



Published in final edited form as:

Neuron. 2009 May 14; 62(3): 413–425. doi:10.1016/j.neuron.2009.03.014.

Spontaneous events outline the realm of possible sensory responses in neocortical populations

Artur Luczak^{1,2}, Peter Barthó¹, and Kenneth D. Harris^{1,*}

¹ Center for Molecular and Behavioural Neuroscience, Rutgers University, 197 University Avenue, Newark NJ 07102, USA

Summary

Neocortical assemblies produce complex activity patterns both in response to sensory stimuli, and spontaneously without sensory input. To investigate the structure of these patterns, we recorded from populations of 40–100 neurons in auditory and somatosensory cortices of anesthetized and awake rats using silicon microelectrodes. Population spike time patterns were broadly conserved across multiple sensory stimuli and spontaneous events. Although individual neurons showed timing variations between stimuli, these were not sufficient to disturb a generally conserved sequential organization observed at the population level, lasting for approximately 100ms with spiking reliability decaying progressively after event onset. Preserved constraints were also seen in population firing rate vectors, with vectors evoked by individual stimuli occupying subspaces of a larger but still constrained space outlined by the set of spontaneous events. These results suggest that population spike patterns are drawn from a limited “vocabulary,” sampled widely by spontaneous events but more narrowly by sensory responses.

Introduction

Single-unit recordings in sensory cortex have revealed much about how the firing of individual neurons is modulated by sensory stimuli. However, any individual neuron functions only as part of a much larger population whose combined activity underlies an animal’s processing of information. Characterizing the structure of neuronal population activity, and the way it is modulated by sensory stimuli, is a necessary step toward understanding the principles of information processing in the cortex.

Much has been learned about the structure of cortical population spike patterns by studying spontaneous activity. Cortical circuits both *in vitro* and *in vivo* during resting and sleep spontaneously produce periods of activity known as “upstates” (Battaglia et al., 2004; Luczak et al., 2007; Massimini et al., 2004; Petersen et al., 2003; Sirota et al., 2003; Steriade et al., 1993a). *In vitro* experiments have shown that neural activity within upstates has a sequential structure, with the order in which neurons fire largely conserved from one upstate to the next, reflecting the interaction of recurrent circuitry and intrinsic cellular dynamics (Cossart et al., 2003; MacLean et al., 2005; Mao et al., 2001). *In vivo*, early evidence for sequentially structured spiking activity came from studies detecting the presence of precisely repeating spike motifs

*Correspondence: E-mail: kdharris@rutgers.edu.

²Present address: CCBN, University of Lethbridge, Lethbridge, AB, Canada

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

(Abeles, 1991), although the statistical methods employed, as well as the long duration and high temporal precision of the detected motifs, have been controversial (Baker and Lemon, 2000; Mokeichev et al., 2007; Oram et al., 2001). Recent analyses of *in vivo* population data have confirmed using straightforward statistical methods that upstates are indeed sequentially patterned for a period of the order of 100ms, with temporal precision decaying as the upstate progresses (Luczak et al., 2007).

The structure of cortical population activity evoked by sensory stimuli, and the way this structure varies with the stimulus, is not yet fully clear. Recordings of single neurons have shown that spike timing relative to stimulus onset, as well as firing rate, can vary with the stimulus presented (Heil, 2004; Nelken et al., 2005; Optican and Richmond, 1987; Oram et al., 2002; Petersen et al., 2002). From this one might predict that, at the population level, sensory stimuli induce activity patterns whose neuronal composition and sequential order both depend on the stimulus (Harris, 2005). On the other hand, *in vitro* studies suggest that the activity patterns evoked by thalamic stimulation show the same sequential order as those occurring spontaneously (MacLean et al., 2005); from this one might infer that sensory-evoked sequences *in vivo* might share the same temporal order as spontaneous events, with the further implication that temporal order should be conserved between stimuli. *In vivo* recordings in hippocampus and neocortex have shown that spontaneous patterns during resting or sleep may mimic those seen in prior behavior (Diba and Buzsaki, 2007; Hoffman and McNaughton, 2002; Ji and Wilson, 2007; Walker and Stickgold, 2006). While this is usually interpreted as replay of patterns evoked during the behavior session, mimicry between spontaneous and sensory-evoked patterns has also been reported in optical imaging studies, in the absence of behavioral paradigms (Kenet et al., 2003; Petersen et al., 2003; Tsodyks et al., 1999).

The ability of sensory stimuli to change the firing rate of individual neurons is well known. Less is known about the way firing rates are coordinated in neural populations. The firing rates of a set of neurons at any moment can be summarized by a vector containing spike counts of all cells in a particular time window. The number of firing rate vectors that can be expressed by even a small number of neurons is in principle astronomical. However, there are reasons to suspect constraints on the patterns that can in practice be expressed by a given circuit. Recordings in several cortical areas have found persistent correlations between simultaneously recorded neuronal pairs (Averbeck and Lee, 2003; Gawne and Richmond, 1993; Jung et al., 2000; Montani et al., 2007; Narayanan et al., 2005; Zohary et al., 1994). Although these correlations are often of modest magnitude, apparently weak pairwise correlations may actually reflect strong constraints on patterns expressible at the population level (Schneidman et al., 2006). Cortical population patterns evoked by thalamic stimulation *in vitro* are subject to similar structural constraints as those occurring spontaneously (MacLean et al., 2005), suggesting that these constraints may arise from the dynamics of the cortical network.

Here, we characterize the fine structure of sensory responses and spontaneous activity bursts in the auditory and somatosensory cortices of anesthetized and unanesthetized rats *in vivo*. We find that activity patterns produced in response to multiple sensory stimuli, as well as occurring during spontaneous upstates, show largely conserved temporal structure. The possible population firing rate vectors observed are also restricted, with sensory responses and spontaneous events being drawn from a common limited “vocabulary.”

Results

Neurons respond to tones with stimulus-dependent firing rates, but stereotyped temporal profiles

We recorded simultaneously from populations of 40–100 neurons in rat neocortex using silicon microelectrodes. The first set of experiments was conducted in auditory cortex under urethane

anesthesia (8 rats) while presenting 500ms long acoustic stimuli (5 pure tones and 5 natural sound snippets; see Experimental Procedures for details). We began the analysis of this data at the single-cell level. The response of individual neurons to tone stimuli was visualized by raster plots, and quantified by the peri-event time histogram (PETH, Figures 1A–D). Although the firing rate evoked in any given neuron varied with tone frequency, PETH shapes were largely conserved across tone frequencies. Between neurons, however, PETH shapes differed considerably. PETH shapes showed strongest consistency in the period of ~100ms following stimulus onset (Supplementary Figure 1). To quantify the preservation of temporal structure across stimuli, we computed for each PETH a measure of mean spike latency (MSL; defined as the mean spike time in the 100 ms after stimulus onset; see Experimental Procedures; for analysis of first spike times see Supplementary Figure 2). Figure 1E plots each neuron's MSL to its preferred tone frequency versus its average MSL to all other tones. For neurons with short latency (MSL below 40ms for best frequency), the majority of points are below the diagonal, confirming that for such neurons preferred stimuli often induce earlier firing ($p < 0.001$, paired Wilcoxon signed rank test; for neurons with $MSL > 40ms$ this effect was present but less robust, $p < 0.05$ in 3 out of 8 animals). Nevertheless, the wide scatter between neurons and robust correlation (mean $R = 0.72 \pm 0.24$; $p < 0.001$ for all 8 datasets individually) confirms that temporal profiles are diverse between neurons, and largely conserved within the responses of each cell to different tones. The mean latency measure does not fully summarize the shape of PETHs; for example a number of PETHs had a bimodal structure (e.g. Figure 1C). To confirm that PETH shapes were preserved between stimuli, beyond the conservation of mean latency, we employed a "PETH consistency measure" (Luczak et al., 2007), which showed that a neuron's PETH is more similar to 83% of its own responses to different tones, than to presentations of the same tone to another neuron (see also Supplementary Figure 3). In cortex, spike width can be used to classify cells into putative fast spiking and pyramidal cells (Bartho et al., 2004; Luczak et al., 2007), however we did not observe a significant difference between MSLs of these cell classes ($p > 0.1$, Figure 1E). We conclude that, if a neuron is driven to fire in response to a given tone, it will do so with a stereotyped cell-specific temporal profile.

Responses to natural sounds show similar temporal organization as responses to tones

The fact that individual neurons have consistent and stereotyped PETHs, and that these differ between neurons, indicates that at the population level, responses have a sequential organization. To visualize this organization, Figure 2B shows in grayscale the mean PETHs of a simultaneously recorded population to all tones, sorted by mean response time with MSL indicated by red dots; raster plots for two of the individual neurons are shown above in Figure 2A. To determine whether this sequential organization was preserved in responses to more complex sensory stimuli, we similarly analyzed the response of the same population to a natural sound stimulus (insect vocalization; Figures 2C,D), with neurons sorted in the same order as in Figure 2B. The sequential structure was largely preserved in response to this stimulus. To statistically confirm this finding, we again performed a correlation analysis of MSLs (Figure 2E; $R = 0.69 \pm 0.21$; $p < 0.001$ individually for all 5 rats to which natural sounds were presented). As before, no significant difference was found in MSL between putative pyramidal cells and interneurons ($p > 0.1$). In addition to latency analysis we also quantified PETH consistency across stimuli (79% similarity; Supplementary Figure 3). We therefore conclude that presentation of natural sounds initiates – at least for the first 100ms – activity patterns whose temporal structure is homologous to those evoked by tones.

The analyses presented above were based on PETHs, which are computed from a neuron's response averaged over multiple stimulus presentations. How closely do responses on single trials match this average picture? To address this, we performed a direct comparison of spike times on each individual trial, to the mean temporal profile represented in the PETHs. For each trial, a mean spike time was computed for each neuron firing, and the sequence in which

neurons fired on that trial compared to the PETH means by rank correlation (see Experimental Procedures and Supplementary Figure 4A). Figures 2F and 2G show histograms of rank correlations comparing single-trial sequences evoked by tones and natural sounds, respectively, to PETHs computed from all tones, indicating that single trial spike sequences showed significant similarity to those predicted from the PETHs (t-test: $p < 0.01$ for each experiment; see Supplementary Figure 4C for an alternative approach yielding similar results). As with PETHs, the match of single trials to the average was strongest in the initial response period (~100ms), but decayed thereafter (Supplementary Figure 4B).

