

# Non-Gaussian Membrane Potential Dynamics Imply Sparse, Synchronous Activity in Auditory Cortex

Michael R. DeWeese and Anthony M. Zador

Cold Spring Harbor Laboratory, Cold Spring Harbor, New York 11724

Many models of cortical dynamics have focused on the high-firing regime, in which neurons are driven near their maximal rate. Here we consider the responses of neurons in auditory cortex under typical low-firing rate conditions, when stimuli have not been optimized to drive neurons maximally. We used whole-cell patch-clamp recording *in vivo* to measure subthreshold membrane potential fluctuations in rat primary auditory cortex in both the anesthetized and awake preparations. By analyzing the subthreshold membrane potential dynamics on single trials, we made inferences about the underlying population activity. We found that, during both spontaneous and evoked responses, membrane potential was highly non-Gaussian, with dynamics consisting of occasional large excursions (sometimes tens of millivolts), much larger than the small fluctuations predicted by most random walk models that predict a Gaussian distribution of membrane potential. Thus, presynaptic inputs under these conditions are organized into quiescent periods punctuated by brief highly synchronous volleys, or “bumps.” These bumps were typically so brief that they could not be well characterized as “up states” or “down states.” We estimate that hundreds, perhaps thousands, of presynaptic neurons participate in the largest volleys. These dynamics suggest a computational scheme in which spike timing is controlled by concerted firing among input neurons rather than by small fluctuations in a sea of background activity.

**Key words:** auditory; cortex; whole-cell recording; rat; synchrony; random walk; neural coding

## Introduction

Most of what we know about the dynamics of neuronal population activity in the cortex comes from recording neurons individually or in small groups. However, what is most relevant to any particular neuron is not the activity of randomly chosen nearby neurons or the activity of the entire cortical population; rather, it is the activity of the specific subpopulation that provides input to that neuron. A typical cortical neuron receives input from ~10,000 other neurons (Braitenberg and Schuz, 1998), distributed within the same cortical column as well as much more distal parts of cortex and elsewhere. Identifying and simultaneously recording from even a modest fraction of these numerous and widely dispersed neurons would be a daunting task using conventional techniques. A more practical approach is to infer the activity of this subpopulation by recording the membrane potential of a single neuron within the intact cortex. Specifically, the dynamics of the membrane potential on individual trials places constraints on the degree to which the presynaptic subpopulation cooperates in driving the recorded neuron to fire.

What can we infer about population activity from the dynamics of the subthreshold membrane potential? One can imagine two extreme limiting cases. At one extreme, the input to the

neuron consists of many small uncorrelated postsynaptic potentials (PSPs) summed together; in this case, the membrane potential is Gaussian distributed and follows a “random walk.” Such models have been widely used to describe the inputs to cortical neurons in both theoretical (Gerstein and Mandelbrot, 1964; Calvin and Stevens, 1967; Softky and Koch, 1993; Shadlen and Newsome, 1994; Shadlen and Newsome, 1995; Tsodyks, 1995; van Vreeswijk and Sompolinsky, 1998; Song et al., 2000; Fellous et al., 2003; Rudolph and Destexhe, 2003) and experimental (Destexhe et al., 2003; Carandini, 2004) studies but have yet to be experimentally tested in auditory cortex. At the other extreme, the presynaptic population is highly correlated; neurons might be silent most of the time, except during brief moments when large groups of them fire in a concerted manner, which would elicit rare large excursions (or “bumps”) of the membrane potential of the postsynaptic neuron.

One cannot distinguish between even these two extreme models by observing the output spike train alone: either model could account for essentially any observed spike train. We therefore used whole-cell patch-clamp recording techniques in both anesthetized and awake animals to measure directly the dynamics of the membrane potential of auditory cortical neurons on a single-trial basis. We focused on neural responses recorded under the low-firing rate conditions typically seen in the awake state (Evans and Whitfield, 1964), when stimuli have not been optimized to drive neurons maximally. We found that membrane potential dynamics were very similar across neurons: under all stimulation conditions tested, fluctuations suggested the second model, in which highly correlated firing within the population of presynaptic afferents manifested as bumps in the postsynaptic voltage.

Received July 2, 2006; revised Sept. 17, 2006; accepted Oct. 17, 2006.

This work was supported by grants from the National Institutes of Health, the Sloan Foundation, the Packard Foundation, and the Mathers Foundation (A.M.Z.) and by a Swartz fellowship (M.R.D.).

Correspondence should be addressed to Michael R. DeWeese, Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724. E-mail: deweese@cshl.edu.

DOI:10.1523/JNEUROSCI.2813-06.2006

Copyright © 2006 Society for Neuroscience 0270-6474/06/2612206-13\$15.00/0

These dynamics are compatible with computational schemes in which spike timing is precisely controlled by concerted firing in other neurons.

## Materials and Methods

To study the single-trial dynamics of membrane potential fluctuations in auditory cortex, we presented brief pure-tone pips as well as long pure tones of different frequencies and fixed intensity while recording from neurons in the auditory cortex of rats; silent epochs were analyzed as well. As described previously (DeWeese and Zador, 2004), we used *in vivo* whole-cell patch-clamp methods (Ferster and Jagadeesh, 1992; Metherate and Ashe, 1993; Nelson et al., 1994; Hirsch et al., 1995; Borg-Graham et al., 1998; Moore and Nelson 1998; Zhu and Connors, 1999) to record the PSPs in the intact animal. The main analyses are based on  $n = 35$  neurons recorded in anesthetized animals; in addition, we recorded  $n = 5$  neurons in awake animals.

### Quantifying bump duration

We quantified the duration of bumps in the membrane potential by measuring the width at half-maximum for every bump that was at least 10 mV tall as measured relative to the rest potential. For the “short-tone protocol,” the rest potential for each trace was estimated by first taking the mean of the membrane potential during the 15 ms preceding each tone onset and then taking the median across all of these pre-tone values. For the “long-tone protocol” and the “silent protocol,” we estimated the rest potential for each 4 s trace by taking the fifth percentile value of the membrane potential across the trace. Before computing bump statistics for the long-tone data, we excised the initial 250 ms from each trace to ensure that they reflected the steady-state behavior of the neuron during prolonged tone presentation rather than the transient stimulus onset response; trials concluded at tone termination so that OFF responses were also excluded from the long-tone protocol. We distinguished between “tone-evoked” and “spontaneous” bumps within the short-tone protocol by designating as “tone-evoked” all bumps crossing threshold between 10 and 100 ms after the onset of any of the 25-ms-duration tones.

The number reported in the text is the maximum bump width for each neuron, averaged across neurons from all four stimulus conditions (maximum width of bumps,  $54.6 \pm 3.3$  ms;  $n = 30$  neurons, 2 stimulus conditions per neuron; all quantities are mean  $\pm$  SE unless otherwise specified). Broken down by stimulation protocol, the maximum width of tone-evoked bumps during the short-tone protocol was  $37 \pm 4$  ms,  $n = 16$  cells; spontaneous bumps during the short-tone protocol was  $48 \pm 8$  ms,  $n = 16$  cells; the long-tone protocol was  $71 \pm 5$  ms,  $n = 14$  cells; and the silent protocol was  $66 \pm 5$  ms,  $n = 14$ .

### Random walk model

We made use of three variants of a “random walk model” that integrated unitary inputs to produce an output voltage trace.

**Excitation only model.** The input for the simplest of these variants (see Fig. 5*b*) consisted of purely excitatory (i.e., positive) unitary events, with a time course  $q_{\text{syn}}(t)$ ; the arrival time of each of these unitary events was statistically independent of all the others, and they occurred at a constant rate so that they obeyed a homogeneous Poisson process. We can write the time-dependent voltage,  $v(t)$ , generated by this model as follows:

$$v = q_{\text{syn}} * z,$$

where  $z(t)$  is an instantiation of a homogeneous Poisson process with rate  $\lambda_e$ , and  $*$  denotes convolution,

$$v(t) = \int q_{\text{syn}}(\tau) z(t - \tau) d\tau.$$

The excitatory event rate,  $\lambda_e$ , is the only free parameter in this model, which we fit to the difference between the mean voltage of the data trace and the rest potential of the neuron. (We used the voltage corresponding to the fifth percentile of the membrane potential histogram as the rest potential of the neuron, although the results were insensitive to this

choice for all variants of the random walk model.) For the instantiation of the model shown in Figure 5*b*, we used unitary events [miniature EPSPs (mEPSPs)] that were 0.4 mV at the peak, with a 10 ms exponential decay. Therefore, each event contributed a finite area of 4 (millivolts)  $\times$  (milliseconds) to the resulting output voltage trace. Thus, to fit the 1.3 mV mean of the recorded trace, we set the event rate to  $\sim 0.33$  events/ms.

**Excitation and inhibition model.** We next considered a model that included both excitatory inputs and inhibitory inputs (see Figs. 5*c*, 6*a,b*):

$$v = q_{\text{syn}} * z_e - q_{\text{syn}} * z_i,$$

where  $z_e(t)$  and  $z_i(t)$  are instantiations of point processes drawn from Poisson processes with the mean rate given by the excitatory and inhibitory event rates,  $\lambda_e$  and  $\lambda_i$ , respectively.

For a Poisson process, the mean of the number of events per time bin is equal to the variance in the event count in the same bin, and we chose inhibitory events that were the same magnitude as the excitatory events but opposite in sign. Thus, we could simultaneously fit the mean,  $\mu$ , of the data trace, which was proportional to the difference between the event rates:

$$\mu = (\lambda_e - \lambda_i)q$$

and the variance of the trace,  $\sigma^2$ , which was proportional to the sum of the event rates:

$$\sigma^2 = (\lambda_e + \lambda_i)q^2,$$

where  $q$  is the integrated contribution from one unitary event,  $q_{\text{syn}}$ , divided by the unit time step. Note that the mean and variance of the measured whole-cell records were measured across time for each trace rather than across trials.

**Rate modulated excitation and inhibition model.** Finally, we allowed the Poisson rate parameters  $\lambda_e$  and  $\lambda_i$  to vary with time (with the timescale,  $\tau$ , specified) (see Fig. 5*d,e*). Essentially, we fit this rate-modulated random walk model to the measured voltage traces by fitting both rate parameters at every time point to the mean and variance of a short segment (of length  $\tau$ ) of the trace centered on that time point. However, occasionally, the variance for a segment was less than what the “excitation only model” would give once the mean was fit perfectly; adding inhibitory events can only decrease the mean and increase the variance, so in these cases, we set the inhibitory rate to 0.

For Figure 5, *d* and *e*, we wanted to fit the model as well as possible to the features in the data, so the algorithm was modified to perform something very close to a proper deconvolution. Specifically, as the algorithm stepped from time point to time point, it accounted for the contribution to the mean of the model trace resulting from the tails of unitary events that began before the current segment, and it discounted the contribution to the mean resulting from the loss of tails passing out of the current segment. Unfortunately, there are (well known) oscillatory instabilities that arise during convolution when an algorithm attempts to correct for mistakes made on previous time steps attributable to overfitting features of the data that are much finer than the unitary events in the model. To avoid this, we only added inhibitory events when the actual mean value (i.e., the mean value before accounting for the contributions to the model trace from tails of events occurring before the current segment) was negative or when the variance of the current segment was greater than the best fit of the excitation only model.

