

Auditory Cortical Receptive Fields: Stable Entities with Plastic Abilities

Mounya Elhilali,¹ Jonathan B. Fritz,¹ Tai-Shih Chi,² and Shihab A. Shamma¹

¹Center for Auditory and Acoustic Research, Institute for Systems Research, University of Maryland, College Park, College Park, Maryland 20742, and

²Department of Communication Engineering, National Chiao Tung University, Taiwan

To form a reliable, consistent, and accurate representation of the acoustic scene, a reasonable conjecture is that cortical neurons maintain stable receptive fields after an early period of developmental plasticity. However, recent studies suggest that cortical neurons can be modified throughout adulthood and may change their response properties quite rapidly to reflect changing behavioral salience of certain sensory features. Because claims of adaptive receptive field plasticity could be confounded by intrinsic, labile properties of receptive fields themselves, we sought to gauge spontaneous changes in the responses of auditory cortical neurons. In the present study, we examined changes in a series of spectrotemporal receptive fields (STRFs) gathered from single neurons in successive recordings obtained over time scales of 30–120 min in primary auditory cortex (A1) in the quiescent, awake ferret. We used a global analysis of STRF shape based on a large database of A1 receptive fields. By clustering this STRF space in a data-driven manner, STRF sequences could be classified as stable or labile. We found that >73% of A1 neurons exhibited stable receptive field attributes over these time scales. In addition, we found that the extent of intrinsic variation in STRFs during the quiescent state was insignificant compared with behaviorally induced STRF changes observed during performance of spectral auditory tasks. Our results confirm that task-related changes induced by attentional focus on specific acoustic features were indeed confined to behaviorally salient acoustic cues and could be convincingly attributed to learning-induced plasticity when compared with “spontaneous” receptive field variability.

Key words: auditory; behavior; cortex; plasticity; receptive field; stability

Introduction

Receptive fields characterize how sensory information is processed, encoded, and mapped to guide perception and behavior (Hubel and Wiesel, 1968; Aertsen and Johannesma, 1981; Eggermont et al., 1981; DeAngelis et al., 1995; Schreiner, 1995; Fitzpatrick, 2000; Ghazanfar et al., 2001; Ringach, 2004; Bair, 2005; Martinez, 2006). Considered as the fundamental building blocks of perception, a plausible assumption was that receptive fields (RFs) should remain stable to maintain perceptual constancy. Although intuitively appealing, the assumption of RF stability was challenged by compelling experimental evidence for RF plasticity, both during development (Hubel and Wiesel, 1963; de Villers-Sidani et al., 2007) in the adult brain (for review, see Fritz et al., 2005b; Weinberger, 2007) or induced by arousal, attention, and stimulus and behavioral context.

A central paradox arises, however, as to how the brain maintains a stable and consistent image of the sensory information, while at the same time adapting to changing behavioral demands

by changing RF properties. Moreover, to heighten the paradox, how can we possibly hope to measure induced RF plasticity in a constantly roiling sea of RF change? In the auditory system, recent studies have addressed this paradox by investigating the nature and magnitude of spontaneous daily variability by measuring variations in frequency tuning parameters of cortical neurons. Galvan et al. (2001) monitored cortical local field potentials (LFPs) over the course of several weeks in A1 of awake guinea pigs and reported small daily changes in frequency tuning (~0.2 octaves). Recent long-term recordings from A1 neurons in the naive owl monkey also indicate stability in RF best frequency (D. Blake and J. Fritz, personal communication). However, previous studies recorded from individual cortical neurons that showed clear variation in RFs as the animal moved between different states of vigilance (wakefulness, slow wave, and paradoxical sleep) (Edeline et al., 2001; Edeline, 2003) although there were no significant population level changes between states. Kisley and Gerstein (2001) also suggested that A1 RFs in anesthetized animals could vary substantially over the course of a week, but the magnitude of the spontaneous RF variation they reported was significantly smaller than the RF changes observed in learning-induced RF plasticity in auditory cortex in their own study (Kisley and Gerstein, 2001) or in other studies (Weinberger et al., 1993).

In this study, we address the question of cortical RF stability, monitored in awake animals over a much finer time scale than the days or weeks of previous studies. Given recent findings of rapid task-related plasticity in RFs in the primary auditory cortex A1

Received April 1, 2007; revised July 10, 2007; accepted Aug. 2, 2007.

This work was supported in part by National Institutes of Health Grants R01DC005779 and R01DC05937 and by National Science Foundation–National Institutes of Health Collaborative Research in Computational Neuroscience Grant R01 AG02757301. We thank Tamar Vardi for assistance with animal training and Shantanu Ray and Mahdvi Jain for technical assistance in electronics and major contributions in customized software design.

Correspondence should be addressed to Mounya Elhilali, Center for Auditory and Acoustic Research, Institute for Systems Research, 2202 A. V. Williams Building, University of Maryland, College Park, College Park, MD 20742. E-mail: mounya@umd.edu.

DOI:10.1523/JNEUROSCI.1462-07.2007

Copyright © 2007 Society for Neuroscience 0270-6474/07/2710372-11\$15.00/0

over the course of minutes or hours (Fritz et al., 2003, 2005a,b), it is critical to evaluate RF stability in behaviorally naive animals over a similar time course to clarify the extent of spontaneous RF variability, which may reflect normal dynamics. Hence, the goal of this research was to examine stability of A1 RFs obtained from consecutive recordings at the same cortical site in awake, quiescent animals so that we could assess how repeatable the measurements were and explore any changes in the structure or shape of these response properties over the course of hours. In the first part of the study, we describe a computational technique that allowed us to perform a global assessment to compare RFs of A1 neurons and then label them as stable or labile. This analysis enabled us to compare RFs in consecutive recordings and showed that, at the population level, the observed variations were unlikely to reflect any systematic shift in tuning or alteration in neuronal selectivity. In the second part of the study, we asked the question of whether spontaneous variations in passive quiescent animals were significantly smaller than behaviorally induced changes observed in RFs of trained and behaving animals. To address this question, we performed a local analysis of RFs by focusing on frequency-specific changes under passive conditions compared with behavioral conditions (with animals performing spectral tasks of either single tone detection or two-tone discrimination).

Materials and Methods

Physiological recordings

All procedures were in accordance with the Institutional Animal Care and Use Committee at the University of Maryland and the Guidelines of the National Institutes of Health for use of animals in biomedical research. We performed extracellular recordings from multiple cortical units in six awake domestic ferrets (*Mustela putorius*). To enhance the stability of the recordings, a stainless-steel headpost was surgically implanted on each animal's skull after behavioral training was completed. The implant procedure is fully described in a previous study (Fritz et al., 2003). After the animal had fully recovered from the implantation procedure, we recorded from the primary auditory cortex (A1) in multiple recording sessions (lasting 6–8 h a day) through small craniotomies (~1 mm in diameter) over A1. Tungsten electrodes (3–8 M Ω) were used to record neural responses from single and multiunits at different depths. The response patterns were stored and processed off-line to sort single-unit activity. Multi-unit records were constructed from spikes triggered by a low threshold level [four SDs (4σ) above baseline]. Single units were derived using a Matlab-based customized manual sorting technique based on spike templates constructed from thresholds defined at multiple time windows. The window thresholds were chosen such that variances from the different sorted classes did not overlap at those chosen points. The variance of each sorted class of units was then estimated and was always well within the threshold windows chosen in the sorting. In addition, we always used two other criteria for the sorted spike classes: (1) the interspike intervals for each class were exponential with a minimum 1 ms spike latency, and the distribution peak was always >2 ms; (2) the spike rate remained stable throughout the recording time.

The ferrets were gradually habituated to lie calmly in an open restraining tube for increasing periods of time period up to 6–8 h. The recordings were performed in awake ferrets in both passive and active conditions. In the passive state (P), the animals were awake and quiescent but were not performing any experimentally defined behavioral task, whereas the active condition (A) required the ferrets to perform an acoustic behavioral task during the recording of their neural activity. The behavioral tasks consisted of a single tone detection or a two-tone discrimination task. In the tone-detection task, the animals were trained to recognize broadband ripples as "safe" sounds, during which they could continuously lick water through a spout, and were trained to refrain from licking during presentation of "warning" target sounds (pure tones) to avoid mild shock (Heffner and Heffner, 1995; Fritz et al., 2003). During two-tone discrimination, the animals were trained to lick during the

presence of ripple noises and reference single tones and refrain from licking during target tones (Fritz et al., 2005a).