Spontaneous upstates have a similar sequential structure to sensory-evoked responses

During sleep, quiet waking, and anesthesia, cortical activity is characterized by an alternation of “downstates” of network silence and “upstates” of generalized spiking and neuronal depolarization, which occur spontaneously in the absence of sensory stimulation (Figure 3A; Supplementary Figure 5; DeWeese and Zador, 2006; Luczak et al., 2007; Steriade et al., 1993a; Steriade et al., 2001). We next asked whether spike patterns accompanying upstates are also temporally homologous to those evoked by sensory stimuli. Figures 3B and 3C show upstate-triggered PETHs of the same neurons as in Figure 2, displayed in the same vertical order. Again, a similar sequential ordering was seen. To statistically confirm this similarity, a slightly different approach was used, as the beginnings of upstates are not experimentally controlled. To measure a cell’s position in the firing sequence accompanying an upstate, without requiring a precise trigger event, we defined a measure μ_{cc} , the center of mass of its cross-correlogram with the summed activity of all other neurons computed in the first 100ms after the onset of each event type (see Experimental Procedures). Values of μ_{cc} were correlated between spontaneous events and stimulus classes, demonstrating that firing order is consistent between sensory stimuli and spontaneous events (Figures 3D,E; $R_{\text{ureth: spont-ton}} = 0.60 \pm 0.14$, $N=8$ rats; $R_{\text{ureth: spont-nat}} = 0.57 \pm 0.18$; $R_{\text{ureth: ton-nat}} = 0.65 \pm 0.07$, $N=5$ rats. $p < 0.001$). Consistency of firing order was again confirmed at the single trial level by rank correlation (Figure 3F; t-test: $p < 0.01$ for each experiment).

Although single-trial responses showed significant homology to the mean, spike timing patterns were not identical across trials, even for repetitions of a single stimulus. Further analyses (Supplementary Figures 6 and 7) suggested that spike timing variability in the initial 100ms period is close to that predicted from the PETHs, under an inhomogeneous Poisson model. After the initial 100ms of tone stimuli, spike timing relative to onset is looser. However, use of the μ_{cc} measure in the late response period (300–500ms after event onset) suggested that temporal relationships between neurons are preserved throughout sensory responses and upstates (Supplementary Figures 5, 8A). With natural sound stimuli, neurons can exhibit stimulus-locked timing later into the stimuli, presumably in response to acoustic features in this period (Supplementary Figure 9)

Preservation of sequential structure in unanesthetized animals

To verify that the above results, obtained under urethane anesthesia, generalize to unanesthetized animals, we performed three further recordings in unanesthetized subjects in a passive listening condition (see Experimental Procedures). As previously reported for resting animals (Luczak et al., 2007; Petersen et al., 2003; Poulet and Petersen, 2008), spontaneous global fluctuations in network activity were seen, although the length and depth of downstates was reduced compared to the anesthetized condition (Figures 4A,B). Consistent sequential activation of neurons between tone responses and upstates was again seen (Figures 4C–E; $R_{\text{unanesth: spont-ton}} = 0.53 \pm 0.17$; $p < 0.001$ for all 3 datasets), with rank correlation analysis indicating a significant homology of single trials to the population mean (Figures 4F,G, t-test: $p < 0.01$ for each experiment). Again, temporal relationships between cells persisted after the

immediate onset period, as shown by correlation of μ_{cc} in the entire stimulated and unstimulated periods (Supplementary Figure 8A).

Sequential structure of sensory-evoked responses in somatosensory cortex

To investigate whether consistent temporal patterns are specific to auditory cortex, or a more general feature of cortical processing, we conducted additional experiments in somatosensory cortex (3 rats, urethane anesthesia). Sensory responses were evoked by whisker stimulation (200ms long air puffs with 1s inter stimulus interval; Figure 5A). As in auditory cortex, consistent temporal profiles of activation were observed between spontaneous and evoked events ($R = 0.57 \pm 0.12$, Figure 5B,C).

Conserved constraints on population firing rate vectors

Population spike patterns thus show constrained temporal structure that is largely conserved across stimuli and spontaneous upstates. However, the number of spikes fired by a given neuron can vary strongly with the stimulus. We next asked whether the possible combinations of neural firing rates, like the temporal pattern of spikes, were subject to conserved constraints across stimuli and spontaneous events. For this analysis we therefore discarded temporal information, and summarized each population pattern by a vector containing the firing rates of each recorded neuron during the first 100 milliseconds after stimulus or upstate onset.

To gain insight into the nature of constraints on firing rates, we initially focused on cell pairs. Figure 6A illustrates, for one pair of cells, the number of spikes fired in individual upstates (black dots), responses to a representative tone (green), and natural sound (red). The region occupied by upstate spike counts has a triangular shape, suggesting the presence of a constraint on the possible spike count combinations: if neuron 2 fired in any given upstate, neuron 1 almost always fired also. The regions occupied by the responses to the two stimuli differed, but both fell within the region outlined by the set of spontaneous events. Figure 6B shows, in outline view, the regions occupied by responses to the two stimuli, spontaneous upstates, and a set of vectors in which spike count correlations had been destroyed by shuffling (spontaneous spike counts were shuffled for each neuron separately, preserving each cell's firing rate distribution but destroying relations between neurons). Part of the realm of shuffled responses (marked in gray) is not occupied by either spontaneous or evoked responses, indicating that these spiking combinations are not produced by the circuit.

We next asked whether a similar phenomenon occurs at the level of the full spike count vectors. Visualization of high-dimensional data requires techniques to map this data into two dimensions. We used multidimensional scaling (MDS; Kruskal and Wish, 1978), a nonlinear method whereby points which are close in the original high-dimensional space will also be placed close by in the 2D projection (Figures 6C,D; see Experimental Procedures). It can be seen that each stimulus produces response vectors that occupy a specific subspace within the realm outlined by spontaneous activity, which is itself contained in the realm outlined by shuffled patterns. To statistically confirm this visual impression, we computed for each sensory response the difference between the Euclidean distances to its closest neighbor in the spontaneous events (E_{spont}), and to its closest neighbor in the shuffled spontaneous events (E_{shuf} ; Figures 6E,F; distances were calculated in the original, high-dimensional space, not the 2D MDS projection). Nearly every evoked event lay closer to a true spontaneous event than a shuffled one ($p < 0.01$ for all experiments, paired Wilcoxon signed rank test). Similar results were seen in the unanesthetized data (Figure 6G; $p < 0.01$ for all experiments), and also when computing count vectors from time bins 300–500ms after onset (Supplementary Figure 8B). Thus, spike count vectors accompanying spontaneous events occupy only a small subspace of the space of possible rate vectors, and responses to individual sensory stimuli occupy subspaces of this “allowed region.”

Analysis of firing rate correlations

The above analyses demonstrated that spontaneous and evoked activity patterns are subject to similar constraints on population firing rates. To gain insight into the character of these constraints, we took a geometrical approach. The firing rates of a set of n neurons can be represented by a point in an n -dimensional vector space. Characterizing the distribution of points in this space is equivalent to characterizing the constraints on population activity. While complete analysis of the structure of high-dimensional distributions is notoriously complex (Agresti, 2002; Scott, 1992), analysis of firing rate correlations between neuronal pairs will allow us to build up a basic geometric intuition for the sizes and orientations of these sets. The correlations between a set of simultaneously recorded neurons may be summarized by a correlation matrix, in which the (i,j) th element gives the correlation between cells i and j . Figure 7A shows a pseudocolor representation of the correlation matrix for spontaneous upstates. For ease of visual analysis, neurons were ordered by a search algorithm so that the most positive correlations were close to the diagonal (see Experimental Procedures). Figure 7B shows the correlation matrix for all sensory responses, with neurons arranged in the same order. The visual similarity of these matrices was confirmed by correlation of their off-diagonal elements (Figure 7D; $p < 0.001$ for all datasets; see Experimental Procedures for how significance was assessed).

The correlations shown in Figure 7B are derived from presentations of all stimuli, without regard to stimulus identity. These correlations could arise from two sources: trial-to-trial correlations in the responses of the population to repetitions of a single stimulus (typically called “noise correlations”); or systematic correlations in the mean responses to multiple stimuli (typically called “signal correlations”). Figure 7C shows the noise correlation matrices for repeated presentations of a single tone and natural sound. The visual similarity of these matrices to Figure 7A was again confirmed by correlation of off-diagonal elements (Figure 7D; $p < 0.001$ for all datasets). The correlation matrix for signal correlation (correlation between average spike counts for each type of stimulus) was also significantly similar to spontaneous correlations ($p < 0.01$ for all datasets, Supplementary Figure 10D). We thus hypothesize that the effect suggested by Figure 6, the constraint of sensory responses to the realm outlined by spontaneous upstates, arose both because the mean responses to individual stimuli were inside this realm, and also because fluctuations around these means were aligned along a similar direction as spontaneous fluctuations. This was further supported by repeating the analyses of Figures 6E and 6F after eliminating noise correlations by shuffling, and directly shifting the mean vectors (Supplementary Figure 10).

To obtain an estimate of the strength of these constraints on population rate vectors, we next asked what were the volumes of the spaces outlined by spontaneous events and sensory responses (note that as multidimensional scaling does not preserve volume, thus this cannot be determined from Figure 6D). Volumes were estimated by calculating the square root determinant of the cell-by-cell spike count covariance matrix (see Supplementary Figure 11). Volumes were expressed as a fraction the volume that would be available to a population displaying the same range of firing rates, in the absence of correlational constraints, as estimated by the volume of the shuffled spontaneous vectors. The volume fraction depended on the size of the population considered. Figure 7E shows volume fractions for the set of responses to a single sensory stimulus, the pooled responses to the 10 stimuli we presented, and the set of spontaneous events, averaged over randomly chosen cell subsets of varying sizes, out of 55 cells recorded in one experiment. In all cases the volume fraction decreased monotonically with the number of cells considered, suggesting that each additional neuron added further constraints at the population level (the distribution of slopes across experiments is shown in Figure 7F).