### Kurtosis analysis

We quantified the agreement between the excitation and inhibition (E&I) variant of the random walk model and each 4 s voltage trace from our dataset by computing the kurtosis of each trace; specifically, we computed the “sample kurtosis excess,” defined as the ratio of the fourth central moment divided by the fourth power of the SD, minus three (the value expected for a Gaussian distribution):

$$\text{kurtosis} = n \sum_i (v_i - v_{\text{mean}})^4 / (\sum_i (v_i - v_{\text{mean}})^2)^2 - 3,$$

where  $n$  is the number of sample points in the trace,  $v_i$  is the voltage measured at the  $i$ th sample point, and  $v_{\text{mean}}$  is the mean value of the

voltage across the trace. Before computing the kurtosis of the long-tone data, we excised the initial 250 ms from each trace to ensure that the kurtosis reflected the steady-state behavior of the neuron during prolonged tone presentation rather than the transient stimulus onset response; trials concluded at tone termination so that OFF responses were also excluded from the long-tone protocol.

Within each stimulation protocol, many factors contributed to the variability in the measured kurtosis from one trace to the next. Unsurprisingly, the predominant source of this variability was whether or not a large, spontaneous bump appeared during that 4 s period. Accordingly, kurtosis rose with increasing trace length, which increased the probability of the occurrence of a large event on any given trace (short tones, kurtosis of  $23.6 \pm 5.8$ ,  $n = 17$  neurons, 8 s traces; silent epochs, kurtosis of  $25.7 \pm 4.7$ ,  $n = 18$  neurons, 9 s traces; all quantities are mean  $\pm$  SE unless otherwise specified). Note that longer records were more susceptible to slow-change electrode drift, slow movement artifacts attributable to pulmonary and cardiac pulsations, recording instabilities, etc., which typically have the effect of decreasing the kurtosis, because they smear out the strong peak in the mean of the voltage histogram near the resting potential of the cell.

#### Estimation of the presynaptic firing rate

We will make an order-of-magnitude estimate of the number of presynaptic neurons responsible for a typical synchronous volley (see Fig. 7). This procedure will involve four steps. First, based on previous voltage-clamp measurements from a similar population of neurons (Wehr and Zador, 2003, their Fig. 1*d*), we will model the relationship between the recorded membrane potential and the waveforms of the (unobserved) underlying excitatory and inhibitory conductances. To a first approximation, conductance and voltage follow very similar time courses, with a conversion factor of  $\sim 1$  nS/mV for both the excitatory and inhibitory contributions; a minor refinement is to delay the onset of the inhibitory conductance by  $\sim 2.5$  ms and slightly hasten the decay of each curve relative to the voltage trace, but these subtleties will not affect our order-of-magnitude estimates. Thus, accepting this relationship by fiat, we can estimate the hidden time course of both excitatory and inhibitory conductances based on our present measurements of membrane potential.

Second, we must model the canonical waveform for each unitary synaptic event. In one study (Stevens and Zador, 1998), spontaneous miniature EPSCs (mEPSCs) recorded in cortical neurons had a mean value of 6.4 pA, consistent with other *in vitro* (Gil et al., 1999) and *in vivo* (DeWeese and Zador, 2004) studies. Accordingly, we will assume mEPSCs to be 6 pA tall, with a 3 ms exponential decay. Note that recording spontaneous mEPSCs in tetrodotoxin (TTX) tends to undercount the number of very small mEPSCs, which fall below the detection threshold and into the noise, suggesting that the actual number of active synapses might be greater than our estimate.

Third, we need a way to relate our conductance waveforms to our canonical mEPSC. For a holding potential of  $-60$  mV, the mean size of an mEPSC is close to 6 pA (Stevens and Zador, 1998; Gil et al., 1999), corresponding to a conductance change of approximately  $(6 \text{ pA}) / (60 \text{ mV}) = 0.1$  nS:

$$\begin{aligned} I &= (V - E) g_{\text{syn}} \\ &= (-60 - 0) g_{\text{syn}} \\ &= -6 \text{ pA}, \end{aligned}$$

where  $I$  is the value of the current at the peak of the mEPSC. We now use this conversion factor,  $(0.1 \text{ nS}) / (6 \text{ pA})$ , to rescale our canonical mEPSC waveform from units of current to units of conductance.

Finally, we compute the number of unitary events required at every time step to sum to our estimate for the full excitatory conductance time course. To convert this to a firing rate of presynaptic neurons, we need to choose a value for the probability of release,  $p$ , which is the probability that a synapse will release neurotransmitter given that an action potential has invaded the nerve terminal. We chose  $p = 1$  for Figure 7*b* and all numerical estimates in the text, which is likely to cause us to underestimate the number of active presynaptic fibers given the much lower values

for  $p$  that are often reported at neocortical (Castro-Alamancos and Connors, 1997) and hippocampal (Dobrunz and Stevens, 1997; Huang and Stevens 1997; Murthy et al., 1997) (see also Castro-Alamancos and Connors, 1996) synapses. Release probability is different at different synapses, with most synapses showing very low probability, particularly during periods of high firing that elicit synaptic depression. In addition, our estimate for the size of the mean mEPSC is probably an overestimate attributable to the difficulty in observing the smallest events, which would also lead to an underestimate of the number of neurons participating in a synchronous volley. It should be noted that our methods cannot distinguish whether input synchrony arises from correlated firing across distinct presynaptic neurons or across synapses driven by a single presynaptic neuron making many contacts, as at the neuromuscular junction. However, connections in cortex are often quite weak (Atzori et al., 2001), although in layer 5 especially, strong connections are not uncommon (Markram et al., 1997; Song et al., 2005).

#### Stimuli

Stimulus delivery followed essentially the same protocol as described previously (DeWeese et al., 2003; DeWeese and Zador, 2004), with the exception that here we presented tones of two different durations; we also analyzed silent epochs in the present study. The stimuli used for the short-tone protocol consisted of 25 ms pure-tone pips of 32 different frequencies (logarithmically spaced between 2 kHz and 46731 Hz) with 5 ms 0–100% cosine-squared ramps applied to the onset and termination of each pip. All 32 tones were repeatedly presented at 65 dB in a fixed pseudorandom order at a rate of 2 tones/s. The long-tone protocol consisted of 4-s-duration pure tones of 56 different frequencies (logarithmically spaced between 1 kHz and 45255 Hz) with 5 ms 10–90% cosine-squared ramps applied to tone onsets and terminations. All 56 tones were presented in a pseudorandom order with silent gaps of either 1 or 13 s duration between them; 4-s-duration periods that occurred at least 1 s after the termination of a tone were used for the “silent” kurtosis analysis. All experiments were conducted in a double-walled sound booth (Industrial Acoustics Company, Bronx, NY). Free-field stimuli were presented using a System II (Tucker-Davis Technologies, Gainesville, FL) running on a host personal computer connected to an amplifier (Stax SRM 313), which drove a calibrated electrostatic speaker (Stax SR303). The speaker was placed 6 cm to the right of, and at the same elevation as, the rat's head. The head was rotated to the right by an angle between 60 and 90° about the rostrocaudal axis, so that the plane of the craniotomy was nearly horizontal. For the unanesthetized recording shown in Figure 2*c*, the rat's head was held upright, so that the animal's right ear directly faced the speaker, which was placed 6 cm away from the head; for this recording, the stimuli followed the short-tone protocol, except that the tones were 100 ms in duration.

#### Surgery

Sprague Dawley rats (17–25 d old) were anesthetized in strict accordance with the National Institutes of Health guidelines as approved by the Cold Spring Harbor Laboratory Animal Care and Use Committee. Pentobarbital (65 mg/kg) was used for the six neurons (from five animals) recorded after TTX application (see below, TTX application) and for the 17 neurons (from nine animals) recorded during the short-tone protocol. Diazepam (5 mg/kg) was also used in three of the later cases, but it did not make a statistically significant difference in the measured kurtosis (pentobarbital with diazepam, kurtosis of  $12.1 \pm 10.4$ ,  $n = 3$  neurons; pentobarbital without diazepam,  $26.1 \pm 6.6$ ,  $n = 14$  neurons). We therefore pooled these data for the group statistics. For the 18 neurons (from five animals) recorded during silence and the long-tone protocol, recordings were performed under ketamine (60 mg/kg) and medetomidine (0.50 mg/kg). We found no significant difference between the kurtosis measured during the three protocols not involving TTX (pentobarbital/short tones, kurtosis of  $15.2 \pm 3.5$ ,  $n = 17$  neurons; ketamine/long tones, kurtosis of  $9.5 \pm 2.2$ ,  $n = 18$  neurons; ketamine/silence, kurtosis of  $14.4 \pm 3.0$ ,  $n = 18$  neurons).

After the animal was deeply anesthetized, it was placed in a custom naso-orbital restraint that left the ears free and clear. Local anesthetic was applied to the scalp, and a  $1 \times 2$  mm craniotomy and duratomy were

performed above the left auditory cortex. A cisternal drain was performed before the craniotomy. Before the introduction of electrodes, the cortex was covered with physiological buffer (in mM: 127 NaCl, 25 Na<sub>2</sub>CO<sub>3</sub>, 1.25 NaH<sub>2</sub>PO<sub>4</sub>, 2.5 KCl, 1 MgCl<sub>2</sub>, and 25 glucose) mixed with 1.5% agar. Temperature was monitored rectally and maintained at 37°C using a feedback-controlled blanket (Harvard Apparatus, Holliston, MA). Breathing and response to noxious stimuli were monitored throughout the experiment, and supplemental dosages of pentobarbital or ketamine were provided when required.

For the cell shown in Figure 2c, which was recorded in the awake condition, a craniotomy and duration were performed under ketamine and medetomidine as described above, except that no stereotaxic frame was used. The area surrounding the craniotomy was then protected with a plastic well with removable screw cap, and the cortical surface was covered with Kwik-Cast (World Precision Instruments, Sarasota, FL) between recording sessions. An aluminum head post was attached to the skull with Relyx Luting Cement (3M ESPE, St. Paul, MN). A silver chloride ground wire was implanted subcutaneously on the back of the animal. The rat was allowed at least 24 h of recovery time before the first recording session. During the recording session, the head post was fixed in the head-post holder and the animal was standing inside a hard plastic tube, which provided a loose restraint for body movements. The plastic cap and Kwik-Cast were removed, and the cortex was covered with physiological buffer as for the anesthetized recordings. Recording sessions lasted for a maximum of 4 h. During the recording sessions, the animal stood fairly motionless some of the time and occasionally moved its limbs, whisked, groomed, etc.

