

Learning Mid-Level Auditory Codes from Natural Sound Statistics

Wiktor Młynarski*, Josh H. McDermott
mlynar@mit.edu

Abstract

Interaction with the world requires an organism to transform sensory signals into representations in which behaviorally meaningful properties of the environment are made explicit. These representations are derived through cascades of neuronal processing stages in which neurons at each stage recode the output of preceding stages. Explanations of sensory coding may thus involve understanding how low-level patterns are combined into more complex structures. To gain insight into such mid-level representations for sound, we designed a hierarchical generative model of natural sounds that learns combinations of spectrotemporal features from natural stimulus statistics. In the first layer the model forms a sparse convolutional code of spectrograms using a dictionary of learned spectrotemporal kernels. To generalize from specific kernel activation patterns, the second layer encodes patterns of time-varying magnitude of multiple first layer coefficients. When trained on corpora of speech and environmental sounds, some second-layer units learned to group spectrotemporal features that occur together in natural sounds. Others instantiate opponency between dissimilar sets of spectrotemporal features. Such groupings might be instantiated by neurons in the auditory cortex, providing a hypothesis for mid-level neuronal computation.

Introduction

The ability to interact with the environment requires an organism to infer characteristics of the world from sensory signals. One challenge is that the environmental properties an organism must recognize are usually not explicit in the sensory input. A primary function of sensory systems is to transform raw sensory signals into representations in which behaviorally important features are more easily recovered. In doing so the brain must generalize across irrelevant stimulus variation while maintaining selectivity to the variables that matter for behavior. The nature of sensory codes and the mechanisms by which they achieve appropriate selectivity and invariance are thus a primary target of sensory system research.

The auditory system is believed to instantiate such representations through a sequence of processing stages extending from the cochlea into the auditory cortex. Existing functional evidence suggests that neurons in progressively higher stages of the auditory pathway respond to increasingly complex and abstract properties of sound [Chechik and Nelken, 2012, Atencio et al., 2012, Carruthers et al., 2015, Bizley et al., 2005, Elie and Theunissen, 2015, Russ et al., 2008, Mesgarani et al., 2008, Obleser et al., 2010, Bendor and Wang, 2005, Overath* et al., 2015, Norman-Haignere et al., 2015]. Yet our understanding of the underlying transformations remains limited, particularly when compared to the visual system.

Feature selectivity throughout the auditory system has traditionally been described using linear receptive fields [Aertsen and Johannesma, 1981, Theunissen et al., 2000, Depireux et al., 2001]. The most common instantiation is the spectrotemporal receptive field

(STRF), which typically characterizes neural activity with a one-dimensional linear projection of the sound spectrogram transformed with a nonlinearity [Sharpee et al., 2011]. As a neural data analysis technique, STRFs are widespread in auditory neuroscience and have generated insight in domains ranging from plasticity to speech coding (e.g. [Fritz et al., 2003, Woolley et al., 2005]).

Despite their utility, it is clear that STRFs are at best an incomplete description of auditory codes, especially in the cortex [Machens et al., 2004, Sahani and Linden, 2003, Williamson et al., 2016]. Experimental evidence suggests that auditory neural responses are strongly nonlinear. As a consequence, auditory receptive fields estimated with natural sounds differ substantially from estimates obtained with artificial stimuli [Theunissen et al., 2000]. STRF descriptions also fail to capture the dimensionality expansion of higher representational stages. In contrast to the brainstem, neurons in the auditory cortex seem to be sensitive to multiple stimulus features at the same time [Atencio et al., 2012, Kozlov and Gentner, 2016, Harper et al., 2016]. The presence of strongly non-linear behavior and multiplexing necessitates signal models more sophisticated than one-dimensional, linear features of the spectrogram such as STRFs.

An additional challenge to characterizing mid-level features of sound is that humans lack strong intuitions about abstract auditory structure. In specific signal domains such as speech, progress has been made by cataloging phonemes and other frequently occurring structures, but it is not obvious how to generalize this approach to broader corpora of natural sounds.

An alternative approach to understanding sensory representations that is less reliant on domain-specific intuition is that of efficient coding [Barlow, 1961, Attneave, 1954]. The efficient coding hypothesis holds that neural codes should exploit the statistical structure of natural signals, allowing such signals to be represented with a minimum of resources. Numerous studies have demonstrated that tuning properties of neurons in early stages of the visual and auditory systems are predicted by statistical models of natural images or sounds [Srinivasan et al., 1982, Olshausen and Field, 1997, Bell and Sejnowski, 1997, van Hateren and van der Schaaf, 1998, Lewicki, 2002, Klein et al., 2003, Smith and Lewicki, 2006, Carlson et al., 2012, Terashima and Okada, 2012, Młynarski, 2015, Młynarski, 2014, Carlin and Elhilali, 2013, Lee et al., 2008, Hosoya and Hyvärinen, 2015]. Although the early successes of this approach engendered optimism, applications have largely been limited to learning a single stage of representation, and extensions to multiple levels of sensory processing have proven difficult. The underlying challenge is that there are many possible forms of high-order statistical dependencies in signals, and the particular dependencies that occur in natural stimuli are typically not obvious. The formulation of models capable of capturing these dependencies requires careful analysis and design [Hyvärinen and Hoyer, 2000, Karklin and Lewicki, 2005, Karklin and Lewicki, 2009, Cadieu and Olshausen, 2008], and perhaps good fortune, and is additionally constrained by what is tractable to implement. In the auditory system in particular, it remains to be seen whether modeling statistical signal regularities can reveal the complex acoustic structures and invariances that are believed to be represented in higher stages of the auditory system.

The primary goal of the present work was to discover such high-order structure in natural sounds and generate hypotheses about not-yet-observed intermediate-level neural representations. To this end we developed a probabilistic generative model of natural sounds designed to learn a novel stimulus representation - a population code of naturally occurring combinations of basic spectrotemporal patterns, analogous to STRFs.

The representations learned by the model from corpora of natural sounds suggest grouping principles in the auditory system. Model units learned to pool sets of similar spectrotemporal features, presumably because they occur together in natural sounds. In addition, some model units encode opponency between different sets of features. Although

not yet described in the auditory system, such tuning patterns may be analogous to phenomena such as end-stopping or cross-orientation suppression in the visual system. The representations learned by our model also resemble some recently reported properties of auditory cortical neurons, providing further evidence that natural-scene statistics can predict neural representations in higher sensory areas.

Methods and Models

Overview of the hierarchical model

To learn mid-level auditory representations, we constructed a hierarchical, statistical model of natural sounds. The model structure is depicted in Fig 1. The model consisted of a stimulus layer and two latent layers that were adapted to efficiently represent a corpus of audio signals. Because our goal was to learn mid-level auditory codes, we did not model the raw sound waveform. Instead, we assumed an initial stage of frequency analysis, modeled after that of the mammalian cochlea. This frequency analysis results in a spectrogram-like input representation of sound, which we term a 'cochleagram,' that provides a coarse model of the auditory nerve input to the brain (Fig 1 A, bottom row). This input representation is an $F \times T$ matrix, where F is the number of frequency channels and T the number of time-points. Our aim was to capture statistical dependencies in natural sounds represented in this way.




Figure 1: **Overview of the hierarchical model.** A) A spectrogram (bottom-row) is encoded by a set of spectrotemporal kernels (middle row). The features learned by the second layer encode temporal patterns of multiple STK activations. B) A graphical model depicting statistical dependencies among variables.

The first layer of the model (Fig 1, middle row) was intended to learn basic acoustic features, analogous to STRFs of early auditory neurons. Through the remainder of the paper, we refer to the features learned by the first layer as spectrotemporal kernels (STKs), in order to differentiate them from neurally derived STRFs. The second layer, depicted in the top row of Fig 1 A, was intended to learn patterns of STK co-activations that frequently occur in natural sounds.

The model specifies a probability distribution over the space of natural sounds, and its parameters can be understood as random variables whose dependency structure is depicted in Fig 1 B. The spectrogram x is represented with a set of spectrotemporal features ϕ convolved with the latent activation time-courses s . The second latent layer encodes the magnitudes of s with the basis functions B convolved with their activation time-courses v .

In the following sections, we present the details of each layer.

First layer of model - convolutional, non-negative sparse coding of spectrograms

The first layer of the model was designed to learn basic spectrotemporal features. One previous attempt to learn sparse spectrotemporal representations of natural sounds [Carlson et al., 2012] produced structures reminiscent of receptive fields of neurons in the auditory midbrain, thalamus, and cortex. Here, we extended this approach by learning a spectrogram representation which is sparse and convolutional.

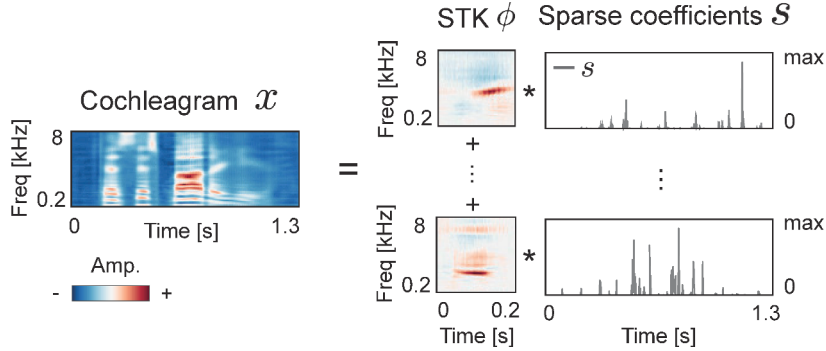


Figure 2: **Explanation of first layer of the model.** A cochleagram x of arbitrary length is represented as a sum of spectrotemporal kernels ϕ convolved with time-courses of corresponding coefficients s . Kernels span the full extent of the frequency dimension, and convolution occurred only in time. Coefficients s are non-negative and have sparse distributions (i.e. remain close to zero most of the time).

A schematic of the first layer is depicted in Fig 2. We modeled the cochleagram $(x_{t,f})$ as a linear combination of spectrotemporal kernels ϕ convolved in time with their activation time-courses s and distorted by additive Gaussian noise $\xi_{t,f}$ with variance σ^2 :

$$\hat{x}_{t,f} = \left[\sum_{i=1}^N \phi_{i,f} * s_i \right]_t \quad (1)$$

$$x_{t,f} = \hat{x}_{t,f} + \xi_{t,f} \quad (2)$$

Kernel activations $s_{i,t}$ are assumed to be independent, i.e. their joint distribution is equal to the product of marginals:

$$p(\mathbf{s}) = \prod_{i=1}^N \prod_{t=1}^T p(s_{i,t} | \lambda_i) \quad (3)$$

We assumed that each spectrotemporal kernel remains inactive for most of the time, i.e. that the distribution of its activations is sparse. Moreover, we imposed a non-negativity constraint on the coefficients s . This facilitates interpretations in terms of neural activity and improves the interpretability of the learned representation. These constraints were embodied in an exponential prior on the coefficients s :

$$p(s_{i,t} | \lambda_i) = \frac{1}{\lambda_i} \exp \left[-\frac{s_{i,t}}{\lambda_i} \right] \quad (4)$$

where λ_i is the scale parameter. The scale of the exponential distribution determines its "spread" along the real line. Small values of λ_i yield distributions tightly peaked at 0.

Large λ_i generate broader distributions allowing $s_{i,t}$ to attain large values with higher probability. When training the first layer we assumed all λ_i to be constant and equal to 1. As we will describe in the following sections, the second layer of the model relaxes that assumption and learns a representation of time-varying scale parameters λ_i .

The first layer of the model specifies the following negative log-posterior probability of the data:

$$E_1 \propto \frac{1}{\sigma^2} \sum_{f=1}^F \sum_{t=1}^T (\hat{x}_{t,f} - x_{t,f})^2 + \sum_{i=1}^N \frac{1}{\lambda_i} \sum_{t=1}^T s_{i,t} \quad (5)$$

This negative log-probability can be viewed as a cost function to be minimized when inferring the value of coefficients s : while maintaining low-reconstruction error (first term on the right-hand side), the sparsity of representation should be maximized (the second term).

Dependencies between coefficients - a signature of mid-level structure

Although the sparse coding strategy outlined above learns features that are approximately independent across the training set, residual dependencies nonetheless remain. In part this is because not all dependencies can be modeled with a single layer of convolutional sparse coding. However, due to the non-stationary nature of natural audio, coefficient dependencies can also be present locally despite not being evident across a large corpus as a whole. Empirically, the learned features exhibit dependencies specific to particular sounds [Karklin and Lewicki, 2009], and thus deviate locally from their (approximately independent) marginal distribution. For example, a spoken vowel with a fluctuating pitch contour would require many harmonic STKs to become activated, and their activations would become strongly correlated on a local time scale. Such local correlations reflect higher-order structure of particular natural sounds. Statistically speaking, this is an example of marginally independent random variables exhibiting conditional dependence (in this case conditioned on a particular point in time or a type of sound).

