

Reliability and Representational Bandwidth in the Auditory Cortex

Review

Michael R. DeWeese, Tomáš Hromádka,
and Anthony M. Zador*

Cold Spring Harbor Laboratory
Watson School of Biological Sciences
Cold Spring Harbor, New York 11724

It is unclear why there are so many more neurons in sensory cortex than in the sensory periphery. One possibility is that these “extra” neurons are used to overcome cortical noise and faithfully represent the acoustic stimulus. Another possibility is that even after overcoming cortical noise, there is “excess representational bandwidth” available and that this bandwidth is used to represent conjunctions of auditory and nonauditory information for computation. Here, we discuss recent data about neuronal reliability in auditory cortex showing that cortical noise may not be as high as was previously believed. Although at present, the data suggest that auditory cortex neurons can be more reliable than those in the visual cortex, we speculate that the principles governing cortical computation are universal and that visual and other cortical areas can also exploit strategies based on similarly high-fidelity activity.

Sensory cortex represents information about the sensory world. According to one common model, sensory processing proceeds by transforming the raw sensory signal through a series of higher level (more complex) representations; the features represented at each later stage are more abstract than at the earlier stage. In the visual system, for example, where this model has been most thoroughly elaborated, it is posited that in the first cortical stage (area V1) the features represented are oriented edges and that after a series of transformations (in areas V2 and V4), very complex features such as faces and objects (area IT) are represented. A similar but less well-understood series of transformations is presumed to occur between the early stages of auditory cortical processing (e.g., areas A1 and AAF) and later stages.

Underlying the processing of acoustic signals are the neural circuits in the subcortical and cortical auditory pathway. A remarkable feature of the anatomy of these pathways is that there are many more neurons in auditory cortex than there are in the auditory nerve. This divergence can be dramatic. For example, a human has about 16,000 hair cells in each cochlea, but at least ten thousand times as many neurons in the primary auditory cortex. The visual system shows a comparable divergence from the retina to the visual cortex.

Why are there so many more neurons in auditory cortex than in the auditory nerve? (The same question—and possibly the same answers—arise about divergence between sensory afferents and sensory cortex in vision and other modalities as well, but we will focus our discussion on the auditory system.) To the extent that the

function of the auditory cortex is to provide a faithful representation—or even a series of rerepresentations—of the auditory signal at the periphery, it appears that there are orders of magnitude more cortical neurons available than are needed. We call this the “excess representational bandwidth” problem. The simple explanation is, of course, that the cortex has so many neurons because it must perform complex computations; the extra neurons are there to subserve these computations. However, without a model of how and what the cortex computes, this simple explanation has little explanatory or predictive value. Indeed, this excess capacity is all the more puzzling because the goal of sensory processing is presumably to compute, and computation usually involves discarding information. For example, computing the sum $1 + 2 + 3 + 4 = 10$ discards the details about the order of the numbers on the left side of the equal sign.

One possible explanation is that the extra neurons are needed to overcome cortical noise. This explanation follows from a common model of cortical computation that posits that cortical representations are necessarily noisy (Shadlen and Newsome, 1998). If neurons in the auditory cortex were much noisier than in the auditory nerve, then many more cortical neurons would be required to represent faithfully the information available at the cochlea. This view is implicit in models of representation that assume that neurons within a cortical column all represent approximately the same information and that downstream neurons must average over such columns in order to “read out” these neuronal populations. In this view, cortical processing proceeds by means of a series of rerepresentations of the same sensory information, and the cortical circuitry underlying the transformation of the sensory signal from one stage to the next necessarily corrupts the signal; only by averaging over large populations can this corruption be overcome.

An alternative or additional function of these extra neurons might be to facilitate computation by providing multiple representations of the same information in somewhat different forms, an approach used in artificial neural networks (Vapnik, 1999). Continuing with the example above, it might be useful to represent not just the sum of the numbers above but their product as well: $1 \times 2 \times 3 \times 4 = 24$. Indeed, the existence of multiple cortical areas (A1, AAF, etc.) may be an example of such multiple representations. There may also be multiple representations within an area.

A third possibility is that the auditory cortex may facilitate computation by representing more than just the auditory stimulus itself; it may also represent conjunctions of “bottom-up” sensory signals from the auditory periphery and “top-down” signals encoding expectations, goals, attention, and other nonsensory information relevant to auditory processing. For example, at a cocktail party one can voluntarily attend to any of several simultaneous voices; to the extent that this attention is reflected in neuronal activity in the auditory cortex (Hubel et al., 1959; Fritz et al., 2003), the signal specifying which voice to attend to must be conveyed in a top-down

*Correspondence: zador@cshl.edu

fashion. Little is known about how this occurs, but it is possible that such top-down signals require substantial representational bandwidth.

How we as experimenters should interpret cortical spikes depends strongly on which of these explanations is correct. Of course, the full explanation may involve some combination of these and other factors. If cortical neurons are very noisy, then it seems reasonable to hypothesize that it is necessary to average over large populations to overcome this noise. If on the other hand cortical neurons are not very noisy, then perhaps the apparent excess representational bandwidth has some other functions. We believe that neuronal noise may not impose an important constraint on cortical computation and that the “extra” neurons in auditory cortex aid computation by providing multiple representations of the auditory signal and conjunctions of auditory and non-auditory signals.

In this review, we will focus on two issues. First, we will summarize some of what is known about the reliability of neurons in the auditory cortex and what this reliability implies for cortical representations. We will emphasize experimental data that bear on these issues and attempt to maintain a clear separation between the data themselves and our speculations about the implications of the data for theories about neural coding. We will show that according to two widely applied measures of neuronal reliability—spike timing and spike count—neurons in auditory cortex can be quite precise. In particular, the spike-count reliability of neurons in auditory cortex can be much greater than has previously been reported in other cortical regions. Second, we will briefly summarize both classical and recent data suggesting that the auditory cortex represents more than simply auditory information. Together, these two lines of data raise the possibility that cortical computation is not limited by neuronal noise and that much of the apparent excess of cortical bandwidth may be used to represent the conjunction of auditory and nonauditory information.

The scope of this review is necessarily limited. We focus discussion mainly on the subset of data relevant to neuronal noise in the auditory cortex, comparing and contrasting occasionally with data from visual cortex, in which reliability has been studied extensively (cf., [Petersen et al., 2002](#)). We will not speculate on what computations (e.g., acoustic source separation and scene analysis [[Micheyl et al., 2005](#)] or representation of auditory objects [[Nelken et al., 2003](#)]) the auditory cortex might be performing.