#### Whole-cell patch-clamp recording

We used standard blind whole-cell patch-clamp recording techniques, modified from brain slice recordings (Stevens and Zador, 1998). Membrane potential was sampled at either 4 or 10 kHz in current-clamp mode ( $I = 0$ ) using an Axopatch 200B amplifier (Molecular Devices, Palo Alto, CA) with no on-line series resistance compensation. Data were acquired using either an Igor (WaveMetrics, Lake Oswego, OR)-based system written by Dr. Bernardo Sabatini (Harvard Medical School, Boston, MA) or a Matlab (MathWorks, Natick, MA)-based system, controlling a National Instruments (Austin, TX) card on a Dell personal computer (Dell Computer Company, Round Rock, TX). For some whole-cell recordings, recording pipettes were filled with an internal solution consisting of the following: 10 mM KCl, 140 mM K-gluconate, 10 mM HEPES, 2 mM MgCl<sub>2</sub>, 0.05 mM CaCl<sub>2</sub>, 4 mM Mg-ATP, 0.4 mM Na<sub>2</sub>-GTP, 10 mM Na<sub>2</sub>-phosphocreatine, 10 mM BAPTA, and 1% biocytin, pH 7.25, diluted to 290 mOsm. The remaining recordings had an additional 1% biocytin but did not include the 10 mM KCl. The fast sodium channel blocker QX-314 [5 mM; *N*-(2,6-dimethylphenyl)-3,3-diethylpentamide], which also blocks some other activity-evoked conductances, was added to the intracellular solution to block sodium channels and therefore prevent spiking. Consequently, spiking was rare in these neurons; when a spike did occur, the 7 ms waveform beginning 1 ms before the onset of the spike was removed from the full record before the kurtosis was calculated. QX-314 was not used for the recordings after TTX application, nor was it used for the unanesthetized recording shown in Figure 2c. Electrodes were pulled from filamented, thin-walled borosilicate glass (1.5 mm outer diameter, 1.17 mm inner diameter; World Precision Instruments) on a vertical Narishige (Tokyo, Japan) two-stage puller. Resistance to bath was 3–5 M $\Omega$  before seal formation. Recordings were performed throughout auditory cortex at electrode depths ranging from 140 to 834  $\mu$ m below the cortical surface.

Of the 43 neurons recorded for the kurtosis analysis, 35 (from 14 animals) passed our criteria for inclusion: recordings had to be very stable for at least one 8 s trial; trials affected by electrode drift or appreciable motion artifacts resulting from pulmonary or cardiac pulsations were not included in the analysis; and the resting potential had to be at or below  $-55$  mV, corrected for the liquid junction potential, which we calculated to be 12 mV for our internal solution. Of the 21 neurons recorded after TTX application, six recordings (from five animals) had sufficiently low recording noise to allow unambiguous detection of mEPSPs.

#### TTX application

For the recordings performed after the pharmacological removal of presynaptic input correlations, a standing pool of 1.0 mM TTX in a physiological buffer (see above, Surgery) was applied to the surface of the cortex while playing 65 dB stimuli (short-tone stimulation protocol; see above, Stimuli) and monitoring the local field potential (LFP) recorded with a patch electrode  $\sim 700$   $\mu$ m below the cortical surface. Whole-cell recordings were only attempted after complete abolition of all evoked and spontaneous LFP responses. As is commonly done when measuring miniature PSPs *in vitro*, we used TTX to prevent spiking in any of the presynaptic fibers of the neuron. Subsequently, the only PSPs we observed were attributable to the stochastic, spontaneous release of individual synaptic vesicles. Spontaneous release events are statistically independent; thus, if the collective spontaneous event rate is high enough across the population of synapses, fluctuations in the membrane potential attributable to the simultaneous release of multiple vesicles results in a random walk, with a Gaussian distributed histogram of membrane potential values.

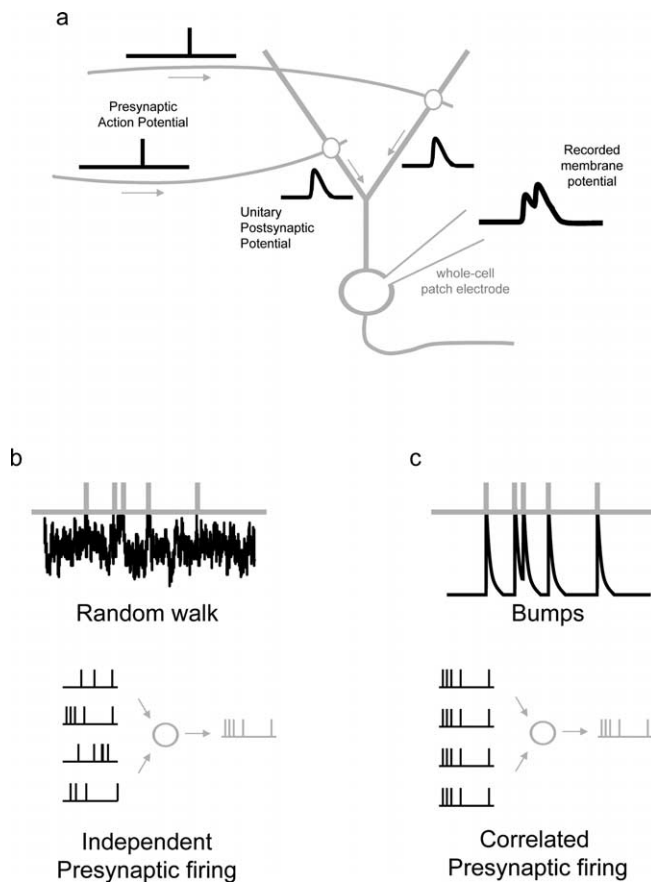
## Results

### Random walk models

The inputs that provide the synaptic drive to any given cortical neuron arise mainly from the spiking outputs of other cortical neurons. In theoretical work, it is often assumed that the statistical properties of these cortical inputs are identical. The appeal of this simple assumption is that it can give rise to analytically tractable models that make quantitative and testable predictions. Specifically, a broad class of such models predict that many small postsynaptic potentials summate to give rise to membrane trajectories that follow a random walk (Tuckwell, 1988; Softky and Koch, 1993).

To make these ideas concrete, consider a simple model of a neuron integrating synaptic inputs from a large population of other neurons. In this model, responses from individual presynaptic inputs are summed together to produce the membrane potential at that moment, as illustrated in Figure 1a for a pair of inputs. How might this look when many inputs are simultaneously active? One can imagine two extreme limiting cases. In the first case, the activity of the presynaptic subpopulation is completely uncorrelated (i.e., Poisson) both in time and across neurons; the timing of synaptic events in this case would resemble the ticks of a Geiger counter. This situation would result in a highly variable (Gaussian) membrane potential that undergoes a random walk (Fig. 1b) (Tuckwell, 1988), provided that the synaptic event rate is sufficiently high for the conditions of the central limit theorem to be met, as would likely be the case if the random fluctuations were sufficiently large to cause spiking. This model and close variants have been widely used to describe the inputs to cortical neurons in both theoretical (Gerstein and Mandelbrot, 1964; Calvin and Stevens, 1967; Softky and Koch, 1993; Shadlen and Newsome, 1994, 1995; Tsodyks 1995; van Vreeswijk and Sompolinsky, 1998; Song et al., 2000; Fellous et al., 2003; Rudolph and Destexhe, 2003) and experimental (Destexhe et al., 2003; Carandini, 2004) studies, but it has yet to be experimentally tested in auditory cortex.

At the other extreme is the possibility that the presynaptic population is highly correlated: neurons might be silent most of the time, except during brief moments when large groups of them fire in a concerted manner, as schematized in Figure 1c. In this case, the postsynaptic membrane potential would sit at rest except for those brief moments of synchronous presynaptic activity that would elicit large excursions (or bumps) of the membrane potential of the postsynaptic neuron. Of course, it may be that auditory cortex operates in an intermediate regime that does not fit neatly into either of these categories.



**Figure 1.** The dynamics of the membrane potential of a given cortical neuron *in vivo* provides a way of inferring activity among the network of neurons presynaptic to that neuron. **a**, In this simplified example, action potentials arriving at two synapses on the dendritic arbor of the recorded neuron each result in synaptic transmission, which in turn evokes a pair of unitary postsynaptic potentials. To a first approximation, the membrane potential of the cell is the sum of these two events, along with any other events that may have occurred. If the two synaptic events are nearly simultaneous, they can add and be seen as a single large event, whereas if they occurred at different times, they would be seen as two separate events. More generally, the degree to which spiking activity of the network of presynaptic afferents is correlated may be reflected in the dynamics of the recorded membrane potential. **b**, **c**, The same spike train can result from very dissimilar membrane potential dynamics. In **b**, the timing of each of the five spikes (top, gray) is determined by random, threshold-crossing fluctuations in the membrane potential (top, black) as it follows a highly stochastic, random walk that hovers just below the spike threshold, as one might expect if the synaptic inputs to the neuron are statistically independent of one another (bottom). In **c**, each spike from an identical spike train as in **b** results from tall bumps in the membrane potential, which would result if the synaptic inputs to the neuron were highly correlated in their activity (bottom).

Previous studies have compared the spike trains predicted from these models with the spike trains observed in the cortex, particularly in the visual cortex (Softky and Koch, 1993; Shadlen and Newsome, 1995, 1998; Troyer and Miller 1997). In the simplest random walk model (see Materials and Methods, Excitation only model), all inputs are assumed to be excitatory neurons. However, the spiking statistics predicted by the excitation only model are inconsistent with the observed statistics in middle temporal area MT and in other brain areas (Softky and Koch, 1993; Shadlen and Newsome, 1998). This inconsistency led to variant models in which inputs consist of both excitatory and inhibitory synapses (Shadlen and Newsome, 1994, 1995; Bell et al., 1995; Troyer and Miller, 1997) (see Materials and Methods, Excitation and inhibition model); the addition of inhibitory inputs permitted the model to fit the observed low-order spike statistics. In one

common variant, excitatory and inhibitory inputs are “balanced,” so that the membrane potential hovers just below threshold (van Vreeswijk and Sompolinsky 1996; Chance et al., 2002; Hertz et al., 2003).