Predicting receptive fields from spontaneous correlations

To further illustrate the conservation of relationships between neurons in spontaneous and evoked conditions, we used a prediction method. If linear relationships provide a good approximation to the restrictions on population rate vectors, and if these relationships are conserved between spontaneous and evoked activity, then it should be possible to predict a neuron's receptive field based only on its correlations with the rest of the population during spontaneous activity, and from the receptive fields of these other cells. We predicted the firing rate r_j of neuron j as a weighted sum of the rates of all other neurons $r_i = \sum r_i * w_i$, with weights fit to optimize the prediction on spontaneous data (Figure 8A; Harris et al., 2003; Itskov et al., 2008; Luczak et al., 2004). Figures 8B and 8C show the original and predicted receptive field of a representative neuron. Repeating this analysis for all cells, we found the mean correlation between original and predicted receptive fields was $R = 0.62 \pm 0.24$ (112 neurons from 3 experiments in which tuning curve stimuli were presented). To ensure that this effect did not simply reflect similarity of receptive fields of neighboring neurons, we repeated the above analyses excluding neurons recorded from the same shank as the predicted neuron, again finding a significant effect ($R = 0.56 \pm 0.25$). This indicated that receptive field predictability reflects a more complex organization of correlations in the population than simple tonotopy (this is also visible in the correlation matrices of Figure 7B, where even neurons at distant shanks can have substantial correlation). Thus, any particular cell's sensory tuning can be largely predicted from the tuning of its peers, and its correlation with them during spontaneous activity.

Geometrical interpretation of firing rate constraints

The geometrical picture painted by the above results is illustrated in Figures 8D–F. Multidimensional scaling analysis (Figure 6) suggested that clusters corresponding to individual stimuli are confined to the same subregion as spontaneous events, but said little about the structure of these clusters. The conservation of correlation matrices and prediction analysis (Figures 7, 8A–C) suggested a simple picture in which constraints on rate vectors can be approximated by linear correlations. Geometrically, a correlation matrix defines the orientation of the corresponding cluster (Figures 8D,E). The fact that both noise and signal correlation matrices were similar to the correlation matrix for spontaneous events (Figures 7A–D & Supplementary Figure 10) thus suggests that the corresponding clusters have a similar orientation. The geometrical interpretation of this is shown in Figure 8F, in which spontaneous rate vectors are shown as occupying a narrow subspace of theoretically possible vectors, with individual stimuli occupying smaller subspaces of similar orientation.

Discussion

To study the structure of neocortical population spiking activity *in vivo*, we recorded the stimulus-driven and spontaneous activity of neural populations of auditory and somatosensory cortices. We found that population patterns occurring spontaneously and in response to sensory stimulation were subject to common constraints, on both the order in which neurons fire, and the possible combinations of neural firing rates. Although it is only experimentally possible to present a finite number of stimuli, the fact that similar constraints applied to each stimulus we (arbitrarily) chose to present suggests that these constraints will likely apply to all stimuli. These results therefore suggest that the population spike patterns expressible by the cortical circuit are restricted to a limited “vocabulary,” with spontaneous events widely sampling this vocabulary, and responses to sensory stimuli sampling smaller subspaces of it.

Precision and variability of population activity patterns

Previous work has shown that individual neurons can exhibit changes in timing depending on the stimulus (Heil, 2004; Nelken et al., 2005; Optican and Richmond, 1987; Oram et al.,

2002;Panzeri et al., 2001;Petersen et al., 2002). Based on these single-cell results, one might hypothesize that different stimuli may evoke completely different firing sequences (in fact, we hypothesized just this in a recent review article: Harris, 2005). Our present results suggest that this hypothesis is inaccurate. Although individual neurons did show changes in spike timing between stimuli, these were not sufficient to disrupt a sequential structure broadly conserved between stimuli, and the order in which active neurons fired on any trial could be predicted even from the mean response to other stimuli.

A number of previous studies have described spike patterns, sometimes several seconds in length, that repeat with millisecond precision more often than expected under shuffling manipulations (Abeles, 1991;Ikegaya et al., 2004;Nadasdy et al., 1999). Our results suggest a different timescale, with robust sequential organization seen for approximately the first 100ms of responses, and spiking reliability decaying progressively after stimulus onset (see also Luczak et al., 2007). This apparent discrepancy most likely reflects a difference of statistical methodology rather than of biology. If a search is conducted for repeating patterns of millisecond precision, these are the only sequences that can be found; if such patterns are found more often than after shuffling, it shows the original and shuffled data are different, but does not indicate that the particular precision searched for is biologically meaningful. Previous work has shown that the structure of precisely repeating patterns can be predicted from individual neurons' temporal responses to behavioral events or upstate onsets (Baker and Lemon, 2000;Luczak et al., 2007;Oram et al., 2001), suggesting that repetition of millisecond-scale spike templates could be a consequence of response stereotypy at slower timescales.

Possible mechanisms of constrained activity

Sensory responses and spontaneous upstates are likely initiated by different mechanisms. While sensory responses reflect the effects of thalamic input, upstates are believed to be of cortical origin (Sanchez-Vives and McCormick, 2000;Steriade et al., 1993b;Timofeev et al., 2000). One might therefore expect that spontaneous and evoked patterns would propagate differently through cortical circuits. By contrast, our analyses show that both types of activity are subject to common constraints within local populations. We therefore suggest that these constraints arise largely from the dynamics of the local cortical circuit.

One can imagine a number of ways in which the physical properties of a neural circuit could impose consistent constraints on the spike patterns it can generate. First, cortical neurons express diverse sets of voltage-gated ion channels, and are diverse in their intrinsic physiological properties (Storm, 2000;Sugino et al., 2006;Vervaeke et al., 2006). This may contribute to the consistent cellular timing we observe with, for example, cells of low threshold firing earliest (Kang et al., 2008). Second, connectivity within cortical circuits is far from homogenous, for example with strong reciprocal connectivity occurring more than expected by chance (Song et al., 2005). Such connectivity patterns may impose constraints on the possible cell groups that can be active at any time. We observed conserved correlations in firing rate not just locally, but also between neuronal pairs recorded from shanks over 1mm apart (c.f. Eggermont, 2007). Examination of cross-correlograms, however, typically did not indicate a functional monosynaptic connection between correlated pairs (data not shown). These correlations may thus reflect larger-scale network interactions, such as the consistent participation of cells in neuronal assemblies spread over wide cortical areas (Harris, 2005). Although sensory response latencies have been shown to vary across cortical layers (Armstrong-James et al., 1992;Wallace and Palmer, 2008), this is unlikely to account for our results: because we used multi-shank probes inserted perpendicular to the cortical surface, all tetrodes were at approximately the same depth; furthermore consistent timing differences were seen even between neurons recorded from a single tetrode.

Repeated presentation of a single stimulus led to variations in response from trial to trial, which were subject to the same constraints as variations in mean responses between stimuli. Although trial-to-trial variability is often called “noise,” this variability may largely reflect systematic coding of variables not under experimental control (Harris, 2005). Only a small fraction of the input to a cortical column arises from primary sensory thalamus, and responses in sensory cortex can be affected by cognitive factors such as reward and attention (Brosch et al., 2005;Fritz et al., 2007;Shuler and Bear, 2006), other sensory modalities (Brosch et al., 2005;Ghazanfar and Schroeder, 2006), and ongoing oscillations (Hasenstaub et al., 2007;Womelsdorf et al., 2007). We hypothesize that on any trial, these factors, together with direct sensory information from the ascending pathways, determine the activity pattern produced by the circuit; but in all cases the dynamics of the local circuit imposes common constraints on the patterns produced.

Relation to memory replay studies

A large number of studies have suggested that spontaneous activity in resting or sleep replays firing patterns seen during prior behavior (Diba and Buzsaki, 2007;Hoffman and McNaughton, 2002;Ji and Wilson, 2007;Walker and Stickgold, 2006). In our data, similarity between spontaneous and evoked activity is unlikely to reflect replay of the specific stimuli we presented; indeed the animal had never heard these sounds before the experiment. Instead, we hypothesize that this similarity arises because of constraints imposed by the cortical circuit on all spike patterns it can possibly produce. Nevertheless, our results do not contradict the possibility that stimulus replay can occur. For example, after a salient experience of a particular sensory stimulus, spontaneous patterns might preferentially occupy regions of the space of possible events close to the pattern evoked by that stimulus. Our results do however suggest that for experimental characterization of replay, firing rate vectors might be as useful as temporal order and correlation (Battaglia et al., 2005).

Significance of constraints for information encoding

Our results indicate conserved constraints on the activity of cortical populations. What is the significance of these constraints for encoding of sensory information? The fact that timing differences between stimuli occur within a broadly conserved sequential pattern does not imply that timing information cannot be extracted by downstream structures – it simply implies that to do so, the downstream structures must detect variations on a common temporal theme. While accurate information-theoretic analysis of large populations requires currently unrealistic amounts of data, previous studies of single neurons in auditory cortex have shown that both spike times and rates can carry information about stimuli, with partial redundancy between them (Nelken et al., 2005). Our data are consistent with these studies, but suggest that spatiotemporal constraints may limit the information carried by both timing and rates in a population. Although the total strength of constraints increases with the number of neurons considered, the size of the possible “vocabulary” also increases, just less rapidly than in a hypothetically unconstrained situation. We thus anticipate that large enough cortical populations could still encode arbitrarily large ensembles of stimuli, but that this would take more neurons (or more spikes) than in the absence of constraints.