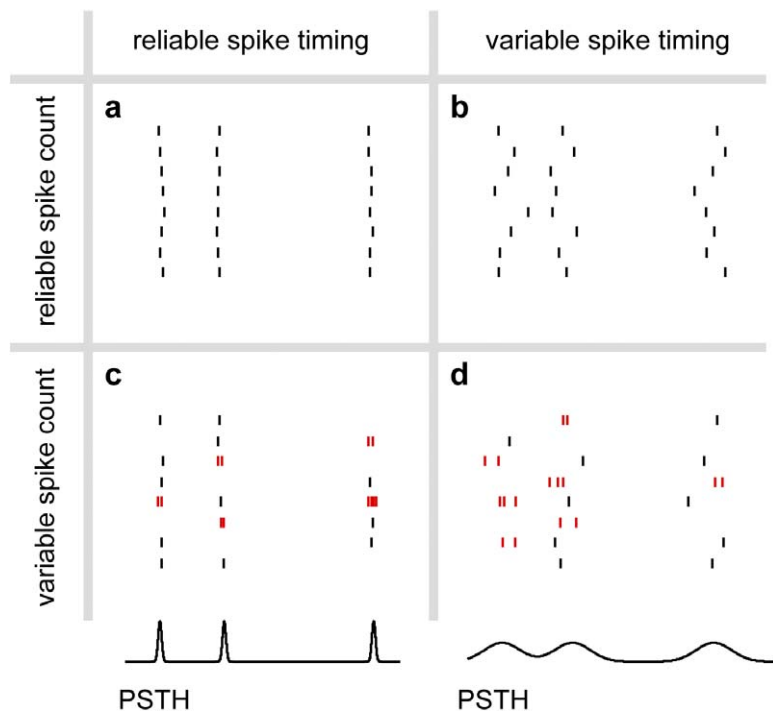
Neuronal Reliability

Sensory-evoked responses in the cortex have generally been found to be highly variable: the same stimulus typically elicits a different response on each presentation. In this respect, the visual cortex differs from the sensory periphery and even from other parts of the central nervous system such as the spinal cord and retina, where responses are more orderly and regular ([Calvin and Stevens, 1968](#); [Berry and Meister, 1998](#)). Although it is often supposed that high variability is a hallmark of all cortical regions, much of the support for this view comes from recordings in visual cortex where neural reliability has been scrutinized intensively.

The high variability of responses in visual cortex has often been used to make inferences about the kinds of codes neurons might use to represent stimuli ([Softky and Koch, 1993](#); [Shadlen and Newsome, 1998](#); [Pouget et al., 2000](#); [Mazurek and Shadlen, 2002](#)). The spikes from a well-driven neuron in visual cortex, when played through an audio monitor, typically sound like the ticks of a Geiger counter (but see [Kara et al. \[2000\]](#)). For some stimuli (e.g., a bar moving in a particular direction), the neuron fires at a higher rate, and for others (e.g., a bar moving in a different direction), at a lower rate; thus, the experimenter can “read out” some stimulus parameter (like the direction of visual motion) by estimating the spike rate. Under these conditions, the experimenter extracts the information about stimulus direction from the neuron’s firing rate rather than from the precise times of the individual spikes; the timing of the spikes themselves, and the number of spikes on a trial, typically vary from trial to trial and so are not useful for estimating the stimulus. The spike trains can thus be modeled as instantiations of a random process, with the stimulus controlling only the statistical parameters of the stochastic process. In the simplest case, the spike train might be modeled as a Poisson process (the name given to the stochastic process governing the random ticks of the Geiger counter in our example), and the direction of the bar would be encoded in the mean spike rate, in the same way that one could estimate one’s distance from a radioactive source by listening to the mean click rate of the Geiger counter. This is the basis for the idea of a classical “rate code.”

Because under many conditions, the experimenter can extract significant information about the stimulus from the firing rate, but not much more from the exact number or arrival times of individual spikes, the residual neuronal variability is often interpreted as “noise.” But in what sense is it really noise? Certainly from the point of view of the experimenter, it is noise; by definition, noise is simply unexplained variability. For researchers interested in understanding cortical circuitry, however, the key question is whether it is also noise from the organism’s point of view. The experimenter has access to only a small number of variables—often just the sensory stimulus she presents—and from these she must attempt to predict the spike train of the neuron from which she is recording. In this regard, the organism is at a distinct advantage, in that he has access to many more internal variables; in principle, he might have access to the activity of all his own neurons.

Another way to think about this is that the cortex might be in a different state, or it might be performing a different computation, from one trial to the next. For example, changes in the animal’s motivation or attention can alter cortical responses throughout cortex, even in sensory areas ([Hocherman et al., 1976](#); [Reynolds and Chelazzi, 2004](#); [Fritz et al., 2003](#)). Thus, trial-to-trial differences in the details of what the entire cortex is trying to compute might masquerade as noise in a neuron’s representation of a sensory stimulus ([Arieli et al., 1995](#); [Buracas et al., 1998](#); [DeWeese and Zador, 2004](#)). This implies that the mere demonstration of high neuronal variability does not necessarily mean that the fluctuating features of the spike train are incapable of conveying signals important for coding or computation from the animal’s perspective.



of rasters are transient, but their high spike-count variability is more consistent with a highly stochastic Poisson process than a binomial process that generates only one or zero spikes per trial. All multispike responses to any of the three stimulus transients are shown in red. As this toy example shows, it is a logical possibility for sensory cortical neurons to exhibit low jitter but high spike-count variability, or vice versa.

In what follows, we will summarize data from auditory cortex showing that spike trains can be much more reliable than previously reported in visual cortex. The fact that a spike train can, under at least some circumstances, be reliable even from the experimenter's point of view raises the possibility that it is reliable from the organism's point of view under other conditions—even when the experimenter does not have access to all the relevant variables. However, the mere demonstration of high neuronal reliability does not in any way imply that this reliability is used by the nervous system. It is quite possible that the high reliability of these auditory spike trains is irrelevant to behavior or coding. (In the same way, the fact that area X contains a neural representation of stimulus property Y does not imply that the organism uses the representation in area X.) A different kind of experiment—for example, some sort of “reliability knockout” experiment—would be needed to demonstrate relevance to behavior or coding.

In quantifying neuronal reliability, it is useful to distinguish between two different forms: timing reliability and spike-count reliability (Figure 1). Timing reliability, or its inverse timing “jitter,” is typically assessed with reference to some well-defined event present in the stimulus such as its onset or termination. Spike-count reliability is assessed by counting the number of spikes elicited by a stimulus in some time window. We will see that fairly precise stimulus-evoked spike timing can coexist with the high spike-count variability expected from a classical rate code (Buracas et al., 1998), but we will also see that the existence of low spike-count variability in both anesthetized (DeWeese et al., 2003) and unanesthetized

Figure 1. A Cartoon Illustrating the Distinction between Two Types of Neural Response Variability: Spike-Timing Variability and Spike-Count Variability

Each of the four panels depicts a set of artificial spike rasters generated in response to eight repeated presentations of a stimulus consisting of three transient events that each evoke an average of one spike per trial; rows represent the time courses of individual trials, with hash marks indicating spikes. Both panels on the left (A and C) contain responses that were generated with high temporal precision from trial to trial, whereas the two panels on the right (B and D) display much greater spike-timing variability, or “jitter.” Idealized peri-stimulus time histograms (PSTHs) for the low temporal jitter cases (bottom left) and the high temporal jitter cases (bottom right) are shown at the bottom of the figure. Grouping the panels another way reveals that the top two rasters (A and B) show perfect spike-count reliability—every trial contains exactly one spike for each of the three stimulus-locked responses—whereas the bottom panels (C and D) exhibit spike counts that are highly variable from one trial to the next. We would classify the responses in each of the top two panels as “binary” because no more than one spike is elicited on any trial for each of the three transient responses. The responses in the bottom pair

(Chimoto et al., 2002; Barbour and Wang, 2003) auditory cortex provides a clear example of neuronal activity pointing beyond a rate code. These highly reliable spiking responses are consistent with the observation that subthreshold responses in auditory cortical neurons can approach the limit imposed by single-neuron biophysics ($[\text{variance in PSP height}] / [\text{mean PSP height}] \sim 0.25 \text{ mV}$) (DeWeese and Zador, 2004).

