Representation of Sounds in Auditory Cortex of Awake Rats

A Dissertation Presented

by

Tomáš Hromádka

Engelhorn Scholar

to

The Watson School of Biological Sciences

at

Cold Spring Harbor Laboratory

in Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

in

Biological Science

March 2007

Contents

Acknowledgments

would probably cease to exist without him. Anita Matthes and Becky Dong deserve my thanks for their excellent assistance with various interesting administrative tasks.

A very special thanks goes to the members of my thesis committee: Carlos Brody, Zach Mainen, Mitya Chklovskii, and Bill Tansey. I have to thank them for their patience and support. Bill Tansey also deserves my deepest gratitude and admiration for being my academic mentor at the Watson School. I would also like to thank Kenneth Harris for taking time out from his busy schedule to serve as my external examiner.

I had enough opportunities during the past several years to realize how *extraordinarily* lucky I am to have Eva, my wife and best friend, and Samo and Timo by my side. Without them, my life would be incomplete. I also need to thank my mother and father who always supported me in everything I was doing.

Finally, I recognize that my work would not be possible without the generous support provided by Curt Engelhorn and the Watson School of Biological Sciences, under the leadership of Lilian Clark, and Winship Herr.

List of Figures

Chapter 1

Introduction

The brain is the most complex computational device known to Man. Not only does it mediate our orientation in both external (physical) and internal worlds, but—even more astonishingly—the brain enables study of itself. Yet, this amazing device is composed of only a limited set of neurons and their connections. Neuroscience—in all its disguises—approaches the device with the belief that knowledge of function and structure of these basic building blocks will eventually help us understand how the brain understands itself.

Brain's connection to the external world is provided by sensory systems. External inputs (light, sound, touch, etc.) are detected by sensory receptors, translated into internal representations, which in turn are interpreted into percepts, and eventually lead to (motor) actions. Such transformations—especially in the beginning of the processing chain—are the main topics of sensory neuroscience.

The auditory branch of sensory neuroscience tries to explain how the brain (i.e. us) makes sense of the acoustic environment. Unfortunately for the reader, despite centuries of interest (Cooper, 1801; Cooper and Home, 1800; Holder, 1668), auditory neuroscience lacks famous clinical cases, which would make for a catchy introduction. Thus, rustling leaves attracting the attention of a barn owl (Winer et al., 2005) cannot match an iron rod shattering the prefrontal cortex of Phineas Gage (Harlow, 1848), or the profound memory loss of the truly unfortunate patient H. M (Scoville and Milner, 1957). Nevertheless, as we are constantly bombarded by sounds, we effortlessly extract necessary information, whether surrounded by fellow colleagues at a cocktail party, or hunting for prey (being a barn owl). When rapid changes in air pressure (sounds) hit our ears, sounds are transformed in cochlea into series of action po-

tentials, which, after several processing steps in the brainstem and thalamus, reach the primary auditory cortex, the acoustic gateway of cerebral cortex.

The electrophysiological approach to study the function of neurons in the auditory system is to record the electrical activity of small numbers of neurons, while probing the sensory receptors with various sounds. The recorded electrical activity then serves as an estimate of the internal representation of the particular acoustic stimulus. The activity of neurons in the primary auditory cortex is the main topic of this thesis. Before the main chapters, however, we step back and offer a brief—necessarily idiosyncratic—summary of some of the key aspects of neuronal activity in auditory cortex.

1.1 Auditory cortex—pariah of sensory cortices

After more than six decades of research (Bremer and Bonnet, 1949; Tunturi, 1944; Woolsey and Walzl, 1942) there is no consensus on how auditory cortex is defined (Read et al., 2002). Although many would agree on a definition based on the pattern of input connections and the type of sensory inputs provided by the thalamus, it is both astonishing and confusing that sixty years of physiological research were not enough for a comprehensive characterization of even the primary auditory cortex (Ehret, 1997).