Stimuli

To derive receptive fields from the recorded units, we used noise-like broadband sounds consisting of ripple mixtures, called temporally-orthogonal ripple combinations (TORCs) (Klein et al., 2000; Depireux et al., 2001). Each TORC is a broadband noise with a logarithmically spaced carrier and a dynamic spectral profile consisting of six superimposed envelopes drifting at different velocities from 4 to 24 Hz. The spectral envelope of each TORC consists of equally spaced peaks from 0 (flat) to 1.4 peaks per octave. A full set of 30 TORCs (at all spectral spacing and upward vs downward dynamic drifting) was required to characterize the receptive field of each unit. The TORC stimuli were 1.25 s (in active conditions) and 3 s (in passive conditions), and the full set was repeated 6–15 times on average. The interstimulus interval was 1.2 s for the active conditions and between 1 and 1.2 s in the passive conditions. The sounds were computer generated and were delivered to the animal's ear through inserted earphones that were calibrated *in situ* at the beginning of each experiment. The amplitude of the TORCs was fixed for a recording session at a level between 55 and 75 dB sound pressure level, depending on the maximally effective level at best frequency (Kowalski et al., 1996).

Receptive field estimation

Determining the spectrotemporal receptive field. Characterizing the receptive field of the cortical units required at least the presentation of one full set of 30 TORCs. We hence defined one stimulus repetition to be equal to one set of 30 TORCs, which were necessary to span the entire stimulus space (Klein et al., 2000). Depending on the signal-to-noise ratio (SNR) for each recording, multiple stimulus presentations were required to obtain a reliable receptive field, with an average recording time of 12–30 min for each recording (i.e., between 6 and 15 repetitions of each stimulus set). The spectrotemporal receptive fields (STRFs) from single unit and multiunit responses in both passive and active conditions were computed using standard reverse correlation techniques (Klein et al., 2000; Depireux et al., 2001; Miller et al., 2002). In the active condition, only the TORC portion of the recording was used to estimate the STRF while discarding the responses to the reference or target tones (see Fig. 1). The reliability of an STRF was measured using a bootstrap technique (Efron and Tibshirani, 1998), which allowed us to quantify the variance of the neural response and hence the overall SNR for each STRF. The bootstrap analysis, as well as SNR measurement procedure, are explained in details by Klein et al. (2006). Most SNRs were >1 , and we excluded all STRFs with an SNR <0.2 from our analysis.

Determining the center of the receptive field. To define the center of the receptive field of a given unit, we extracted the Hilbert envelope (Oppenheim and Schaffer, 1999) of the spectral receptive field and defined the peak of the magnitude of this Hilbert envelope as indicative of the center of the receptive field of the cell. This chosen center usually corresponded to the maximal excitatory or inhibitory response of the cell. However, this analysis was not possible in ~15% of the cases where the STRF exhibited a complex spectral pattern (e.g., multi-peaked excitatory responses). In those cases, we chose the center of the receptive field of the unit to be the maximal excitatory or inhibitory response in the STRF.

Data selection

Spike-matching algorithm. We were interested in characterizing the stability of the neural response over a sequence of consecutive recordings. To ensure that the recording was stable over the course of the experiment, we investigated the persistence of the action potential waveform shape. Hence, we used a spike-matching algorithm to test the similarity of spike shape clusters across sequences of consecutive recordings (both for single unit and multiunit responses). We used a clustering technique based on a Fisher discriminant analysis (FDA) to project the data on a two-dimensional space (Duda et al., 2001). An FDA projection minimized the scatter within each recording and maximized the discrimination across recordings, making it a very conservative selection criterion. Recording sequences that achieved a high overlap in this projected plane [with statistical significance quantified by a permutation test (Efron and Tibshirani, 1998)] were labeled as belonging to the same unit or group of

units and included in the analysis. Recordings with spikes not significantly overlapping in the FDA projection were rejected from our analysis.

Stability criterion

Tree-structured vector quantization algorithm. We used a tree-structured vector quantization (TSVQ) technique to project the STRFs into a space where we can define a stability criterion. The hierarchical TSVQ algorithm has been successfully used to analyze multiscale data (Breiman et al., 1984; Gersho and Gray, 1991).

TSVQ is a clustering algorithm that constructs a binary-search-tree data structure by recursively dividing the data space into two subspaces at each resolution level. The division at each level is performed by minimizing the cost (distance) between a “standard” STRF, the VQ encoder γ , and all STRFs in the space. We define the distortion as $D(\gamma) = \|f - \gamma(f)\|^2$, where f is the individual STRF. Using an L2 distance is akin to the use of correlation coefficient measures to quantify STRF similarities (DeAngelis et al., 1999).

The solution to the minimization problem is given as follows: the “standard” STRFs (S_i) should be the average of all those training STRFs that are in the encoding region R_i , where:

$$R_i = \{f: \|f - S_i\|^2 \leq \|f - S_j\|^2, \forall j\}.$$

In other words, the TSVQ divides the data space at different levels. At each resolution, the space is partitioned into different clusters or cells, which are determined by repeated application of the Linde–Buzo–Gray algorithm (Abut, 1990). The algorithm finds the optimal clustering of the data at a given level, as defined by the cost function above. The procedure is first applied to the coarsest resolution of the data vectors; then, clustering is performed successively at finer resolutions to yield additional insights about the structure of the data space.

Multiscale analysis of STRFs

The different resolution levels were defined by approximations from singular value decomposition (SVD) of the STRFs at different ranks. SVD is a factorization technique that can be applied to any finite dimensional matrix (in this case, the STRF) by writing it in the form $S = U\tilde{E}V^T$ (Haykin, 1996; Hansen, 1997). The columns of U and V form an orthonormal basis of left and right, respectively, singular vectors, whereas \tilde{E} is diagonal matrix with entries that correspond to the singular values (λ) of the matrix S . Singular value decomposition can be viewed as reformulating the STRF matrix S into a sum of separable matrices, where the columns u_1, \dots, u_m of U and v_1, \dots, v_n of V correspond to spectral and temporal cross sections of these separable transfer functions.

Typically, the number of nonzero singular values is equal to the rank r of the matrix S . However, because of the presence of noise in the measurement, the λ values are all expected to be nonzero with their values decreasing monotonically to a noise floor, which depends on the level of the noise (Depireux et al., 2001). Depending on the spectrotemporal interactions in each neuron, its receptive field can either be fully separable (rank 1), partially separable (rank 2), or of higher order. In the case of higher-order STRF, adding more terms to the SVD summation allows to capture more of the receptive field features. To illustrate this concept, Figure 2a (right panel) shows an example of an STRF with tuning that can be captured by its rank 1 approximation, but its orientation (tuning to upward moving spectrotemporal patterns) can only be captured starting from its rank 2 estimate.

Based on previous studies (Depireux et al., 2001; Simon et al., 2007), it has been shown that most STRFs obtained in both anesthetized and

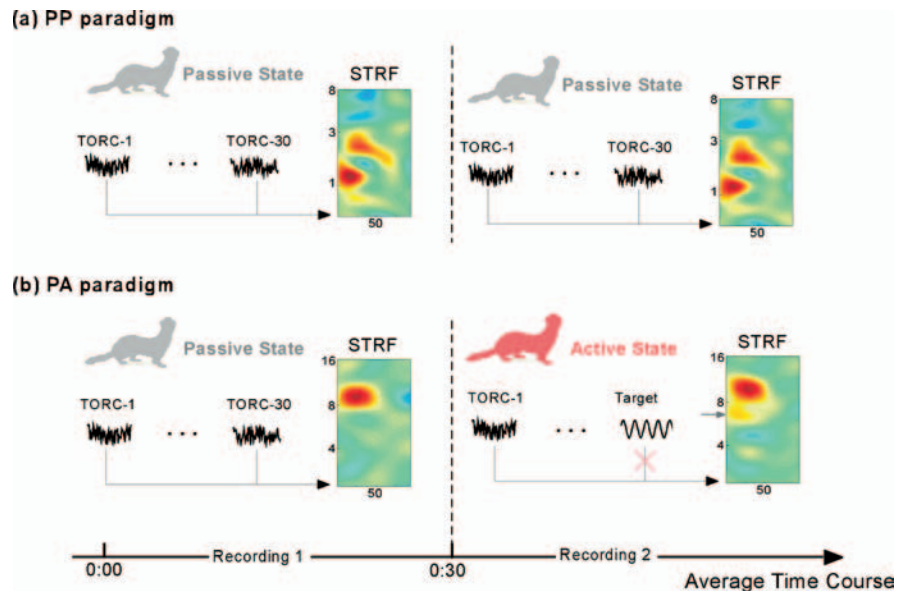


Figure 1. Experimental design for PP and PA paradigms. **a**, PP paradigm. Sequences of receptive field recordings were performed in awake, nonbehaving animals (P). The stimulus set consisted of broadband noise-like stimuli known as TORCs, which were used to derive the spectrotemporal receptive field of each neuron using standard reverse correlation techniques. The successive recordings were performed over the course of 30 min to 2 h. **b**, PA paradigm. Similar receptive field sequences were measured using the same stimulus design but changing the behavioral state of the animal. Initially, a P receptive field was obtained from a nonbehaving animal. Next, the animal was engaged in a behavioral task (A) and a similar receptive recording was obtained again. The behavioral paradigm has been described previously (Fritz et al., 2003, 2005) and included the use of a set of TORC and single tone stimuli. Only the TORC stimuli were used for deriving the receptive fields in the active state, allowing us to compare receptive fields that are obtained from the same stimulus set but under different behavioral conditions.

awake recordings in A1 tend to be of rank 1 and 2, with few exceptions at higher ranks. Following these results, we based our multiscale clustering on gradual levels of SVD approximations of the STRF, up to rank 4. In the first level, we estimated the rank-1 approximation of all STRFs in the training set and used the VQ algorithm for clustering. As we move to the next level, we added an additional rank to the approximation and clustered the data within each group into two subsets.