Timing Reliability

Since the earliest experiments, it has been evident that the cortical response to the onset of a stimulus can be quite temporally precise. The most reliable stimulus-locked spikes in visual cortex have a jitter of less than five milliseconds (Bair and Koch, 1996; Buracas et al., 1998; Bair, 1999), and the most reliable spikes in auditory cortex have a jitter of less than 1 ms (DeWeese et al., 2003; Heil, 2004). Much of the information present in a spike train about a stimulus can often be extracted from the timing of the first few spikes (Bair and Koch, 1996; Buracas et al., 1998; Petersen et al., 2002; Stecker and Middlebrooks, 2003; Heil, 2004; Osborne et al., 2004; Nelken et al., 2005).

Are the spikes that occur well after the initial transient spikes also laid down precisely? In experiments in which a stimulus (such as a long tone or an oriented grating moving across a receptive field) is turned on abruptly and then maintained at a constant level for hundreds of milliseconds or even seconds, it certainly appears—at least from the point of view of the experimenter—that the answer is no, that the timing of spikes after the onset is quite random. Such experiments are, however, not definitive; this negative result may reflect merely the

limitations of our experiment. In the absence of a stimulus transient, there might exist some “internal clock”—for example, a theta or gamma rhythm—that neurons use as a timing reference (O’Keefe and Recce, 1993; Lisman and Idiart, 1995; deCharms and Merzenich, 1996; Hopfield and Brody, 2001). Unless the experimenter has access to the clock, she would remain unaware that the spikes were timed precisely. Nevertheless, although there is ample evidence for synchrony and rhythms in auditory cortex (Eggermont, 2000; Sukov and Barth, 2001), it remains unclear whether these rhythms play the role of a master clock.

It might be argued that such static stimuli are highly unnatural; animals in nature are rarely confronted with static stimuli. In the acoustic domain, typical sounds—the chatter of birds, the rustling of leaves, speech, traffic, hoof beats, thunder—are characterized by their rich temporal structure. Even when the stimulus itself does not vary in time, as, for example, with static visual scene, active sensation (eye movements, sniffing, whisking) can impose temporal structure on the raw sensory signal. Temporal structure might conceivably also arise from internal processes (such as suddenly making a decision), but such internal processes are more difficult to monitor experimentally.

Because dynamic stimuli effectively consist of a series of transients, they can elicit a corresponding series of precisely timed onset responses. However, the fact that these responses can be modulated with millisecond precision is compatible with a natural extension of a classical rate code. Returning to the previous analogy of the Geiger counter, imagine attaching the probe to a swinging pendulum and placing a radioactive source directly at its nadir. The click rate attains a maximum at the nadir when the probe is close to the source, and a minimum at the zenith when it is far; the instantaneous rate is at all times precisely governed by the distance of the probe from the source. The rapid modulation of the Geiger counter’s click rate in this example reveals something about how quickly the probe moves to and from the source but does not change our view of radioactive decay; the mechanism governing the click times remains fundamentally random no matter how fast the probe is moved. Similarly, the fact that spike rates can be modulated on a fast time scale—in which fast is defined relative to the typical interspike interval—tells us little about whether rates alone matter.

Precisely timed spikes elicited by dynamic stimuli thus do not shed light on whether spikes evoked long after any stimulus transients have faded may also have precise timing from the animal’s point of view. In order to examine whether spike times are well described by the conventional view in which they are drawn randomly from a stimulus-modulated rate or if instead they are being laid down precisely through the activity of neurons in the network, we turn now to spike-count reliability.

Spike-Count Reliability

A hallmark of any Poisson process is that the number of events in any specified time window is drawn from a Poisson distribution. This relation holds even for a time-varying, or rate-modulated, Poisson process, such as that governing the Geiger probe affixed to the pendulum. If a “trial” is defined as one cycle of the pendulum, then the distribution of the number of clicks produced within

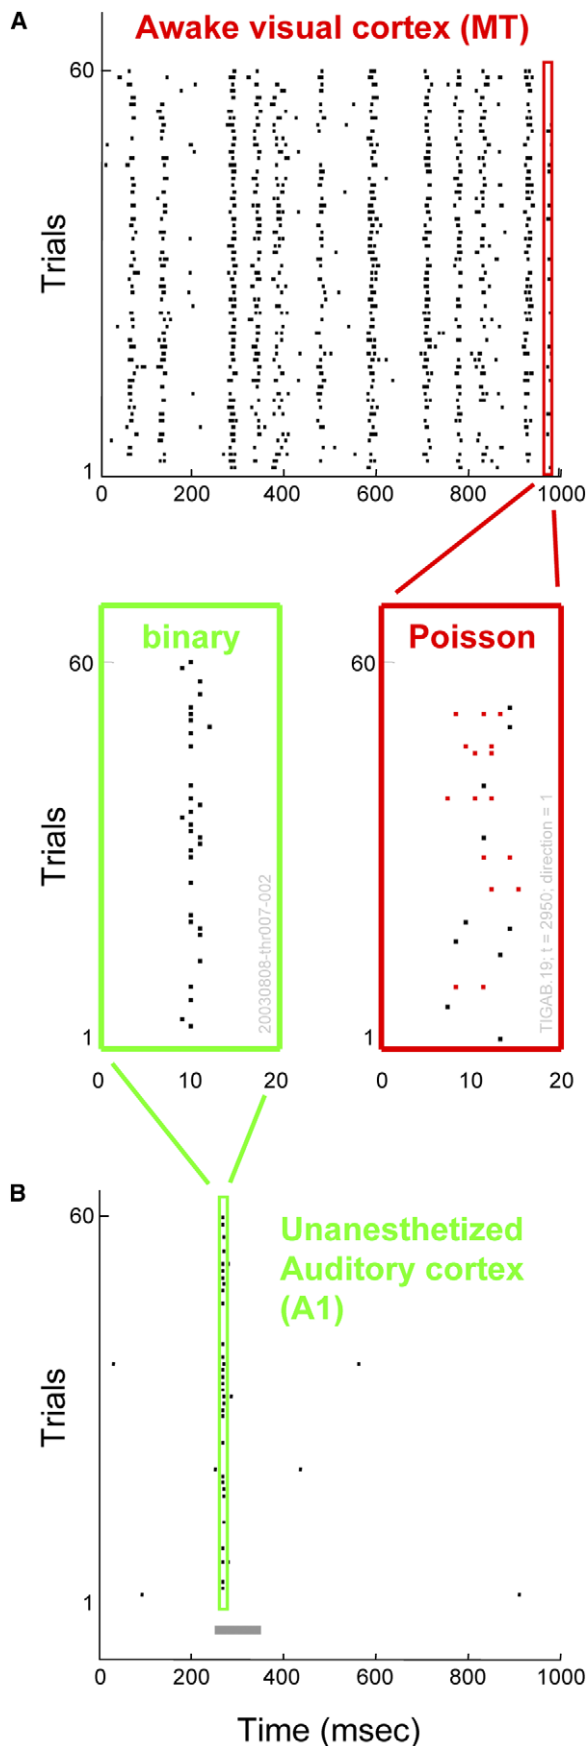
a trial will obey a Poisson distribution, even though the Poisson event rate within each trial varies as the pendulum swings. Accordingly, if the same stimulus is presented on multiple trials, then the distribution of spike counts generated by a model Poisson neuron will follow a Poisson distribution, even if the stimulus—and the resulting stimulus-locked spike rate—vary in time.