Yet, the neural activity in the auditory system seems to be "simple" enough to study. Small variations in sound pressure are transformed at the cochlea into spikes that traverse the auditory nerve, pass several subcortical auditory stations (cochlear nucleus, superior olive, inferior colliculus, thalamus), and eventually reach auditory cortex. To study neural processing in this pathway, one might begin by studying the transformation of sound into spikes at the auditory nerve, and then follow the neural representation of sound through the various subcortical stations. Eventually, one would have to grapple with the more complex representations emerging at the higher stations, but the whole process seems straightforward.

An analogous research paradigm dominated the study of the *visual cortex* for several decades since Hubel and Wiesel's seminal work on the primary visual cortex (Hubel and Wiesel, 1977), and remains important even today. Hubel and Wiesel showed that many neurons in the primary visual area could be driven to fire at high rates in response to oriented bars appropriately positioned in the visual field. This eventually led to an appealing general model of sensory processing in which visual recognition is achieved by constructing an appropriate series of representations of the visual world

from the simpler representations at the previous stage, a view supported by anatomical evidence for an anatomical hierarchy of connections in visual cortex (Felleman and Essen, 1991).

Since the early start of single-unit recordings in the auditory system (Davies et al., 1956; Galambos and Davis, 1943) many acoustic stimuli have been used to probe neurons in auditory cortex, including tones (Moshitch et al., 2006), sweeps (Orduña et al., 2001; Zhang et al., 2003), ripples (Kowalski et al., 1996a,b), TORCS (Elhilali et al., 2004) and natural sounds (Bar-Yosef et al., 2002; Machens et al., 2004). However, no acoustic feature analogous to the oriented bar has emerged (Nelken, 2004; Nelken et al., 2003).

The universally agreed upon property of primary auditory cortex is its tonotopical organization, which replicates decomposition of sounds into separate frequency components at the cochlea (Merzenich et al., 1973, 1975; Nicholls et al., 2001; Schreiner, 1995; Schreiner et al., 2000), and is preserved at the subcortical processing stations as well. Because tonotopy is such a robust feature, many studies have tried to relate responsiveness to various features to tonotopical organization. Among features which have been mapped relative to isofrequency columns are response latency (Mendelson et al., 1997), intensity threshold (Phillips et al., 1994; Sutter and Schreiner, 1995), operating range (Heil et al., 1992a, 1994), frequency modulation rate (Heil et al., 1992b; Mendelson et al., 1993), binaural response type (Imig and Adrián, 1977), sharpness of frequency tuning (Heil et al., 1992b; Schreiner and Mendelson, 1990), etc.; see also (Ehret, 1997; Read et al., 2002; Schreiner et al., 2000) and references therein for more details. Overall, the overlap of various parameter mappings parallel and orthogonal to the tonotopic axis is usually described as the modular organization of auditory cortex (Schreiner et al., 2000), lacking the elegance and effectiveness of visual cortex organization.

Auditory cortex thus remains a pariah among her (non-chemical) sensory sisters, visual and somatosensory cortices. Our knowledge of basic processing in auditory cortex still does not match the detailed descriptions of processing (understanding of structure and function) of visual and somatosensory (especially barrel) cortices (Callaway, 1998; Douglas and Martin, 2004). Although the majority of auditory studies mentioned in this section were conducted in anesthetized animals, the situation in awake preparations seems to be even more complicated.

1.2 Awake auditory cortex

Most of what is known about the activity of neurons in the auditory cortex comes from studies in *anesthetized* preparations. However, the presence of anesthetic drugs (Gaese and Ostwald, 2001; Populin, 2005), as well as depth of anesthesia (Armstrong-James and George, 1988) influence neuronal activity, with different drugs having different effects on neuronal activity (Cheung et al., 2001; Wehr and Zador, 2005). In addition, studying neuronal activity during behavior, which requires active participation of the animal is, of course, impossible in anesthetized preparations.

Studies of neuronal activity in the auditory cortex of *awake* (unanesthetized) animals have a long history (Evans and Whitfield, 1964; Gerstein and Kiang, 1964; Goldstein et al., 1968; Hubel et al., 1959; Katsuki et al., 1962, 1958; Whitfield, 1957), and were recently resurrected by several groups (see for example Barbour and Wang, 2003; Chimoto et al., 2002; deCharms et al., 1998; Elhilali et al., 2004; Gaese and Ostwald, 2003; Recanzone, 2000; Wang et al., 2005).