Fig 3 depicts such dependencies via cross-correlation functions of selected STK activations for a white noise sample and for three different natural sounds. Coefficient correlations (black curves in each subplot) vary from sound to sound, but in all cases deviate from those obtained with noise (gray bars within subplots), revealing dependence. These dependencies are indicative of "mid-level" auditory features, perhaps analogous to the correlations between oriented Gabor filters induced by an elongated edge. In this work, we exploited the fact that intermediate level representations can be learned by modeling dependencies among first-layer features [Karklin and Lewicki, 2009, Karklin and Lewicki, 2005, Cadieu and Olshausen, 2008].

Variability of STK activations

A second phenomenon evident in STK activations is that particular patterns of co-activation occur with some variability. This is visible in Fig 4, which depicts activations of selected STKs when encoding multiple exemplars of the same sound - the word "one" spoken twice by the same speaker (Fig 4A) and two exemplars of water being poured into a cup (Fig 4B).

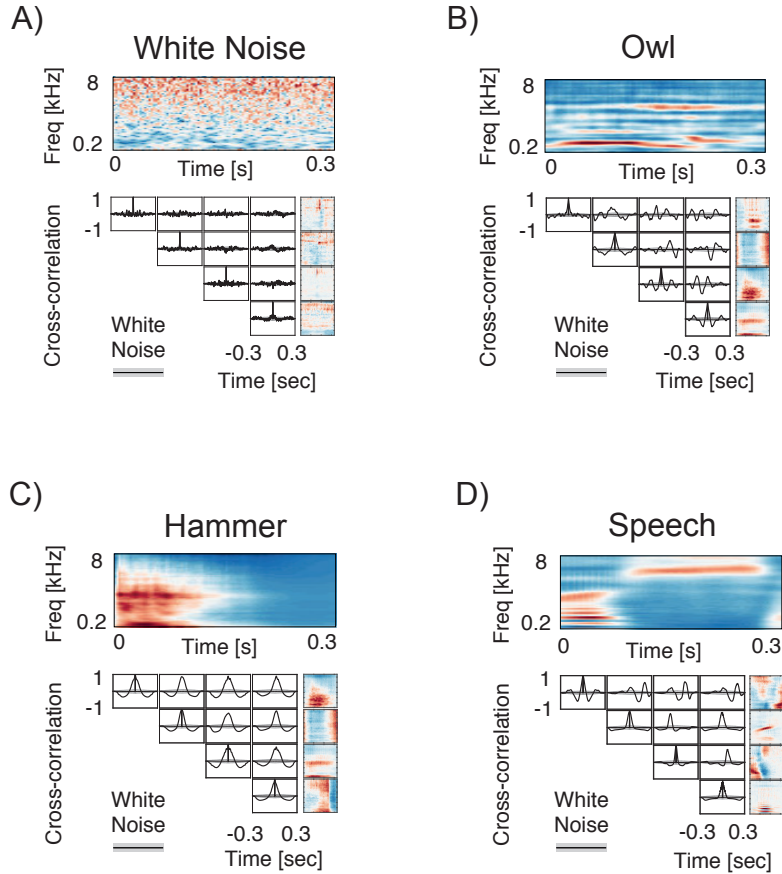


Figure 3: **Dependencies among spectrotemporal feature activations.** A) When encoding white noise (cochleagram depicted on top), sparse coefficients are uncorrelated on short time-scales. This is visible in the shape of the coefficient cross-correlation functions (black lines) corresponding to the four most strongly activated STKs (bottom, right column). Due to the lack of temporal structure, the cross-correlation between STKs is flat, and the autocorrelation of individual STKs is a Dirac delta function at 0. Here and in other panels, the mean was subtracted from the (non-negative) coefficient trajectories prior to computing the cross-correlation. B) When encoding structured stimuli such as an owl vocalization, STK activations reveal strong local correlations - the cross-correlations deviate from those for white noise (thick gray lines). C) Same as B, for a hammer hit. D) Same as B, for a speech excerpt.




Figure 4: **Variability of spectrotemporal feature activations.** A) Coefficient trajectories of spectrotemporal features ϕ are depicted for two utterances of word "one" spoken by the same speaker. Although both coefficient trajectories (gray and black lines) exhibit the same global structure, they are not identical. B) Coefficient variability visualized in the same way for two different examples of water pouring into a cup.

It is apparent that the STK coefficient trajectories for the two exemplars in each case (black and gray lines) reflect the same global pattern even though they differ somewhat from exemplar to exemplar. The similarity suggests that the trajectories could be modeled as different samples from a single time-varying distribution parameterized by a non-stationary coefficient magnitude. When the magnitude increases, the probability of a strong STK activation increases. Retaining the (inferred) time-varying magnitude instead of precise values of STK coefficients would yield a representation more invariant to low-level signal variation, potentially enabling the representation of abstract regularities in the data. Such a representation bears an abstract similarity to the magnitude operation used to compute a spectrogram, in which each frequency channel retains the time-varying energy in different parts of the spectrum. Here we are instead estimating a scale parameter of the underlying distribution, but the process similarly discards aspects of the fine detail of the signal.

Second layer of model - encoding of STK combinations

The second layer of the model was intended to exploit the two statistical phenomena detailed in the previous section: conditional dependencies between STKs and their variation across exemplars. Similarly to the first layer, the second layer representation is formed by a population of sparsely activated basis functions. These basis functions capture local dependencies among STK magnitudes by encoding the joint distribution of STK activations rather than exact values of STK coefficients. The resulting representation is thus more specific than the first-layer code - instead of encoding single features independently, it signals the presence of particular STK combinations. It is also more invariant, generalizing over specific coefficient values.

Because the proposed representation is a population code of a distribution parameter, it bears conceptual similarity to previously proposed hierarchical models of natural stimuli that encoded patterns of variance [Karklin and Lewicki, 2005, Bumbacher and Ming, 2012], covariance [Karklin and Lewicki, 2009] or complex amplitude [Cadieu and Olshausen, 2008, Hyvärinen and Hoyer, 2000]. The novelty of our model structure lies in being convolutional (i.e., it can encode stimuli of arbitrary length using the same representation) and in parameterizing distributions of non-negative STK coefficients, increasing the inter-

pretability of the learned spectrogram features. The novelty of the model’s application is to learn hierarchical representations of sound (previous such efforts have largely been restricted to modeling images; though see [Lee et al., 2009]).

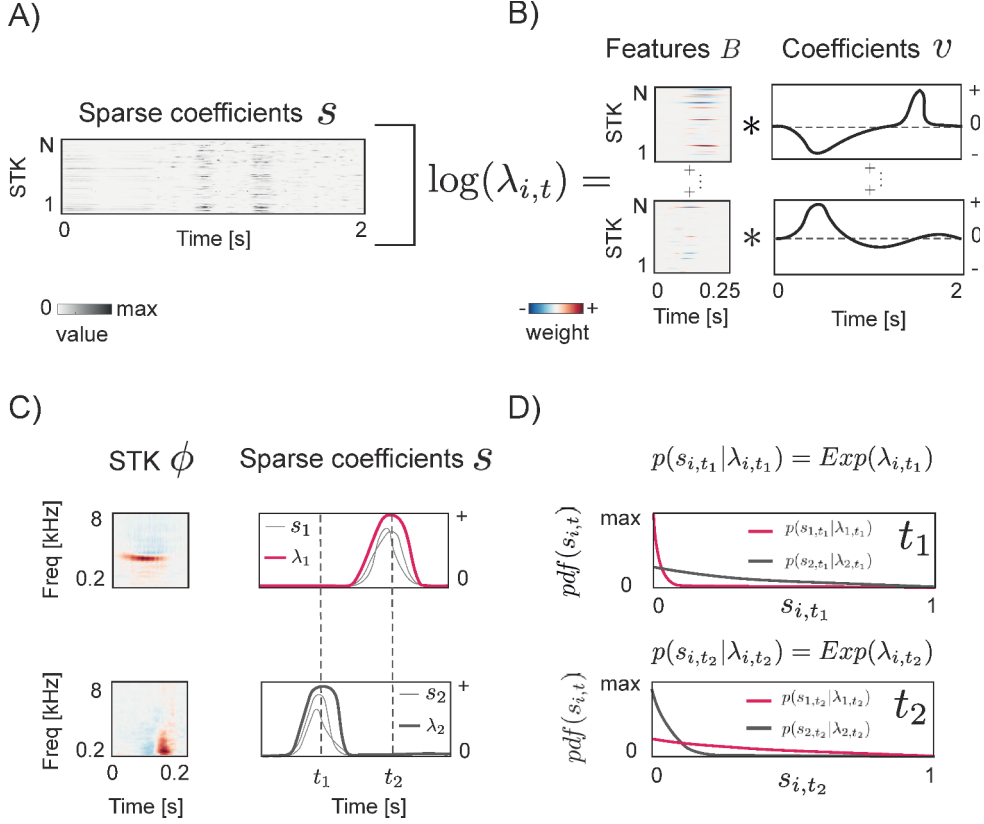


Figure 5: **Explanation of second layer of the model.** A) An array of STK activations s (a "STK-gram") serves as an input to the second layer. Rows correspond to first layer features ϕ_i and columns to time points. B) The second layer uses a population of features B to encode the logarithm of STK activation magnitudes. C) Coefficient trajectories s (thin grey lines) and their magnitudes λ (thick red and black lines) for two example STKs. D) Distributions of s at time points t_1 and t_2 are depicted in the right column.

The second layer of the model is depicted schematically in Fig 5A. We assume that STK activations s are samples from a non-stationary exponential distribution with time-varying scale parameter $\lambda_{i,t}$, relaxing the assumption of stationary λ_i made in learning the first layer:

$$p(s_{i,t}) = \text{Exp}(\lambda_{i,t}) = \frac{1}{\lambda_{i,t}} \exp\left[-\frac{s_{i,t}}{\lambda_{i,t}}\right] \quad (6)$$

When $\lambda_{i,t}$ is high, the distribution of $s_{i,t}$ becomes broader (Fig 5C, black line in the top row, red line in the bottom). This allows the coefficient s_i to attain large values. For small values of the scale parameter, the probability density is concentrated close to 0 (Fig 5D), red line in the top row, black line in the bottom), and coefficients $s_{i,t}$ become small. To model the magnitudes λ (which are non-negative, akin to variances), we took their logarithm, mapping their values onto the entire real line so that they could be represented by a sum of real-valued basis functions.

Patterns of STK magnitudes are represented in the second layer by a population of

features B convolved with coefficients v :

$$\lambda_{i,t} = \exp \left[\sum_{j=1}^M B_{j,i} * v_j + \rho_i \right]_t \quad (7)$$

where ρ is a bias vector. Each second-layer basis function B represents a particular temporal pattern of co-activation of first-layer STKs. Their corresponding coefficients v are assumed to be sparse and independent:

$$p(\mathbf{v}) = \prod_{j=1}^M \prod_{t=1}^T p(v_{j,t}) \quad (8)$$

$$p(v_{j,t}) \propto \exp \left(-\alpha |v_{j,t}| \right) \quad (9)$$

where α controls the degree of sparsity.

As with the first layer, learning and inference are performed by gradient descent on the negative log posterior. Because the second layer units encode combinations of sparse first-layer coefficients, we placed a sparse prior on the L_1 norm of the basis functions B . The overall cost function to be minimized during learning in the second layer is then:

$$E_2 \propto \sum_{i=1}^N \sum_{t=1}^T \frac{s_{i,t}}{\lambda_{i,t}} + \log(\lambda_{i,t}) + \alpha \sum_{j=1}^M \sum_{t=1}^T |v_{j,t}| + \beta \sum_{j=1}^M \sum_{i=1}^N \sum_{t_b=1}^{T_b} |B_{j,i,t_b}| \quad (10)$$

where β controls the strength of the sparse prior on B , and T_b is the temporal extent of each second-layer basis function. As in the first layer cost function (Eq 5), the first term on the right hand side of Eq 10 enforces a match of the representation to data (the magnitudes λ are pushed away from zero towards the observed coefficients s), while the second and third terms promote sparsity of second-layer coefficients and basis functions, respectively. Sparsity of second layer units is a reasonable assumption given the sparsity of first-layer coefficients, and we found that using a sparse prior on second-layer units (i.e. setting β to be larger than 0) was necessary to achieve convergence. Others have found that the addition of sparse priors can enable faster learning with less data without substantially altering the obtained solution [Hyvärinen and Raju, 2002].