In each clustering level, we first aligned all STRFs by their best frequencies. This allowed us to avoid a trivial outcome that could result in a classification based on spectral tuning. We were primarily interested in a classification based on the spectrotemporal features (e.g., tuning pattern and temporal dynamics) of the receptive field; whereas a division of units based on high-frequency/low-frequency cells is of little interest in this analysis.

Localized analysis of receptive field changes

To quantify the local changes in receptive fields from one STRF measurement to the next, we subdivided the spectral axis into bands (Δf_i) of 0.5 octave width ($1 \leq i \leq 10$, over five octaves) and extracted localized measures of spectral change in each recording pair. This procedure was performed by first normalizing each STRF by its Euclidean norm and then calculating the difference between the two STRFs in the sequence (STRF_{diff}). Next, we extracted a measure of change at each frequency band ΔA_i , defined as the local maximum difference for that band i . This point is taken as the one spectrotemporal bin with maximal change over the entire spectral band Δf_i (± 0.25 octave around the tone of interest). The values of ΔA_i were reported as percentages relative to the maximum value of the first STRF (Fritz et al., 2005a). This analysis yielded 10 different ΔA_i values for each STRF pair. In the case of the passive STRF pairs [passive–passive (PP) recordings], we reported the average value of ΔA_i for each STRF pair. We also ran a statistical analysis of the ΔA_i distributions, without averaging the values of ΔA_i to ensure that we were not minimizing or washing away the effects by taking the average. In the passive–active (PA) recording sequences, we distinguished between two kinds of spectral bands (Δf_i): (1) ΔA_{NB} (nonbehavioral), corresponding to the average ΔA_i obtained from spectral bands that did not coincide

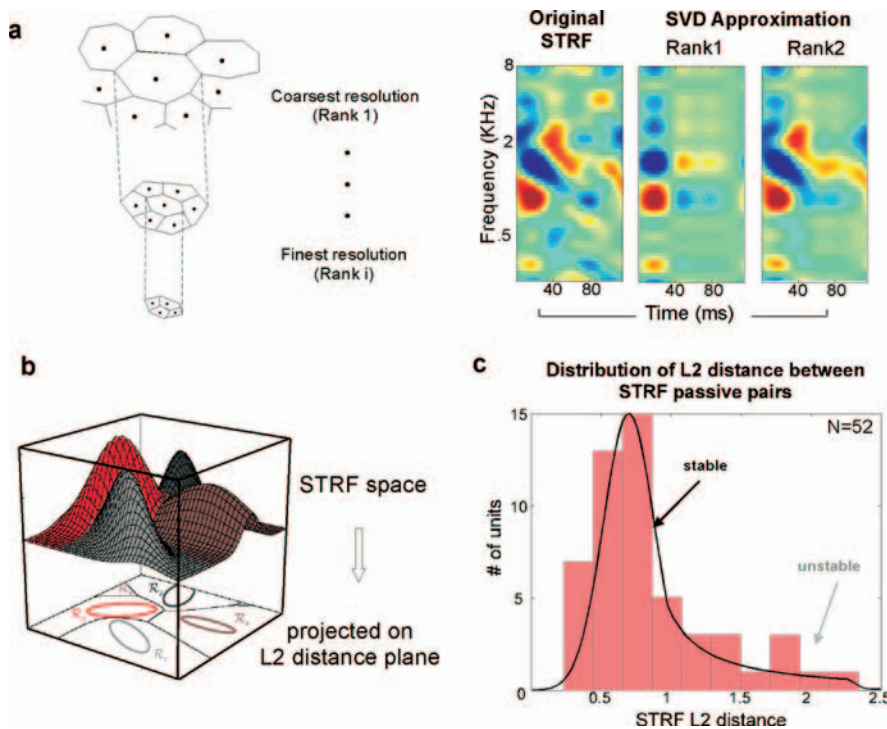


Figure 2. Stability analysis for receptive fields based on TSVQ. *a*, TSVQ analysis. Tree-structure vector quantization is a hierarchical clustering algorithm that organizes data at increasing levels of resolution (based on multiple scales). The right panel depicts the multiple resolutions of analysis used to cluster the STRF space. Each original receptive field matrix can be approximated by STRFs with different ranks using SVD. Each additional rank captures more specific features of the STRF. The rank-1 STRF fails to capture the orientation of the STRF and rather denotes the general tuning of the receptive field. By adding an additional degree (Rank-2), the original receptive field is now better approximated as an oriented function. *b*, Schematic of TSVQ clustering of STRF space [reproduced from Duda et al. (2001)]. Functions in the STRF space are projected on a two-dimensional plane by computing Euclidean distance (L2) between each pair. STRFs that have smaller L2 distances are deemed closer, and effectively belonging to a same cluster, and vice versa. *c*, Distribution of L2 distances in PP pairs. On average, most STRF pairs in the PP sequences tend to have a small L2 distance, making them effectively stable receptive fields (based on our stability criterion). The distribution, however, has a longer tail, indicating that a smaller subset of cells does in fact exhibit a wider variation in receptive field shapes over successive recordings.

with a behavioral tone frequency; and (2) ΔA_B (behavioral), which corresponded to the spectral change observed in the vicinity of a behaviorally relevant tone (i.e., target tone in a detection task, or reference, and target tones in a discrimination task).

Because the spectral effects captured by ΔA_i could be either “facilitative” (positive) or “depressive” (negative), we compared the populations of ΔA_{PP} , ΔA_{NB} , and ΔA_B by taking the absolute magnitude only and ignoring the sign of the spectral change. This analysis allowed us to focus on how “small” or “significant” the effects were and, hence, how much change was induced when the animal engaged in behavior relative to the spontaneous changes observed in receptive fields. For population analysis, we ran a two-sample Kolmogorov–Smirnov test and *t* test (Lindgren, 1993) to compare how statistically different or similar these populations were.

Results

We examined changes in receptive field properties in primary auditory cortex of six awake ferrets by performing sequences of extracellular recordings at multiple sites. In the first set of experiments, we obtained at least two consecutive recordings at each cortical site during a passive (nonbehavioral) state of the animal. Sequences of passive recordings were labeled PP and were obtained from 52 cortical sites that conformed to our selection criteria (see Materials and Methods). In a second set of experiments, the receptive field sequences were obtained from a passive followed by an active behavioral state of the animal (Fig. 1). The active recordings were acquired while ferrets, previously trained

on various acoustic paradigms, were performing a behavioral task of single tone detection (Fritz et al., 2003) or two-tone discrimination (Fritz et al., 2005a). In tone detection, animals were trained to lick water during the presentation of safe broadband sounds (TORCs) and refrain from licking during a target single frequency tone. In tone discrimination paradigm, animals could safely lick during the TORC sounds as well as a safe reference tone and were trained to stop licking during a target different frequency tone. Sequences of interleaved passive and active STRF measurements were labeled PA. We recorded 101 such PA sequences.

The STRF estimation was performed following standard reverse correlation techniques (see Materials and Methods), where only the responses to TORC sounds were used to measure the receptive field, in both passive and active states. We note that in the active behavioral states, the response to the reference or target tones was not included in the STRF measurements. Hence, the same stimulus set was used to characterize receptive fields in both passive and active settings. The only major difference between the PP versus PA conditions was the uniform versus changing behavioral state of the animal during the recording of the sequence.

Clustering of STRF shapes

To quantify the stability of STRF sequences, we aimed to define a criterion by which we could label our data set into stable versus labile units. We opted for a clustering algorithm that would span the space of receptive fields naturally observed in the primary auditory cortex of awake ferrets. By dividing this STRF space in a data-driven manner, we could organize STRFs into different groups based on their spectrotemporal properties. We could then label two STRFs in a sequence as belonging to a stable unit if all receptive fields in the sequence belonged to the same cluster in our STRF classification. Similarly, a unit was labeled as labile if the receptive fields recorded within minutes/hours at the same site fell in different clusters.

We used a TSVQ technique, using a Euclidean distance (L2), to cluster STRFs into different groups (see Materials and Methods). We analyzed 794 single-unit STRFs measured in previous studies from awake ferrets (Elhilali et al., 2004; Fritz et al., 2003, 2005a) as our training set to define the optimal projection from STRF space to an L2 plane (Fig. 2*a*). The algorithm organized the pool of STRFs into subgroups at different levels of resolution starting from a coarse level (rank 1 STRF approximation) to a finer level (rank 4).