A convenient scalar measure of deviation from a Poisson process is the Fano factor, defined as the variance divided by the mean of the spike-count distribution. The Fano factor is 1 for a Poisson process, and 0 for a very regular sequence such as the ticks of a clock. In general, the Fano factor can take any value greater than or equal to zero. The classical rate model predicts that the spike-count distribution, recorded over trials, should be approximately Poisson distributed, with a Fano factor near 1. Of course, the Poisson model is just an idealization; we expect deviations. In fact, some visual cortical spike trains have been shown to be well fit by a mixture of several Poisson processes with different rates, which actually results in a Fano factor even greater than 1 (Wiener and Richmond, 2003).

How well can real neurons be modeled as a Poisson spike generators? Spike-count variability has been studied extensively in visual cortex. The consistent finding is that Fano factors are high (Heggelund and Albus, 1978; Dean, 1981; Tolhurst et al., 1983; Shadlen and Newsome, 1998), although sub-Poisson Fano factors have been observed under conditions in which the firing rate approached the upper limit set by the refractory period (Kara et al., 2000). Even dynamic stimuli specifically designed to elicit rapid modulation of the spike rate (Figure 2A) did not lower the average spike-count variability appreciably (Buracas et al., 1998). Neural responses from visual cortex appeared compatible with the classical rate model, and high spike-count variability seemed to be a universal feature of cortical responses.

Although, of course, universality can never be proven conclusively—the absence of evidence is not evidence of absence—the failure to find counterexamples to high cortical spike-count variability seemed to present a serious challenge to intriguing models of cortical computation that relied on very precise neuronal interactions (Diesmann et al., 1999; Thorpe et al., 2001). Instead, it seemed reasonable to focus on computing in the face of unavoidable neuronal noise.

Surprisingly, it has recently become clear that spike-count variability in auditory cortex can be much lower than had previously been suspected based on reports from other sensory areas (DeWeese et al., 2003). In that study, roughly half of the responses from the recorded population in the auditory cortex of anesthetized rats were found to consist almost entirely of 1 or 0 spikes, with very few multispike responses. In several extreme examples, neurons fired exactly one spike on every trial. In such cases, the Fano factor was thus zero, reflecting the total lack of spike-count variability. For the more typical case in this population—in which the neuron responded with 0 or 1 spikes on every trial, but not more—the Fano factor could be greater than zero, but it was always less than unity and still as low as mathematically possible given the spike rate. Because these neurons fired either 0 or 1 spikes per stimulus, but not more, we have called this “binary spiking” (Figures 2B



and 3). These binary responses thus rule out the Poisson model as a universal description of cortical responses.

However, not all auditory cortical responses exhibit such high reliability. In fact, the same neuron can show binary firing for some stimuli and high variability firing for others. Figure 3 shows an example of this from an unanesthetized animal. Thus, low spike-count variability is not a neuron-specific property but appears to depend on the stimulus as well.

Binary spiking is more common in the anesthetized than the unanesthetized animal, in part because transient spiking is more common in the anesthetized preparation (Evans and Whitfield, 1964; DeWeese et al., 2003; Wang et al., 2005). Binary spiking is not, however, an artifact of anesthesia. It is clear that some neurons in the unanesthetized cortex can fire in a binary fashion (Figures 2B and 3). Other clear examples of binary spiking can be found in the transient responses from a neuron recorded in the unanesthetized marmoset monkey (Figure 4A from Barbour and Wang [2003]) and in the alert cat (Figure 4A from Chimoto et al. [2002]). Although not highlighted by the authors, a close examination (e.g., by enlarging the figure in PDF format) of these figures, which each depict stimulus-evoked rasters from a well-isolated single unit in the primary auditory cortex, reveal that responses consist of exactly 0 or 1—but never more than 1—spikes on each trial. These data demonstrate that binary spiking is not restricted to the anesthetized state and that it can be observed in the awake preparation.

The theoretical significance of binary spiking is that it provides an existence proof—a counterexample to the presumed universality of high spike-count variability in cortex. It is clear that not all stimuli elicit low-variability responses, just as not all stimuli elicit high-firing rates. In the same way that stimuli can be specially selected—

Figure 2. Temporal Jitter Can Be Low in Both Visual and Auditory Cortex, but Binary Spiking Has Only Been Reported in Auditory Cortex

(A) An example motion-sensitive neuron (from Buracas et al. [1998]) recorded in middle temporal (MT) area of a monkey’s visual cortex exhibits relatively low temporal jitter (5–10 ms) but high spike-count variability. The neuron was responding to 60 presentations of the same visual stimulus while the alert monkey fixated the direction of its gaze. The expanded view of the stimulus-locked response (red rectangle) reveals several two and three spike responses among the 60 trials (multispike responses are indicated by red hash marks). The spike-count variability was high for this example (Fano factor ~ 1.4 ; see text), even greater than one would expect from a highly stochastic Poisson process.

(B) An example of a sound-sensitive neuron from the primary auditory cortex (area A1) of an unanesthetized rat exhibits both low temporal jitter (2–3 ms) and low spike-count variability (“binary spiking”). On each trial, a 100 ms, 60 dB tone (gray bar) between 8 and 13 kHz was presented. The expanded view shows that every stimulus-locked response (green rectangle) consists of either one or zero spikes, but not more (note that there is no choice for the duration or placement of the red rectangle in [A] that would have resulted in a binary set of responses with more than a small fraction of trials containing a spike). Spike-count variability for a binary set of responses such as this is as low as mathematically possible (in this case, Fano factor = 0.49) given the firing rate (0.52 spikes/trial). Other examples of binary spiking, recorded in the auditory cortex of the unanesthetized marmoset and the alert cat, can be found in Figure 4A from Barbour and Wang [2003] and Figure 4A from Chimoto et al. [2002], respectively.