Viewed from one perspective, such redundancy might appear inefficient. Elimination of redundancy has been suggested to be an important function of sensory systems: in populations of primary receptors, activity is likely to be highly correlated, due to statistical regularities in patterns of sensory stimulation, as well as overlapping receptive fields; it has been proposed that higher processing centers remove these correlations, thereby achieving an “economy of impulses” (Barlow, 1962;Chechik et al., 2006). Our data suggest that in auditory cortex, this process is at best incomplete. Another possibility however is that constraints impose a degree of redundancy that is actually beneficial. While redundant codes might require a larger number

of neurons, they can be “robust,” in the sense that the code can still be read if the activity of component cells is unavailable or corrupted. In written language, it is precisely because only a small set of letter combinations form meaningful words that we can understand words in which letters are missing or misprinted. A certain amount of redundancy in cortical coding may thus be beneficial, particularly given that any downstream neuron will only sample a subset of the neurons in the auditory cortex. The fact that constraints are consistent across multiple stimuli and spontaneous patterns might allow downstream structures to consistently utilize this redundancy, to correctly interpret the activity of cortical populations.

Experimental Procedures

Surgery and recording

For anesthetized experiments, Sprague-Dawley rats (300–500g) were anaesthetized with urethane (1.5g/kg) and held with a custom naso-orbital restraint. After preparing a 3mm square window in the skull over auditory cortex (8 rats) or somatosensory cortex (3 rats), the dura was removed and silicon microelectrodes (Neuronexus technologies, Ann Arbor MI) inserted. Probes had 8 or 4 shanks spaced by 200 μ m, with a tetrode recording configuration on each shank (25 μ m spacing between sites; a shank typically yielded 5–15 well isolated units). For awake head-fixed experiments we used an incremental training procedure (see e.g. de Kock and Sakmann, 2008; Hromadka et al., 2008; Luczak et al., 2007; Robbe et al., 2006). A headpost was implanted on the skull of the animal under ketamine-xylazine, and a well drilled above the auditory cortex and covered with wax and dental acrylic. After recovery the animal was trained daily to remain motionless in the restraining apparatus for increasing periods. Excessive movement, signs of stress or discomfort were used to indicate the end of the training session; in some cases chocolate milk reward was given during training. Typically 6–8 days of training were required to reach the target of 1 hour fixation. On the day of the surgery, the animal was briefly anesthetized with isoflurane and the dura resected; after a 1 h recovery period, recording began. Only experiments where the animal stayed motionless for at least an hour, indicated by stable, clusterable recorded units, were included in this study (3/7 rats). The location of the recording sites was estimated to be primary auditory cortex by stereotaxic coordinates, vascular structure (Doron et al., 2002; Rutkowski et al., 2003; Sally and Kelly, 1988) and tonotopic variation of frequency tuning across recording shanks; for somatosensory experiments by stereotaxic coordinates and robust whisker responses. Electrodes were estimated to be in deep layers by field potential reversal (Kandel and Buzsaki, 1997), most likely layer V due to electrode depth and the presence of broadly tuned units of high background rate (Sakata and Harris, SFN Abstract #389.14, 2007). Units were isolated by a semiautomatic algorithm (<http://klustakwik.sourceforge.net>) followed by manual clustering (<http://klusters.sourceforge.net>). To ensure accurate estimation of PETHs, only neurons with firing rates higher than 2 Hz were used in further analysis ($n_{\text{ureth}}=274$ cells, $n_{\text{unanesth}}=81$ cells, $n_{\text{somatosens}}=131$ cells).

Stimuli

As stimuli we used pure tones (3, 7, 12, 20, 30 kHz at 60dB) and in 5 anesthetized experiments also 5 different natural sounds (extracted from the CD “Voices of the Swamp”, Naturesound Studio, Ithica, NY). Each stimulus had duration of 500ms followed by 1500ms of silence (for unanesthetized animals 1s tones were presented followed by 1s silence). All stimuli were tapered at beginning and end with a 5ms cosine window. Experiments took place in single-walled sound isolation chamber (IAC, Bronx, NY) with sounds presented free field (RP2/ES1, Tucker-Davis, Alachua, FL). Spontaneous activity was recorded for 10min before and after presenting stimuli; during stimulus presentation, activity occurring >300ms after stimulus offset and before the next stimulus onset was also regarded as spontaneous. For somatosensory experiments stimuli consisted of 200ms air puffs directed at contralateral whiskers.

Upstate detection

Upstate onsets were identified from the spiking activity of all recorded cells, as the time of the first spike marking a transition from a period of global silence (30ms with at most one spike from any cell) to a period of activity (60ms with at least 15 spikes from any cells; Luczak et al., 2007). To ensure our results did not depend on these particular criteria, analyses were repeated with different parameter values (#spikes in pre-upstate window between 0 and 2, #spikes in following window between 5 and 25). These changes resulted in a different number of detected upstates (within $\pm 20\%$) but did not affect our conclusions. The number of analyzed upstates per experiment (\pm SD) was 701 ± 246 . Note that measures used for statistical analysis (μ_{cc} and population vectors) are by design unaffected by the precise upstate onset time.

Mean spike latency

For each neuron stimulus- and upstate-triggered PETHs were computed using a 10ms Gaussian smoothing kernel. Mean latency was defined as the center of mass of the PETH in the 0–100ms period (equivalent to mean spike time in that time window). To ensure that our results did not depend on these specific definitions, the time of PETH maximum was used as an alternative timing criterion; both procedures resulted in similar conclusions.

μ_{cc}

As upstate onsets can only be determined approximately, accurate estimation of firing sequences during upstates requires a measure insensitive to exact onset time. We therefore computed for each neuron a measure μ_{cc} , defined as the center of mass of the cross-correlogram of this neuron with the summed activity of all other simultaneously recorded cells, computed in the first 100ms after the approximately determined onset of each event type (see Figure 3D and text below Supplementary Figure 8). Analyses were repeated by computing μ_{cc} using multiunit activity taken only from each neuron's local recording shank; for all datasets this resulted in changes in μ_{cc} of less than 5%, indicating that consistent sequential activation was found within local populations, rather than just reflecting spatial spread of activity across multiple shanks (c.f. Luczak et al., 2007).

Single-Trial Rank Correlation Measure

To analyze the preservation of firing sequences on single trials, we compared the sequence of firing evoked by a single stimulus presentation to the mean response to this or other stimuli. Because neurons may fire more than one spike in response to a stimulus, we first calculated for each trial the mean spike time for each firing neuron, in the 100ms window after stimulus onset. Similarity to the mean response sequence was assessed by Spearman's rank correlation of single-trial firing times with the mean spike times computed from the same neurons' PETHs (Supplementary Figure 4A). Only trials with at least three neurons active in the 100ms window were considered. To confirm that the distribution of rank correlations across events was significantly different to zero, a t-test was used. In addition to this approach, a measure based on pattern matching gave similar results (Supplementary Figure 4C).

Multidimensional scaling

Multidimensional scaling was performed in MATLAB with Euclidian metric and Kruskal's normalized stress1 criterion (Kruskal and Wish, 1978). Due to the increase in computational demands with the number of data points, analysis was performed on a random subset of the data (150 points of each class in Figure 6C; the precise number of points chosen did not affect results).

Matrix reordering

To allow visual comparison of correlation matrices (Figure 7A–C), neurons were reordered so that the most positive correlations were placed close to the diagonal. This was achieved with a greedy stepwise search algorithm, optimizing the Frobenius inner product of the reordered correlation matrix with a Toeplitz matrix whose entries decayed exponentially with distance from the main diagonal.

Statistical assessment of matrix similarity

The similarity between correlation matrices was assessed using the correlation coefficient of off-diagonal elements as an intuitive measure. Nevertheless, because these elements are not statistically independent, the significance of this correlation cannot be assessed by standard linear regression. We therefore assessed significance with a randomization method, comparing against a null distribution obtained by randomly reassigning cell identities within each group separately.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This study was supported by NIH grants MH073245 and DC009947. K.D.H. is an Alfred P. Sloan fellow.

Reference List

1. Abeles, M. *Corticonics: neural circuits of the cerebral cortex*. Cambridge: Cambridge University Press; 1991.
2. Agresti, A. *Categorical Data Analysis*. Wiley; 2002.
3. Armstrong-James M, Fox K, Das-Gupta A. Flow of Excitation Within Rat Barrel Cortex on Striking A Single Vibrissa. *Journal of Neurophysiology* 1992;68:1345–1358. [PubMed: 1432088]
4. Averbach BB, Lee D. Neural noise and movement-related codes in the macaque supplementary motor area. *Journal of Neuroscience* 2003;23:7630–7641. [PubMed: 12930802]
5. Baker SN, Lemon RN. Precise spatiotemporal repeating patterns in monkey primary and supplementary motor areas occur at chance levels. *J Neurophysiol* 2000;84:1770–1780. [PubMed: 11024069]
6. Barlow, HB. Possible principles underlying the transformations of sensory messages. In: Rosenzweig, WA., editor. *Sensory communication*. MIT Press; Cambridge, MA: 1962.
7. Bartho P, Hirase H, Monconduit L, Zugaro M, Harris KD, Buzsaki G. Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J Neurophysiol* 2004;92:600–608. [PubMed: 15056678]
8. Battaglia FP, Sutherland GR, Cowen SL, McNaughton BL, Harris KD. Firing rate modulation: A simple statistical view of memory trace reactivation. *Neural Networks* 2005;18:1280–1291. [PubMed: 16257176]
9. Battaglia FP, Sutherland GR, McNaughton BL. Hippocampal sharp wave bursts coincide with neocortical “up-state” transitions. *Learn Mem* 2004;11:697–704. [PubMed: 15576887]
10. Brosch M, Selezneva E, Scheich H. Nonauditory events of a behavioral procedure activate auditory cortex of highly trained monkeys. *J Neurosci* 2005;25:6797–6806. [PubMed: 16033889]
11. Chechik G, Anderson MJ, Bar-Yosef O, Young ED, Tishby N, Nelken I. Reduction of information redundancy in the ascending auditory pathway. *Neuron* 2006;51:359–368. [PubMed: 16880130]
12. Cossart R, Aronov D, Yuste R. Attractor dynamics of network UP states in the neocortex. *Nature* 2003;423:283–288. [PubMed: 12748641]
13. de Kock CPJ, Sakmann B. High frequency action potential bursts (≥ 100 Hz) in L2/3 and L5B thick tufted neurons in anaesthetized and awake rat primary somatosensory cortex. *Journal of Physiology-London* 2008;586:3353–3364.