The clustering obtained from our training set is shown in Figure 3. Each horizontal level represents a clustering of the STRFs at a given resolution. We examine the receptive field features that arose in each branch and note the following. Level 1: as expected, receptive fields were naturally clustered into an “excit-

tatory branch” (Fig. 3, left) and an “inhibitory branch” (Fig. 3, right). Two-thirds of the data set (498 U) fell in the “excitatory” group, and one-third (296 U) was classified as “inhibitory.” Level 2: the “excitatory branch” appears to subdivide further based on temporal criteria such as the latency of the excitatory peak of the STRF and its temporal extent, both significantly longer in the left branch (slow STRFs) compared with the right (fast STRFs). The “inhibitory branch” segregates based on spectral (and not temporal) criteria. Thus, the temporal properties of STRFs in the two branches at this level do not show significant differences, whereas the spectral properties are noticeably diverse. For example, the right branch of this inhibitory cluster (Fig. 3) groups spectrally symmetric units, in contrast to the left branch that has mostly asymmetric STRFs with side-band excitatory fields. Level 3: on the “excitatory branch,” another temporal subdivision occurs between faster and slower cells, whereas the “inhibitory branch” subdivides based on spectral criteria (STRFs with different degrees of asymmetry). Finally, in the final step (level 4), STRFs in both branches subdivide based on mixed spectrotemporal criteria.

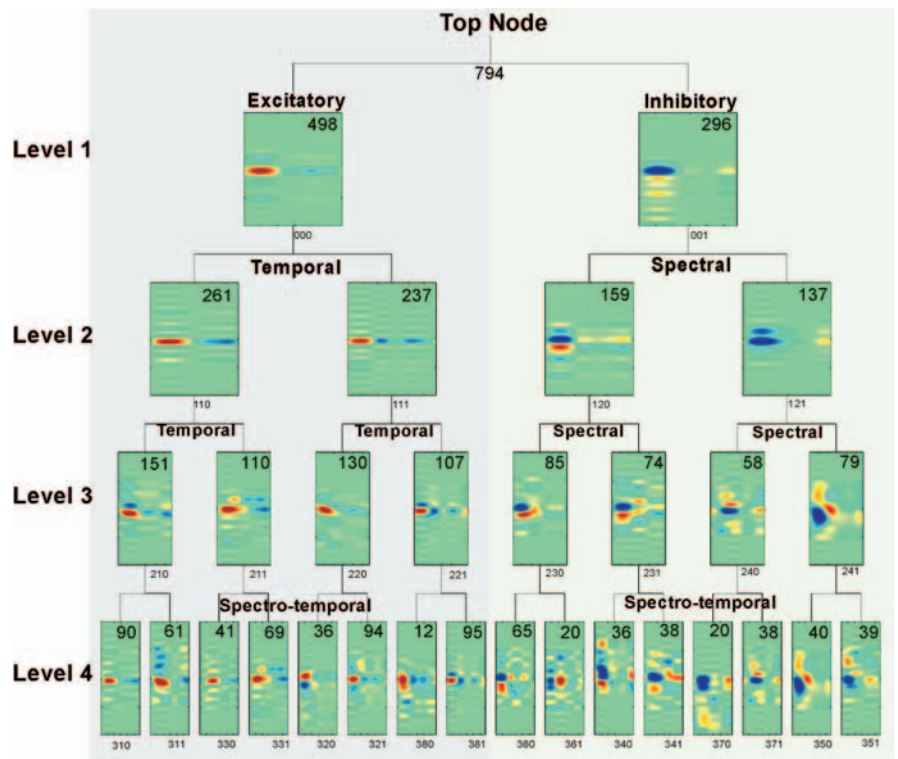


Figure 3. Clustering of receptive field space. The tree-structured diagram depicts the classification of STRFs at different levels of resolution. The gray versus green vertical branches illustrate a natural division of receptive fields into excitatory versus inhibitory units. The vertical direction reveals a more refined division of each branch by adding increasingly more information going from levels 1 to 4 (see description in Materials and Methods).

STRF stability

Using the TSVQ classification as our stability criterion, we found that 73% (38 of 52) of multiunit passive recording pairs (PP) could be labeled as stable (i.e., recordings of successive responses in the same site yielded receptive fields that shared the same branch in the tree, up to the fourth level). Additionally, of those 52 multiunit PP pairs, 72% had at least one single unit that could be labeled as stable. These numbers suggest that the majority of passive cortical neurons in awake nonbehaving ferrets had consistent spectral and temporal features that allowed them to maintain stable receptive fields over the course of minutes to hours.

Figure 2*c* illustrates the population results of L2 distances between consecutive STRFs in a PP sequence. As expected, we obtained a unimodal distribution with a long tail. The majority of units yielded a small L2 distance between consecutive pairs (mean, 0.6), hence allowing them to fall in the same cluster. A smaller subset (belonging to the tail of the distribution) appeared to have a bigger L2 difference yielding a labile STRF population.

There was no apparent correlation between a particular branch or node, and the ratio of stable versus labile units belonging to that node. Hence, the stability or instability of A1 cells does not appear to be obviously related to the spectrotemporal properties of their receptive fields. In addition, given that the TSVQ algorithm aligns STRFs according to their best frequency, we checked whether there was any obvious change in the center of the receptive field of each sequence. This analysis confirmed that the STRF center in all PP sequences did not change significantly from one STRF to the other (mean difference between PP pairs, 0.047 ± 0.04 octaves for multiunits, $n = 52$, and 0.095 ± 0.13 octaves for single units, $n = 94$).

Stable versus labile STRFs

To illustrate the variations in the properties of a stable versus labile sequence, Figure 4 shows examples of four receptive field sequences. In Figure 4*a*, the sequence of stable STRFs was measured over the course of ~ 2 h (time stamps of the recordings shown above each STRF). Despite some inherent spontaneous changes in their spectral and temporal properties, these receptive fields exhibited low variability and could be labeled as stable according to our TSVQ classification. To confirm this assessment, we plot separately the spectral receptive field and temporal impulse response of the STRF (right panels). These plots confirm that no significant variability was observed in the properties of this receptive field sequence. Similarly, the second sequence in Figure 4*a* illustrates a series of six stable receptive fields, measured over the course of 2.5 h. Again, whereas we observed some small variability in the background features of these receptive fields, the key properties of each STRF remained unchanged.

In contrast, the sequences shown in Figure 4*b* are receptive fields measured from a labile site. Both sequences exhibited noticeable variability in their spectral and temporal properties. In the first sequence, the right panel of Figure 4*b* shows a strengthening of the excitatory field and weakening of the inhibitory field in the receptive field. Temporally, we note a 10 ms shift in the latency of the excitatory peak. This variability leads to a significant difference between the STRFs in this sequence, causing them to fall in different TSVQ leaves, and hence were labeled as labile. In the second example of Figure 4*b*, the STRFs exhibit variations in both of their spectral and temporal features in time. Starting from a receptive field tuned at 500 Hz, with a slight side-band inhibition at ~ 800 Hz, both excitatory and inhibitory fields

