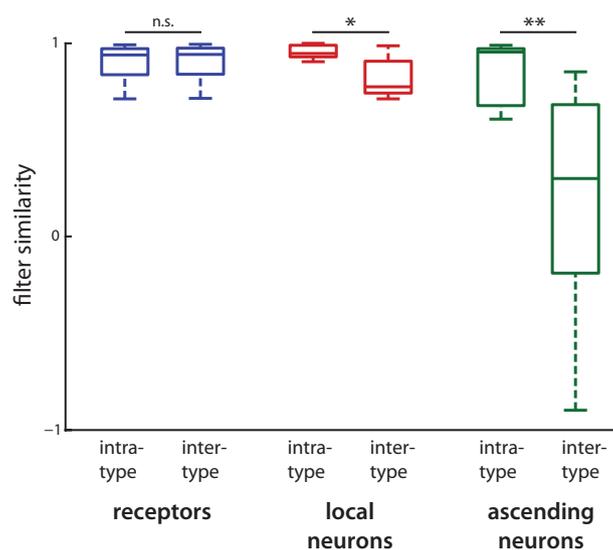


Fig. S3. Cell-type specificity of STA filters. Shown is the similarity of the STA filters of different specimens of the same cell type (intratype) and the similarity of the STA filters of different cell types (intertype; same as Fig. 3C in the main text). n.s., nonsignificant. $P > 0.05$, $*P < 0.05$, $**P < 0.01$, rank-sum test.

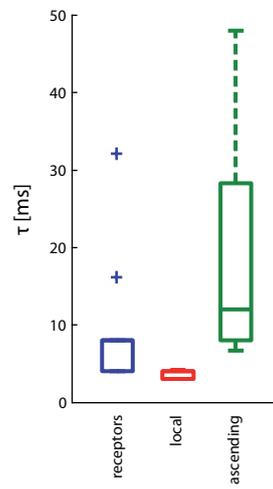


Fig. 56. Optimal timescales for decoding. Box plots show the τ s that maximized the mutual information for each cell type. These determine the width of the α functions with which spike trains were convolved in the decoding procedure and indicate the timescale at which the decoder operated.