14. DeWeese MR, Zador AM. Non-Gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. *Journal of Neuroscience* 2006;26:12206–12218. [PubMed: 17122045]
15. Diba K, Buzsaki G. Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience* 2007;10:1241–1242.
16. Doron NN, LeDoux JE, Semple MN. Redefining the tonotopic core of rat auditory cortex: Physiological evidence for a posterior field. *Journal of Comparative Neurology* 2002;453:345–360. [PubMed: 12389207]
17. Eggermont JJ. Correlated neural activity as the driving force for functional changes in auditory cortex. *Hearing Research* 2007;229:69–80. [PubMed: 17296278]
18. Fritz JB, Elhilali M, David SV, Shamma SA. Auditory attention - focusing the searchlight on sound. *Current Opinion in Neurobiology* 2007;17:437–455. [PubMed: 17714933]
19. Gawne TJ, Richmond BJ. How Independent Are the Messages Carried by Adjacent Inferior Temporal Cortical-Neurons. *Journal of Neuroscience* 1993;13:2758–2771. [PubMed: 8331371]
20. Ghazanfar AA, Schroeder CE. Is neocortex essentially multisensory? *Trends in Cognitive Sciences* 2006;10:278–285. [PubMed: 16713325]
21. Harris KD. Neural signatures of cell assembly organization. *Nat Rev Neurosci* 2005;6:399–407. [PubMed: 15861182]
22. Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsaki G. Organization of cell assemblies in the hippocampus. *Nature* 2003;424:552–556. [PubMed: 12891358]
23. Hasenstaub A, Sachdev RNS, McCormick DA. State changes rapidly modulate cortical neuronal responsiveness. *Journal of Neuroscience* 2007;27:9607–9622. [PubMed: 17804621]
24. Heil P. First-spike latency of auditory neurons revisited. *Current Opinion in Neurobiology* 2004;14:461–467. [PubMed: 15321067]
25. Hoffman KL, McNaughton BL. Coordinated reactivation of distributed memory traces in primate neocortex. *Science* 2002;297:2070–2073. [PubMed: 12242447]
26. Hromadka T, DeWeese MR, Zador AM. Sparse representation of sounds in the unanesthetized auditory cortex. *Plos Biology* 2008;6:124–137.
27. Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, Yuste R. Synfire chains and cortical songs: temporal modules of cortical activity. *Science* 2004;304:559–564. [PubMed: 15105494]
28. Itskov V, Curto C, Harris KD. Valuations for spike train prediction. *Neural Computation* 2008;20:644–667. [PubMed: 18045025]
29. Ji DY, Wilson MA. Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience* 2007;10:100–107.
30. Jung MW, Qin Y, Lee D, Mook-Jung I. Relationship among discharges of neighboring neurons in the rat prefrontal cortex during spatial working memory tasks. *Journal of Neuroscience* 2000;20:6166–6172. [PubMed: 10934266]
31. Kandel A, Buzsaki G. Cellular-synaptic generation of sleep spindles, spike-and-wave discharges, and evoked thalamocortical responses in the neocortex of the rat. *J Neurosci* 1997;17:6783–97. [PubMed: 9254689]
32. Kang S, Kitano K, Fukai T. Structure of spontaneous UP and DOWN transitions self-organizing in a cortical network model. *Plos Computational Biology* 2008;4
33. Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A. Spontaneously emerging cortical representations of visual attributes. *Nature* 2003;425:954–956. [PubMed: 14586468]
34. Kruskal, JB.; Wish, M. *Multidimensional Scaling*. Beverly Hills, California: Sage; 1978.
35. Luczak A, Bartho P, Marguet SL, Buzsaki G, Harris KD. Sequential structure of neocortical spontaneous activity in vivo. *Proceedings of the National Academy of Sciences of the United States of America* 2007;104:347–352. [PubMed: 17185420]
36. Luczak A, Hackett TA, Kajikawa Y, Laubach M. Multivariate receptive field mapping in marmoset auditory cortex. *Journal of Neuroscience Methods* 2004;136:77–85. [PubMed: 15126048]
37. MacLean JN, Watson BO, Aaron GB, Yuste R. Internal dynamics determine the cortical response to thalamic stimulation. *Neuron* 2005;48:811–823. [PubMed: 16337918]
38. Mao BQ, Hamzei-Sichani F, Aronov D, Froemke RC, Yuste R. Dynamics of spontaneous activity in neocortical slices. *Neuron* 2001;32:883–898. [PubMed: 11738033]

39. Massimini M, Huber R, Ferrarelli F, Hill S, Tononi G. The sleep slow oscillation as a traveling wave. *J Neurosci* 2004;24:6862–6870. [PubMed: 15295020]
40. Mokeichev A, Okun M, Barak O, Katz Y, Ben Shabar O, Lampl I. Stochastic emergence of repeating cortical motifs in spontaneous membrane potential fluctuations in vivo. *Neuron* 2007;53:413–425. [PubMed: 17270737]
41. Montani F, Kohn A, Smith MA, Schultz SR. The role of correlations in direction and contrast coding in the primary visual cortex. *Journal of Neuroscience* 2007;27:2338–2348. [PubMed: 17329431]
42. Nadasdy Z, Hirase H, Czurko A, Csicsvari J, Buzsaki G. Replay and time compression of recurring spike sequences in the hippocampus. *J Neurosci* 1999;19:9497–507. [PubMed: 10531452]
43. Narayanan NS, Kimchi EY, Laubach M. Redundancy and synergy of neuronal ensembles in motor cortex. *Journal of Neuroscience* 2005;25:4207–4216. [PubMed: 15858046]
44. Nelken I, Chechik G, Mscis-Flogel TD, King AJ, Schnupp JWH. Encoding stimulus information by spike numbers and mean response time in primary auditory cortex. *Journal of Computational Neuroscience* 2005;19:199–221. [PubMed: 1613819]
45. Optican LM, Richmond BJ. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex. III Information theoretic analysis. *J Neurophysiol* 1987;57:162–178. [PubMed: 3559670]
46. Oram MW, Hatsopoulos NG, Richmond BJ, Donoghue JP. Excess synchrony in motor cortical neurons provides redundant direction information with that from coarse temporal measures. *J Neurophysiol* 2001;86:1700–1716. [PubMed: 11600633]
47. Oram MW, Xiao D, Dritschel B, Payne KR. The temporal resolution of neural codes: does response latency have a unique role? *Philos Trans R Soc Lond B Biol Sci* 2002;357:987–1001. [PubMed: 12217170]
48. Panzeri S, Petersen RS, Schultz SR, Lebedev M, Diamond ME. The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron* 2001;29:769–777. [PubMed: 11301035]
49. Petersen CC, Hahn TT, Mehta M, Grinvald A, Sakmann B. Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. *Proc Natl Acad Sci USA* 2003;100:13638–13643. [PubMed: 14595013]
50. Petersen RS, Panzeri S, Diamond ME. Population coding in somatosensory cortex. *Curr Opin Neurobiol* 2002;12:441–447. [PubMed: 12139993]
51. Poulet JFA, Petersen CCH. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature*. 2008advance online publication 16 July 2008
52. Robbe D, Montgomery SM, Thome A, Rueda-Orozco PE, McNaughton BL, Buzsaki G. Cannabinoids reveal importance of spike timing coordination in hippocampal function. *Nature Neuroscience* 2006;9:1526–1533.
53. Rutkowski RG, Miasnikov AA, Weinberger NM. Characterisation of multiple physiological fields within the anatomical core of rat auditory cortex. *Hear Res* 2003;181:116–130. [PubMed: 12855370]
54. Sally SL, Kelly JB. Organization of auditory cortex in the albino rat: sound frequency. *J Neurophysiol* 1988;59:1627–1638. [PubMed: 3385476]
55. Sanchez-Vives MV, McCormick DA. Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nat Neurosci* 2000;3:1027–1034. [PubMed: 11017176]
56. Schneidman E, Berry MJ, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 2006;440:1007–1012. [PubMed: 16625187]
57. Scott, D. *Multivariate Density Estimation: Theory, Practice, and Visualization*. Wiley; 1992.
58. Shuler MG, Bear MF. Reward timing in the primary visual cortex. *Science* 2006;311:1606–1609. [PubMed: 16543459]
59. Sirota A, Csicsvari J, Buhl D, Buzsaki G. Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences of the United States of America* 2003;100:2065–2069. [PubMed: 12576550]
60. Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB. Highly nonrandom features of synaptic connectivity in local cortical circuits. *Plos Biology* 2005;3:507–519.
61. Steriade M, Nunez A, Amzica F. A novel slow (< 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *J Neurosci* 1993a;13:3252–65. [PubMed: 8340806]

62. Steriade M, Nunez A, Amzica F. Intracellular analysis of relations between the slow (< 1 Hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. *J Neurosci* 1993b; 13:3266–83. [PubMed: 8340807]
63. Steriade M, Timofeev I, Grenier F. Natural waking and sleep states: a view from inside neocortical neurons. *J Neurophysiol* 2001;85:1969–1985. [PubMed: 11353014]
64. Storm JF. K⁺ channels and their distribution in large cortical pyramidal neurones. *Journal of Physiology-London* 2000;525:565–566.
65. Sugino K, Hempel CM, Miller MN, Hattox AM, Shapiro P, Wu CZ, Huang ZJ, Nelson SB. Molecular taxonomy of major neuronal classes in the adult mouse forebrain. *Nature Neuroscience* 2006;9:99–107.
66. Timofeev I, Grenier F, Bazhenov M, Sejnowski TJ, Steriade M. Origin of slow cortical oscillations in deafferented cortical slabs. *Cerebral Cortex* 2000;10:1185–1199. [PubMed: 11073868]
67. Tsodyks M, Kenet T, Grinvald A, Arieli A. Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science* 1999;286:1943–1946. [PubMed: 10583955]
68. Vervaeke K, Hu H, Graham LJ, Storm JF. Contrasting effects of the persistent Na⁺ current on neuronal excitability and spike timing. *Neuron* 2006;49:257–270. [PubMed: 16423699]
69. Walker MP, Stickgold R. Sleep, memory, and plasticity. *Annu Rev Psychol* 2006;57:139–166. [PubMed: 16318592]
70. Wallace MN, Palmer AR. Laminar differences in the response properties of cells in the primary auditory cortex. *Experimental Brain Research* 2008;184:179–191.
71. Womelsdorf T, Schoffelen JM, Oostenveld R, Singer W, Desimone R, Engel AK, Fries P. Modulation of neuronal interactions through neuronal synchronization. *Science* 2007;316:1609–1612. [PubMed: 17569862]
72. Zohary E, Shadlen MN, Newsome WT. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 1994;370:140–3. [PubMed: 8022482]