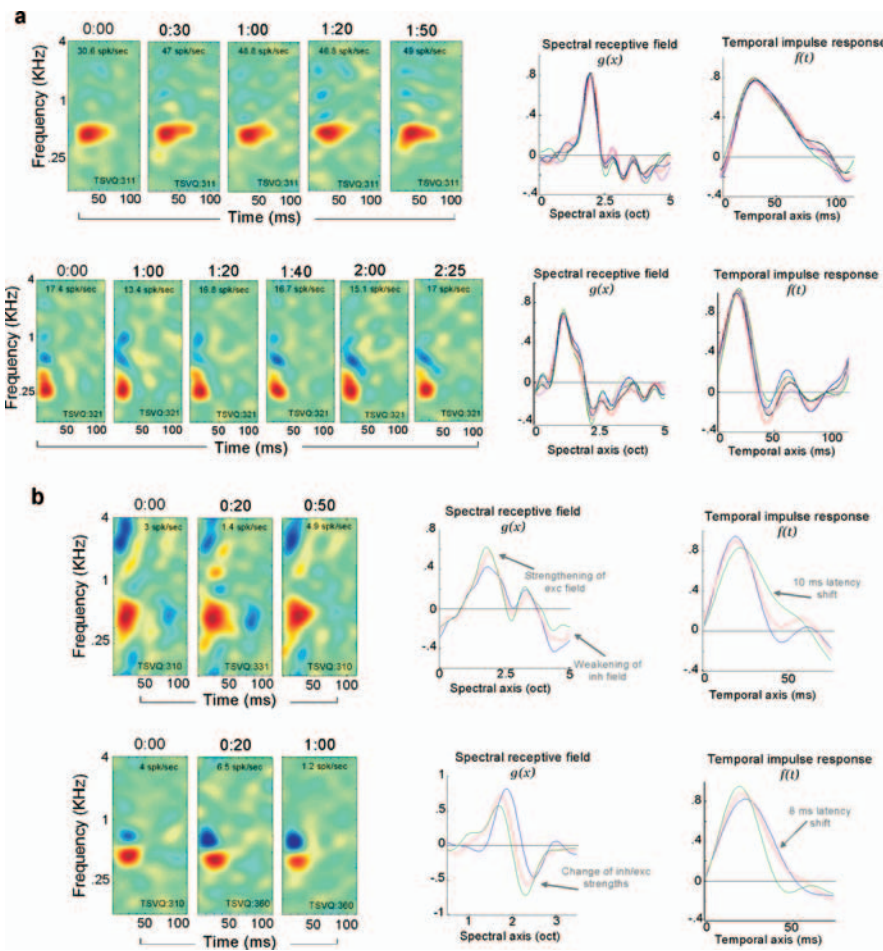


Figure 4. Examples of stable versus labile units. **a**, Stable units. Each row shows a sequence of receptive fields obtained from successive recordings over the course of 2–2.5 h. The receptive fields show a remarkable stability and hardly change their main tuning features. The right panels depict the spectral receptive fields and temporal impulse responses extracted from these receptive fields and confirm the stability of the shape and tuning of the response properties of these units. The number shown on top of each STRF plot indicates the average firing rate of the cell during that particular recording. The L2 distances between the consecutive STRFs in the top example are 0.51, 0.52, 0.65, and 0.72 and in the second example are 0.38, 0.26, 0.35, 0.37, and 0.3. **b**, Labile units. The panels show examples of two labile units. The receptive fields obtained over the course of 1 h exhibit changes in the spectral and temporal tuning, as depicted in the right panels. The L2 distances between the consecutive STRFs in the first labile example are 0.94 and 1.05 and in the second labile example are 0.93 and 0.94.

change in strength and location. The temporal dynamics also exhibit a change in latency of ~ 8 ms.

Stability and STRF convergence

What properties distinguish stable and labile populations? Although we could not observe any apparent correlation between the STRF shapes and their stability, we investigated a possible correspondence between the receptive field stability and their rate of convergence during measurement. Each receptive field was obtained by presenting multiple repetitions of the TORC stimuli. On average, we collected seven repetitions (or ~ 15 min) worth of data in each recording site to measure a receptive field. However, units varied in how fast they yield a convergent receptive field. An STRF^N obtained from N stimulus repetition was called convergent if the variation between STRF^N and STRF^{N-1} conforms to our stability criterion (i.e., had an L2 difference less than a SD, $\sigma = 0.85$). In other words, we define convergence to mean the number of stimulus repetitions needed for the L2 criterion to be met. Figure 5 shows four examples of STRFs obtained as we accumulated additional data from multiple presentations

of the stimulus (the stimulus presentation number is indicated above each STRF). In Figure 5, **a** and **b** illustrate examples of neurons that yielded a convergent receptive field after two or three stimulus repetitions; whereas **c** and **d** depict examples of slowly converging receptive fields that minimally vary only after six or seven repetitions of the stimulus. Overall, STRFs labeled as stable converged much faster than labile ones. Stable units required $\sim 3 \pm 0.5$ stimulus repetitions (which is equivalent to ~ 10 min). In contrast, the labile population showed a great variability in the convergence rate, with an average of 4.8 repetitions and a SD of ± 2 repetitions. Hence, we can infer that stable units yielded receptive fields converged significantly faster than the labile neurons.

We also explored any dependence of stability of the units on the temporal separation between recordings. In most cases, our STRF recordings were obtained within 10–15 min of each other. Of our 52 PP recordings, 50 were performed in < 60 min of each other, with an average inter-recording time of 12.8 min (minimum of 2 min, median of 11.2 min). Two recordings were done > 1 h apart (one 65.1 min and a second one 2.3 h apart). There was no correlation between this recording pattern and the likelihood of stability/labability of the units. The two recordings that were more than 1 h apart yielded 1 stable and 1 labile multiunit. In addition, there was no correlation between the temporal separation between recordings and the unit stability in the remaining 50 U.

Stability versus plasticity of STRFs

In the second part of this study, we contrasted the spontaneous changes in receptive fields with the behaviorally induced plastic changes in cortical neurons. We compared the receptive field properties of the PP versus PA sequences. Our active set of units includes neural responses measured during the performance of spectral tasks (tone detection and discrimination), and hence mostly induced localized plastic changes correlated with the frequency of the target and/or reference tones. Our recent results show that during performance of a tone detection task, receptive fields exhibit an enhanced response at the frequency of the target tone (Fritz et al., 2003). Tone discrimination tasks induce a distinct pattern of STRF changes, by selective enhancement of the target frequency and depression of the reference frequency (Fritz et al., 2005a). These results emphasize that differing predictive roles for the same stimulus during task performance in distinct behavioral contexts induces a differential response in cortical neurons.

We derived a localized change measure of receptive fields over specific frequency bands. We quantified the average changes in two consecutive passive STRFs over spectral bands of 0.5 octave width (Fig. 6a). These changes lead to a distribution over our

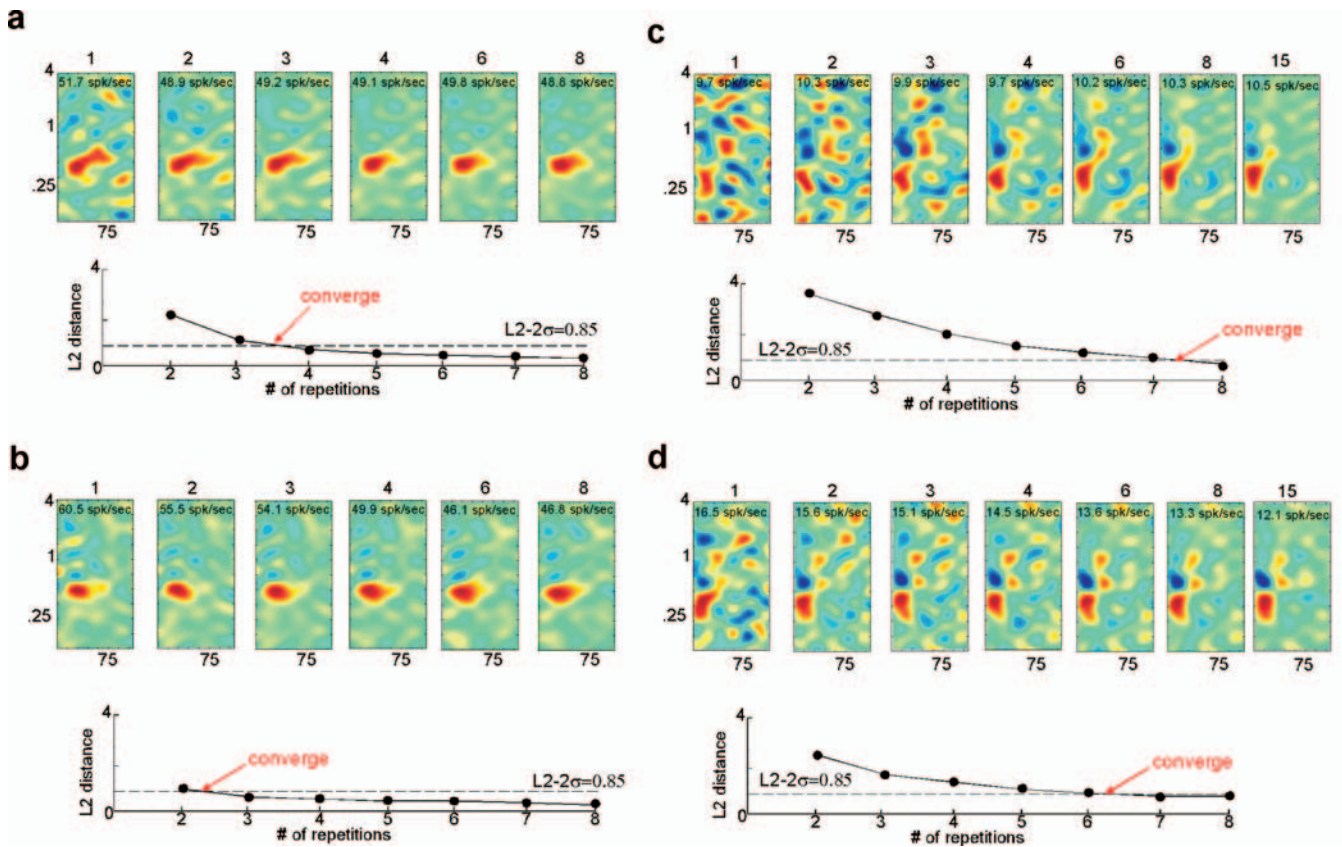


Figure 5. Convergence of receptive field estimates over time. Examples in *a* and *b* depict two units with rapidly converging STRFs. The top panels show the STRF estimate after increasing numbers of repetitions of the stimulus sequence (1, 2, up to 8 repetitions). The plot below the STRFs measures the L2 distance between every receptive field and its previous estimate. Examples in *c* and *d* illustrate two cases of slowly converging receptive fields. As reflected in the L2 distance plots, the STRFs keep changing after each repetition of the stimulus and take approximately six and seven stimulus iterations before reaching the convergence criterion.