Learning procedure

The two layers of the model were trained separately, i.e. the training of the second layer occurred after the first layer training was completed. In the first layer, training was performed with an EM-like procedure that iteratively alternated between inferring STK coefficients and updating STK features [Olshausen and Field, 1997, Cadieu and Olshausen, 2008]. Spectrotemporal features ϕ were initialized with Gaussian white noise. For each excerpt in the training set, coefficients $s_{i,t}$ were inferred via gradient descent on the energy function (5). Because inference of all coefficients s is computationally expensive, we adopted an approximate inference scheme [Blumensath and Davies, 2006]. Instead of inferring values of all coefficients for each excerpt, we selected only a subset of them to be minimized. This was done by computing the cross-correlation between a sound excerpt and features ϕ_i and selecting a fixed number of the largest coefficients $s_{i,t}$. The inference step adjusted only this subset of coefficients while setting the rest to 0. Given the inferred coefficients, a gradient step on the spectrotemporal features ϕ was performed.

Each learning iteration therefore consisted of the following steps:

1. Draw a random sound excerpt from the training data set. Excerpts were 403 ms in length (129 time samples of the spectrogram, sampled at 320 Hz).

2. Compute the cross-correlation of all basis functions ϕ_i with the sound excerpt. Select the 1024 pairs of coefficient indices and time-points (i, t) that yield the highest correlation values.
3. Infer the values of the selected coefficients by minimizing Eq. 5 with respect to $s_{i,t}$ via gradient descent. Set the rest of coefficients to 0.
4. Compute the gradient step on the basis functions as the derivative of Eq. 5 with respect to basis functions ϕ using inferred coefficient values \hat{s} . Update basis functions according to the gradient step.
5. Normalize all basis functions to unit norm.

This procedure was terminated after 200000 iterations.

The second layer was then learned via the same procedure used for the first layer. In each iteration a 528 ms long (169 samples at 320 Hz) randomly drawn sound excerpt was encoded by the first layer, and the resulting matrix of coefficients s served as an input to the second layer. A subset of coefficients v was selected for approximate inference by computing the cross correlation between features B_i and the logarithm of the first-layer coefficients s (analogous to step 2 in the procedure described above for the first layer). The energy function E_2 was first minimized with respect to coefficients v followed by a gradient update to the basis functions B (analogous to steps 3 and 4 for the first layer). Entries in the bias vector ρ corresponding to each coefficient s_i were set to the expectation of the coefficient across the entire training set: $\rho_i = \mathbb{E}_t[s_{i,t}]$ (i.e., the estimate of the marginal scale parameter λ_i for the corresponding STK). Learning was again terminated after 200000 iterations. To validate the learning algorithm, we ran it on a toy data set generated using known second-layer units. The algorithm successfully recovered the parameters of the generating model, as desired (see Appendix C for these results).

Training data and spectrogram parameters

We trained the model on two different sound corpora. The first corpus was the TIMIT speech database [Garofolo et al., 1993]. The second corpus combined a set of environmental sounds (the Pitt sound database [Lewicki, 2002]) and a number of animal vocalizations downloaded from freesound.org. The environmental sounds included both transient (breaking twigs, steps, etc.) and ambient (flowing water, wind, etc.) sounds; the animal vocalizations were mostly harmonic.

We computed cochleagrams by filtering sounds with a set of 65 bandpass filters intended to mimic cochlear frequency analysis. Filters were equally spaced on an equivalent rectangular bandwidth (ERB) scale [Glasberg and Moore, 1990]), with parameters similar to that from a previous publication [McDermott and Simoncelli, 2011]. Center frequencies ranged from 200 Hz to 8 kHz. We computed the Hilbert envelope of the output of each filter and raised it to the power of 0.3, emulating cochlear amplitude compression [Robles and Ruggero, 2001]. To reduce dimensionality, each envelope was downsampled to 320 Hz.

We set the number of features in the first layer to 128 and in the second layer to 100. Pilot experiments yielded qualitatively similar results for alternative feature dimensionalities. Each first layer feature encoded a 203 ms interval (65 time samples of the spectrogram).

Results

First layer: Basic spectrotemporal features of natural sounds

The first-layer features learned from each of the two sound corpora are shown in Fig 6 A and B. These features could be considered as the model analogues of neural STRFs.

The vast majority of features are well localized within the time-frequency plane, encoding relatively brief acoustic events. The STKs learned from speech included single harmonics and harmonic "stacks" (Fig 6 A - features numbered 1 and 2), frequency sweeps (feature 3), and broadband clicks (feature 4). The features learned from environmental sounds also included single harmonics and clicks (Fig 6 B - features 1, 2 and 4). In contrast to the results obtained with speech, however, harmonic stacks were absent, and a number of high-frequency hisses and noise-like features were present instead (feature 3).

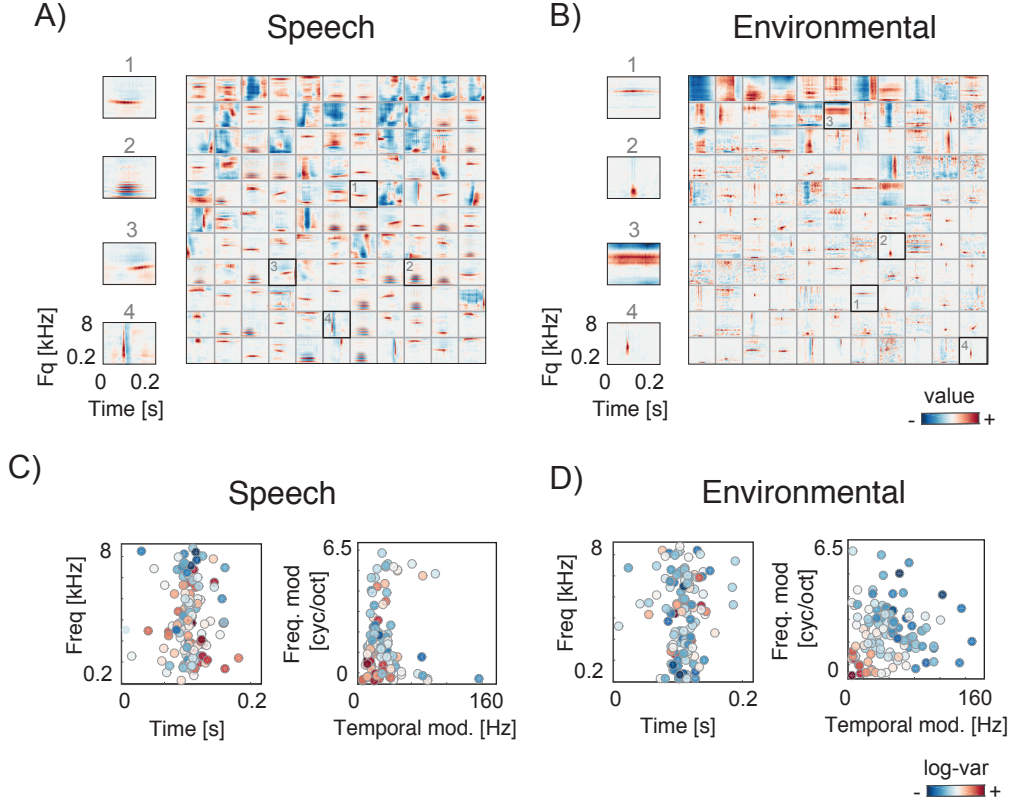


Figure 6: **Spectrotemporal kernels learned by the first layer.** A) A population of STKs learned from a speech database. Representative STKs for each corpus are magnified and numbered from 1-4 for ease of reference. B) Population of STKs learned from environmental sounds. C) Speech-trained STK population plotted on time-frequency (left) and spectral-temporal modulation (right) planes. Each dot corresponds to a single STK, and its color encodes the log of its mean coefficient (averaged over entire training dataset). D) The STK population trained on environmental sounds, represented as in (C).

The properties of the learned dictionaries are also reflected in distributions of feature locations in the time-frequency and spectrotemporal modulation planes, as visible in Fig 6C and D. Each dot position denotes the center of mass of a single STK, while its color signals the feature's average coefficient value over the stimulus set. For both speech and environmental sounds, the learned STKs uniformly the audio frequency spectrum (Fig 6 C and D, left panels). Due to the convolutional nature of the code, the energy of each feature is concentrated near the middle of the time axis. The modulation spectra of the learned STKs (Fig 6 C and D, right panels) are somewhat specific to the sound-corpus. Features trained on a speech corpus were more strongly modulated in frequency, while environmental sounds yielded STKs with faster temporal modulations. STKs learned from both datasets exhibit a spectrotemporal modulation tradeoff: if a STK is strongly temporally modulated its spectral modulation tends to be weaker. This tradeoff is an inevitable consequence of time-frequency conjugacy [Singh and Theunissen, 2003], and is also found in the STRFs of

the mammalian and avian auditory systems [Miller et al., 2002, Woolley et al., 2005].

Second layer: Combinations of spectrotemporal features

STKs captured by the first layer of the model reflect elementary features of natural sounds. By contrast, the features learned by the second layer capture how activations of different STKs cluster together in natural sounds, and thus reflect more complex acoustic regularities. We first present several ways of visualizing the multi-dimensional nature of the second-layer representation, then make some connections to existing neurophysiological data, and then derive some neurophysiological predictions from the model.

Visualizing second-layer features

The second-layer features encode temporal combinations of STK log-magnitudes. Each feature can be represented as a $N \times T_b$ dimensional matrix, where rows correspond to STKs and columns to time-points. An example second-layer unit is shown in Fig 7A (left panel). A positive value in the i -th row and t -th column of a feature B_j encodes a local increase in the magnitude of the i -th STK. A negative value encodes a decrease in magnitude.

An alternative visualization is to examine the spectrotemporal structure of the STKs that have large weights in the second-layer feature. The same feature B_j is depicted in this way in the center panel of Fig 7A, which displays the four STKs with highest average absolute weights for this particular feature. To the right of each STK are its weights (i.e., the corresponding row of the B_j matrix). It is apparent that the weights increase and decrease in a coordinated fashion, and thus likely encode particular dependencies between the STKs.

To summarize the full distribution of STKs contributing to a second-layer unit, we adopted the visualization scheme illustrated in the right panel of Fig 7A. We plot the center of mass of each STK in the modulation (top row) and time-frequency (middle row) planes, as in Fig 6C and D. The dot for a first-layer STK is colored red or blue, depending on the sign of their time-averaged weight, with the average absolute value of the weight signaled by the intensity of the color. The bottom row of the panel depicts the temporal pattern of STK magnitudes - line colors correspond to dots in the top and middle rows of the panel. Although the weights of most STKs maintain the same sign over the temporal support of the second-layer unit, they were not directly constrained in this regard, and in some cases the weight trajectories cross zero.




Figure 7: **Second layer model features.** A) Feature visualizations. Left panel - the first layer STK weights for an example second layer feature (representing magnitudes of each STK over a time window). Middle panel - STKs whose weights in the same second-layer feature deviate most strongly from 0 are displayed along with their weight profile over time. Right panel - STKs are plotted as dots in the modulation and time-frequency planes, with the dot location indicating the STK center of mass in the plane, and the color indicating the weight sign and magnitude (red denoting positive and blue denoting negative). STKs are divided into those with positive and negative weights for clarity. Bottom row visualizes temporal trajectories of STK weights for the feature. B) Examples of learned second-layer features. First two columns (labeled 1 and 2) depict units with positive ("excitatory") weights only. Last two columns (labeled 3 and 4) depict "opponent" units that pool features with both positive and negative weights.

Representative examples of second-layer basis functions are depicted in this way in Fig 7B. We separated them into two broad classes - "excitatory" units (columns 1 and 2 in Fig 7), which pool STKs using weights of the same sign, and therefore encode a pattern of coordinated increase in their magnitudes, and "excitatory-inhibitory" units (columns 3 and 4 in Fig 7) which pool some STKs with positive average weights and others with negative

average weights. We note that excitatory-only and inhibitory-only units are functionally interchangeable in the model, because the encoding is unaffected if the sign of both the STK weights and coefficients are reversed. From inspection of Fig 7B it is evident that some second-layer units pool only a few STKs (e.g. A1, C1, D1) while others are more global and influence activations of many first-layer units (e.g. A2, A3, C3).

Second layer features encode patterns of STK dependencies

To better understand the structure captured by the second-layer units, we considered their relationship to the two kinds of dependencies in STK activations that initially motivated the model. In Methods and Models section we observed that STK activations to natural sounds exhibit strong local cross-correlations as well as variations of particular temporal activation patterns. The second layer of the model was formulated in order to capture and encode these redundancies.