Spiking data provide only an indirect means to assess such models, because, by definition, the spike-generating mechanism discards all of the information about the subthreshold membrane potential preceding the spike. The fact that a particular simple random walk model is compatible with an observed spike train is suggestive but not conclusive; many subthreshold trajectories are compatible with the same observed spike train (Fig. 1*b,c*). Distinguishing among different models thus requires other methods (Azouz and Gray, 1999; Carandini, 2004)

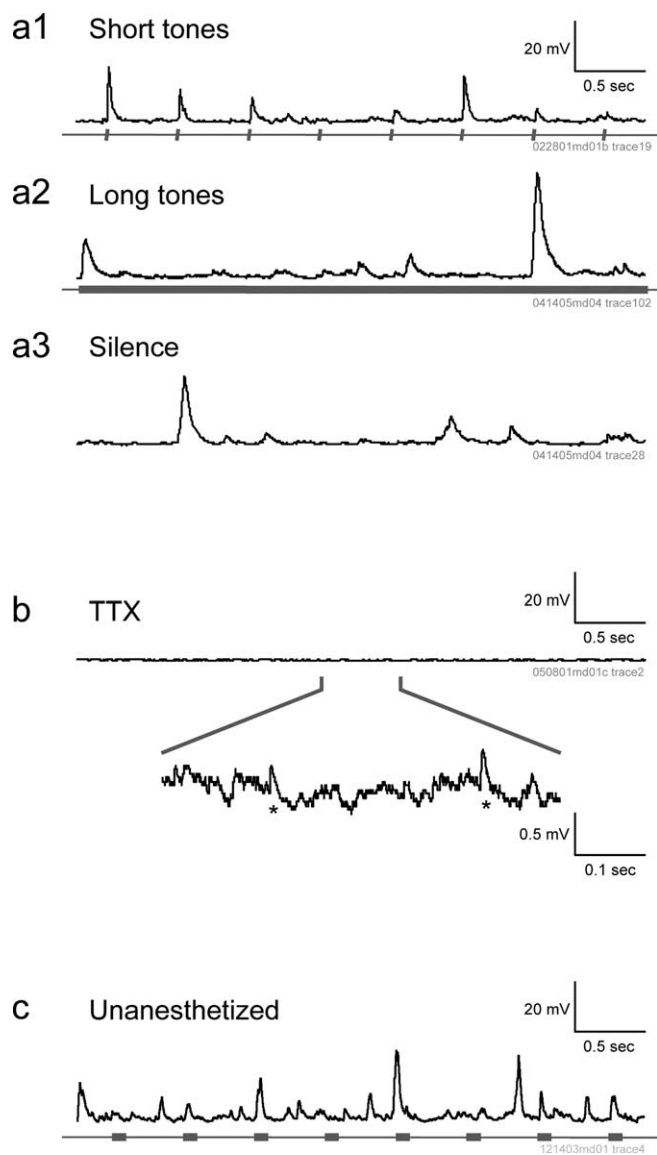
### Subthreshold recordings in auditory cortex

To gain insight into the network activity driving individual cortical neurons, we used *in vivo* whole-cell patch-clamp methods to record the membrane potential from neurons in the rat auditory cortex. Because we were interested in the synaptic inputs to the neuron under study rather than its spiking output, we blocked action potentials in the target neuron, but not other neurons in the circuit, by adding the sodium channel blocker QX-314 to the recording pipette. We considered both spontaneous and sound-evoked activity. To study the activity of the network on a single trial rather than just its mean activity, we analyzed traces from individual trials rather than averaging across multiple stimulus presentations.

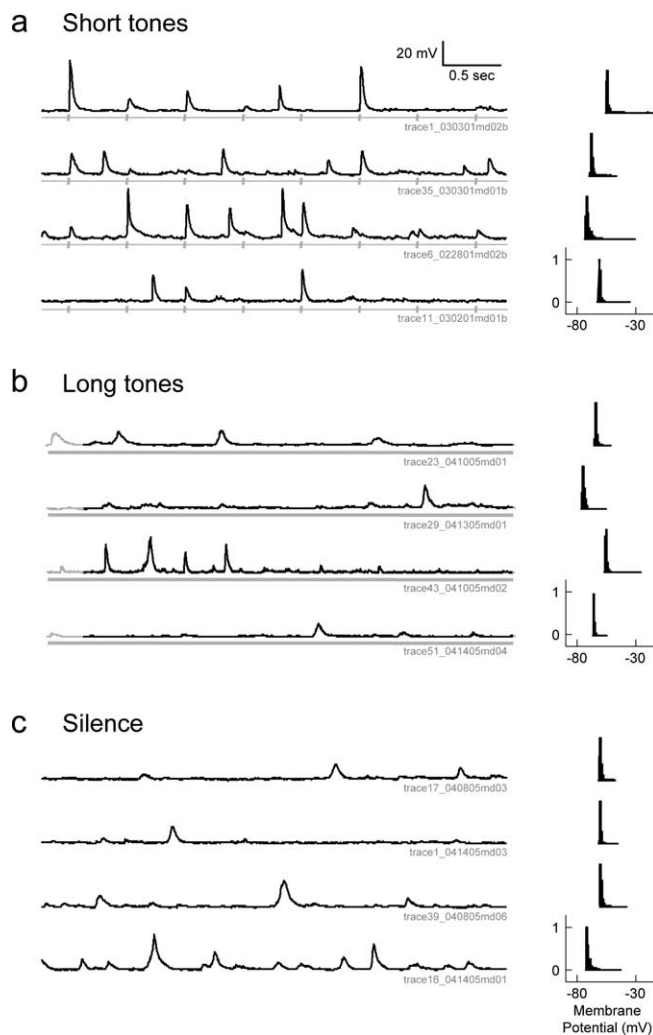
We began by recording responses to a series of short tones, consisting of 25-ms-duration pure tones presented twice per second (see Materials and Methods, Stimuli). As expected from previous results (Wehr and Zador, 2003; Zhang et al., 2003; DeWeese and Zador, 2004; Tan et al., 2004; Las et al., 2005), tone-evoked responses were brief and often occurred with a short latency immediately after the tone (Figs. 2*a1*, 3*a*). The response was well characterized as essentially flat, except for infrequent, but substantial, bumps. The overall response was consistent with a picture in which the neurons were typically silent except for specific moments when many inputs became active at once. We presumed that these moments of activity represented mainly some combination of thalamocortical and intracortical inputs.

We wondered whether the structure we observed (occasional bumps superimposed on a quiet background) was merely the result of the pulsatile stimulus protocol we had tested, which consisted of brief infrequent tone pips. We therefore tested two other stimulus protocols (see Materials and Methods, Stimuli): long tones (Figs. 2*a2*, 3*b*), consisting of 4-s-duration pure tones, and silence (Figs. 2*a3*, 3*c*). The resulting traces looked qualitatively similar to the responses elicited by brief tone pips. Specifically, they often contained well isolated bumps that rose slightly less steeply than the short-tone-evoked bumps but which were comparable in height. Thus, bumps were not solely attributable to stimulus transients but rather appeared to represent some manifestation of network dynamics.

We interpret these bumps as arising mainly from a concerted barrage of presynaptic activity rather than from postsynaptic nonlinearities such as dendritic sodium or calcium spikes, for at least three reasons. First, in many neurons, previous work has demonstrated that the synaptic current–voltage relationship is linear, and the inferred synaptic conductance can be used to predict the observed membrane voltage (Wehr and Zador, 2003). Second, the presence of QX-314 in the internal solution blocked not only the fast sodium channels but possibly other voltage-sensitive channels as well (Talbot and Sayer, 1996; Deisz et al.,



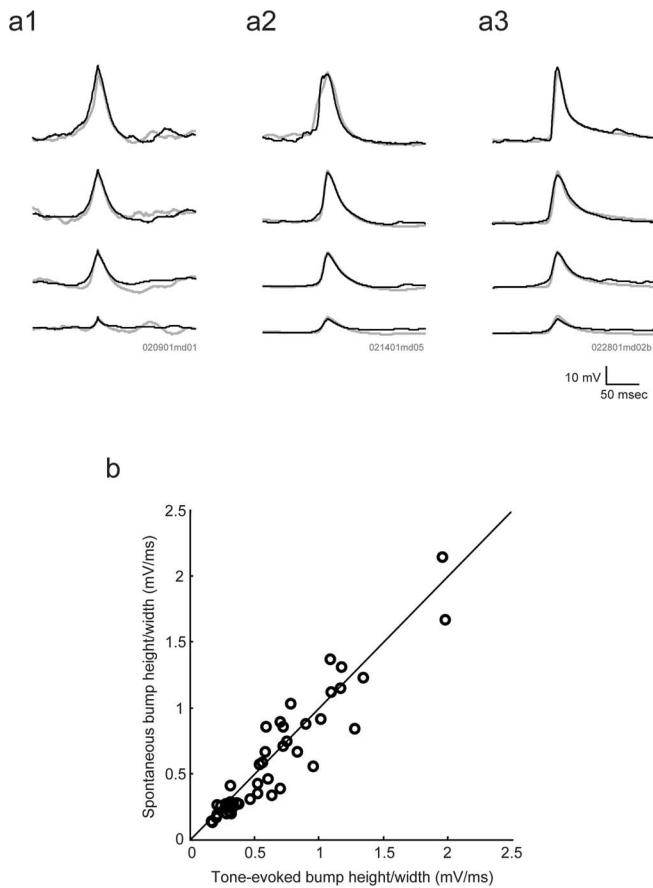
**Figure 2.** Membrane potential dynamics in auditory cortex do not resemble a random walk. **a1**, This 4 s example of a whole-cell patch-clamp recording from an auditory cortical neuron *in vivo* clearly exhibits the bumpy appearance ubiquitous in our dataset; QX-314, an intracellular fast sodium-channel blocker, was included in the patch pipette to prevent spiking as well as some other nonlinearities that can distort the relationship between synaptic activity and membrane potential. Throughout the trace, steeply rising peaks in the postsynaptic potentials follow most of the 65 dB, 25 ms tone pips (gray hash marks below the voltage trace; for stimulus protocol, see Materials and Methods), consistent with the occurrence of synchronous volleys of synaptic input. Aside from these narrow bumps, the membrane voltage remained close to the resting potential of the neuron. **a2**, These large excursions from rest were not restricted to stimulus transients. The stimulus here consisted of a 4-s-duration tone (gray bar beneath voltage trace), which began 15 ms following the far left of the trace. **a3**, Even in the absence of any auditory stimulation, the membrane potential displayed the same bumpy appearance, even long after the onset of the tone. **b**, As a control, we repeated these experiments after the topical application of TTX, a fast sodium channel blocker, to the cortical surface so as to abolish all presynaptic spiking and thus ensure the independence of synaptic events (see Materials and Methods, TTX application). As the expanded view of the trace demonstrates (bottom), the membrane potential does resemble a random walk in the absence of input correlations. We indicated two putative mEPSPs with asterisks. **c**, An example trace from an unanesthetized, head-restrained rat shows the same bumpy appearance as the records from anesthetized animals. The stimulus consisted of 100-ms-duration pure tones of 65 dB presented every 500 ms (gray hash marks below the voltage trace). Because QX-314 was not included in the patch pipette for this recording, the neuron sometimes fired action potentials. To allow comparison with the previous figures, we therefore median filtered [filter duration of 3 ms (Jagadeesh et al., 1997)] the trace to remove three spikes.



**Figure 3.** Throughout the neuronal population, membrane potential time courses looked “bumpy” under all stimulus conditions tested. Four-second-duration whole-cell records obtained during the presentation of 25-ms-duration tones (**a**, gray hash marks below traces), 4-s-duration tones (**b**; gray bar below traces), and silence (**c**) all display well isolated bumps superimposed on otherwise flat traces sitting at the resting potential of each neuron. For each record, the frequency histogram of membrane potential appears to the right of the corresponding trace; histograms are normalized to unit height, and the membrane potential values are uncorrected for the junction potential. Note the long tail and sharp peak of every distribution. Each of the 12 traces was recorded from a different neuron.

1997), presumably reducing the contribution of postsynaptic nonlinearities, but bumps were observed whether or not QX-314 was present. Finally, the bumps recorded intracellularly often co-occur with large network events, as confirmed previously by simultaneously recording the local field potential far ( $\sim 0.5$  mm) from the whole-cell recording site (DeWeese and Zador, 2004).