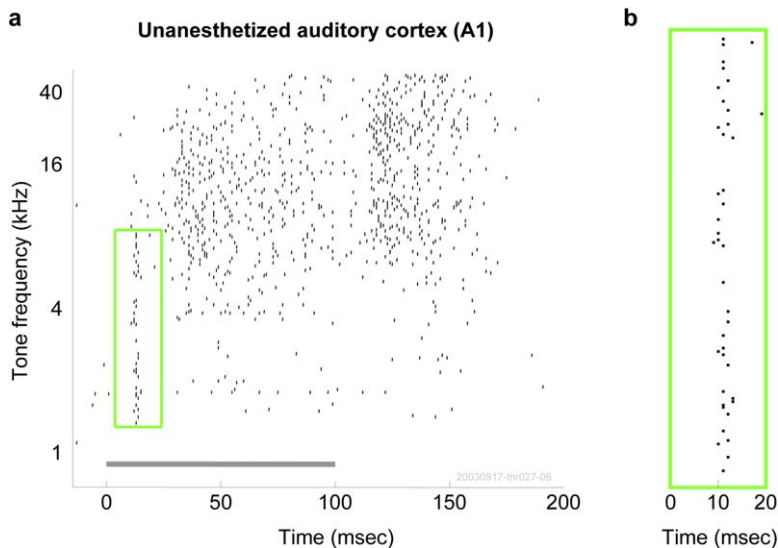


Figure 3. Some Neurons in the Unanesthetized Auditory Cortex Can Respond Reliably and Transiently to Some Acoustic Stimuli but in a Sustained and More Variable Fashion to Other Sounds

(A) In this example, the same neuron responds with low temporal jitter and low spike-count variability to the onset of 100 ms duration tones (gray bar at bottom) in the low-to-middle frequency range. These rasters consist of responses to four repeats each from 64 different frequency tones logarithmically spaced between 1 and 40 kHz; all tones were 60 dB.

(B) As is evident in the expanded view shown here, the transient response after the onset of each tone is binary—it consists of one or zero spikes, but never more. For mid-to-high frequencies, the neuron produces sustained firing, both in response to the onset and the termination of the tone. Thus, we use the descriptors “binary” or “sustained” to refer to a given neuron’s response to a particular stimulus, rather than as a general attribute of the neuron.

“optimized”—to maximize the firing rate (deCharms et al., 1998; Wang et al., 2005), so too could they be optimized to maximize response reliability. The fact that we as experimenters can find examples of stimulus-locked binary spiking raises the possibility that from the animal’s point of view, spike count is much more precisely controlled than has usually been assumed.

Rate versus Temporal Coding

We have so far drawn a sharp distinction between rate coding—the view that information about a stimulus is only represented in a neuron’s firing rate—and “temporal coding”—a less precisely defined term often used to encompass most of the alternatives, but particularly those in which the timing of individual spikes, or the temporal correlations among spikes, encodes information. The rate-coding model is often treated as the null hypothesis; the question is usually formulated in terms of whether there is additional information about the stimulus available beyond that conveyed by the spike rate. From the point of view of the experimenter, if cortical responses were always well described by Poisson processes, then all the information would be available in the rate (but at a cost: estimating the rate from the noisy spike train requires averaging over time or over a population of neurons). However, we have seen (Figures 2B and 3) that auditory cortex responses can be more reliable than Poisson, opening up the possibility that information is represented by features of the spike train not fully captured by the time-varying firing rate.

Does this mean that the auditory cortex necessarily uses a temporal code? No. The notion that these data alone can resolve the nature of the neural code is based on a fallacy: as experimenters, we can determine with certainty whether information about a stimulus is available in a neural spike train. It is much more difficult to establish whether and how this information is used by other neurons (but see Cohen and Newsome [2004]) and by the organism. Indeed, these issues sometimes verge on the philosophical; for example, if the activity of an inhibitory neuron can be used to reconstruct a stimulus faithfully, but if that neuron has only local projec-

tions, in what sense is it accurate to say that that neuron’s representation of the stimulus is used by the organism? Analyzing the statistics of neural responses to repeated presentations of stimuli can place some constraints on the possible coding schemes employed by the cortex, but the definitive tests will most likely require a combined approach involving both cortical physiology and behavior.

Representations in Auditory Cortex

In the discussion so far, we have repeatedly stressed the extent to which the “experimenter’s point of view” colors the interpretation of physiological data. The cortical physiologist traditionally manipulates sensory stimuli and records neuronal responses. From these responses, she sets herself the task of inferring the representation of acoustic stimuli in the auditory cortex. In this paradigm, the independent variable under the experimenter’s control—the stimulus—must carry the burden of explaining everything about the observed neuronal response. In this paradigm, other factors controlling the neuronal response remain inaccessible to the experimenter, even if they are accessible and important from the animal’s point of view. These other factors include: information carried from other auditory and nonauditory cortical areas, state-dependent modulation because of variable levels of arousal and fear, task-dependent modulation because of attention and reward expectation, and many others.

In what follows, we will briefly discuss the representation of auditory stimuli from the experimenter’s point of view and then conclude by discussing one form of nonsensory modulation that has been studied experimentally.

Optimal Stimuli

One of the most powerful approaches to studying representations in the visual cortex has been to search for “optimal stimuli” that drive neurons to fire at high rates. By using this approach, optimal stimuli have been identified for neurons in areas V1, MT, and other visual and nonvisual areas. The search for optimal stimuli in auditory

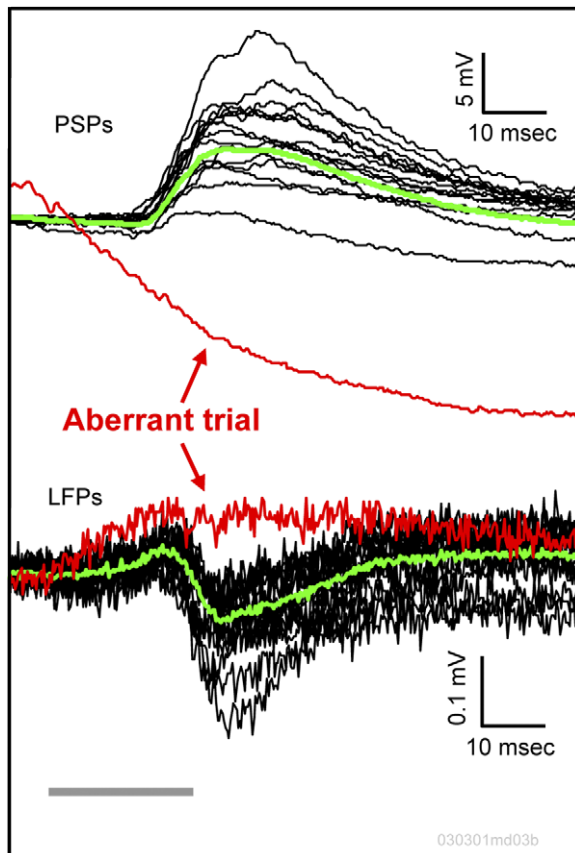


Figure 4. Stimulus-Evoked Responses in Auditory Cortical Neurons Can Be Strongly Affected by Stimulus-Independent, Circuit-Wide Activity