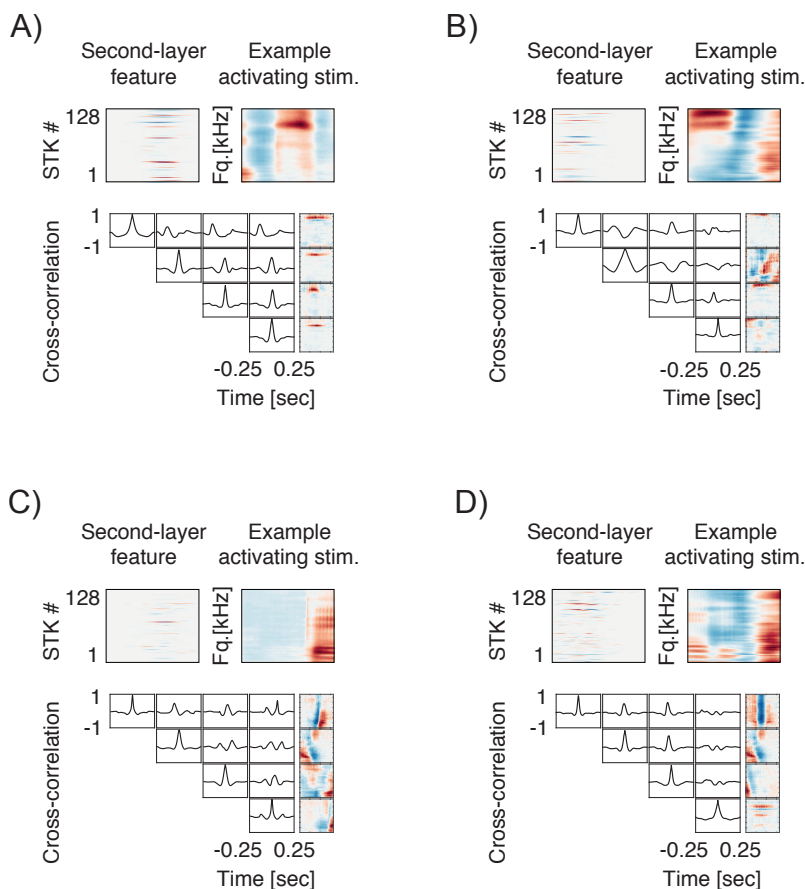


Figure 8: **Second-layer units respond to specific STK cross-correlation patterns.** A) A second-layer unit (top row, left column) and an example sound excerpt eliciting a strong positive response in the unit (top row, right column). Bottom shows cross-correlation functions of coefficient trajectories for four STKs. The STKs selected were those with the largest weights for this second-layer unit. Cross-correlations were averaged across 25 stimuli eliciting the strongest positive response of the second-layer unit across a large subset of the TIMIT corpus. B, C, D) same as A for three other second-layer units.

To first test whether the second layer captures the sorts of residual correlations evident in the first layer output, we measured cross-correlations between STK activations conditioned on the activation of particular second-layer units. Fig. 8A-D depicts four example second-layer units (top row, left column) along with an example stimulus that produced a

strong positive response in the unit (selected from the TIMIT corpus). The bottom section of each panel depicts cross-correlation functions of activations of four STKs, averaged over 25 stimulus epochs that produced a strong positive response of the second-layer unit. The cross-correlation functions deviate substantially from 0, as they do when conditioned on excerpts of natural sounds (Fig. 3). These correlations reflect the temporal pattern of STK coefficients that the second-layer unit responds to.

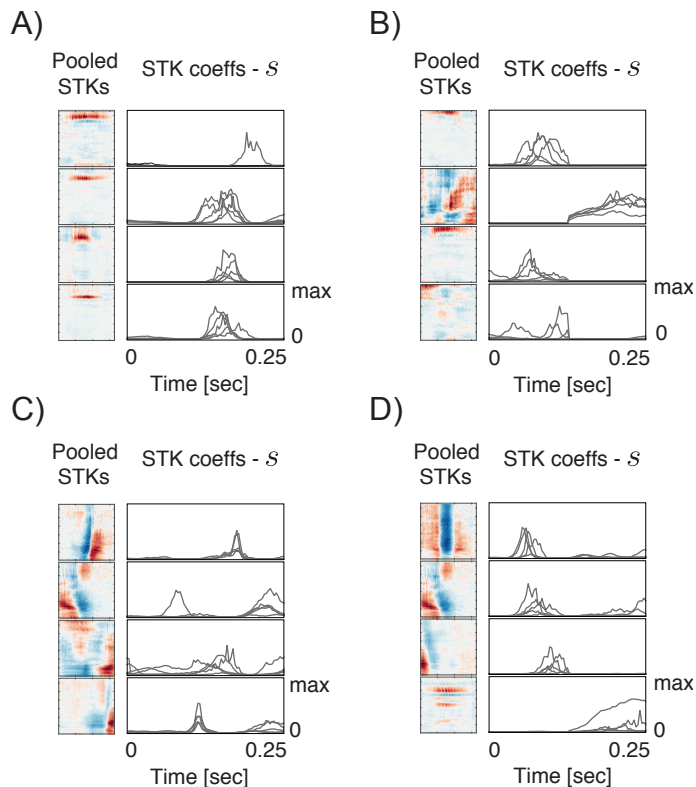


Figure 9: **Second-layer units generalize across variation in STK coefficients.** Panels A-D correspond to second-layer units depicted in Fig. 8. Each panel shows four STKs pooled with strongest weights (left column) by the corresponding second-layer unit. Next to each STK are their activation patterns (right column) for the 5 stimuli eliciting the strongest response in the second-layer unit. Despite some variability they share a global trend.

We next examined whether the second-layer units respond to STK activations fluctuating around particular global patterns (as depicted in Fig. 4). Because the second layer of the model represents the magnitude, rather than the precise values of first-layer coefficients, it should be capable of generalizing over minor STK coefficient variation. Fig. 9 plots STK coefficient trajectories for stimuli eliciting a strong response in the second-layer features from Fig. 8. The STK activation traces reveal variability in each case, but nonetheless exhibit a degree of global consistency, as we saw earlier for natural sound exemplars (Fig. 4). These results provide evidence that the second layer is capturing the dependencies it was intended to model.

Experimental predictions

Pooling of similar spectrotemporal patterns

What do the model results suggest about mid-level auditory structure? Our visualizations of second layer features in Fig 7B revealed that many represent the concurrent activation of many STKs. Examination of Fig 7B (columns 1 and 2) suggests that the first layer STKs

that such units pool are typically similar either in their spectrotemporal or modulation properties. For example, the unit depicted in panel A2 encodes joint increases in the magnitude of high-frequency STKs, while that in panel C2 encodes joint increases in low-frequency STKs.

In order to quantitatively substantiate these observations, we measured the spread of the center-of-mass of each of the STKs pooled by "excitatory-only" second-layer units (those whose STK weights exceeding 5% of the maximum absolute weight were all of the same sign). Specifically, we computed the standard deviation of the center-of-mass on each dimension of the time-frequency and modulation planes, for all STKs with weights exceeding 5% of the maximum weight for the unit, and compared this to the standard deviation of centers of mass for random samples of STKs (matched in size to the mean number of STKs pooled by the excitatory-only units). These standard deviations measure the spread of STKs in spectrotemporal and modulation domains; if the pooled STKs are similar, the spread should be small.

As shown in Fig.10, the STKs pooled with the same sign were more similar in each of the four analyzed dimensions than random groups of STKs, for both sound corpora (red bar vs. gray bar, t-tests yielded $p < .05$ in all cases). This result suggests the hypothesis that STRFs that are pooled by downstream auditory cortical neurons with excitatory weights [Atencio et al., 2012, Kozlov and Gentner, 2016] should also tend to be more similar than expected by chance.

Opponency patterns in mid-level audition

Fig 7B also shows examples of a distinct set of second-layer units in which increased activation of one group of STKs (again, typically similar to each other) is associated with a decrease in activity of another group of STKs. For example, the unit in panel B3 encodes an increase in magnitude of temporally modulated features (clicks) together with a simultaneous decrease of activation of a strongly spectrally modulated (harmonic) STK. Such "opponency" was evident in a large subset of second layer units, and represents the main novel phenomenon evident in our model.

To our knowledge no such opponent tuning has been identified in auditory neuroscience, but qualitatively similar opponent behavior is evident in visual neurons exhibiting end-stopping or cross-orientation inhibition. The results raise the possibility that coordinated excitation and inhibition could be a feature of central auditory processing, and we thus examined this model property in detail.

To quantitatively substantiate the observation of opponency, we again examined the variation in STKs pooled by the second-layer units, in this case those that pooled STKs with both positive and negative weights. We analyzed all units that had at least one STK weight of each sign whose absolute value exceeded 5% of the maximum weight. For each such second-layer unit we identified centers of mass of the STKs pooled with weights exceeding 5% of the maximum and again computed their standard deviation along time-frequency and modulation coordinates. We compared the standard deviation for STKs pooled with the same sign (separately computed for positive and negative signs) to that of STKs pooled with both signs. If opponent units contrast different types of features we expect the standard deviation of STKs pooled with both signs to be higher than that of STKs pooled with a single sign.

As shown in Fig.10, even for STKs pooled with the same sign, the spread was typically higher for opponent units than for excitatory-only units, for both speech and environmental sounds (blue vs. red bars in Fig.10A and B respectively). However, the average spread of STKs pooled with the both signs by opponent units (Fig.10, dark blue bars) was significantly higher than the spread of STKs pooled with the same sign (Fig.10, light blue bars). This finding held for all dimensions and both corpora ($p < .05$ in all cases, via

t-test).). Moreover, the spread of STKs pooled with the same sign by opponent units were also significantly lower than the spread of random groups of STKs (light blue bars vs. gray bars; $p < .05$ in all cases, via t-test). These results provide quantitative evidence that second-layer units pool features in a structured way and that opponent units assign opposite signs to to groups of STKs that are similar within groups but differ across groups.

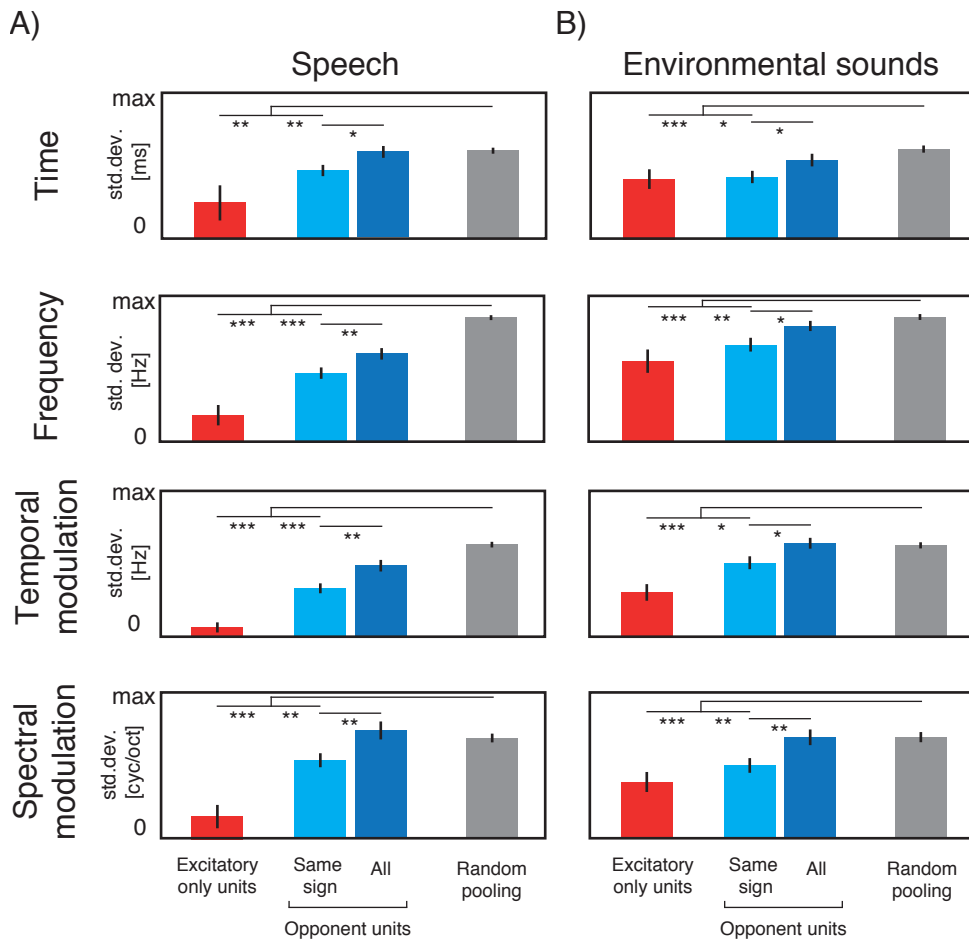


Figure 10: **Statistics of STK pooling by second layer units.** Average standard deviation (spread) of centers of mass of STKs pooled by second layer units. Red bars plot spread of STKs pooled by excitatory-only units. Blue bars plot spread of STKs pooled by opponent units. Light and dark blue bars plot the spread of STKs pooled by opponent units with the same sign (light blue) and with either sign (dark blue). Gray bars plot the spread of random subsets of STKs equal in size to the average number of STKs pooled by second layer units. Error bars plot the standard error of the mean. A) Pooling statistics for model trained on speech. B) Pooling statistics for model trained on environmental sounds.