For all three protocols, the duration of bumps was consistently brief. We quantified this by thresholding the membrane potential at 10 mV above resting potential and measuring the width at half the peak value for every bump that crossed threshold. The longest bump duration for every neuron was typically on the order of tens of milliseconds (maximum width of bumps,  $54.6 \pm 3.3$  ms;  $n = 30$  neurons; all quantities are mean  $\pm$  SE unless otherwise specified) (see Materials and Methods, Quantifying bump duration). Thus, these auditory cortex neurons did not show the bistable “two-state” behavior reported in other areas (Steriade et al., 1994; Wilson and Kawaguchi, 1996; Anderson



**Figure 4.** Stimulus-evoked postsynaptic potentials (or bumps) were very similar in shape to bumps that occurred spontaneously. **a1**, For this neuron, each trace corresponds to the membrane potential averaged across either stimulus-evoked (gray traces) or spontaneous (black traces) bumps with peak heights falling between 5 and 10 mV (bottom pair of traces), 10 and 20 mV (second from bottom), 20 and 30 mV (third from bottom), or 30 and 40 mV (top pair of traces). The gray and black traces are nearly identical for each of the four pairs. Bumps were identified as every excursion of the membrane potential exceeding a 5 mV threshold above the resting potential of a neuron. Bumps with peaks occurring between 10 and 100 ms after stimulus onset were classified as stimulus evoked, and all others were classified as spontaneous; this neuron was recorded during the short-tone protocol, which consisted of 25-ms-duration tone pips. Bumps were aligned horizontally based on the times of their peaks. **a2**, **a3**, Same format as **a1** for two other neurons. **b**, For each neuron in the population, stimulus-evoked bumps were similar in shape to spontaneous bumps, as indicated by the high correlation between the ratio of height to width (full-width at half-maximum) of spontaneous versus tone-evoked bumps (correlation coefficient of 0.94), and the fact that all points lie close to the diagonal line indicating equality. Each point in the scatter plot corresponds to 1 of 17 neurons and 1 of the 4 peak height categories defined in **a1** [ $n = (17 \text{ neurons})(1\text{--}4 \text{ peak height categories}) = 44 \text{ points}$ ]. For any given neuron, only those peak height categories containing at least one spontaneous and at least one tone-evoked bump were included in the analysis.

et al., 2000; Sanchez-Vives and McCormick, 2000; Cossart et al., 2003; Petersen et al., 2003).

In addition to being brief, the shapes of stimulus-evoked bumps were surprisingly similar to those of spontaneous bumps in the same neuron. Mean tone-evoked and spontaneous bumps were nearly identical in several neurons (Fig. 4a). Across the population, the shapes of evoked and spontaneous bumps were highly correlated within neurons (Fig. 4b).

As a control for the possibility that the apparent bumpiness of these records was attributable to some artifact of our recording methodology rather than a reflection of network dynamics, we repeated these experiments under conditions in which synaptic events were guaranteed to be uncorrelated. Specifically, we re-

corded the membrane potential after topical application of TTX to the cortical surface (Fig. 2b) (see Materials and Methods, TTX application). TTX is an extracellular fast sodium channel blocker that prevents spike propagation in axons, thus ensuring that each unitary synaptic input to the neuron was attributable to the spontaneous release of a synaptic vesicle, independent of the activity of all the other synapses. (Note that spontaneous synaptic release, unlike spontaneous spiking activity, is independent of action potentials.) In contrast to the first three examples, the time course of this trace (Fig. 2b, expanded view) closely resembles a random walk, as one would expect once correlations among different synaptic inputs have been pharmacologically removed.

Most of our recordings were obtained from anesthetized animals. However, in a few cases ( $n = 5$  neurons), we obtained whole-cell recordings from the auditory cortex of head-fixed unanesthetized animals. An example of a response from an unanesthetized animal (Fig. 2c) shows bumps that are qualitatively similar to those observed in the anesthetized preparation. Because intracellular recording in unanesthetized animals is more technically challenging (Wilson and Groves, 1981; Covey et al., 1996; Fee, 2000; Steriade et al., 2001; Margrie et al., 2002), and because sound-evoked responses in auditory cortex show more heterogeneity in the unanesthetized preparation (Evans and Whitfield, 1964), these recordings must be considered anecdotal; although the majority (four of five) of neurons showed bumps, we cannot yet make any definitive statement about the prevalence of such responses in the unanesthetized animal. Although in this small sample we did not encounter neurons that responded with high firing rates to the pure-tone stimuli we used, such responses do exist in the unanesthetized preparation (Evans and Whitfield, 1964; Wang et al., 2005) and may show different subthreshold membrane dynamics. Nevertheless, we conclude from these recordings that the bumpiness of the membrane potential we observed is not solely an anesthesia artifact.

### Up and down states

Interestingly, subthreshold dynamics seem to be qualitatively less bumpy, and more like a random walk, in some studies of non-auditory cortical areas (Ferster and Jagadeesh, 1992; Destexhe et al., 2003; Carandini, 2004). In many cases, subthreshold dynamics have been characterized as two state, in which the membrane potential toggles back and forth between a “down” state and an “up” state. During the down state, the neuron sits quietly at its rest potential, whereas in the up, state it is depolarized, with fluctuations resembling the random walk behavior depicted in Figure 1b. Each state can last for periods as long as several seconds. These two-state dynamics have been observed in a wide range of *in vivo* preparations, including the visual cortex (Anderson et al., 2000), other cortical areas (Steriade et al., 1994; Destexhe and Pare, 1999; Petersen et al., 2003; Leger et al., 2005), the neostriatum (Wilson and Groves, 1981; Wilson and Kawaguchi, 1996), in some (Petersen et al., 2003) but not other (Wilent and Contreras, 2005; Bruno and Sakmann, 2006) barrel cortex recordings, as well as *in vitro* (Sanchez-Vives and McCormick, 2000; Cossart et al., 2003).

It is unclear why we did not observe up states in the auditory cortex *in vivo*. One possibility is that the issue is merely one of nomenclature and that the synchronized volleys of activity we observed in auditory cortex represent extremely brief up states. In this view, the duration of up states can range over two orders of magnitude, from the tens of milliseconds we typically observed (maximum bump duration,  $54.6 \pm 3.3$  ms) to seconds, but arise from similar network mechanisms. Indeed, it has been suggested

(Benucci et al., 2004) that the existence of pairwise correlations among neurons can facilitate up-state stability, despite the random walk appearance of the membrane potential during prolonged up states. However, the network dynamics responsible for bumps and for up states have not yet been fully elucidated, so it remains an open question whether or not their underlying mechanisms are the same.

### Comparison between observed subthreshold fluctuations and random walk models

The large excursions in membrane potential seen in the data are grossly incompatible with the simplest excitation only or E&I random walk models. To illustrate this, we computed the best fit of each model to a representative trace shown in Figure 2*a1* (see Materials and Methods, Random walk model). For the excitation only model, there was only one free parameter, the rate  $\lambda_e$  at which excitatory presynaptic events arrived; this parameter was used to fit the mean membrane potential. For the E&I model, there were two free parameters  $\lambda_e$  and  $\lambda_i$ , corresponding the arrival rate of excitatory and inhibitory events, respectively; these were used to fit the mean and variance of the membrane potential. The excitation model is thus a special case of the E&I model. Other model parameters, such as the resting potential of the neuron, were measured directly from the recording.


Simulated sample traces generated from these models bear little resemblance to the sample trace to which they were fit (Fig. 5*a*). For the excitation only model (Fig. 5*b*), fluctuations were small and steady, reminiscent of (but larger than) the minute fluctuations arising from spontaneous synaptic vesicle release seen in the TTX control experiment, in which all spiking input was blocked (Fig. 2*b*). For the E&I model (Fig. 5*c*), fluctuations were larger, but, as expected, the fluctuations are steady and do not resemble the large, well isolated bumps so prominent in the data. These simulated traces illustrate that an uncorrelated steady mixture of excitation and inhibition is unable to account for the observed fluctuations in membrane potential, consistent with previous *in vitro* results (Stevens and Zador, 1998).

### Time-varying random walk models

Why did the E&I model fail? As noted above, even casual inspection of the sample traces in Figure 2*a* reveals large, infrequent bumps in membrane potential. Because these bumps occur only rarely, any model must fail to account for these data if it predicts that membrane potential fluctuations arise from a constant rain of excitatory and inhibitory inputs. The problem, then, rests in the assumption of time invariance, or stationarity.

We can attempt to salvage the random walk models by generalizing them, at the cost of extra parameters, to include time-varying inputs. We therefore tested a model in which the excitatory and inhibitory event rates varied with time, on a timescale  $\tau$ , thus inducing correlations among the different synaptic inputs (see Materials and Methods, Rate modulated excitation and inhibition model). To fit this model, the mean and variance at each point in time were fit to a window of length  $\tau$  centered at that time, to produce time-varying population rates  $\lambda_e(t)$  and  $\lambda_i(t)$ , respectively.

For long windows ( $\tau = 200$  ms) corresponding to a relatively slowly varying input, the model fit remained poor (Fig. 5*d*). Not surprisingly, the fit improved as the window length decreased; decreasing the window length corresponds to increasing the number of model degrees of freedom. For sufficiently fast modulations ( $\tau = 10$  ms) (Fig. 5*e*) of the excitatory and inhibitory event rates, the model was quite good; however, despite the effective



**Figure 5.** The random walk model does not capture the key features of the data unless we introduce correlations among the presynaptic inputs. **b** shows an example trace generated by the excitation only model (see text and Materials and Methods, Random walk model) that takes only excitatory input that is fit to the mean of the data trace shown in **a** (replotted from Fig. 2*a1*). The model trace looks nothing like the original. **c**, By including inhibitory inputs to the model (E&I model), we can fit both the mean and the variance of the data trace, but the fluctuations in the model do not have the same bumpy character as the data. **d**, Allowing the model parameters to slowly vary on a 200 ms timescale (Rate Modulated E&I model) still does not capture the structure evident in the data. **e**, Fitting the mean and variance of the model to the data trace on a fast, 10 ms timescale greatly improves the performance of the model, but it still makes errors during abrupt changes in the membrane potential, such as the downward swings that often occur at the onsets of steeply rising bumps in the data (expanded view, bottom), unless we impose additional correlations between the inhibitory and excitatory inputs (data not shown).

addition of a new pair of free parameters every 10 ms, the model still made qualitative errors, especially near abrupt changes in membrane potential such as occur at the onsets of bumps at which the model often made large negative swings (Fig. 5*e*, expanded view). To fix this, we would need to impose additional correlations between the excitatory and inhibitory inputs so that excitation leads inhibition by a few milliseconds during sharp increases in synaptic activity, as has previously been shown experimentally for these neurons (Wehr and Zador, 2003).

We conclude that, although it is true that we can improve the fit to data with the addition of an ever increasing number of free parameters, the time-varying E&I model does not capture the essence of the measured membrane potential unless we fit it to the data on such a fine timescale that it no longer follows anything resembling a random walk trajectory. These extra parameters are manifestations of correlations among the population of presynaptic inputs feeding into the recorded neuron. Unlike a random walk, which crosses threshold at random times determined by the chance coincidence of many small uncorrelated events, the measured membrane potential is only great enough to reach threshold at specific moments when many correlated synaptic inputs are active at once. Thus, the time-varying random walk model,



although formally adequate, does not appear to be as useful a description as our initial characterization that activity consists of extended periods of silence punctuated by brief synchronous bouts of activity (bumps).