In this example (from DeWeese and Zador [2004]), the membrane potential of an auditory cortical neuron was recorded by whole-cell patch clamp methods in the intact, anesthetized rat during repeated presentations of a 25 ms duration pure tone (65 dB, 31 kHz tone; gray bar at bottom). On most trials, the postsynaptic potentials (PSPs; black traces at top) were well described as rescaled versions of the mean response (upper green trace), but one PSP (top red trace) was qualitatively different. Simultaneously recorded local field potentials (LFPs; bottom) from a second nearby (~0.5 mm) electrode followed the same pattern; the LFP (lower red trace) corresponding to the aberrant PSP was also an outlier, indicating a source of variability shared between the neuron recorded in whole-cell mode and the population of neurons contributing to the LFP. Shared fluctuations such as this contributed significantly to the total neural variability of the majority of neurons in this study (DeWeese and Zador, 2004). Because these fluctuations were shared across many neurons, they could reflect trial-to-trial differences in what computation the cortex was performing, rather than an unstructured degradation of the sensory signal. More concretely, the fact that these aberrant trials were shared by a population of neurons implied that the fluctuations in membrane potential were not due to sources of noise inside the neuron. Thus, nonauditory influences over the synaptic drive to the neuron can result in trial-to-trial fluctuations that might represent a “signal” from the perspective of the animal, rather than noise resulting from biophysical processes private to the neuron.

cortex has been hampered by an important difference between auditory and visual cortex. In the visual cortex, sustained high-rate (“optimal”) cortical responses can be elicited in both the awake and anesthetized preparation, whereas in the auditory cortex, only transient responses are observed under most forms of anesthesia;

Table 1. Comparison of Response Types Observed in Anesthetized and Unanesthetized Auditory and Visual Cortex

Preparation	Sustained	Transient	
		Nonbinary	Binary
Visual cortex (unanesthetized)	Observed	Observed	Not observed
Visual cortex (anesthetized)	Observed	Observed	Not observed
Auditory cortex (unanesthetized)	Observed	Observed	Observed
Auditory cortex (anesthetized)	Not observed	Observed	Observed

sustained responses in auditory cortex are typically observed only when the subject is unanesthetized. Note that we are distinguishing between sustained and transient responses here. This is quite separate from the distinction between transient and binary responses; all binary responses are transient, but not all transient responses are binary. Both binary and nonbinary transient responses are observed in both the awake and anesthetized auditory cortex (see Table 1).

Many of the earliest recordings in auditory cortex were performed in the unanesthetized preparation (Hubel et al., 1959; Evans and Whitfield, 1964; Whitfield and Evans, 1965; Goldstein et al., 1968). From these recordings, it was clear that at least some neurons could respond in a sustained fashion to simple stimuli such as tones and sweeps. In one study (Evans and Whitfield, 1964), for example, tones elicited a sustained responses in about 24% of neurons in primary auditory cortex of the unanesthetized cat, whereas they elicited a transient response in about 9% of units; in some of the remainder, they elicited more complex responses such as suppression, and in 45%, tones elicited no response at all. (About 3% of neurons in auditory cortex responded only to visual stimulation in this study; cf., Brosch et al. [2005]). Unfortunately, the authors did not define these categories quantitatively. However, recent results with more rigorous definitions (T.H., M.R.D., and A.M.Z., unpublished data) support the general view that auditory neurons in the unanesthetized preparation have highly heterogeneous response properties to tones.

With the recent resurgence of work in the unanesthetized preparation, a wider variety of stimuli have been identified that can drive neurons to fire in a sustained fashion. In some cases, researchers have used stimulus ensembles designed as auditory analogs of visual stimuli—edges and drifting gratings—that have been successful driving neurons in area V1 (Kowalski et al., 1996a, 1996b; deCharms et al., 1998; Fishbach et al., 2001, 2003). In addition, online stimulus optimization is sometimes used to adapt the stimulus properties in real time to find stimuli that drives neurons best (deCharms et al., 1998; Barbour and Wang, 2003; O’Connor et al., 2005; Wang et al., 2005). These recent studies greatly clarify and extend the earlier evidence showing that tones and other stimuli can elicit sustained activity in the auditory cortex of unanesthetized animals.

When interpreting these optimal stimulus experiments, three issues must be distinguished: (1) Is there an optimal stimulus for every neuron in auditory cortex? This is a difficult question to address experimentally

because it can never be answered definitively in the negative—failure to find a stimulus capable of driving a neuron at a high rate does not prove that such a stimulus does not exist. Notwithstanding this difficulty, an exciting recent study (Wang et al., 2005) provides an indication that there may indeed exist an optimal stimulus for every neuron: as many as 72% fired in a sustained fashion to the preferred temporally modulated stimulus, almost three times as high a fraction as, for example, in earlier work (Evans and Whitfield, 1964). However, sustained firing in this study was defined on the basis of an absolute level—five spikes/second—rather than relative to spontaneous firing. Moreover, it is not clear whether a stimulus eliciting a firing rate of five spikes/second should be deemed “optimal” (the authors use the more conservative term “preferred stimulus”). Thus, although it may be appealing from a theoretical point of view to posit the existence of an optimal stimulus for every neuron, it must for now be considered an open question.

(2) For any given stimulus, is there at least one neuron that fires vigorously? And (3), is an optimal (i.e., high-firing) response necessary for perception, decisions, and behavior? These questions go to the heart of the representation issue. If there is at least one well-driven neuron in auditory cortex for any acoustic stimulus—if optimal responses cover the auditory space fully—then it raises the possibility that these well-driven responses are sufficient for perception and behavior. An extreme view is that only the neurons responding optimally are important. However, if there exist acoustic stimuli that do not elicit sustained firing in any cortical neuron, but the organism can nevertheless perceive and act upon it, then this would imply that such sustained firing is not necessary. Unfortunately, little is known at present about the role of sustained firing in perception and behavior.

Nonsensory Representations in Auditory Cortex

We close the discussion of auditory representations by asking more broadly what it is that activity in the auditory cortex represents. Of course, the auditory cortex must represent information about acoustic stimuli. But the majority of the projections to the auditory cortex are not feed-forward connections from the thalamus carrying acoustic information from the periphery; most of the inputs originate in other cortical areas. To the extent that the activity in these other areas is structured rather than random, then the activity in auditory cortex must in some sense represent the activity—albeit in a processed form—from these other sensory and nonsensory brain areas.