One might imagine that examination of the STKs pooled by each unit would be sufficient to determine the sort of stimuli eliciting strong positive or negative responses. But because the activation of each unit is the result of a non-linear inference process in which units compete to explain the stimulus pattern [Olshausen and Field, 1997], it is often not obvious what a set of STKs will capture. Thus to understand which stimuli 'excite' or 'inhibit' second-level features, we inferred coefficients v by encoding the entire training dataset and selected two sets of 25 sound epochs that elicited the strongest positive and strongest negative responses, respectively, in each unit. Examples of opponent stimulus patterns encoded by second-layer features are visualized in Fig 11. These positive and negative stimuli are depicted in the second and third columns from the left, respectively.

In the last two columns, the center of mass of each of these stimuli is plotted on the time-frequency plane (fourth column) and modulation plane (fifth column). Although the center of mass of each stimulus is admittedly a crude summary, the simplicity of the representation facilitates visualization and analysis of the clustering of positive and negative stimuli. We note also that because second-layer coefficients and weights can be sign-reversed without changing the representation, the designation of stimuli (and STK weights) as 'excitatory' or 'inhibitory' is arbitrary.

In some cases, some natural function can be ascribed to the unit, particularly upon listening to stimuli eliciting positive and negative responses. For instance, positive activations of the unit shown in Fig 11B appear to encode onsets of voiced speech, while its negative activations encode voicing offsets. By contrast, the unit shown in Fig 11H appears to code speaker gender. For this unit, representations of positive and negative stimuli were mixed on the time-frequency and modulation planes, but the excitatory stimuli exhibited somewhat coarser spectral modulation (red circles in the right column are more concentrated in the lower part of the plane than blue circles). Listening to the stimuli revealed a clear difference in pitch/gender (23 out of 25 of the negative stimuli were vowels uttered by a female speaker, while all of the positive stimuli were vowels uttered by a male speaker, as confirmed by the first author upon listening to the stimuli). This observation underscores the fact that the time-frequency and modulation planes are but two representational spaces in which to assess stimuli, and are used here primarily because they are standard. Moreover, the centroids in these planes are but one way to summarize the STK. Although opponent sets of STKs and positive/negative stimuli often exhibited separation in one of the two planes when analyzed in terms of centroids, they need not do so in order to be distinct.

Taken together, our results demonstrate opponent units that encode STK activation patterns which are mutually exclusive and presumably do not co-occur in natural signals. The phenomenon is a natural one to investigate in the auditory cortex.

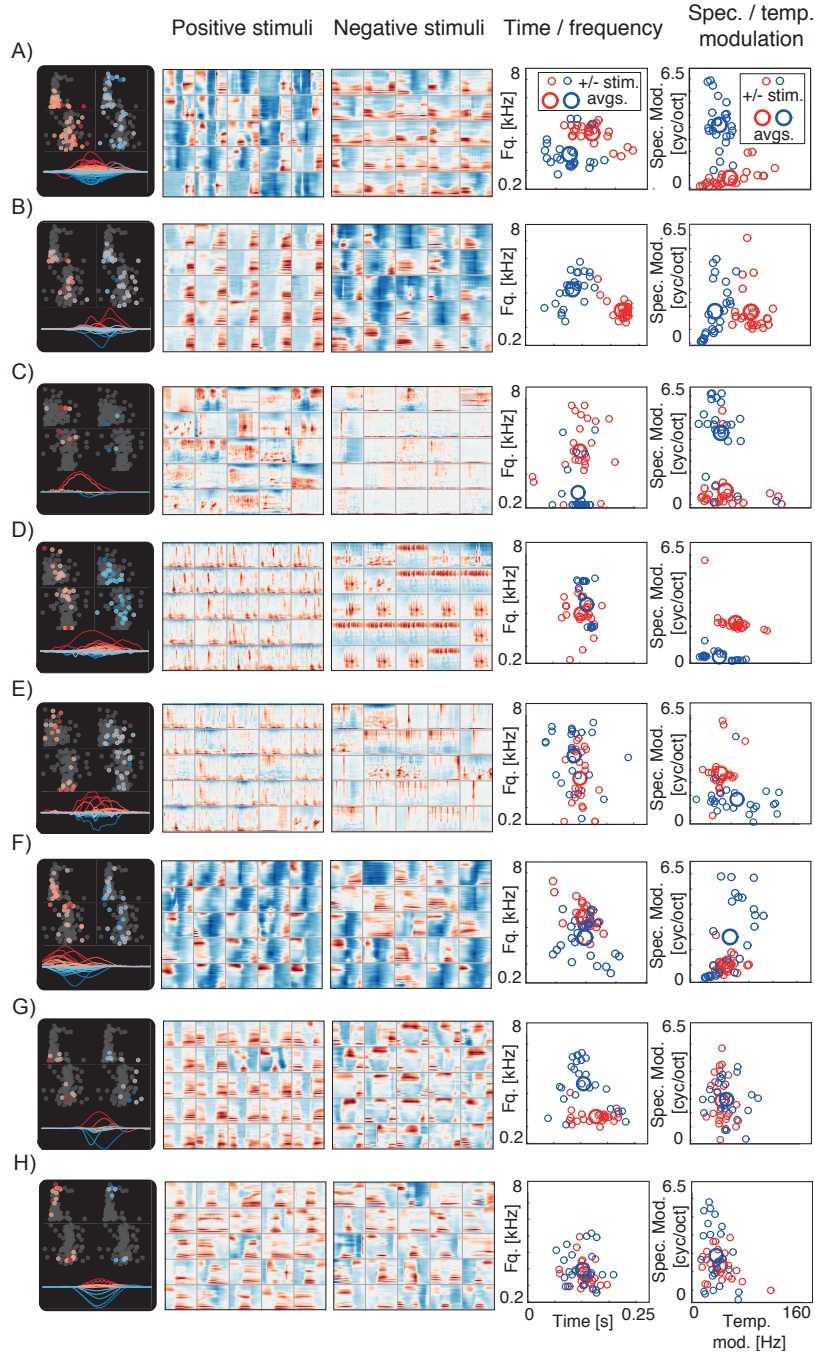


Figure 11: **Opponent stimulus selectivity in second layer units.** Each row corresponds to a particular second-layer unit. The leftmost column plots the center of mass of each first-layer STK in the modulation and time-frequency planes, along with the time courses of their weights in the second-layer unit (as in Fig 7B). The second and third columns from the left depict 25 stimuli eliciting strong positive and negative responses in the corresponding second-layer unit. In the fourth and fifth columns, positive and negative stimuli are visualized as red and blue circles, respectively, in time-frequency and modulation planes (the circle is located at the center of mass of the stimulus). Large circles correspond to centroids of positive and negative stimulus clusters.

Comparison with neurophysiological data

Although our primary goal was to generate predictions of not-yet observed neural representations of sound, we also sought to test whether our model would reproduce known findings from auditory neuroscience. The first layer STKs replicated some fairly standard findings in the STRF literature, as discussed earlier. To compare the results from the second layer to experimental data, we examined their receptive field structure and the specificity of their responses, and compared each to published neurophysiology data.

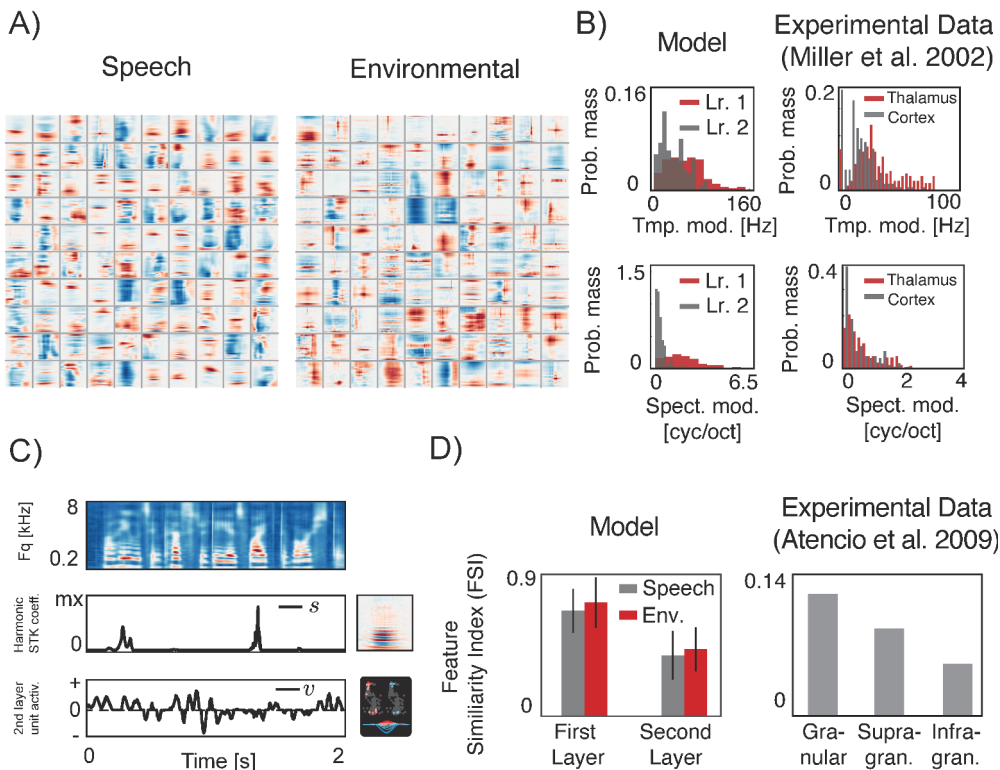


Figure 12: **Comparison with neurophysiological data.** A) Spectrotemporal receptive fields estimated for second layer units trained on speech and environmental sounds. B) Comparison of spectral and temporal modulation tuning of first and second layer receptive fields (panels in left column) to experimental measurements in auditory thalamus and cortex of the cat (panels in right column were provided by Lee Miller [Miller et al., 2002]). Higher processing stages both in the model and the auditory system exhibit tuning for coarser modulations in both frequency and time. Sub-panels with experimental data reprinted with permission of original author. C) Example activation trajectories of first (middle row) and second layer (bottom row) units to an excerpt of speech (top row). The activation of the first layer unit is tightly locked to presence of a preferred stimulus. Responses of the second layer unit are less specifically locked to particular spectrotemporal structures. D) Comparison of tuning specificity, as measured with the Feature Selectivity Index (FSI), for first and second layers of the model trained either on speech or environmental sounds (gray and red bars, respectively), and for different layers of the auditory cortex of the cat (right panel replotted from [Atencio et al., 2009]).

Differences in modulation tuning between first and second model layers

First, we estimated spectrotemporal receptive fields for units in both layers of the model. The receptive fields of the first-layer units are simply the STK of the unit. To estimate receptive fields of a second-layer unit, we drew inspiration from the spike-triggered average, generating a number of cochleagram samples from each basis function and averaging them. To generate samples from the j -th basis function, we set a coefficient v_j to 1 with all other $v_{i \neq j}$ set to zero. We then sampled STK activation trajectories from the distribution dictated by the second-layer feature's coefficient and weights, convolved them with the corresponding STKs and summed the results. We then averaged multiple such samples together. Although we could have computed something more directly analogous to a spike-triggered average, the average sample (which we can compute only because we have the underlying generative model, unlike when conducting a neurophysiology experiment) has the advantage of alleviating the influence of stimulus correlations on the signature of the receptive field. See Appendix B for an illustration of this analysis.