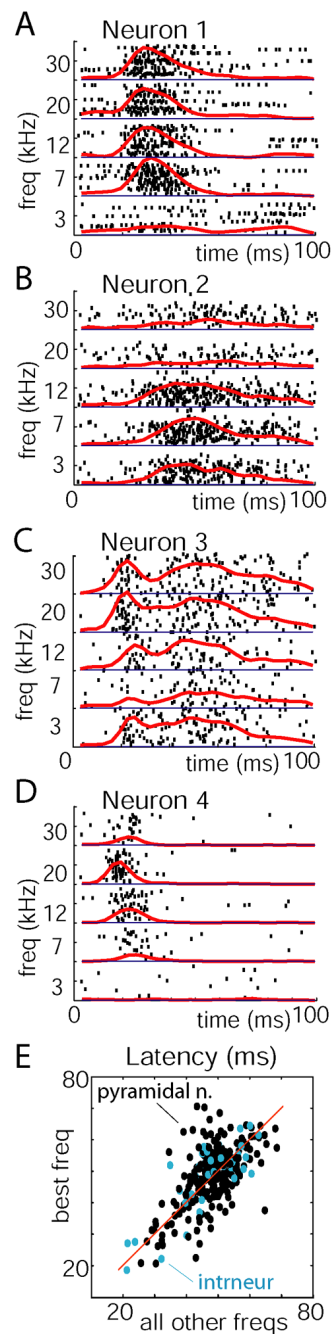


Figure 1.

Individual neurons respond to different tones with stereotyped temporal profiles, but varying firing rates. (A–D) Raster plots showing responses of representative neurons to presentations of 5 pure tones (100 trials for each tone). Red lines represent peri-stimulus time histograms. (E) Scatter plot showing each neuron’s mean spike latency (MSL) to its preferred tone frequency vs. to all other tones. The red line corresponds to equal latencies. Blue dots indicate putative interneurons as defined by spike width. While neurons typically show earlier firing to their preferred tone, this difference is an order of magnitude smaller than the differences between cells.

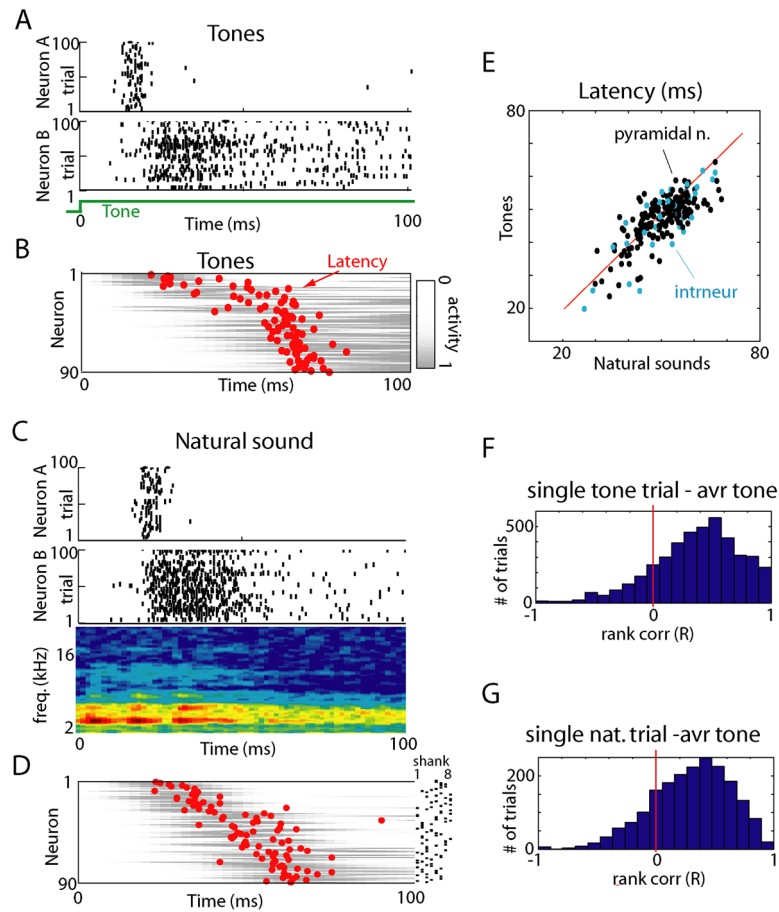


Figure 2.

Similar temporal activity patterns initiated by presentation of tones and natural sounds. (A) Raster plots showing spike times for two representative neurons to repeated presentations of a pure tone stimulus. (B) Average activity of 90 simultaneously recorded neurons to tone stimuli. Grey bars show pseudocolor representations of each neuron's perievent time histogram (PETH), red dots denote each neuron's mean spike latency in the 100ms after tone onset. Neurons are ordered vertically by the mean latency over all stimuli, to illustrate sequential spread of activity. (C) Response of the same two neurons as in (A) to a natural sound (insect vocalization; sound spectrogram shown below rasters), illustrating similar temporal response profiles as to the tone. (D) Response of the same population as (B), displayed in the same vertical order, indicating that the sequential order of firing is preserved. The dots on the right indicate at which shank neurons were recorded. (E) Scatter plot showing each neuron's mean spike latency for tones and natural sounds with putative interneurons marked in blue. The distribution of points along the diagonal indicates preservation of sequential structure across conditions. (F) Histogram of rank correlations between mean spike times for individual tone presentations and mean response profile across all tones (see Supplementary Figure 4A). (G) Histogram of rank correlations between mean spike times for single natural sound presentation and average across all tones. The prevalence of positive correlations indicates that for the majority of trials, the sequence of neuronal activation was preserved.