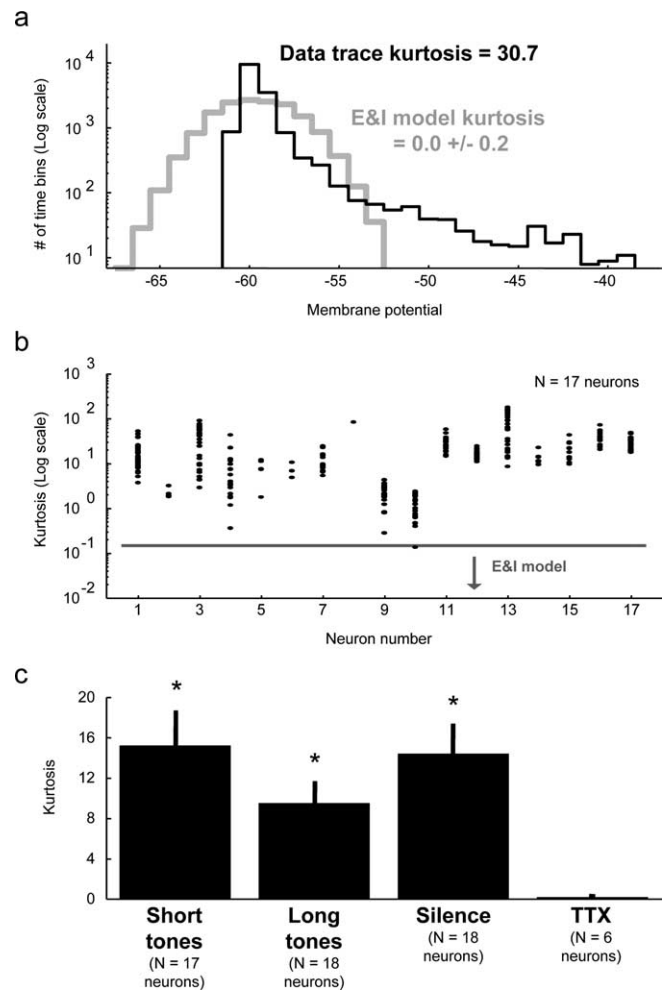
### Kurtosis of membrane potential distribution

The simple time-varying models we considered represent a qualitative improvement over time-invariant models. Their success stems from the correlations induced among the inputs by the time-varying population rate; these correlations give rise to the relatively infrequent periods of elevated firing required to drive the membrane potential to tens of millivolts above its resting level. However, these simple models represent only a small subset of the large class of models in which the input neurons have correlations. There are many possible ways in which the input neuron activity could be correlated. Just as the spiking output places constraints on (but does not fully specify) the underlying subthreshold membrane potential (Fig. 1*b,c*), so too do subthreshold membrane potential fluctuations constrain (without fully specifying) the network input to the neuron. The mean and the variance of the membrane potential, used above to fit the simple time-invariant models, represent two such constraints.

If the distribution of membrane potential were Gaussian, the mean and the variance would fully characterize the distribution. Because these two parameters fully specify the time-invariant models, they determine how good a fit the excitation only and E&I random walk models provide to the data. To relate the observed distribution with a Gaussian, we considered the entire distribution of membrane potentials from each full trace. Figure 6*a* compares the membrane potential histogram for the sample data trace and the E&I model from Figure 5*c*. When the synaptic event rate is high enough to match the experimentally observed membrane potential mean, the E&I model generates distributions of membrane potentials that are nearly Gaussian, as expected from the central limit theorem (Fig. 6*a*, gray line). In contrast, the experimentally observed distribution is far from Gaussian (Fig. 6*a*, black line; note logarithmic scale); the data have many more small values (near the resting potential of the neuron) and many more large values (attributable to the tall bumps) than expected from the Gaussian. The poor fit of the Gaussian to the observed distribution provides additional intuition about the failure of the random walk models.

This overrepresentation of outliers in the data are a hallmark of sparseness. Sparse signals have received much attention recently in the signal processing and natural scenes community (Simoncelli and Olshausen, 2001). A convenient way to quantify sparsity is the kurtosis, a function of the fourth central moment of a distribution (see Materials and Methods, Kurtosis). The kurtosis is a measure of the shape of a distribution compared with a Gaussian; it is greater for distributions with more outliers, or more values very close to the mean, making it well suited for quantifying the “bumpiness” of a trace.

The kurtosis provides a convenient way to characterize the sparseness of the data across the population (Fig. 6*b*). When applied to our data, we found that the kurtosis was high for all conditions (short tones,  $15.2 \pm 3.5$ ,  $n = 17$  neurons; long tones,  $9.5 \pm 2.2$ ,  $n = 18$  neurons; silence,  $14.4 \pm 3.0$ ,  $n = 18$  neurons); as expected, only the recordings performed after the application of TTX to block presynaptic spiking (included as a control) had a low kurtosis (kurtosis,  $0.2 \pm 0.3$ ;  $n = 6$  neurons). Thus, the high kurtosis observed experimentally is inconsistent with the most basic random walk models and consistent with a time-varying



**Figure 6.** The random walk model is inconsistent with membrane potential dynamics across the neural population. *a*, A histogram of membrane potential values (black line) taken from the example trace shown in Figures 2*a1* and 5*a* exhibits much more weight both at its peak (approximately  $-60$  mV) and in its tail (more than approximately  $-50$  mV) than the Gaussian-distributed histogram (thick gray line) corresponding to the example trace plotted in Figure 5*c* generated by the E&I model fit to the same mean and variance (note the logarithmic scale of the ordinate). We quantified the difference in histogram shape with the kurtosis (see Materials and Methods, Kurtosis), which is always 0 for the E&I model, and  $>0$  for otherwise flat traces containing tall, well isolated bumps. For this example, the data trace has a kurtosis of 30.7, whereas for traces drawn from the random walk model, kurtosis was  $0.0 \pm 0.2$  (mean  $\pm$  SD). *b*, Across the population of 17 neurons responding to the short-tone protocol, the kurtosis of individual traces (black points) were large and positive (note logarithmic scale) compared with the range of values corresponding to the random walk model; the gray line indicates 1 SD above 0, which is the mean value for the model. *c*, Across the population, the kurtosis was high for short tones ( $15.2 \pm 3.5$ ;  $n = 17$  neurons; all quantities are mean  $\pm$  SE unless otherwise specified), long tones ( $9.5 \pm 2.2$ ;  $n = 18$  neurons), and even silence ( $14.4 \pm 3.0$ ;  $n = 18$  neurons), but it was consistent with the random walk model when input correlations were removed through the application of TTX to the cortical surface (kurtosis of  $0.2 \pm 0.3$ ;  $n = 6$  neurons; asterisks denote mean values significantly different from zero according to a single sample Student's *t* test for  $p < 0.01$ ) (see Materials and Methods, TTX application).

extension of the model only when the model parameters vary on a fast ( $\tau$  of approximately tens of milliseconds) timescale.

### Estimating the number of synaptic inputs participating in a bump

Because we are searching for a simple and compact model of the data, we neglected the impact of many potentially important biophysical phenomena, including dendritic integration and synaptic saturation (Reyes, 2001; Benucci et al., 2004; Kuhn et al., 2004;

Rudolph et al., 2004), all of which would be expected to influence the observed distribution of membrane potential. Moreover, we stress that the particular time-varying model we considered is by no means uniquely specified by the data. We therefore return here to our central, robust finding, evident in all of the sample traces and indeed in all of the neurons from which we recorded: fluctuations are dominated by stereotyped bumps reflecting an abrupt, synchronized increase in network activity. It thus seems natural to focus on the structure of the bumps themselves.

How many unitary synaptic events participate in a bump? In previous experiments using voltage-clamp methods, the excitatory and inhibitory conductance changes underlying a typical 15 mV stimulus-evoked event were estimated to be comparable, each rising  $\sim 15$  nS at their peaks (Wehr and Zador, 2003), with each following a time course similar to the resulting membrane potential. Using a rough estimate of 0.1 nS per synaptic vesicle (Stevens and Zador, 1998; Gil et al., 1999), we estimate that underlying a typical 15 mV bump are at least 1000 EPSPs and 1000 IPSPs, concentrated over a brief period and arriving at a peak rate of at least 50 PSPs/ms (Fig. 7) (for details of this calculation, see Materials and Methods, Estimation of the presynaptic firing rate). This suggests that a substantial fraction of the synaptic inputs to a neuron may participate in each network event.

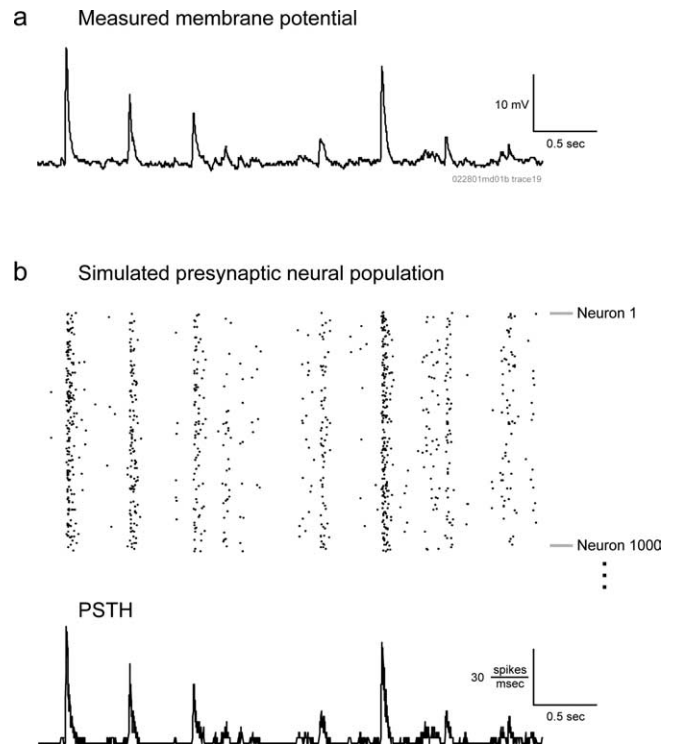
## Discussion

We used single-trial records of spontaneous and evoked fluctuations of membrane potential in auditory cortex neurons to make inferences about network dynamics underlying spikes. Membrane potential fluctuations were characterized by occasional large excursions (bumps, sometimes tens of millivolts high) that were much larger than the small fluctuations predicted by simple random walk models. We infer that these bumps are the manifestation of large synchronized volleys of action potentials from the presynaptic population, with  $\sim 1000$  PSPs contributing to the larger events. These dynamics suggest that spike timing is controlled by concerted firing among input neurons rather than by small fluctuations in a sea of background activity.

Membrane potential dynamics in the auditory cortex of intact animals have been examined previously by several groups (De Ribaupierre et al., 1972; Volkov and Galazyuk 1992; Ojima and Murakami, 2002; Wehr and Zador, 2003; Zhang et al., 2003; DeWeese and Zador, 2004; Tan et al., 2004; Las et al., 2005), and the present data appear consistent with that body of work. In particular, examples of subthreshold activity presented in those reports typically appeared to consist of brief excursions from rest, similar to the bumps that dominate the dynamics we observed. The subthreshold dynamics we see also appear consistent with at least some data from barrel cortex (Wilson and Groves, 1981; Ferster and Jagadeesh, 1992; Steriade et al., 1994; Wilson and Kawaguchi, 1996; Destexhe and Pare, 1999; Anderson et al., 2000; Sanchez-Vives and McCormick, 2000; Cossart et al., 2003; Destexhe et al., 2003; Petersen et al., 2003; Benucci et al., 2004; Carandini, 2004; Leger et al., 2005; Wilentz and Contreras, 2005).