population of PP single-unit sequences (Fig. 6*b*, green curve) with mean value 10.73 ± 8.22 . In contrast, we display the average amplitude change in the PA paradigm during detection and discrimination tasks. We derived two values from each PA sequence: (1) ΔA_B , or average behavioral change, which quantifies the receptive field difference at a behaviorally relevant frequency band (for instance, for a PA recording during a detection task of a 6 KHz tone, ΔA_B computes the receptive field change in the vicinity of 6 KHz); (2) ΔA_{NB} , which is a similar measure but away from the behaviorally relevant tone. We averaged the change values of all 0.5 octave bands outside the behaviorally relevant tone, yielding the ΔA_{NB} value for each receptive field pair (Fig. 6*a*). The population of ΔA_{NB} is shown in gray in Figure 6*b*, yielding a Gaussian fit with mean value 17.74 ± 16.48 . The distribution of ΔA_B is shown in red, with an average value of 53.14 ± 32.99 . The same analysis can be made for the multiunit clusters, as shown in Figure 6*c*. The Gaussian fits yielded a comparable trend to the single unit data, with ΔA_{PP} fitting a Gaussian with mean 10.59 ± 8.7 . The fit for the multiunit ΔA_{NB} was 15.73 ± 14.16 ; whereas for ΔA_B , it was 50.82 ± 29.8 .

Based on these population results, the passive distribution gives us a benchmark against which we can compare the behaviorally induced plastic changes. Our results clearly suggest that the behaviorally induced plastic changes are beyond the normal spontaneous variations in cortical receptive fields. To confirm that the passive and behavior populations were different, we ran statistical tests comparing these distributions. Using a two-

sample Kolmogorov–Smirnov test (Lindgren, 1993), we confirmed that the PP and PA-behave (ΔA_B) distributions were different with confidence p value $< 10^{-6}$. Also, we verified that the PA-nonbehave (ΔA_{NB}) and PA-behave (ΔA_B) distributions were different with confidence p value $< 10^{-7}$. Using the same statistical test, we found that the PP and PA-nonbehave (ΔA_{NB}) distributions were statistically similar. These same findings were also confirmed using a two-sample t test and applied to both the single and multiunit data.

To ensure that using ΔA averages in the PP and PA-nonbehave populations did not dilute any possible strong effects at specific spectral locations, we ran a permutation test comparing these distributions, without averaging the values of ΔA obtained at each spectral band. Instead, we compared the entire population for PP and PA-nonbehave (ΔA_{PP} and ΔA_{NB}) and tested the null hypothesis that these two groups were drawn from the same underlying probability distribution. The permutation test statistic confirms that these two populations were the same. In comparing these distributions with the PA-behave population (ΔA_B), the permutation test yielded an achieved significance level < 0.01 , which is considered very strong evidence for rejecting the null hypothesis (Efron and Tibshirani, 1998). This statistical analysis of population data confirms that task-related plastic changes of cortical receptive fields were significantly above the spontaneous variability in STRFs.

Finally, we address the question of whether the likelihood that an STRF would change during a behavioral task was in any way

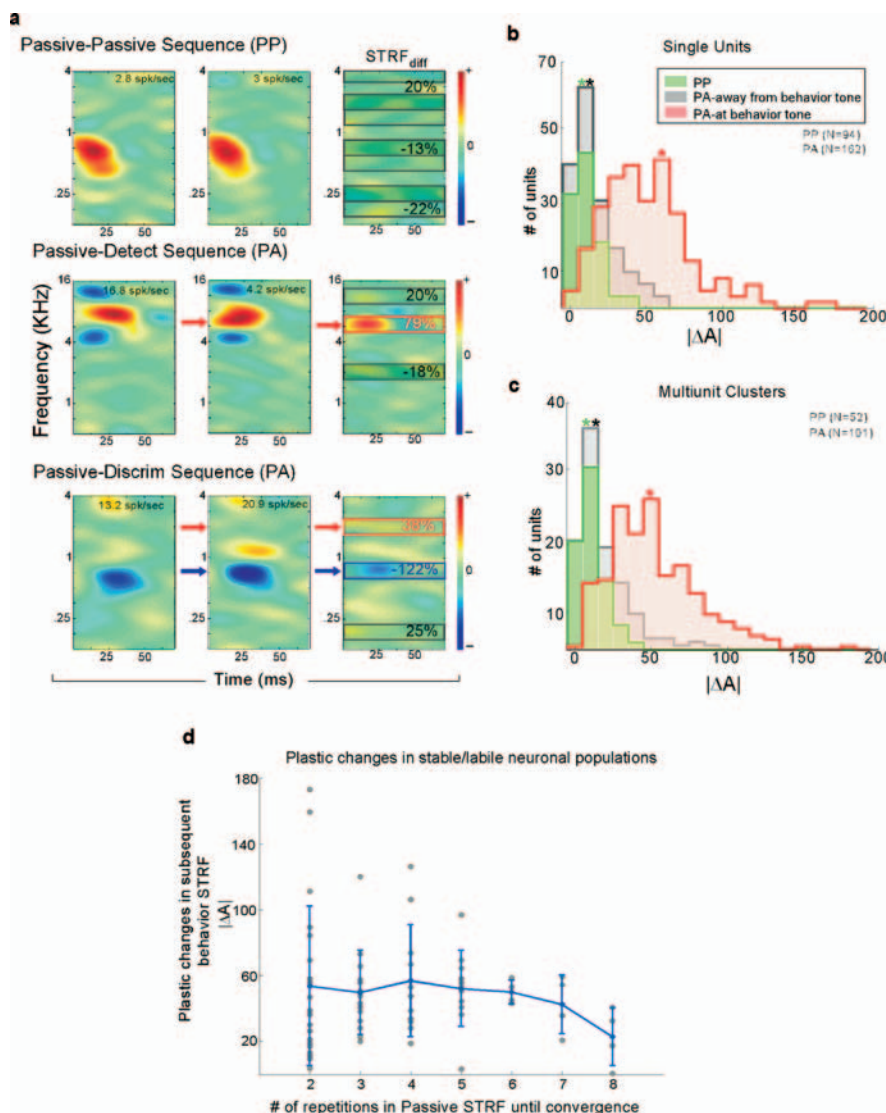


Figure 6. Local spectral changes in receptive fields in both behavioral and passive conditions. *a*, STRF sequences and difference in one PP and two PA examples. Each row depicts a sequence of two consecutive STRF measurements. The third panel shows the difference in STRF obtained by subtracting the second receptive field from the first one. The highlighted spectral bands in the STRF_{diff} illustrate the amount of change (ΔA) at that given band. *b*, Local spectral change in PP and PA populations in single units and multiunit clusters. Each panel shows the distribution of ΔA values at spectral bands across the tonotopic axis in passive STRF sequences in green, away from any behaviorally relevant spectral bands in gray (in PA sequences), and at the behavior tone in red (in PA sequences). The summary histograms (shown as smooth curves) depict a Gaussian curve to the underlying distributions. They reveal a significant difference between the three populations, with a mean value of 10.73% for PP and 17.74% for nonbehavioral bands in PA, in contrast with 53.14% for the behaviorally relevant tone locations. The same trend holds for multiunit data, with mean values of 10.59% for PP, 15.73% for nonbehavioral bands in PA, and 50.82% for behavioral bands.

For each sequence, we compared the firing rate of the first recording (R1) to that of the second sequence (R2). Firing rates are computed as the total number of spikes throughout the stimulation period divided by the total stimulus duration. Figure 7 shows the average firing rate change, defined by $\Delta r = (R2 - R1) / (R2 + R1)$. The distributions for the PP and PA sequences were quite symmetrical at ~ 0 , indicating no systematic increase or decrease of firing rate over the course of the recording. These distributions are statistically identical to Gaussian function with zero mean, and hence we cannot claim any systematic change of response gain over successive recordings. Nonetheless, there was a noticeably big variance observed in PA sequences, reflected in their Gaussian-fit SD of -0.06 ± 0.47 for multiunits and -0.05 ± 0.42 for single units. This variance was much greater than that from PP sequences (0.0 ± 0.27 for multiunits and -0.02 ± 0.29 for single units) and was likely caused by the changing levels of arousal in the behaving animal. Arousal or lick-related influences during behavior may induce overall changes in responsiveness of cortical cells beyond the auditory-related behaviorally induced changes in receptive fields, and hence confound our ability to correlate any variations in overall firing rate with behavior. Hence, the lack of any systematic bias in gain changes confirms that normalizing the rate of the STRFs is an appropriate way to analyze our data, allowing us to focus on STRF shape variations instead.