"Receptive fields" obtained in this way are depicted in Fig 12A. To compare the model units to neurophysiological data, we generated histograms of average spectral and temporal modulation frequency (center of mass in the modulation plane) of first- and second-layer receptive fields and plot them next to distributions of preferred modulation frequencies of neurons in the auditory thalamus and cortex of the cat [Miller et al., 2002] (Fig 12B). The same trend is evident in the model and the auditory system: the second-layer prefers features with slower/coarser spectral and temporal modulations relative to the first-layer, mirroring the difference seen between the cortex and thalamus. This analysis used features trained on speech, but environmental sounds yielded qualitatively similar results. In the model as well as the brain, lower modulation frequencies may result in part from combining multiple distinct STKs in downstream units.

Model units reflect neural hierarchical trends of response specificity

Neuronal tuning in early and late stages of the auditory system also tends to differ in specificity [Chechik and Nelken, 2012, Kozlov and Gentner, 2016]. Compared to the auditory brainstem, cortical neurons are less selective and respond to multiple features of sound [Kozlov and Gentner, 2016, Atencio et al., 2012], consistent with an increase in abstraction of the representation [Chechik and Nelken, 2012]. Suggestions of similar behavior in our model are apparent in the activations of first- and second-layer units to sound, an example of which is shown in Fig 12C. The first layer feature (middle row) becomes activated only when it is strongly correlated with the stimulus. In contrast, activations of a typical second layer feature (bottom row) deviate from zero during many, seemingly different parts of the stimulus.

We quantified the specificity of tuning with the feature selectivity index (FSI), a measure introduced previously to quantify how correlated a stimulus has to be with a neuron's spike-triggered average to evoke a response [Miller et al., 2001] (see Appendix B for definition). An FSI equal to 1 implies that a neuron spikes only when a stimulus is precisely aligned with its STRF (defined as the spike-triggered average), whereas an FSI equal to 0 means that the stimuli triggering neural firing are uncorrelated with the STRF. We computed the FSI using the 25 cochleagram excerpts that most strongly activated each of the first- and second-layer units. The average FSI of each model layer is plotted in Fig 12D. Second-layer features are substantially less specific than first-layer features. A similar effect occurs across different cortical layers [Atencio et al., 2009] (Fig 12D, right). Analogous differences seem likely to occur between thalamus and cortex as well, although we are not aware of an explicit prior comparison. The decrease in response specificity in our model can be explained by the fact that second layer units can become activated when any of the

pooled first-layer features (or their combination) appears in the stimulus.

Discussion

Natural sounds are highly structured. The details of this structure and the mechanisms by which it is encoded by the nervous system remain poorly understood. Progress on both fronts is arguably limited by the shortage of signal models capable of explicitly representing natural acoustic structure. We have proposed a novel statistical model that captures an unexplored type of high-order dependency in natural sounds – correlations between the activations of basic spectrotemporal features. Our model consists of two layers. The first layer learns a set of elementary spectrotemporal kernels and uses them to encode sound cochleagrams. The second layer learns a representation of co-occurrence patterns of the first layer features.

We adopted a generative modelling approach, inspired by its previous successes in the domain of natural image statistics. Previous hierarchical, probabilistic models of natural images were able to learn high-order statistical regularities in natural images [Karklin and Lewicki, 2009, Karklin and Lewicki, 2005, Cadieu and Olshausen, 2008, Hyvärinen et al., 2009, Hoyer and Hyvärinen, 2002, Lee et al., 2008, Hosoya and Hyvärinen, 2015, Garrigues and Olshausen, 2008, Berkes et al., 2009]. In many cases the representations learned by these models exhibit similarity to empirically observed neural codes in the visual system.

The hierarchical representations learned by our model provide predictions about neural representations of mid-level sound structure. Moreover, the model reproduces certain aspects of the representational transformations found through the thalamus and cortex. The results suggest that principles of efficient coding could shape mid-level processing in the auditory cortex in addition to the auditory periphery [Lewicki, 2002].

Statistical dependencies in natural sounds

Our model exploits statistical dependencies between first-layer spectrotemporal kernels. These dependencies have two likely causes. First, dependencies almost surely remain from limitations of the convolutional sparse coding model, in that first-layer coefficients are never completely marginally independent even after learning (due to insufficient expressive power of the code). Second, even if the first-layer coefficients were fully marginally independent they could exhibit local dependencies when conditioned on particular sound excerpts. Observations of such conditional dependence have been made previously, mostly in the context of modelling natural image statistics [Karklin and Lewicki, 2009, Karklin and Lewicki, 2005, Hyvärinen et al., 2009, Cadieu and Olshausen, 2008, Karklin et al., 2012]. We similarly observed strong non-zero cross-correlations between STKs for particular natural sound excerpts, and found that the second-layer units captured these sorts of dependencies (Fig. 3,8).

We also observed that different instances of the same acoustic event (such as the same word uttered twice by the same speaker) yield similar STK coefficient trajectories (Fig. 4), which can be thought of as samples from a nonstationary distribution with a time-varying magnitude parameter. By modeling this magnitude the model learned representations with some degree of invariance (Fig. 9).

Model results - Convolutional spectrotemporal features

The model first learned a sparse convolutional decomposition of cochleagrams. Convolutional representations are less redundant than “patch-based” codes and can represent signals of arbitrary length, and have been successfully applied to spectrogram analysis in engineering (e.g. [Blumensath and Davies, 2006, Grosse et al., 2012]). The features learned

with convolutional sparse coding improved sound classification accuracy, musical note extraction and source separation in prior work [Blumensath and Davies, 2006, Grosse et al., 2012]. By contrast, the model presented here is, to our knowledge, the first use of convolutional sparse coding of audio in a neuroscience context. Accordingly, we provide an in-depth analysis of feature shapes learned by the first layer and relate them to known spectrotemporal tuning properties of auditory neurons.

The learned spectrotemporal kernels span clicks, harmonics, combinations of harmonics, bandpass noise, frequency sweeps, onsets, and offsets. Some of these structures are present in previous learned codes of spectrograms, but due to the convolutional nature of the learned code, the model learns only a single version of each feature rather than replicating it at different time points. Although difficult to quantify, our impression is that the code is more diverse than that obtained with previous patch-based approaches [Carlson et al., 2012, Klein et al., 2003]. As in previous neurophysiological measurements of cortical STRFs, the model STKs tile the spectrotemporal modulation plane, subject to the constraints of the tradeoff between time and frequency [Singh and Theunissen, 2003].

As with features learned from sound waveforms [Lewicki, 2002, Smith and Lewicki, 2006], we found the first-layer spectrogram features to depend on the training corpus (speech or a set of environmental sounds). Although certain spectrotemporal patterns (e.g. single harmonics or clicks) appeared in both features sets, others were corpus-specific. Corpus-specific structure was also evident in the distributions of features in the modulation plane. This corpus dependence contrasts with the relatively consistent occurrence of Gabor-like features in sparse codes of images. These results raise the possibility that auditory cortical neuronal tuning might exhibit considerable heterogeneity, given that the full range of natural audio must be encoded by cortical neurons.

Model results – Second-layer features

The second layer of the model encodes a probability distribution of spectrotemporal feature activations. By jointly modelling the magnitudes of the first layer responses, the second layer basis functions capture patterns of spectrotemporal kernel covariation. Magnitude modeling was essential to learning additional structure – we found in pilot experiments that simply applying a second layer of convolutional sparse coding did not produce comparable results. This observation is consistent with previous findings that nonlinear transformations of sparse codes often help to learn residual dependencies [Karklin and Lewicki, 2009, Shan et al., 2007, Hyvärinen et al., 2009].

The second layer of our model learned a representation of spectrotemporal feature co-activations that frequently occur in natural sounds. Typically, the spectrotemporal features pooled by a second-layer unit shared some property, often evident in at least one of the time-frequency or modulation planes - e.g. frequency content, temporal pattern or modulation characteristics (Fig. 7), quantified by the similarity in these planes. Our model also identified 'opponent' patterns – sets of spectrotemporal features that are pooled with opposite sign, presumably because they are rarely active simultaneously in natural sounds (Fig. 11). The opponent features for a second-layer unit were often at least partially separated in at least one of the spectrotemporal modulation or time-frequency planes (10).

Relation to auditory neuroscience

The structure learned by the model from natural sounds replicates some known properties of the auditory system. The modulation frequencies preferred by units dropped from the first to the second layer (Fig. 12B), as has been observed between the thalamus and cortex. Unit “tuning” specificity (i.e. the similarity between stimuli eliciting a strong response) also decreased from the first to the second layer (Fig. 12C,D). A similar specificity decrease has

been observed between granular and supra- and infra-granular layers of the cortex [Atencio et al., 2009]. Although we do not suggest a detailed correspondence between layers of our model and particular anatomical structures, these similarities indicate that some of the principles underlying hierarchical organization of the auditory pathways may derive from the natural sound statistics and architectural choices that constrain our model.

The combinations of spectrotemporal features learned by our model provide hypotheses for neural tuning that might be present in the auditory system. Some units combined only STKs with similar acoustic properties (Fig. 7, 10). Others encode activations of 'opponent' sets of first-layer kernels. Such opponent kernel sets are presumably those that do not typically become active simultaneously in the training corpus. The results are somewhat analogous to excitation-inhibition phenomena in visual neurophysiology (such as end-stopping, length and width suppression or cross-orientation inhibition) emerging in models of natural images [Karklin and Lewicki, 2009, Coen-Cagli et al., 2012, Zhu and Rozell, 2013, Hosoya and Hyvärinen, 2015]. The second-layer units of our model provide candidates of potentially analogous phenomena in the auditory cortex.

Our results are also relevant to recent evidence that central auditory neurons in some cases are driven by more than one stimulus feature [Atencio et al., 2009, Sharpee et al., 2011, Kozlov and Gentner, 2016]. Our model reproduces that trend, pooling up to dozens of single layer features with strong weights, and predicts that distinct dimensions may be combined in opponent fashion.

Relation to other modeling approaches

The model described here represents one of several approaches to construction of hierarchical signal representations. The representations in our model are learned from natural sounds, and thus contrast with hand-engineered models that seek to replicate known or hypothesized features of sensory coding [Chi et al., 2005, McDermott and Simoncelli, 2011]. One advantage of learning models from natural signals is that any similarities with known neural phenomena provide candidate normative explanations for these phenomena, e.g. that they arise from the demands of efficient coding or some other optimality constraint imposed by the model [Olshausen and Field, 1997, Lewicki, 2002]. Another potential advantage is that one might hope to learn structures that have not yet been observed neurally, but that could provide hypotheses for future experiments. The latter was the main motivation for our modeling approach.

The model is also unsupervised, and generative, specifying a joint probability distribution over the data and coefficients in both latent layers. An alternative approach to learning hierarchical representations of data is to use discriminative models, which optimize a representation for performance of a particular task. Prominent recent examples of such discriminative learning come from the field of deep neural networks [Yamins et al., 2014]. Models of this class are comparatively straightforward to optimize and can yield high performance on classification tasks, but typically require large numbers of labeled training examples. The requirement of labeled training data is sometimes a practical limitation, and raises questions about the extent to which the learning procedure could relate to the brain. Generative models currently have the disadvantage of requiring custom optimization procedures. However, their generative nature facilitates certain applications (e.g. denoising) and allows samples to be straightforwardly generated, potentially for use in experiments. Moreover, they do not require labeled data, and learn representations that are independent of any particular task, much like the unsupervised learning believed to occur in sensory systems.

Conclusion

We have presented a hierarchical model of natural sounds. When trained on natural sound corpora, the model learned a representation of spectrotemporal feature combinations. The properties of the model layers resemble aspects of hierarchical transformations previously observed in the brain, suggesting that efficient coding could shape such transformations throughout the auditory system. The learned mid-level features provide hypotheses for auditory cortical tuning as well as a means to parameterize stimuli with which to probe mid-level audition.

Appendix A

Gradients for learning and inference in the first layer Learning and inference in the first layer was achieved via gradient descent on negative log-posterior (Eq 5). Below we provide expressions of E_1 gradient with respect to $s_{i,t}$ and $\phi_{i,f,\tau}$ respectively

The non-negativity constraint on sparse coefficients s complicates the optimization process in learning and inference. To alleviate this complication, we introduced auxiliary coefficients z_i and assumed that non-negative sparse coefficients s_i are equal to squares of z_i :

$$s_i = z_i^2, z_i \in \mathbb{R} \quad (11)$$

We replace s_i with z_i^2 in equations below.

Gradients of first layer energy function (Eq 5) with respect to $z_{i,t}$ and $\phi_{i,f,\tau}$ was respectively:

$$\frac{\partial E_1}{\partial z_{i,t}} \propto -\frac{4z_{i,t}}{\sigma^2} \sum_{f=1}^F (\phi_{i,f} \odot e_f)_t + 2\lambda_i z_{i,t} \quad (12)$$

$$\frac{\partial E_1}{\partial \phi_{i,f,\tau}} \propto -2 \left(s_i \odot e_f \right)_\tau \quad (13)$$

where \odot denotes cross-correlation and $e_{f,t} = x_{f,t} - \hat{x}_{f,t}$ is the reconstruction error.