### Sparse firing regime

Much of the previous work on cortical dynamics has focused on the high-firing regime (Shadlen and Newsome, 1998), or high-conductance regime (Destexhe et al., 2003), in which neurons fire at a high rate and the input conductance of a neuron is dominated by synaptic inputs. Under the conditions of our recordings, however, firing rates in auditory cortex are low to moderate (DeWeese et al., 2003); we refer to this as the “sparse-firing regime.”



**Figure 7.** How many presynaptic action potentials give rise to a typical synchronous volley? To get an order-of-magnitude estimate of the collective firing rate of the presynaptic population, we used a simple model that relied on previous measurements of the relationship between membrane potential and synaptic conductance (Wehr and Zador, 2003) and mEPSC size (Stevens and Zador 1998; Gil et al., 1999) (see Materials and Methods, Estimation of the presynaptic firing rate). To a first approximation, we estimate that the collective firing rate of the excitatory presynaptic population approximately follows the same time course as the recorded membrane potential itself, with  $\sim 3.3$  mEPSCs occurring every millisecond for every millivolt above rest in the whole-cell record. For example, the 15-mV-tall bump in the membrane potential (**a**; same trace as in Figs. 2a1, 5a) occurring  $\sim 1.5$  s before the end of the trace, corresponds to  $\sim 50$  mEPSCs per millisecond at its peak. Assuming that the probability of vesicle release,  $p$ , is close to 1, this is consistent with  $\sim 50$  spikes/ms from excitatory presynaptic fibers, which corresponds to  $\sim 1000$  spikes over the entire volley. Assuming that these spikes are equally distributed across 10,000 presynaptic neurons (Braitenberg and Schuz, 1998), we simulated spike rasters for 1000 of these neurons (**b**, top) and plotted the peristimulus time histogram (PSTH) for the full population (bottom). Note that this does not include any inhibitory inputs, which play as significant a role as the excitatory inputs near the peaks of the larger bumps (Wehr and Zador, 2003), and that these estimated spike rates may well be underestimates if  $p$  is actually  $< 1$ , as it is often reported to be (Castro-Alamancos and Connors, 1997; Dobrunz and Stevens, 1997; Huang and Stevens, 1997; Murthy et al., 1997).

The present results reflect the mechanisms of the underlying cortical networks *in vivo* operating in this regime.

How relevant are neurons in the sparse-firing regime to perception and behavior? One might question their relevance for at least two reasons. First, our data were recorded mainly in anesthetized animals. It has long been known that, in the unanesthetized auditory cortex, spiking responses can be either sustained or transient (Evans and Whitfield, 1964), whereas under anesthesia, transient responses dominate. Anesthesia does not, however, simply cause a gross reduction of activity; in the awake preparation, the typical firing rate for neurons not driven by optimal stimuli is  $\sim 2$ – $4$  spikes/s (Wang et al., 2005), comparable with that observed under anesthesia. Indeed, direct comparison of the same neurons before and after anesthesia reveals only a modest anesthesia-induced decrease in spontaneous firing rate (Talwar and Gerstein, 2001). Thus, the low typical firing rates observed in our preparation are not restricted solely to the anesthetized con-

dition. Moreover, similar behavior was observed among the few recordings obtained in unanesthetized animals.

The second concern is more subtle. In the awake animal, it is sometimes possible to find the “optimal stimulus” that maximizes the firing rate of a particular neuron (deCharms et al., 1998; Barbour and Wang, 2003; Wang et al., 2005). Although the stimuli we used, pure tones, are not optimal for most neurons in the primary auditory cortex, one might imagine that it is precisely the small subpopulation of highly activated neurons (those for which these tones are the optimal stimuli) that drives perception and behavior, as has been suggested for decisions about visual motion direction in area MT (Britten et al., 1996). A strong interpretation of this view would maintain that only those stimuli capable of driving at least some neurons optimally can be perceived or acted on. Although this strong interpretation cannot be ruled out on first principles, it has not been proven either; given how large the space of possible auditory stimuli is and how difficult it is to drive neurons in auditory cortex optimally, it is not clear that for every auditory stimulus capable of eliciting a behavioral response there exist well driven neurons. Thus, we provisionally adopt a conservative position and remain open to the possibility that neurons in the sparse-firing regime may play some role in perception and behavior. It is clear that additional experiments, beyond the scope of the present study, will be needed to understand the relative importance to behavior of sparse- and high-firing regimes under various conditions.

### Computing with bumps

We have shown that what drives auditory cortical neurons to spike in the sparse-firing regime is a tightly synchronized volley of inputs. What does the existence of such events tell us about cortical computation? Previous studies have shown that, at least under some conditions, stimulus transients can trigger spikes with a precision of a few milliseconds [visual cortex (Bair and Koch, 1996; Buracas et al., 1998)] or even one millisecond [auditory cortex (DeWeese et al., 2003; Heil 2004)]. This was demonstrated by measuring spike timing jitter across multiple presentations of the same stimulus. Experimental observations of patterned neural activity involving synchronous spiking have now been made in a variety of cortical preparations (Schwartz et al., 1998; Ikegaya et al., 2004). However, the fact that an experimenter can extract information about a stimulus from the precise timing or pattern of spikes (Bialek et al., 1991; Buracas et al., 1998; Furukawa and Middlebrooks, 2002) does not necessarily imply that this timing is used by the organism to encode stimulus features or perform computations. Moreover, it is uncertain whether such precisely timed spikes are exclusively stimulus locked (e.g., to stimulus transients) or whether such precision can occur in response to internally generated events such as inputs from another (nonsensory) cortical region.

One might suppose that the strong correlations we infer from membrane potential dynamics would be evident from paired extracellular recordings (Eggermont and Smith, 1995). However, such elevated correlations across even a few percentage of neurons might not be detectable, particularly if the correlations were transient, as might be expected if activity were organized into “cell assemblies” (Hebb, 1949; Harris et al., 2003). Perhaps we can detect these synchronized volleys because, by recording intracellularly, we are sampling just the relevant subpopulation that contributes to the cell assembly. Moreover, by simultaneously recording the local field potential with a second electrode, we have shown previously (DeWeese et al., 2003) that fluctuations in the whole-cell record are correlated with the collective synaptic

activity shared by large groups of distant ( $\sim 0.5$  mm) auditory cortical neurons.

From a computational point of view, it might seem puzzling that cortical neurons should behave in such a correlated manner. Indeed, a literal reading of the diagram in which all inputs fire at once during every volley (Fig. 1c) would seem to have little promise for computations beyond merely transmitting an exact copy of the spike train shared by the full presynaptic subpopulation to the next stage of processing. Fortunately, with up to 10,000 excitatory inputs feeding into a typical cortical neuron (Braitenberg and Schuz, 1998), an enormous number of distinct subsets of synapses could potentially be responsible for spike production on different volleys, so that the output spike train of the neuron need not be an exact copy of a spike train from any of its presynaptic inputs.

Our observations suggest that the same network dynamics responsible for laying down stimulus-evoked spikes with high temporal precision underlie other spikes as well. This raises the possibility that, in the sparse-firing regime, precise spike timing may play a role in computation. It is sometimes suggested that computational schemes that depend on precise spike timing must necessarily be sensitive to noise. Our data do not, however, suggest that “every spike matters,” because the synchronous inputs underlying each spike consist of hundreds or thousands of PSPs. The timing of spikes is thus robust to the addition or loss of presynaptic inputs.

Our study raises a number of questions. For example, what network properties could give rise to such dynamics? Theoretical work has shown that synchronized bursts of activity across the network (Tsodyks et al., 2000) and even concerted volleys of activity propagating from one subpopulation to the next (Litvak et al., 2003) might be necessary consequences of cortical connectivity. Additionally, specific proposals for how cortical processing could proceed with these synchronized volleys have been advanced (Aertsen et al., 1996; Diesmann et al., 1999; Loebel and Tsodyks, 2002; Beggs and Plenz, 2003; Reyes, 2003; Shu et al., 2003). In future work, it will be interesting to test whether membrane potential dynamics measured in the awake auditory cortex show the same degree of input correlation as we found in both our anesthetized population and our preliminary recordings from unanesthetized animals. Most exciting will be to observe membrane potential dynamics within the context of well controlled behavioral paradigms, so that we may determine how these input correlations depend on task contingencies and attentional state.

### References

- Aertsen A, Diesmann M, Gewaltig MO (1996) Propagation of synchronous spiking activity in feedforward neural networks. *J Physiol (Paris)* 90:243–247.
- Anderson J, Lampl I, Reichova I, Carandini M, Ferster D (2000) Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex. *Nat Neurosci* 3:617–621.
- Atzori M, Lei S, Evans DJ, Kanold PO, Phillips-Tansey E, McIntyre O, McBain CJ (2001) Differential synaptic processing separates stationary from transient inputs to the auditory cortex. *Nat Neurosci* 4:1230–1237.
- Azouz R, Gray CM (1999) Cellular mechanisms contributing to response variability of cortical neurons *in vivo*. *J Neurosci* 19:2209–2223.
- Bair W, Koch C (1996) Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput* 8:1185–1202.
- Barbour DL, Wang X (2003) Auditory cortical responses elicited in awake primates by random spectrum stimuli. *J Neurosci* 23:7194–7206.
- Beggs JM, Plenz D (2003) Neuronal avalanches in neocortical circuits. *J Neurosci* 23:11167–11177.
- Bell AJ, Mainen ZF, Tsodyks M, Sejnowski TJ (1995) Balancing conduc-