Discussion

The current study investigated the degree of intrinsic receptive field variability in the awake, quiescent, nonbehaving animal. Our results indicate that most receptive fields in A1 (approximately three-fourths of cells) are stable over the course of 30 min to 2 h. Although not previously investigated at such short time scales, this experimental validation of receptive field stability is not surprising by itself. It complements previous studies confirming that receptive fields in the auditory cortex remain unchanged over longer time periods extending to several days when the animal is in the same state of vigilance (Williams et al., 1999; Witte et al., 1999; Galvan et al., 2001). It also supports the notion of a receptive field as an intrinsic property of neurons, which reflects their fixed tuning to spectrotemporal sound features.

related to its characterization as stable or labile according to the convergence criterion discussed previously. To do so, we determined the number of repetitions needed for convergence for each passive STRFs in the PA-behave sequences. We also determined for each PA pair the ΔA changes in the STRF during the active epoch. The two are plotted against each other in Figure 6*d*. The results clearly show that there was no dependence between the two factors (i.e., both stable and labile STRFs could potentially exhibit rapid plasticity effects).

Changes in firing rates

Given that our measures of STRF plasticity normalize by the spike rate of each unit, we examined the possibility of any systematic changes in firing rates as we collected both PP and PA sequences.

Unlike previous studies, however, this paper adopts a new approach to analyzing receptive field stability by taking advantage of years of recordings of cortical receptive fields in our laboratory. This large database of STRFs acquired over the course of 7–8 years offers a rich sampling of the space of

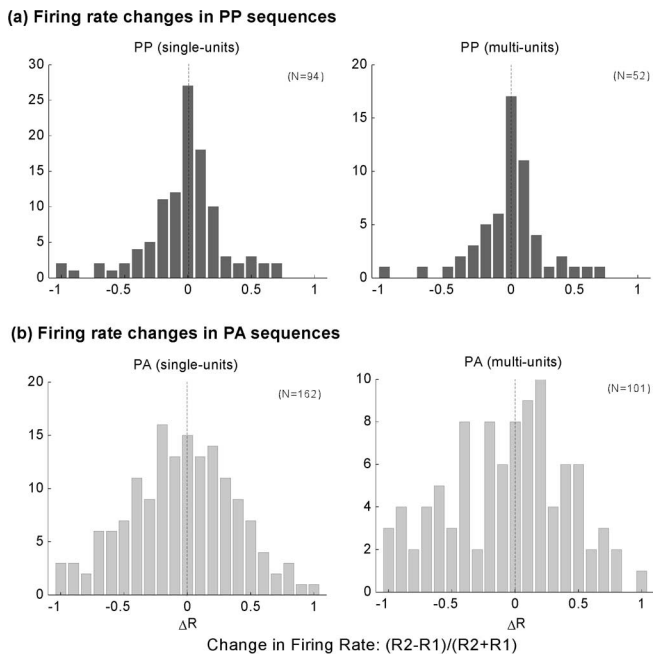


Figure 7. Change in firing rate in PP versus PA sequences. Each panel shows the distribution of rate changes in each sequence, measured as $(R2 - R1)/(R2 + R1)$, where $R1$ is the average firing rate for the first measurement in a sequence, and $R2$ is the rate in the second measurement. The histograms reveal that that gain changes in both PP and PA cases are mostly symmetrical at ~ 0 . The main difference between the two cases is the noticeable increased variability in the PA recordings, possibly because of the changed attentional state, alertness level, and behavioral demands in the active condition.

receptive fields in A1 of awake animals (almost 800 single units) and allowed us to define a nonparametric criterion by which to delineate stable receptive fields. By organizing this space in a data-driven manner, we allow receptive fields to fall into different categories based on how closely matched their spectral and temporal features were. Clearly, radically different methods could be used to study neuronal and receptive field stability, ranging from trial-to-trial variability in latency or amplitude of response to the same acoustic stimuli or changes in frequency response fields in response to pure tones. An advantage of the present approach is that it frees us from defining any specific dimensions along which to analyze receptive field variability (e.g., latency, bandwidth, or best frequency). Instead, the clustering algorithm simultaneously explores all of these dimensions at once by treating the STRF shape as a full feature vector. Finally, we point out (as an aside) that the results of the clustering of A1 STRFs (Fig. 3) reveal an interesting insight into the structure of cortical receptive fields, namely that this variety of STRFs originates primarily from dynamic and spectral variations in the excitatory inputs. Specifically, it suggests that inhibition provides a relatively stable “broad” focus around which excitation distributes both temporally and spectrally.

Utility of the STRFs

A question that emerges from the current stability investigation is the underlying meaning of the possible distinction between the stable and labile neuronal populations. The present study demonstrates that stable units tend to yield stationary and reliable STRFs at a much faster rate (i.e., requiring fewer number of stimulus repetitions) compared with the labile

ones. This distinction might be confounded by the use of a linear model (the STRF model) to capture the tuning properties of these neurons, and the convergence rate could be a reflection of the degree of nonlinearity of the labile population. Nevertheless, we only included units that yielded an STRF estimate that satisfied our selection criterion, and hence our analysis was at the outset dealing with neurons with a certain degree of linearity.

STRFs are quantitative measures of neuronal receptive fields that offer a straightforward quantitative linear description of the selectivity of the neuron to specific stimulus patterns and, hence, provide a deeper understanding of cortical processing and neural encoding of sensory information. They have, however, been criticized for their shortfalls, particularly, linearity. When probed with more behaviorally relevant natural sound ensembles, the linear STRF model proves to be an incomplete description of response properties of nonlinear auditory neurons (Theunissen et al., 2000) and fails to successfully predict responses to many natural stimuli. A success prediction rate of 10% (Machens et al., 2004) to 40% (Sahani and Linden, 2003) was typically reported for classes of natural sounds. A second shortfall is lack of generalization. Whereas the STRF model seems to give satisfactory results for a large set of stimulus ensembles (e.g., ripples, modulated noise, random tone pips, and classes of natural sounds), it appears that comparisons of receptive fields obtained from different stimulus bases lead to a striking difference between the derived kernels (Theunissen et al., 2000, Elhilali et al., 2004). STRFs have also been reported to lack robustness relative to stimulus perturbations, such as use of background noise with natural stimuli (Bar-Yosef et al., 2002).

Despite these known limitations, STRFs remain powerful tools to explore the functional architecture of the auditory cortex. The current study tries to work within these mathematical limitations to compare sequences of STRF measurements. Given the use of the same stimulus ensemble in all recordings, the STRF technique is behaving as a piece-wise linearization of the system over time and, hence, operates within the same regime at different time instants. Therefore, any time-dependent shift of receptive field properties is not a reflection of the constraints of the experimental tool but rather an inherent change in the tuning properties of cell. The population of neurons exhibiting response variations over time may reveal an intrinsic malleability of the sensory system. A subset of A1 cells may be spontaneously dynamic, with receptive fields that are in a state of active flux, or oscillating in a random walk around an attractor state (Hopfield, 1982). From a theoretical and functional perspective, this variability may provide the system with a mechanism for robust sensory encoding at the neuronal population level. Carefully designed studies are needed to elucidate the source of this spontaneous variability. It also remains to be shown how the system may benefit from its dynamic structure to robustly encode sensory information and reliably interpret it to guide perception.

Learning-induced receptive field plasticity

Finally, we used our stability findings as a benchmark of spontaneous natural receptive field variation against which to compare previous results of task-related plasticity of cortical receptive fields (Fritz et al., 2003, 2005a,b, 2007). Establishing a baseline for spontaneous changes in cortical receptive fields is necessary to properly and confidently interpret any changes induced by learning or experimental manipulations. Our anal-

ysis confirms that task-driven rapid plasticity cannot be attributed to simple intrinsic changes in receptive fields of cortical neurons but is significantly bigger and more systematic than inherent receptive field variations. The plastic changes are chiefly observed at particular frequencies, representing behaviorally relevant spectral regions. As has been argued in our previous work, the comparison between active and passive receptive fields is also in agreement with the comparison between active and naive or poorly behaving animals (Fritz et al., 2003, 2005a). Thus, in our studies, systematic changes in receptive fields are observed only during task-related behavior.