Although we know very little about how such information is represented, there is ample evidence that a variety of nonsensory processes can affect neuronal activity in the auditory cortex. For example, many neurons in auditory cortex are strongly affected by stimulus-independent fluctuations in the synaptic input they receive because of concerted activity involving large populations of neurons across the auditory cortex (Figure 4) (DeWeese and Zador, 2004), similar to what has been reported in the visual cortex (Arieli et al., 1995; Buracas et al., 1998). Moreover, in the unanesthetized animal, activity in the auditory cortex is different in different states of arousal (Edeline et al., 2001) and is modulated by emotional states such as fear (Quirk et al., 1997) and by non-

auditory cues and task contingencies (Brosch et al., 2005), and stimulus representations show learning-dependent plasticity (King et al., 2000; Ohl et al., 2001; Sakai and Suga, 2001; Rutkowski and Weinberger, 2005).

One powerful form of nonsensory influence in the auditory cortex is the modulation of responses by attention. Indeed, in one of the earliest single unit studies of auditory cortex, Hubel and colleagues proposed attention as an explanation for why in previous studies, so many neurons could not be driven by acoustic stimulation: “It is not easy to understand why the auditory cortex, in the anesthetized or intact cat, should be populated with so many cells that fail to respond to auditory stimuli. Perhaps these cells become activated only when certain other conditions are simultaneously met. Thus, from our data one may conclude that the neural processes responsible for attention play an important role in determining whether or not a given acoustic stimulus proves adequate” (Hubel et al., 1959).

More recent studies in monkeys (Hoehnerman et al., 1976) and ferrets (Fritz et al., 2003) under better behavioral control confirm the importance of attention in modulating responses. In some cases, attention can completely gate the neural response in primary auditory cortex (Hoehnerman et al., 1976). Based on the profound influence that attention has been shown to hold over auditory cortex in the few published studies of this type, it appears that the auditory cortex could prove ideal for the study of the mechanisms underlying the dynamical control of information flow during cortical sensory processing.

The importance of attention and other nonauditory signals in modulating neural responses in auditory cortex provides a possible resolution of the representational-bandwidth conundrum. Perhaps a substantial fraction of the representational bandwidth in the auditory cortex is devoted to the representation of nonauditory signals. Although we currently tend to think about the representation of such nonauditory signals as exerting only relatively coarse modulation of sensory responses, it is possible that these representations are actually quite rich.

Conclusion

Why does the auditory cortex have so many more neurons than the auditory nerve? Although many neuroscientists consider the answer so obvious—“for computation”—that the question is not worth asking, the consensus breaks down once the answer is specified in more detail. One hypothesis holds that the extra neurons are there to overcome the neuronal noise that necessarily arises in cortical representations. A second hypothesis holds that the extra neurons aid computation by providing multiple convenient rerepresentations of the same signal in different forms. A third hypothesis holds that the auditory cortex must integrate both auditory and nonauditory information.

In this review, we have shown that neurons in auditory cortex can be very reliable, raising the possibility that neuronal noise is not a necessary consequence of cortical circuitry. This leads us, in our experimental work, to join those focusing on the second and third hypotheses—and in particular, on characterizing the powerful modulation that nonauditory signals can exert on responses in auditory cortex. Ultimately, resolving the full structure of these nonsensory signals will

require experiments measuring neural activity within the context of well-controlled behavioral paradigms.

Acknowledgments

The authors would like to thank Adam Kepecs as well as all of the members of the Zador laboratory for many useful discussions and comments on the manuscript. T.H. is an Engelhorn Scholar of the Watson School of Biological Sciences. We gratefully acknowledge support from the National Institutes of Health, the Swartz Foundation, the Mathers Foundation, and the Sloan Foundation.