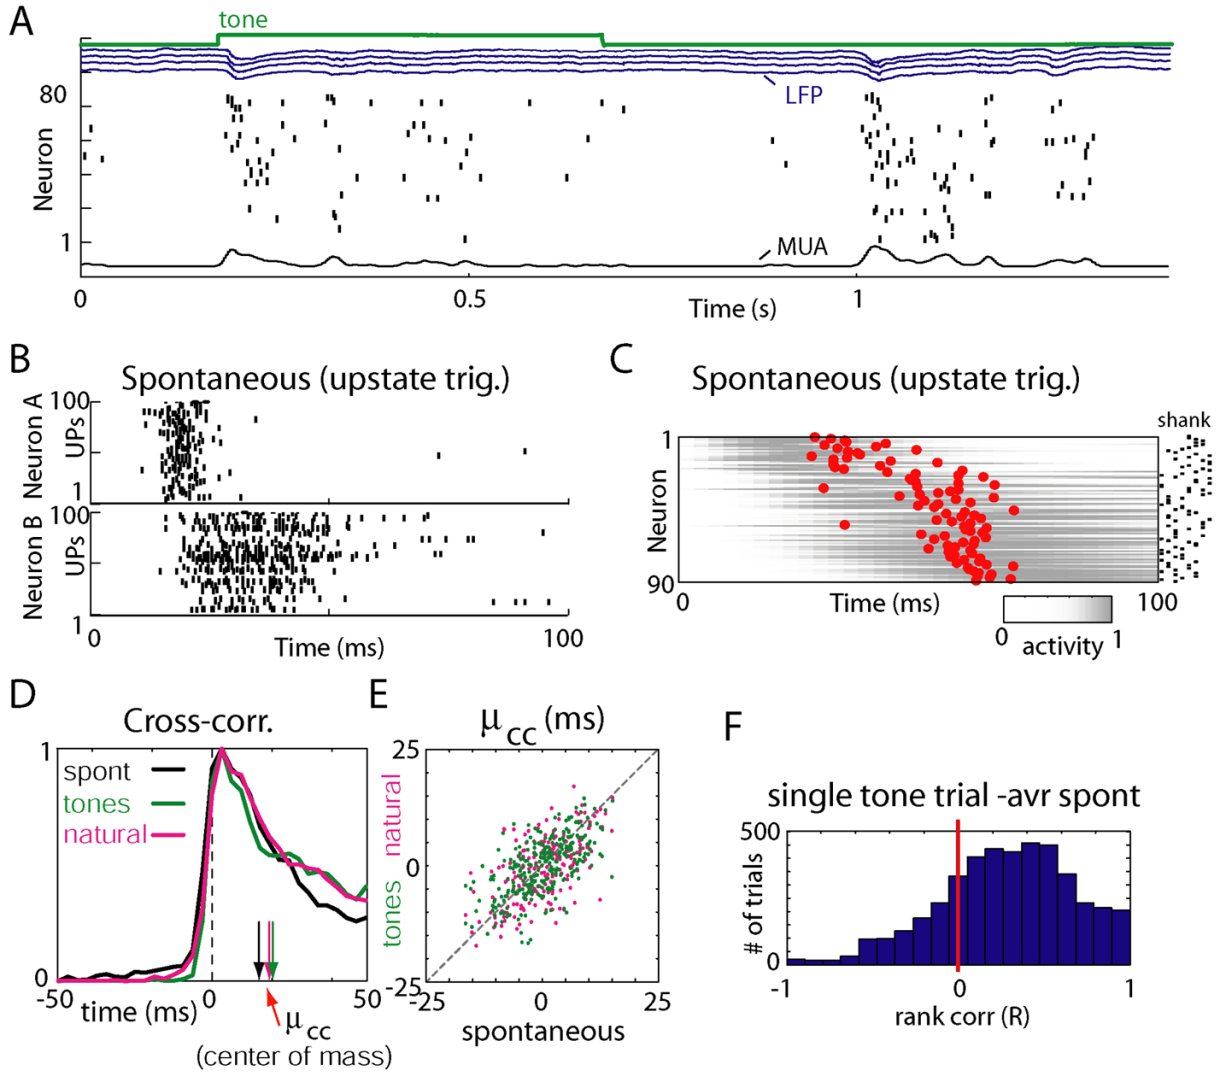


Figure 3. Spontaneous upstates initiate sequential patterns homologous to evoked responses. (A) Representative raw data plot showing a tone response and spontaneous firing event. The green trace is a synchronization pulse indicating the duration of a tone stimulus; blue traces show local field potentials (LFP) from four separate recording shanks; underneath is a raster plot showing the spike trains of simultaneously recorded neurons. At bottom is the multiunit firing rate (MUA) computed by averaging all neurons. Neurons are sorted by average spontaneous mean spike latency to facilitate visual examination of temporal patterns. (B) Raster plots showing spike times for the same neurons as in Figures 2A and 2C, triggered by upstate onsets. Note the similar temporal pattern to Figure 2. (C) Average upstate-triggered activity of all neurons, sorted in the same order as in Figures 2B and D. (D) Cross-correlograms of one neuron's spike times with the summed activity of all other cells, during different experimental conditions. Vertical arrows indicate the center of mass (mean spike time) of correlograms (μ_{cc}). Cross-correlograms are normalized between 0 and 1 to facilitate comparison. (E) Conservation of μ_{cc} across different stimuli and spontaneous events, indicating preservation of sequential order. Each point represents the values of μ_{cc} for a given cell in the conditions

indicated on the axes. (F) Histogram of rank correlations between mean spike times for single-trial tone presentations and average mean spike times for spontaneous events.

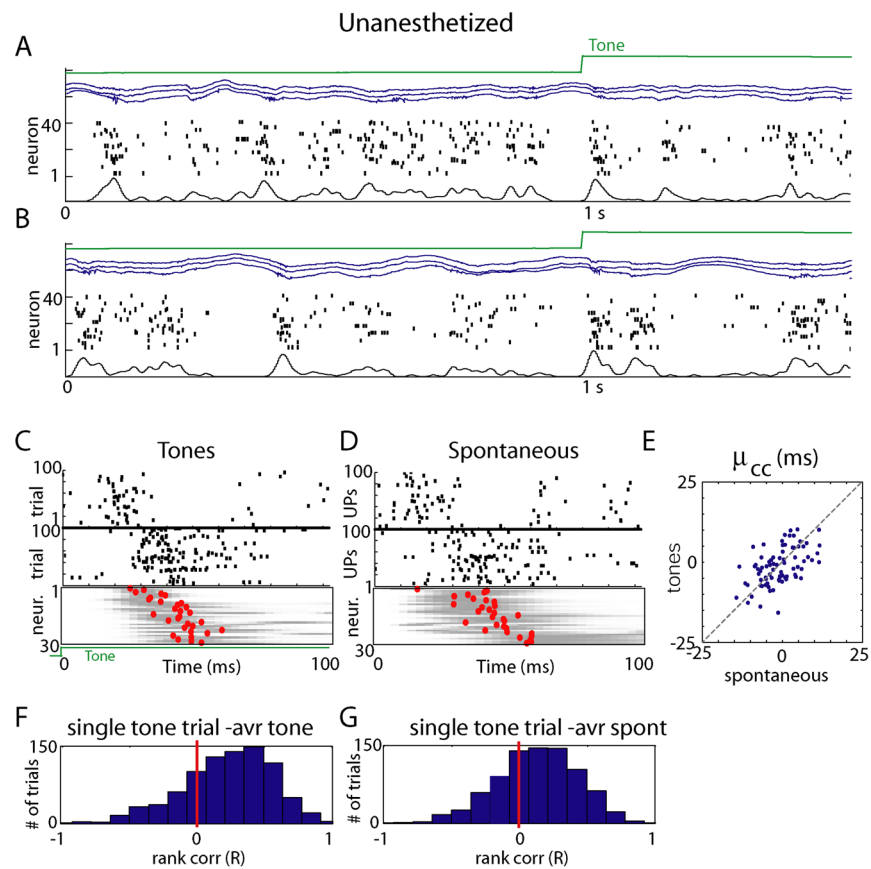


Figure 4.

Preservation of sequential structure between sensory-evoked and spontaneous events in unanesthetized animals. (A, B) Representative raw data plots from an unanesthetized, head-fixed subject in a passive listening paradigm. Again, global fluctuations in activity are seen, although downstates are typically shorter than under anesthesia. (C,D) Similar analysis as in Figure 3B&C, showing preservation of individual neurons' PETH, and conservation of sequential structure. (E) Conservation of μ_{cc} across tones and spontaneous events in unanesthetized animals (similar analysis to Figure 3E). (F,G) Histograms of rank correlations between mean spike times for single tone presentation and average across all tones and spontaneous events respectively.

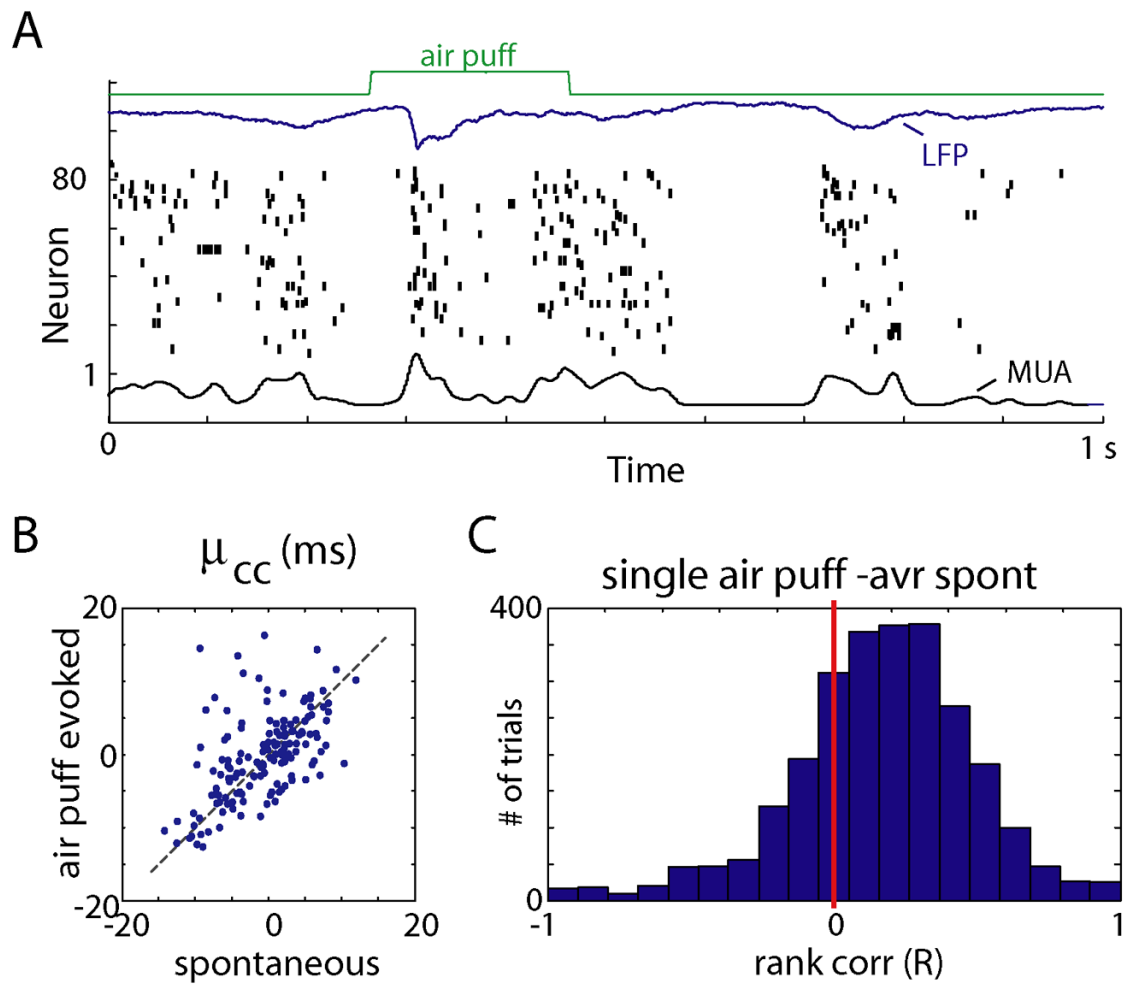


Figure 5. Preservation of temporal structure between sensory-evoked and spontaneous events in somatosensory cortex. (A) Representative raw data plots of spontaneously occurring upstates and air-puff evoked activity. (B) Conservation of μ_{CC} across different stimuli and spontaneous events, indicating preservation of sequential order. (C) Histogram of rank correlations between mean spike times for single air puff and average mean spike times for spontaneous events.