In summary, we described stability of neuronal receptive fields in awake quiescent ferrets. Our results suggest that receptive fields in A1 exist along a stability continuum, where a majority (three-fourths) of cells exhibit stable receptive field properties, and one-fourth of the population varies in the amount of its intrinsic fluctuations. This classification of stability has to be taken with the evident caveats of choice of stability criterion (i.e., TSVQ clustering technique with an L2 measure), use of receptive field measure (i.e., STRF model), as well time window of investigation (i.e., of the order of 30 min to 2 h). With that in mind, the current analysis allows us to emphasize the importance of learning-induced plasticity in cortical neurons against this background of stability. Our data clearly show that behaviorally driven receptive field changes are task specific and perceptible above and beyond any spontaneous variability. However, it remains an open question as to how this stability in the quiescent state and plasticity in the behavioral state of the animal functionally relate to each other. This question invokes the so-called “stability-flexibility” paradox (Liljenstrom, 2003), which describes the constraints on the functional organization of cortical maps to both reliably parse an acoustic scene and at the same time adapt to changing behavioral demands. In this regard, one could entertain multiple possible hypotheses that could offer a resolution to this paradox; such as state-dependent receptive field stability, in which sensory receptive fields would remain unchanged in a given state, and alter only if the state changed, or if new behavioral demands appear. Another hypothesis is that the inherent variability of receptive fields reflects normal dynamics of the neural machinery and provides the system with a robust mechanism for reliably representing sensory information and guiding behavior and perception. Yet, a third hypothesis posits the presence of a mixture of different types of cortical sensory neurons, some with stable and others with labile receptive field properties. Additional studies are required to explore these possibilities and integrate these findings of population neuronal stability with the remarkable degree of plasticity in the auditory cortex of adult animals.

References

- Abut H (1990) Vector quantization. Piscataway, NJ: IEEE.
- Aertsen A, Johannesma P (1981) The spectro-temporal receptive field: a functional characteristic of auditory neurons. *Biol Cybernetics* 42:133–143.
- Bair W (2005) Visual receptive field organization. *Curr Opin Neurobiol* 15:459–464.
- Bar-Yosef O, Rotman Y, Nelken I (2002) Responses of neurons in cat primary auditory cortex to bird chirps: effects of temporal and spectral context. *J Neurosci* 22:8619–8632.
- Breiman L, Friedman JH, Olshen RA, Stone CJ (1984) Classification and regression trees. Boca Raton, FL: Chapman & Hall/CRC.
- DeAngelis GC, Ohzawa I, Freeman RD (1995) Receptive-field dynamics in the central visual pathways. *Trends Neurosci* 18:451–458.
- DeAngelis GC, Ghose GM, Ohzawa I, Freeman RD (1999) Functional micro-organization of primary visual cortex: receptive field analysis of nearby neurons. *J Neurosci* 19:4046–4064.
- Depireux DA, Simon JZ, Klein DJ, Shamma SA (2001) Dynamics of neural responses in ferret primary auditory cortex: I. Spectrotemporal response characterization by dynamic ripple stimuli. *J Neurophysiol* 85:1220–1234.
- de Villers-Sidani E, Chang EF, Bao S, Merzenich M (2007) Critical period window for spectral tuning defined in the primary auditory cortex (A1) in the rat. *J Neurosci* 27:180–189.
- Duda RO, Hart PE, Strok DG (2001) Pattern classification, Ed 2. New York: Wiley.
- Edeline JM, Dutriex G, Manunta Y, Hennevin E (2001) Diversity of receptive field changes in the auditory cortex during natural sleep. *Eur J Neurosci* 14:1865–1880.
- Edeline JM (2003) Thalamo-cortical auditory receptive fields: regulation by the states of vigilance, learning and the neuromodulatory systems. *Exp Brain Res* 153:554–572.
- Efron B, Tibshirani RJ (1998) An introduction to the bootstrap. Boca Raton, FL: Chapman and Hall/CRC.
- Eggermont J, Aertsen A, Hermes D, Johannesma P (1981) Spectrotemporal characterization of auditory neurons: redundant or necessary? *Hear Res* 5:109–121.
- Elhilali M, Fritz JB, Klein DJ, Simon JZ, Shamma SA (2004) Dynamics of precise spike timing in the primary auditory cortex. *J Neurosci* 24:1159–1172.
- Fitzpatrick D (2000) Seeing beyond the receptive field in primary visual cortex. *Curr Opin Neurobiol* 10:438–443.
- Fritz JB, Shamma SA, Elhilali M, Klein DJ (2003) Rapid task-related plasticity of spectrotemporal receptive fields in primary auditory cortex. *Nat Neurosci* 6:1216–1223.
- Fritz JB, Elhilali M, Shamma S (2005a) Differential dynamic plasticity of A1 receptive fields during multiple spectral tasks. *J Neurosci* 25:7623–7635.
- Fritz JB, Elhilali M, Shamma SA (2005b) Active listening: task-dependent plasticity of receptive fields in primary auditory cortex. *Hear Res* 206:159–176.
- Fritz JB, Elhilali M, David SV, Shamma SA (2007) Does attention play a role in dynamic receptive field adaptation to changing acoustic salience in A1? *Hear Res* 229:186–203.
- Galvan VV, Chen J, Weinberger NM (2001) Long-term frequency tuning of local field potentials in the auditory cortex of the waking guinea pig. *JARO* 2:199–215.
- Ghazanfar AA, Nicolelis MA (2001) Feature article: the structure and function of dynamic cortical and thalamic receptive fields. *Cereb Cortex* 11:183–193.
- Gersho A, Gray RM (1991) Vector quantization and signal compression. Norwell, MA: Kluwer AP.
- Hansen PC (1997) Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion. Philadelphia: SIAM.
- Haykin SS (1996) Adaptive filter theory. Upper Saddle River, NJ: Prentice-Hall.
- Heffner HE, Heffner RS (1995) Conditioned avoidance. In: *Methods in comparative psychoacoustics* (Klump G, Dooling R, eds), pp 79–94. Basel: Birkhauser Verlag.
- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci USA* 79:2554–2558.
- Hubel DH, Wiesel TN (1963) Receptive fields of cells in striate cortex of very young, visually inexperienced kittens. *J Neurophysiol* 26:994–1002.
- Hubel DH, Wiesel TN (1968) Receptive fields and functional architecture of monkey striate cortex. *J Physiol (Lond)* 195:215–243.
- Kisley MA, Gerstein GL (2001) Daily variation and appetitive-induced plasticity of auditory cortex receptive fields. *Eur J Neurosci* 13:1993–2003.
- Klein DJ, Depireux DA, Simon JZ, Shamma SA (2000) Robust spectrotemporal reverse correlation for the auditory system: optimizing stimulus design. *J Comp Neurosci* 9:85–111.
- Klein DJ, Simon JZ, Depireux DA, Shamma SA (2006) Stimulus-invariant processing and spectrotemporal reverse correlation in primary auditory cortex. *J Comp Neurosci* 20:111–136.
- Kowalski N, Depireux DA, Shamma SA (1996) An analysis of dynamic spec-

- tra in ferret auditory cortex. I. Characteristics of single unit responses to moving ripple spectra. *J Neurophysiol* 76:3503–3523.
- Liljenstrom H (2003) Neural stability and flexibility: a computational approach. *Neuropsychopharmacology* 28:S64–S73.
- Lindgren BW (1993) *Statistical theory*, Ed 4. Boca Raton, FL: Chapman and Hall/CRC.
- Machens C, Wehr M, Zador A (2004) Linearity of cortical receptive fields measured with natural sounds. *J Neurosci* 24:1089–1100.
- Martinez LM (2006) The generation of receptive-field structure in cat primary visual cortex. *Prog Brain Res* 154:73–92.
- Miller LM, Escabi MA, Read HL, Schreiner CE (2002) Spectrotemporal receptive fields in the lemniscal auditory thalamus and cortex. *J Neurophysiol* 87:516–527.
- Oppenheim AV, Schaffer RW (1999) *Discrete time signal processing*, Ed 2. Upper Saddle River, NJ: Prentice Hall.
- Ringach DL (2004) Mapping receptive fields in primary visual cortex. *J Physiol (Lond)* 558:717–728.
- Sahani M, Linden JF (2003) How linear are auditory cortical responses? In: *Advances in neural information processing systems*, Ed 15, pp 125–132. Cambridge, MA: MIT.
- Schreiner CE (1995) Order and disorder in auditory cortical maps. *Curr Opin Neurobiol* 5:489–496.
- Simon JZ, Depireux DA, Klein DJ, Fritz JB, Shamma SA (2007) Temporal symmetry in primary auditory cortex: implications for cortical connectivity. *Neural Comput* 19:583–638.
- Theunissen FE, Sen K, Doupe AJ (2000) Spectro-temporal receptive fields of nonlinear auditory neurons obtained with natural sounds. *J Neurosci* 20:2315–2331.
- Weinberger NM (2007) Auditory associative memory and representational plasticity in the primary auditory cortex. *Hear Res* 229:54–68.
- Weinberger NM, Javid R, Lapan B (1993) Long-term retention of learning-induced receptive-field plasticity in the auditory cortex. *Proc Natl Acad Sci USA* 90:2394–2398.
- Williams JC, Rennaker RL, Kipke DR (1999) Stability of chronic multichannel neural recordings: implications for a long-term neural interface. *Neurocomputing* 26:1069–1076.
- Witte RS, Otto KJ, Williams JC, Kipke DR (1999) Pursuing dynamic reorganization in auditory cortex using chronic multichannel unit recordings in awake, behaving cats. *Neurocomputing* 26:593–600.