References

- Arieli, A., Shoham, D., Hildesheim, R., and Grinvald, A. (1995). Coherent spatiotemporal patterns of ongoing activity revealed by real-time optical imaging coupled with single-unit recording in the cat visual cortex. *J. Neurophysiol.* *73*, 2072–2093.
- Bair, W. (1999). Spike timing in the mammalian visual system. *Curr. Opin. Neurobiol.* *9*, 447–453.
- Bair, W., and Koch, C. (1996). Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comput.* *8*, 1185–1202.
- Barbour, D.L., and Wang, X. (2003). Auditory cortical responses elicited in awake primates by random spectrum stimuli. *J. Neurosci.* *23*, 7194–7206.
- Berry, M.J., 2nd, and Meister, M. (1998). Refractoriness and neural precision. *J. Neurosci.* *18*, 2200–2211.
- Brosch, M., Selezneva, E., and Scheich, H. (2005). Nonauditory events of a behavioral procedure activate auditory cortex of highly trained monkeys. *J. Neurosci.* *25*, 6797–6806.
- Buracas, G.T., Zador, A.M., DeWeese, M.R., and Albright, T.D. (1998). Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* *20*, 959–969.
- Calvin, W.H., and Stevens, C.F. (1968). Synaptic noise and other sources of randomness in motoneuron interspike intervals. *J. Neurophysiol.* *31*, 574–587.
- Chimoto, S., Kitama, T., Qin, L., Sakayori, S., and Sato, Y. (2002). Tonal response patterns of primary auditory cortex neurons in alert cats. *Brain Res.* *934*, 34–42.
- Cohen, M.R., and Newsome, W.T. (2004). What electrical microstimulation has revealed about the neural basis of cognition. *Curr. Opin. Neurobiol.* *14*, 169–177.
- Dean, A.F. (1981). The variability of discharge of simple cells in the cat striate cortex. *Exp. Brain Res.* *44*, 437–440.
- deCharms, R.C., and Merzenich, M.M. (1996). Primary cortical representation of sounds by the coordination of action-potential timing. *Nature* *381*, 610–613.
- deCharms, R.C., Blake, D.T., and Merzenich, M.M. (1998). Optimizing sound features for cortical neurons. *Science* *280*, 1439–1443.
- DeWeese, M.R., and Zador, A.M. (2004). Shared and private variability in the auditory cortex. *J. Neurophysiol.* *92*, 1840–1855.
- DeWeese, M.R., Wehr, M., and Zador, A.M. (2003). Binary spiking in auditory cortex. *J. Neurosci.* *23*, 7940–7949.
- Diesmann, M., Gewaltig, M.O., and Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature* *402*, 529–533.
- Edeline, J.M., Dutriex, G., Manunta, Y., and Hennevin, E. (2001). Diversity of receptive field changes in auditory cortex during natural sleep. *Eur. J. Neurosci.* *14*, 1865–1880.
- Eggermont, J.J. (2000). Sound-induced synchronization of neural activity between and within three auditory cortical areas. *J. Neurophysiol.* *83*, 2708–2722.
- Evans, E.F., and Whitfield, I.C. (1964). Classification of unit responses in the auditory cortex of the unanaesthetized and unrestrained cat. *J. Physiol.* *171*, 476–493.
- Fishbach, A., Nelken, I., and Yeshurun, Y. (2001). Auditory edge detection: a neural model for physiological and psychoacoustical responses to amplitude transients. *J. Neurophysiol.* *85*, 2303–2323.
- Fishbach, A., Yeshurun, Y., and Nelken, I. (2003). Neural model for physiological responses to frequency and amplitude transitions uncovers topographical order in the auditory cortex. *J. Neurophysiol.* *90*, 3663–3678.
- Fritz, J., Shamma, S., Elhilali, M., and Klein, D. (2003). Rapid task-related plasticity of spectrotemporal receptive fields in primary auditory cortex. *Nat. Neurosci.* *6*, 1216–1223.
- Goldstein, M.H., Jr., Hall, J.L., 2nd, and Butterfield, B.O. (1968). Single-unit activity in the primary auditory cortex of unanesthetized cats. *J. Acoust. Soc. Am.* *43*, 444–455.
- Heggelund, P., and Albus, K. (1978). Response variability and orientation discrimination of single cells in striate cortex of cat. *Exp. Brain Res.* *32*, 197–211.
- Heil, P. (2004). First-spike latency of auditory neurons revisited. *Curr. Opin. Neurobiol.* *14*, 461–467.
- Hoehnerman, S., Benson, D.A., Goldstein, M.H., Jr., Heffner, H.E., and Hienz, R.D. (1976). Evoked unit activity in auditory cortex of monkeys performing a selective attention task. *Brain Res.* *117*, 51–68.
- Hopfield, J.J., and Brody, C.D. (2001). What is a moment? Transient synchrony as a collective mechanism for spatiotemporal integration. *Proc. Natl. Acad. Sci. USA* *98*, 1282–1287.
- Hubel, D.H., Henson, C.O., Rupert, A., and Galambos, R. (1959). Attention units in the auditory cortex. *Science* *129*, 1279–1280.
- Kara, P., Reinagel, P., and Reid, R.C. (2000). Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* *27*, 635–646.
- King, A.J., Parsons, C.H., and Moore, D.R. (2000). Plasticity in the neural coding of auditory space in the mammalian brain. *Proc. Natl. Acad. Sci. USA* *97*, 11821–11828.
- Kowalski, N., Depireux, D.A., and Shamma, S.A. (1996a). Analysis of dynamic spectra in ferret primary auditory cortex. I. Characteristics of single-unit responses to moving ripple spectra. *J. Neurophysiol.* *76*, 3503–3523.
- Kowalski, N., Depireux, D.A., and Shamma, S.A. (1996b). Analysis of dynamic spectra in ferret primary auditory cortex. II. Prediction of unit responses to arbitrary dynamic spectra. *J. Neurophysiol.* *76*, 3524–3534.
- Lisman, J.E., and Idiart, M.A. (1995). Storage of 7 +/- 2 short-term memories in oscillatory subcycles. *Science* *267*, 1512–1515.
- Mazurek, M.E., and Shadlen, M.N. (2002). Limits to the temporal fidelity of cortical spike rate signals. *Nat. Neurosci.* *5*, 463–471.
- Micheyl, C., Tian, B., Carlyon, R.P., and Rauschecker, J.P. (2005). Perceptual organization of tone sequences in the auditory cortex of awake macaques. *Neuron* *48*, 139–148.
- Nelken, I., Fishbach, A., Las, L., Ulanovsky, N., and Farkas, D. (2003). Primary auditory cortex of cats: feature detection or something else? *Biol. Cybern.* *89*, 397–406.
- Nelken, I., Chechik, G., Msrac-Flogel, T.D., King, A.J., and Schnupp, J.W. (2005). Encoding stimulus information by spike numbers and mean response time in primary auditory cortex. *J. Comput. Neurosci.* *19*, 199–221.
- O'Connor, K.N., Petkov, C.I., and Sutter, M.L. (2005). Adaptive stimulus optimization for auditory cortical neurons. *J. Neurophysiol.*, in press. Published online August 31, 2005. 10.1152/jn.00046.2005
- Ohl, F.W., Scheich, H., and Freeman, W.J. (2001). Change in pattern of ongoing cortical activity with auditory category learning. *Nature* *412*, 733–736.
- O'Keefe, J., and Recce, M.L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* *3*, 317–330.
- Osborne, L.C., Bialek, W., and Lisberger, S.G. (2004). Time course of information about motion direction in visual area MT of macaque monkeys. *J. Neurosci.* *24*, 3210–3222.
- Petersen, R.S., Panzeri, S., and Diamond, M.E. (2002). Population coding in somatosensory cortex. *Curr. Opin. Neurobiol.* *12*, 441–447.

- Pouget, A., Dayan, P., and Zemel, R. (2000). Information processing with population codes. *Nat. Rev. Neurosci.* *1*, 125–132.
- Quirk, G.J., Armony, J.L., and LeDoux, J.E. (1997). Fear conditioning enhances different temporal components of tone-evoked spike trains in auditory cortex and lateral amygdala. *Neuron* *19*, 613–624.
- Reynolds, J.H., and Chelazzi, L. (2004). Attentional modulation of visual processing. *Annu. Rev. Neurosci.* *27*, 611–647.
- Rutkowski, R.G., and Weinberger, N.M. (2005). Encoding of learned importance of sound by magnitude of representational area in primary auditory cortex. *Proc. Natl. Acad. Sci. USA* *102*, 13664–13669.
- Sakai, M., and Suga, N. (2001). Plasticity of the cochleotopic (frequency) map in specialized and nonspecialized auditory cortices. *Proc. Natl. Acad. Sci. USA* *98*, 3507–3512.
- Shadlen, M.N., and Newsome, W.T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* *18*, 3870–3896.
- Softky, W.R., and Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J. Neurosci.* *13*, 334–350.
- Stecker, G.C., and Middlebrooks, J.C. (2003). Distributed coding of sound locations in the auditory cortex. *Biol. Cybern.* *89*, 341–349.
- Sukov, W., and Barth, D.S. (2001). Cellular mechanisms of thalamically evoked gamma oscillations in auditory cortex. *J. Neurophysiol.* *85*, 1235–1245.
- Thorpe, S., Delorme, A., and Van Rullen, R. (2001). Spike-based strategies for rapid processing. *Neural Netw.* *14*, 715–725.
- Tolhurst, D.J., Movshon, J.A., and Dean, A.F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* *23*, 775–785.
- Vapnik, V. (1999). *The Nature of Statistical Learning Theory* (New York, NY: Springer-Verlag).
- Wang, X., Lu, T., Snider, R.K., and Liang, L. (2005). Sustained firing in auditory cortex evoked by preferred stimuli. *Nature* *435*, 341–346.
- Whitfield, I.C., and Evans, E.F. (1965). Responses of auditory cortical neurons to stimuli of changing frequency. *J. Neurophysiol.* *28*, 655–672.
- Wiener, M.C., and Richmond, B.J. (2003). Decoding spike trains instant by instant using order statistics and the mixture-of-Poissons model. *J. Neurosci.* *23*, 2394–2406.