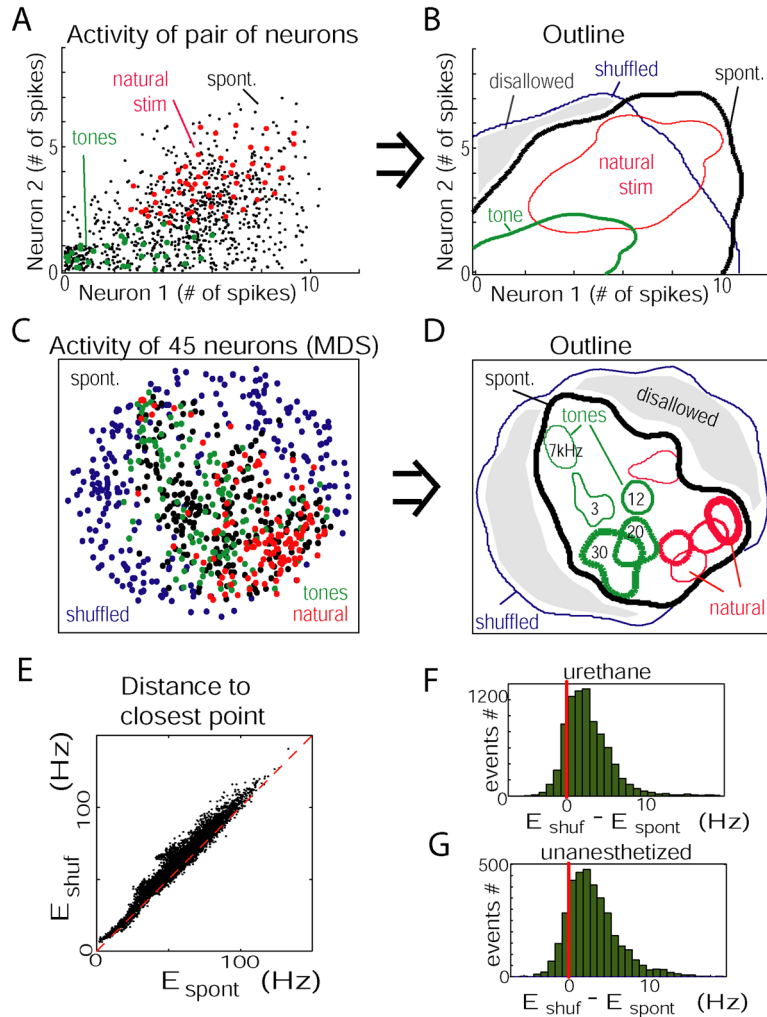


Figure 6. Combinatorial constraints on population firing rate vectors. (A) Spike counts of two neurons (recorded from separate tetrodes) during the first 100ms of spontaneous upstates (black), responses to a tone (green), and natural sound (magenta). Data were jittered to show overlapping points. Note that regions occupied by responses to the sensory stimuli differ, but are both contained in the realm outlined by spontaneous patterns. (B) Contour plot showing regions occupied by points from (A). The blue outline is computed from spike counts shuffled between upstates, indicating the region that would be occupied in the absence of spike count correlations. (C) Firing rate vectors of entire population, visualized using multidimensional scaling; each dot represents the activity of 45 neurons, nonlinearly projected into two-dimensional space. (D) Contour plot derived from multidimensional scaling data, with responses to individual stimuli marked separately. Sensory-evoked responses again lie within the realm outlined by spontaneous events. (E) Scatter plot showing the Euclidean distances from each evoked event to its closest neighbor in the spontaneous events (E_{spont}), and in the shuffled spontaneous events (E_{shuf}). Dashed red line shows equality. (F, G) Histogram showing the difference between distances to shuffled and spontaneous events ($E_{shuf} - E_{spont}$). Top and bottom: data from all anesthetized and unanesthetized experiments, respectively. Almost every evoked event was closer to a true spontaneous vector than to a shuffled vector.

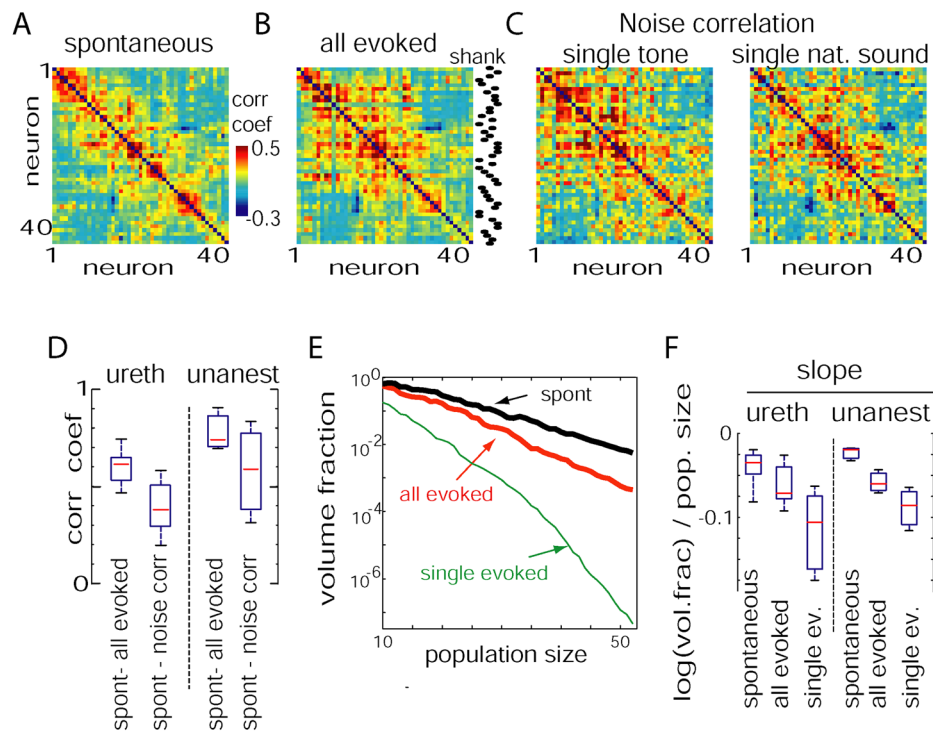


Figure 7.

Analysis of pair-wise correlations during evoked and spontaneous conditions. (A) Correlation matrix between spike counts of 45 neurons calculated during upstates. For ease of visualization, neurons were ordered so that the highest correlations are close to the diagonal. (B) Spike count correlation matrix for responses to all sensory stimuli (5 tones and 5 natural sounds), with neurons ordered the same as in A. The dots on the right indicate at which shank neurons were recorded. (C) Correlation matrices for repeated presentations of a single tone and a single natural sound (“noise correlations”), with neurons again ordered the same as in A; note the similar appearance of all matrices. (D) Boxplots showing distribution across experiments of elementwise correlation coefficients between correlation matrices (diagonal excluded). (E) The volume of response space occupied by spontaneous events (black), all evoked activity (red), and the responses to a single stimulus (green), was estimated as a fraction of the volume occupied by shuffled spontaneous events as the square root ratio of covariance matrix determinants (see Supplementary Figure 11). Note the monotonic decrease with population size. (F) Box plots showing distribution of slopes of log-volume fraction as a function of population size for anesthetized (left) and unanesthetized (right) experiments.

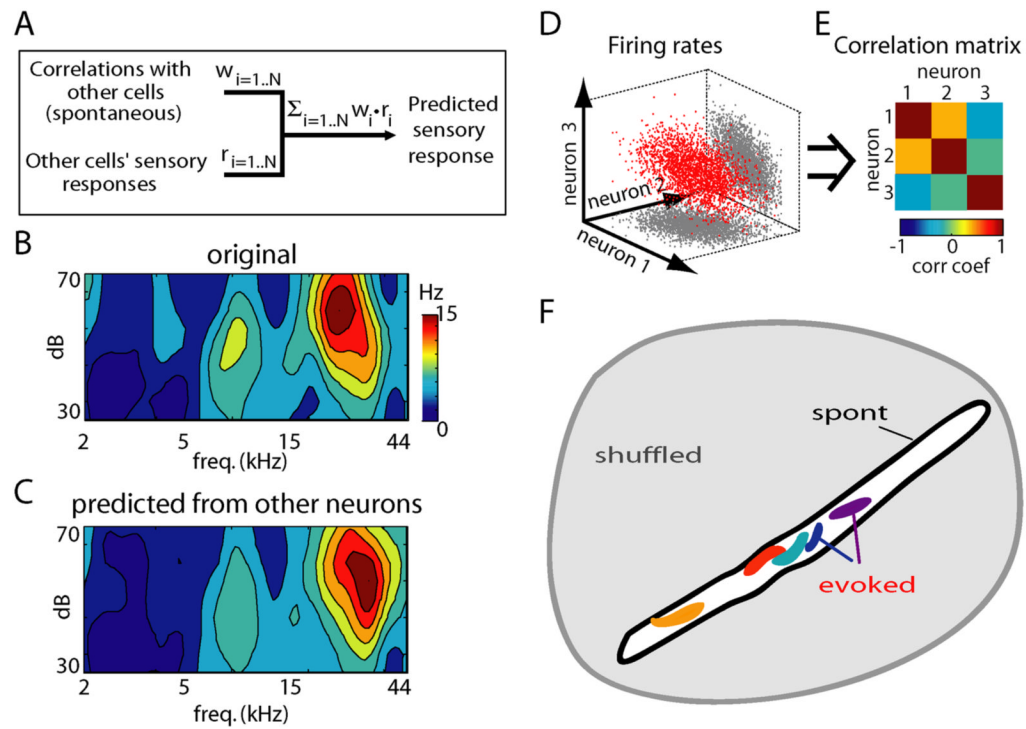


Figure 8.

Preserved constraints on firing rate vectors allow prediction of a neuron's sensory tuning based on spontaneous activity, and a geometrical interpretation of these constraints. (A) Prediction of a neuron's receptive field from correlations with simultaneously recorded neurons during spontaneous activity, and from the receptive fields of these other neurons. Here, w denotes a vector of weights, optimized to maximize prediction of the target neuron's activity during spontaneous activity, and r denotes the receptive field vectors of N other neurons. (B,C) Actual and predicted receptive fields for a representative neuron. On average, predicted and actual receptive fields showed a correlation of 0.62 (see text for details). (D,E) General relationship between correlation matrix and cluster orientation. (D) shows a set of simulated spike count vectors with the correlation matrix shown in (E); the values of the correlation matrix provide information about the orientation of the cluster relative to the coordinate axes. (F) Cartoon illustrating the geometrical interpretation of our findings. The gray area illustrates the space of all rate vectors theoretically possible in the absence of relationships between neurons. The black outline represents the space of spontaneous events; this is shown elongated and of small volume to illustrate strong constraints at the population level. Responses to individual stimuli occupy smaller subsets within this (colored blobs; the irregular shape illustrates possible non-Gaussianity of these clusters). The orientations of the spaces for individual stimuli (corresponding to noise correlation matrices) are approximately aligned with the space of spontaneous events. The mean response to each stimulus also lies within the space of spontaneous events (see Supplementary Figure 10).