Gradients for learning and inference in the second layer Similarly, learning and inference in the second layer was achieved via gradient descent on the corresponding energy function (equation 10). Gradient expressions took the following form:

$$\frac{\partial E_2}{\partial v_{i,t}} \propto \sum_{n=1}^N \psi_n \odot B_{i,n} + \alpha \operatorname{sgn}(v_{i,t}) \quad (14)$$

$$\frac{\partial E_2}{\partial B_{j,i,t_b}} \propto (1 - \psi_i \odot v_j)_{t_b} + \beta \operatorname{sgn}(B_{j,i,t_b}) \quad (15)$$

where $\hat{\lambda}_{n,t} = \lambda_n \exp \left[\sum_j^M v_j * B_j \right]_t$ is the reconstruction of instantaneous scale parameter and $\psi_{n,t} = \frac{s_{n,t}}{\hat{\lambda}_{n,t}}$.

Appendix B

In this appendix we describe the methods used to estimate and analyze receptive fields of model units.

Receptive field estimation

Fig. B1 A illustrates spectrotemporal receptive field estimation of model units through averaging of strongly activating stimuli (akin to computing spike-triggered average), and the biases that result.

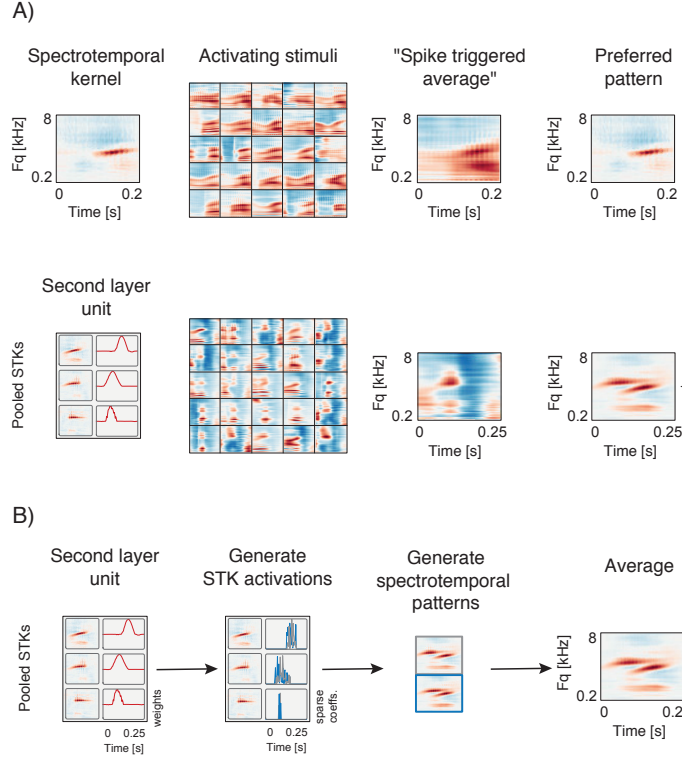


Figure B1: **Estimation of spectrotemporal receptive fields of model units.** A) Top: an example first layer STK (leftmost column) and stimuli which strongly activate it (second column from the left). The "spike triggered average" i.e., the mean of strongly activating stimuli (third column from the left) differs strongly from the true underlying kernel (rightmost column) due to stimulus correlations. Bottom: same as top row for an example second layer unit. Second column from the left depicts strongly activating positive stimuli. Spectrotemporal pattern in rightmost column is derived from the procedure shown in (B). B) Generative estimation of second layer spectrotemporal receptive fields. To avoid the influence of stimulus correlations, we generate a number of STK activation trajectories (second column from the left; gray and blue lines plot two different sets of sampled trajectories) from a distribution encoded by a second-layer unit of interest (leftmost column). For each set of sampled trajectories, we convolve the trajectories with the corresponding STK and sum the results to yield a sampled spectrotemporal pattern (2 example patterns are shown in the third column from the left, resulting from the gray and blue samples, respectively). The set of such spectrotemporal patterns are then averaged to estimate the "true" receptive field.

The "true" spectrotemporal pattern encoded by a STK is simply the STK itself. To identify average spectrotemporal pattern encoded by a second-layer unit we take advantage of the fact that our model is generative (Fig. B1). For each unit we generated a number of STK coefficient trajectories (second panel from the left, Fig. B1 B) from the distribution encoded by that particular unit (e.g. leftmost panel, Fig. B1 B). We then generated spectrotemporal patterns by convolving the trajectories with the corresponding STKs and summing the results together (see Fig.2). Due to the stochasticity of the generated STK trajectories, some variability is visible among the generated spectrotemporal patterns (third panel from the left, Fig. B1 B). We then averaged these patterns to form an estimate of a spectrotemporal receptive field of a second layer unit (right panel, Fig. B1 B). Because the estimation process does not involve averaging stimuli, we avoid biasing the estimate of the average preferred spectrotemporal pattern with stimulus correlations.

It is apparent that "spike-triggered averages" (third column, Fig. B1 A and B) differ from the true stimulus representation in the model for both layers (as embodied by the STK or the average spectrotemporal pattern generated by a second-layer unit). The discrepancy is due to presence of strong correlations in natural stimuli.

Feature Selectivity Index

For comparison with neural data, we computed the feature selectivity index (FSI) proposed in [Miller et al., 2001]. The FSI is a number lying in the $[0, 1]$ interval. FSI values close to 1 imply that stimuli eliciting a response of a neuron (or, in our case, a model unit) are similar to each other (specifically, the stimuli are close to the mean stimulus eliciting a response). When the FSI value is close to 0, the corresponding neuron spikes at random i.e. stimuli preceding spikes are uncorrelated with the spike-triggered average. The relevance of the FSI for our purposes is that a neuron or model unit that exhibits invariance to some type of stimulus variation should have an FSI less than 1. By comparing the FSI across layers we hoped to quantify differences in the degree of representational abstraction.

The FSI computation procedure are described in detail in [Miller et al., 2001]. The only discrepancy between the use of the FSI here and its prior use in neurophysiological studies is that units in our model are continuously active, and do not discretely spike. We emulated the selection of stimuli eliciting spikes by selecting the stimulus excerpts yielding the highest activation of model layer units. In the first layer, we selected the 25 stimuli yielding highest positive activation. In the second layer, we separately computed FSI indices for stimuli eliciting positive and negative responses, using 25 stimuli per unit in each case. We then averaged the indices obtained for the two sets of stimuli for each unit.

Computing FSI for an i -th unit consists of the following steps:

1. Compute average of strongly activating stimuli (analogous to spike-triggered average - STA) - separately for positive and negative stimuli in the case of second-layer units.
2. Compute correlations $c_{S,i,n}$ between each strongly activating stimulus n and its respective STA.
3. Compute correlations $c_{R,i,n}$ between the STA and a randomly selected subset of stimuli n .
4. Compute the area $A_{S,i}$ under the empirical cumulative distribution function of correlations $c_{S,i}$, $A_{S,i} = \int_{-1}^1 ECDF(c_{S,i})dc_{S,i}$

5. Compute area $A_{R,i}$ under the empirical cumulative distribution function of $A_{R,i} = \int_{-1}^1 ECDF(c_{R,i})dc_{R,i}$
6. The FSI of each unit is defined as: $FSI_i = \frac{(A_{R,i}-A_{S,i})}{A_{R,i}}$

Appendix C

In order to verify the correctness of the learning algorithm, we generated a toy dataset using 9 second-layer kernels, out of which 3 were "opponent" (Fig.C1 A, left panel). We generated sample STK activations by randomly superimposing the second-layer kernels to create a "variance map" (Fig.C1 B), from which we sampled STK trajectories (Fig.C1 C). Each iteration of the learning procedure made a gradient step on the model parameters using a single sample of STK trajectories (from a single sampled variance map). After 1500 iterations (Fig.C1 D), the model converged on the solution visualized in the right panel of Fig.C1 A. The generative kernels were recovered up to permutation, sign changes and in some cases, a temporal shift. These latter differences are expected due to the inherent arbitrariness of positive and negative signs in the model and the shift-invariance inherent to a convolutional code. These results demonstrate that the learning algorithm is capable of converging to the correct solution when the data is well-described by the model.

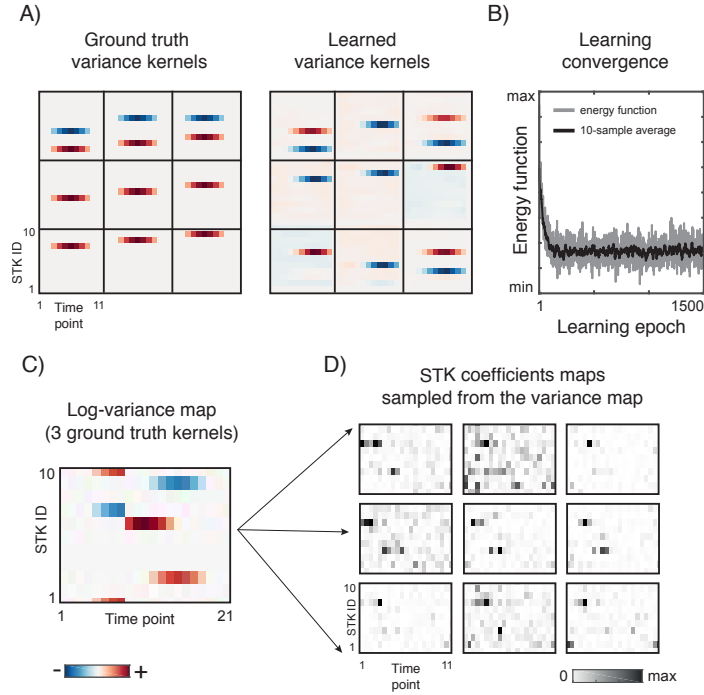


Figure C1: **The model recovers ground truth second-layer kernels from a training dataset.** A) Left panel - nine ground truth second-layer kernels used to generate the toy dataset. The three kernels in the top row are "opponent". Right panel - the generative second-layer kernels recovered by the model. They match the ground truth kernels up to permutation, sign change and time shifts. B) Example variance map generated by the model from the second-layer kernels. Such maps result from superimposing second-layer kernels with coefficients sampled from their sparse prior. This example was generated from three second-layer kernels from panel (A). The variance map encodes temporal changes in the distribution of STK activations. The toy training dataset was generated from 1500 such sampled maps (one per iteration of the learning algorithm). C) STK activations sampled from the variance map in panel B. In the actual training procedure, a single sample was used for each sampled variance map, but here we show 9 examples to illustrate the stochastic nature of the procedure. D) The convergence of the learning algorithm over 1500 learning epochs. The gray line depicts values of the second-layer energy function (see Eq.10). The black line is a 10-epoch moving average.

Acknowledgments

This material is based upon work supported by the Center for Brains, Minds and Machines (CBMM), funded by NSF STC award CCF-1231216, and by a McDonnell Scholar Award to JHM. The authors thank the members of the McDermott lab for helpful comments on earlier drafts of the manuscript.

References

- [Aertsen and Johannesma, 1981] Aertsen, A. and Johannesma, P. (1981). The spectro-temporal receptive field. *Biological cybernetics*, 42(2):133–143.
- [Atencio et al., 2009] Atencio, C. A., Sharpee, T. O., and Schreiner, C. E. (2009). Hierarchical computation in the canonical auditory cortical circuit. *Proceedings of the National Academy of Sciences*, 106(51):21894–21899.
- [Atencio et al., 2012] Atencio, C. A., Sharpee, T. O., and Schreiner, C. E. (2012). Receptive field dimensionality increases from the auditory midbrain to cortex. *Journal of neurophysiology*, 107(10):2594–2603.
- [Attneave, 1954] Attneave, F. (1954). Some informational aspects of visual perception. *Psychological review*, 61(3):183.
- [Barlow, 1961] Barlow, H. B. (1961). Possible principles underlying the transformations of sensory messages.
- [Bell and Sejnowski, 1997] Bell, A. J. and Sejnowski, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vision research*, 37(23):3327–3338.
- [Bendor and Wang, 2005] Bendor, D. and Wang, X. (2005). The neuronal representation of pitch in primate auditory cortex. *Nature*, 436(7054):1161–1165.
- [Berkes et al., 2009] Berkes, P., Turner, R. E., and Sahani, M. (2009). A structured model of video reproduces primary visual cortical organisation. *PLoS Comput Biol*, 5(9):e1000495.
- [Bizley et al., 2005] Bizley, J. K., Nodal, F. R., Nelken, I., and King, A. J. (2005). Functional organization of ferret auditory cortex. *Cerebral Cortex*, 15(10):1637–1653.
- [Blumensath and Davies, 2006] Blumensath, T. and Davies, M. (2006). Sparse and shift-invariant representations of music. *Audio, Speech, and Language Processing, IEEE Transactions on*, 14(1):50–57.
- [Bumbacher and Ming, 2012] Bumbacher, E. and Ming, V. (2012). Pitch-sensitive components emerge from hierarchical sparse coding of natural sounds. pages 219–229.
- [Cadieu and Olshausen, 2008] Cadieu, C. and Olshausen, B. A. (2008). Learning transformational invariants from natural movies. In *Advances in neural information processing systems*, pages 209–216.
- [Carlin and Elhilali, 2013] Carlin, M. A. and Elhilali, M. (2013). Sustained firing of model central auditory neurons yields a discriminative spectro-temporal representation for natural sounds. *PLOS Comput Biol*, 9(3):e1002982.
- [Carlson et al., 2012] Carlson, N. L., Ming, V. L., and DeWeese, M. R. (2012). Sparse codes for speech predict spectrotemporal receptive fields in the inferior colliculus. *PLoS Comput Biol*, 8(7):e1002594.