- tances may explain irregular cortical firing. Technical report. San Diego: Institute for Neural Computation.
- Benucci A, Verschure PF, König P (2004) Two-state membrane potential fluctuations driven by weak pairwise correlations. *Neural Comput* 16:2351–2378.
- Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D (1991) Reading a neural code. *Science* 252:1854–1857.
- Borg-Graham LJ, Monier C, Fregnac Y (1998) Visual input evokes transient and strong shunting inhibition in visual cortical neurons. *Nature* 393:369–373.
- Braitenberg V, Schuz A (1998) *Cortex: statistics and geometry of neuronal connectivity*. Berlin: Springer.
- Britten KH, Newsome WT, Shadlen MN, Celebri S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13:87–100.
- Bruno RM, Sakmann B (2006) Cortex is driven by weak but synchronously active thalamocortical synapses. *Science* 312:1622–1627.
- Buracas GT, Zador AM, DeWeese MR, Albright TD (1998) Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 20:959–969.
- Calvin WH, Stevens CF (1967) Synaptic noise as a source of variability in the interval between action potentials. *Science* 155:842–844.
- Carandini M (2004) Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biol* 2:e264.
- Castro-Alamancos MA, Connors BW (1996) Cellular mechanisms of the augmenting response: short-term plasticity in a thalamocortical pathway. *J Neurosci* 16:7742–7756.
- Castro-Alamancos MA, Connors BW (1997) Distinct forms of short-term plasticity at excitatory synapses of hippocampus and neocortex. *Proc Natl Acad Sci USA* 94:4161–4166.
- Chance FS, Abbott LF, Reyes AD (2002) Gain modulation from background synaptic input. *Neuron* 35:773–782.
- Cossart R, Aronov D, Yuste R (2003) Attractor dynamics of network UP states in the neocortex. *Nature* 423:283–288.
- Covey E, Kauer JA, Casseday JH (1996) Whole-cell patch-clamp recording reveals subthreshold sound-evoked postsynaptic currents in the inferior colliculus of awake bats. *J Neurosci* 16:3009–3018.
- deCharms RC, Blake DT, Merzenich MM (1998) Optimizing sound features for cortical neurons. *Science* 280:1439–1443.
- De Ribaupierre F, Goldstein MH, Yeni-Komshian G (1972) Intracellular study of the cat's primary auditory cortex. *Brain Res* 48:185–204.
- Deisz RA, Billard JM, Zieglansberger W (1997) Presynaptic and postsynaptic GABAB receptors of neocortical neurons of the rat in vitro: differences in pharmacology and ionic mechanisms. *Synapse* 25:62–72.
- Destexhe A, Pare D (1999) Impact of network activity on the integrative properties of neocortical pyramidal neurons in vivo. *J Neurophysiol* 81:1531–1547.
- Destexhe A, Rudolph M, Pare D (2003) The high-conductance state of neocortical neurons in vivo. *Nat Rev Neurosci* 4:739–751.
- Deweese MR, Zador AM (2004) Shared and private variability in the auditory cortex. *J Neurophysiol* 92:1840–1855.
- DeWeese MR, Wehr M, Zador AM (2003) Binary spiking in auditory cortex. *J Neurosci* 23:7940–7949.
- Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–533.
- Dobrunz LE, Stevens CF (1997) Heterogeneity of release probability, facilitation, and depletion at central synapses. *Neuron* 18:995–1008.
- Eggermont JJ, Smith GM (1995) Rate covariance dominates spontaneous cortical unit-pair correlograms. *NeuroReport* 6:2125–2128.
- Evans EF, Whitfield IC (1964) Classification of unit responses in the auditory cortex of the unanaesthetized and unrestrained cat. *J Physiol (Lond)* 171:476–493.
- Fee MS (2000) Active stabilization of electrodes for intracellular recording in awake behaving animals. *Neuron* 27:461–468.
- Fellous JM, Rudolph M, Destexhe A, Sejnowski TJ (2003) Synaptic background noise controls the input/output characteristics of single cells in an in vitro model of in vivo activity. *Neuroscience* 122:811–829.
- Ferster D, Jagadeesh B (1992) EPSP–IPSP interactions in cat visual cortex studied with *in vivo* whole-cell patch recording. *J Neurosci* 12:1262–1274.
- Furukawa S, Middlebrooks JC (2002) Cortical representation of auditory space: information-bearing features of spike patterns. *J Neurophysiol* 87:1749–1762.
- Gerstein GL, Mandelbrot B (1964) Random walk models for the spike activity of a single neuron. *Biophys J* 71:41–68.
- Gil Z, Connors BW, Amitai Y (1999) Efficacy of thalamocortical and intracortical synaptic connections: quanta, innervation, and reliability. *Neuron* 23:385–397.
- Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsáki G (2003) Organization of cell assemblies in the hippocampus. *Nature* 424:552–556.
- Hebb DO (1949) *The organization of behavior*. New York: Wiley.
- Heil P (2004) First-spike latency of auditory neurons revisited. *Curr Opin Neurobiol* 14:461–467.
- Hertz J, Richmond BJ, Nilsen K (2003) Anomalous response variability in a balanced cortical network model. *Neurocomputing* 52–54:787–792.
- Hirsch JA, Alonso JM, Reid RC (1995) Visually evoked calcium action potentials in cat striate cortex. *Nature* 378:612–616.
- Huang EP, Stevens CF (1997) Estimating the distribution of synaptic reliabilities. *J Neurophysiol* 78:2870–2880.
- Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, Yuste R (2004) Synfire chains and cortical songs: temporal modules of cortical activity. *Science* 304:559–564.
- Jagadeesh B, Wheat HS, Kontsevich LL, Tyler CW, Ferster D (1997) Direction selectivity of synaptic potentials in simple cells of the cat visual cortex. *J Neurophysiol* 78:2772–2789.
- Kuhn A, Aertsen A, Rotter S (2004) Neuronal integration of synaptic input in the fluctuation-driven regime. *J Neurosci* 24:2345–2356.
- Las L, Stern EA, Nelken I (2005) Representation of tone in fluctuating maskers in the ascending auditory system. *J Neurosci* 25:1503–1513.
- Leger JF, Stern EA, Aertsen A, Heck D (2005) Synaptic integration in rat frontal cortex shaped by network activity. *J Neurophysiol* 93:281–293.
- Litvak V, Sompolinsky H, Segev I, Abeles M (2003) On the transmission of rate code in long feedforward networks with excitatory–inhibitory balance. *J Neurosci* 23:3006–3015.
- Loebel A, Tsodyks M (2002) Computation by ensemble synchronization in recurrent networks with synaptic depression. *J Comput Neurosci* 13:111–124.
- Margrie TW, Brecht M, Sakmann B (2002) In vivo, low-resistance, whole-cell recordings from neurons in the anaesthetized and awake mammalian brain. *Pflügers Arch* 444:491–498.
- Markram H, Lübke J, Frotscher M, Roth A, Sakmann B (1997) Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. *J Physiol (Lond)* 500:409–440.
- Metherate R, Ashe JH (1993) Nucleus basalis stimulation facilitates thalamocortical synaptic transmission in the rat auditory cortex. *Synapse* 14:132–143.
- Moore CI, Nelson SB (1998) Spatio-temporal subthreshold receptive fields in the vibrissa representation of rat primary somatosensory cortex. *J Neurophysiol* 80:2882–2892.
- Murthy VN, Sejnowski TJ, Stevens CF (1997) Heterogeneous release properties of visualized individual hippocampal synapses. *Neuron* 18:599–612.
- Nelson S, Toth L, Sheth B, Sur M (1994) Orientation selectivity of cortical neurons during intracellular blockade of inhibition. *Science* 265:774–777.
- Ojima H, Murakami K (2002) Intracellular characterization of suppressive responses in supragranular pyramidal neurons of cat primary auditory cortex in vivo. *Cereb Cortex* 12:1079–1091.
- Petersen CC, Hahn TT, Mehta M, Grinvald A, Sakmann B (2003) Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. *Proc Natl Acad Sci USA* 100:13638–13643.
- Reyes A (2001) Influence of dendritic conductances on the input-output properties of neurons. *Annu Rev Neurosci* 24:653–675.
- Reyes AD (2003) Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro. *Nat Neurosci* 6:593–599.
- Rudolph M, Destexhe A (2003) Characterization of subthreshold voltage fluctuations in neuronal membranes. *Neural Comput* 15:2577–2618.
- Rudolph M, Piwkowska Z, Badoual M, Bal T, Destexhe A (2004) A method to estimate synaptic conductances from membrane potential fluctuations. *J Neurophysiol* 91:2884–2896.
- Sanchez-Vives MV, McCormick DA (2000) Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nat Neurosci* 3:1027–1034.
- Schwartz TH, Rabinowitz D, Unni V, Kumar VS, Smetters DK, Tsiola A, Yuste R (1998) Networks of coactive neurons in developing layer 1. *Neuron* 20:541–552.

- Shadlen MN, Newsome WT (1994) Noise, neural codes and cortical organization. *Curr Opin Neurobiol* 4:569–579.
- Shadlen MN, Newsome WT (1995) Is there a signal in the noise? *Curr Opin Neurobiol* 5:248–250.
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896.
- Shu Y, Hasenstaub A, McCormick DA (2003) Turning on and off recurrent balanced cortical activity. *Nature* 423:288–293.
- Simoncelli EP, Olshausen BA (2001) Natural image statistics and neural representation. *Annu Rev Neurosci* 24:1193–1216.
- Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13:334–350.
- Song S, Miller KD, Abbott LF (2000) Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nat Neurosci* 3:919–926.
- Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB (2005) Highly non-random features of synaptic connectivity in local cortical circuits. *PLoS Biol* 3:e68.
- Steriade M, Contreras D, Amzica F (1994) Synchronized sleep oscillations and their paroxysmal developments. *Trends Neurosci* 17:199–208.
- Steriade M, Timofeev I, Grenier F (2001) Natural waking and sleep states: a view from inside neocortical neurons. *J Neurophysiol* 85:1969–1985.
- Stevens CF, Zador AM (1998) Input synchrony and the irregular firing of cortical neurons. *Nat Neurosci* 1:210–217.
- Talbot MJ, Sayer RJ (1996) Intracellular QX-314 inhibits calcium currents in hippocampal CA1 pyramidal neurons. *J Neurophysiol* 76:2120–2124.
- Talwar SK, Gerstein GL (2001) Reorganization in awake rat auditory cortex by local microstimulation and its effect on frequency-discrimination behavior. *J Neurophysiol* 86:1555–1572.
- Tan AY, Zhang LI, Merzenich MM, Schreiner CE (2004) Tone-evoked excitatory and inhibitory synaptic conductances of primary auditory cortex neurons. *J Neurophysiol* 92:630–643.
- Troyer TW, Miller KD (1997) Physiological gain leads to high ISI variability in a simple model of a cortical regular spiking cell. *Neural Comput* 9:971–983.
- Tsodyks M (1995) Rapid switching in balanced cortical network models. *Network* 6:1–14.
- Tsodyks M, Uziel A, Markram H (2000) Synchrony generation in recurrent networks with frequency-dependent synapses. *J Neurosci* 20:RC50(1–5).
- Tuckwell HC (1988) Introduction to theoretical neurobiology. Cambridge, UK: Cambridge UP.
- van Vreeswijk C, Sompolinsky H (1996) Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science* 274:1724–1726.
- van Vreeswijk C, Sompolinsky H (1998) Chaotic balanced state in a model of cortical circuits. *Neural Comput* 10:1321–1371.
- Volkov IO, Galazyuk AV (1992) Peculiarities of inhibition in cat auditory cortex neurons evoked by tonal stimuli of various durations. *Exp Brain Res* 91:115–120.
- Wang X, Lu T, Snider RK, Liang L (2005) Sustained firing in auditory cortex evoked by preferred stimuli. *Nature* 435:341–346.
- Wehr M, Zador AM (2003) Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* 426:442–446.
- Wilentz WB, Contreras D (2005) Dynamics of excitation and inhibition underlying stimulus selectivity in rat somatosensory cortex. *Nat Neurosci* 8:1364–1370.
- Wilson CJ, Groves PM (1981) Spontaneous firing patterns of identified spiny neurons in the rat neostriatum. *Brain Res* 220:67–80.
- Wilson CJ, Kawaguchi Y (1996) The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. *J Neurosci* 16:2397–2410.
- Zhang LI, Tan AY, Schreiner CE, Merzenich MM (2003) Topography and synaptic shaping of direction selectivity in primary auditory cortex. *Nature* 424:201–205.
- Zhu JJ, Connors BW (1999) Intrinsic firing patterns and whisker-evoked synaptic responses of neurons in the rat barrel cortex. *J Neurophysiol* 81:1171–1183.