- [Carruthers et al., 2015] Carruthers, I. M., Laplagne, D. A., Jaegle, A., Briguglio, J. J., Mwilambwe-Tshilobo, L., Natan, R. G., and Geffen, M. N. (2015). Emergence of invariant representation of vocalizations in the auditory cortex. *Journal of neurophysiology*, 114(5):2726–2740.
- [Chechik and Nelken, 2012] Chechik, G. and Nelken, I. (2012). Auditory abstraction from spectro-temporal features to coding auditory entities. *Proceedings of the National Academy of Sciences*, 109(46):18968–18973.
- [Chi et al., 2005] Chi, T., Ru, P., and Shamma, S. A. (2005). Multiresolution spectrotemporal analysis of complex sounds. *The Journal of the Acoustical Society of America*, 118(2):887–906.
- [Coen-Cagli et al., 2012] Coen-Cagli, R., Dayan, P., and Schwartz, O. (2012). Cortical surround interactions and perceptual salience via natural scene statistics. *PLoS Comput Biol*, 8(3):e1002405.
- [Depireux et al., 2001] Depireux, D. A., Simon, J. Z., Klein, D. J., and Shamma, S. A. (2001). Spectro-temporal response field characterization with dynamic ripples in ferret primary auditory cortex. *Journal of neurophysiology*, 85(3):1220–1234.
- [Elie and Theunissen, 2015] Elie, J. E. and Theunissen, F. E. (2015). Meaning in the avian auditory cortex: neural representation of communication calls. *European Journal of Neuroscience*, 41(5):546–567.
- [Fritz et al., 2003] Fritz, J., Shamma, S., Elhilali, M., and Klein, D. (2003). Rapid task-related plasticity of spectrotemporal receptive fields in primary auditory cortex. *Nature neuroscience*, 6(11):1216–1223.
- [Garofolo et al., 1993] Garofolo, J. S., Lamel, L. F., Fisher, W. M., Fiscus, J. G., Pallett, D. S., Dahlgren, N. L., and Zue, V. (1993). Timit acoustic-phonetic continuous speech corpus. *Linguistic data consortium, Philadelphia*, 33.
- [Garrigues and Olshausen, 2008] Garrigues, P. and Olshausen, B. A. (2008). Learning horizontal connections in a sparse coding model of natural images. In *Advances in Neural Information Processing Systems*, pages 505–512.
- [Glasberg and Moore, 1990] Glasberg, B. R. and Moore, B. C. (1990). Derivation of auditory filter shapes from notched-noise data. *Hearing research*, 47(1):103–138.
- [Grosse et al., 2012] Grosse, R., Raina, R., Kwong, H., and Ng, A. Y. (2012). Shift-invariance sparse coding for audio classification. *arXiv preprint arXiv:1206.5241*.
- [Harper et al., 2016] Harper, N. S., Schoppe, O., Willmore, B. D., Cui, Z., Schnupp, J. W., and King, A. J. (2016). Network receptive field modeling reveals extensive integration and multi-feature selectivity in auditory cortical neurons. *PLOS Comput Biol*, 12(11):e1005113.
- [Hosoya and Hyvärinen, 2015] Hosoya, H. and Hyvärinen, A. (2015). A hierarchical statistical model of natural images explains tuning properties in v2. *The Journal of Neuroscience*, 35(29):10412–10428.
- [Hoyer and Hyvärinen, 2002] Hoyer, P. O. and Hyvärinen, A. (2002). A multi-layer sparse coding network learns contour coding from natural images. *Vision research*, 42(12):1593–1605.

- [Hyvärinen and Hoyer, 2000] Hyvärinen, A. and Hoyer, P. (2000). Emergence of phase- and shift-invariant features by decomposition of natural images into independent feature subspaces. *Neural computation*, 12(7):1705–1720.
- [Hyvärinen et al., 2009] Hyvärinen, A., Hurri, J., and Hoyer, P. O. (2009). *Natural Image Statistics: A Probabilistic Approach to Early Computational Vision.*, volume 39. Springer Science & Business Media.
- [Hyvärinen and Raju, 2002] Hyvärinen, A. and Raju, K. (2002). Imposing sparsity on the mixing matrix in independent component analysis. *Neurocomputing*, 49(1):151–162.
- [Karklin et al., 2012] Karklin, Y., Ekanadham, C., and Simoncelli, E. P. (2012). Hierarchical spike coding of sound. In *Advances in neural information processing systems*, pages 3032–3040.
- [Karklin and Lewicki, 2005] Karklin, Y. and Lewicki, M. S. (2005). A hierarchical bayesian model for learning nonlinear statistical regularities in nonstationary natural signals. *Neural computation*, 17(2):397–423.
- [Karklin and Lewicki, 2009] Karklin, Y. and Lewicki, M. S. (2009). Emergence of complex cell properties by learning to generalize in natural scenes. *Nature*, 457(7225):83–86.
- [Klein et al., 2003] Klein, D. J., König, P., and Körding, K. P. (2003). Sparse spectrotemporal coding of sounds. *EURASIP Journal on Advances in Signal Processing*, 2003(7):1–9.
- [Kozlov and Gentner, 2016] Kozlov, A. S. and Gentner, T. Q. (2016). Central auditory neurons have composite receptive fields. *Proceedings of the National Academy of Sciences*, page 201506903.
- [Lee et al., 2008] Lee, H., Ekanadham, C., and Ng, A. Y. (2008). Sparse deep belief net model for visual area v2. In *Advances in neural information processing systems*, pages 873–880.
- [Lee et al., 2009] Lee, H., Pham, P., Largman, Y., and Ng, A. Y. (2009). Unsupervised feature learning for audio classification using convolutional deep belief networks. In *Advances in neural information processing systems*, pages 1096–1104.
- [Lewicki, 2002] Lewicki, M. S. (2002). Efficient coding of natural sounds. *Nature neuroscience*, 5(4):356–363.
- [Machens et al., 2004] Machens, C. K., Wehr, M. S., and Zador, A. M. (2004). Linearity of cortical receptive fields measured with natural sounds. *The Journal of neuroscience*, 24(5):1089–1100.
- [McDermott and Simoncelli, 2011] McDermott, J. H. and Simoncelli, E. P. (2011). Sound texture perception via statistics of the auditory periphery: Evidence from sound synthesis. *Neuron*, 71(5):926–940.
- [Mesgarani et al., 2008] Mesgarani, N., David, S. V., Fritz, J. B., and Shamma, S. A. (2008). Phoneme representation and classification in primary auditory cortex. *The Journal of the Acoustical Society of America*, 123(2):899–909.
- [Miller et al., 2002] Miller, L. M., Escabí, M. A., Read, H. L., and Schreiner, C. E. (2002). Spectrotemporal receptive fields in the lemniscal auditory thalamus and cortex. *Journal of neurophysiology*, 87(1):516–527.

- [Miller et al., 2001] Miller, L. M., Escabi, M. A., and Schreiner, C. E. (2001). Feature selectivity and interneuronal cooperation in the thalamocortical system. *The Journal of Neuroscience*, 21(20):8136–8144.
- [Mlynarski, 2014] Mlynarski, W. (2014). Efficient coding of spectrotemporal binaural sounds leads to emergence of the auditory space representation. *Name: Frontiers in Computational Neuroscience*.
- [Mlynarski, 2015] Mlynarski, W. (2015). The opponent channel population code of sound location is an efficient representation of natural binaural sounds. *PLoS Comput Biol*, 11(5):e1004294.
- [Norman-Haignere et al., 2015] Norman-Haignere, S., Kanwisher, N. G., and McDermott, J. H. (2015). Distinct cortical pathways for music and speech revealed by hypothesis-free voxel decomposition. *Neuron*, 88(6):1281–1296.
- [Obleser et al., 2010] Obleser, J., Leaver, A., VanMeter, J., and Rauschecker, J. P. (2010). Segregation of vowels and consonants in human auditory cortex: evidence for distributed hierarchical organization. *Frontiers in psychology*, 1:232.
- [Olshausen and Field, 1997] Olshausen, B. A. and Field, D. J. (1997). Sparse coding with an overcomplete basis set: A strategy employed by v1? *Vision research*, 37(23):3311–3325.
- [Overath* et al., 2015] Overath*, T., McDermott*, J. H., Zarate, J. M., and Poeppel, D. (2015). The cortical analysis of speech-specific temporal structure revealed by responses to sound quilts. *Nature Neuroscience*, 18:903–911.
- [Robles and Ruggero, 2001] Robles, L. and Ruggero, M. A. (2001). Mechanics of the mammalian cochlea. *Physiological reviews*, 81(3):1305–1352.
- [Russ et al., 2008] Russ, B. E., Ackelson, A. L., Baker, A. E., and Cohen, Y. E. (2008). Coding of auditory-stimulus identity in the auditory non-spatial processing stream. *Journal of neurophysiology*, 99(1):87–95.
- [Sahani and Linden, 2003] Sahani, M. and Linden, J. (2003). How linear are auditory cortical responses? 15:125.
- [Shan et al., 2007] Shan, H., Zhang, L., and Cottrell, G. W. (2007). Recursive ica. *Advances in neural information processing systems*, 19:1273.
- [Sharpee et al., 2011] Sharpee, T. O., Atencio, C. A., and Schreiner, C. E. (2011). Hierarchical representations in the auditory cortex. *Current opinion in neurobiology*, 21(5):761–767.
- [Singh and Theunissen, 2003] Singh, N. C. and Theunissen, F. E. (2003). Modulation spectra of natural sounds and ethological theories of auditory processing. *The Journal of the Acoustical Society of America*, 114(6):3394–3411.
- [Smith and Lewicki, 2006] Smith, E. C. and Lewicki, M. S. (2006). Efficient auditory coding. *Nature*, 439(7079):978–982.
- [Srinivasan et al., 1982] Srinivasan, M. V., Laughlin, S. B., and Dubs, A. (1982). Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London B: Biological Sciences*, 216(1205):427–459.

- [Terashima and Okada, 2012] Terashima, H. and Okada, M. (2012). The topographic unsupervised learning of natural sounds in the auditory cortex. In *Advances in Neural Information Processing Systems*, pages 2312–2320.
- [Theunissen et al., 2000] Theunissen, F. E., Sen, K., and Doupe, A. J. (2000). Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds. *The Journal of Neuroscience*, 20(6):2315–2331.
- [van Hateren and van der Schaaf, 1998] van Hateren, J. H. and van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society of London B: Biological Sciences*, 265(1394):359–366.
- [Williamson et al., 2016] Williamson, R. S., Ahrens, M. B., Linden, J. F., and Sahani, M. (2016). Input-specific gain modulation by local sensory context shapes cortical and thalamic responses to complex sounds. *Neuron*, 91(2):467–481.
- [Woolley et al., 2005] Woolley, S. M., Fremouw, T. E., Hsu, A., and Theunissen, F. E. (2005). Tuning for spectro-temporal modulations as a mechanism for auditory discrimination of natural sounds. *Nature neuroscience*, 8(10):1371–1379.
- [Yamins et al., 2014] Yamins, D. L., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., and DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624.
- [Zhu and Rozell, 2013] Zhu, M. and Rozell, C. J. (2013). Visual nonclassical receptive field effects emerge from sparse coding in a dynamical system. *PLoS Comput Biol*, 9(8):e1003191.