With the resurgence of work in the awake preparations in the last decade, many researchers have emphasized the rich repertoire of neuronal responses in awake animals, including sustained responses to sounds (Barbour and Wang, 2003; Gaese and Ostwald, 2003; Liang et al., 2002; Recanzone, 2000; Wang et al., 2005), which are a sharp contrast to the typically transient sound-evoked responses observed in anesthetized preparations (DeWeese et al., 2003; Doron et al., 2002; Heil, 1997; Phillips and Irvine, 1981). Thus, part of the auditory community emphasises responses observed mostly in awake preparations: "Under anesthetized conditions, phasic responses are the typical discharge patterns in the auditory cortex, even for responses to complex vocalizations. A distinct characteristic of neural responses in *unanesthetized auditory cortex is sustained firing,* a property that stands in sharp contrast to barbiturate-anesthetized cortex." (Liang et al., 2002, emphasis added). Another part of auditory community argues that neuronal activity is almost completely described as seen in anesthetized preparations: "This is not to dispute the existence of sustained excitatory responses to tone or noise burst stimuli as far rostral in the central auditory system as the auditory cortex, especially in unanesthetized animals...*The onset response is, however, arguably the dominant feature* of central neural responses, especially in the forebrain." (Phillips et al., 2002, emphasis added).

Such debate has surprisingly been present (and unresolved) in the auditory community for the past fifty years, since Robert Galambos first stated: "The search for correlates for the steady state [re-

sponses toward tones] at the cortical level has to a large extent been unsuccessful, and for such information as we do possess we have to thank those who have used unanesthetized preparations." (Galambos, 1954). However, the fact that a particular neuron can be driven at high rates when probed with a particular sound tells us little about how that particular sound (or sound in general) is represented across a population of neurons. Indeed, focusing on the rather special types of responses, one is almost guaranteed a biased view of population activity, given the often reported wide range of response properties (response zoo) present in unanesthetized auditory cortex (Chimoto et al., 2002; Evans and Whitfield, 1964; Recanzone, 2000), including a substantial proportion of unresponsive cells (Hubel et al., 1959).

1.3   Beyond sound in auditory cortex

The simple framework for studying visual (or analogously auditory) processing described above assumes that each stage of the processing hierarchy can be studied independently. Such framework is appealing because it posits a series of feedforward, possibly nonlinear, filters organized into feedforward networks, which are much easier to study than recurrent networks.

However, imagine that you hear a shout: "Fire!" If you are part of a firing squad, you might pull a trigger; and if you are on the receiving end of the firing squad, you might merely brace yourself for the inevitable. If you are in your hotel room watching TV, you might decide to put on your slippers and see what the commotion is about; if you are, however, reading this thesis, it would be very hard to distract your attention by any outcry. In all cases, your auditory system will process the same acoustic stimulus, yet the neural activity (which precedes your motor response) will be different. It is clear that how you respond to the exclamation will depend on the behavioral or cognitive context.

Unfortunately for scientists, the mammalian brain is not a simple feedforward network. One cannot simply model the flow of information as undergoing a unidirectional transformation. Particularly in the cortex, it is clear that the representation of sensory information depends also on the animal's behavioral and/or cognitive state, and on how that information is to be used, and not only on the activity at the ear (cochlea).

Since the earliest studies of auditory processing, it has been clear that an animal's behavioral or attentional state can play a crucial role in shaping the response characteristics (Hernandez-Peon et al., 1956), even at the level of single neurons (Hubel et al., 1959). Follow-up studies on the cortical

correlates of auditory attention over the next decades were scant (Hocherman et al., 1976; Miller et al., 1972, 1980), until in recent years a series of studies have renewed interest in the field (Brosch et al., 2005; Fritz et al., 2005a, 2003, 2005b).

Factors other than attention can also influence activity of single neurons in the primary auditory cortex: various behavioral contingencies (Beaton and Miller, 1975; Miller et al., 1972), eye position (Fu et al., 2004; Werner-Reiss et al., 2003), and somatosensory stimulation (Brosch et al., 2005; Fu et al., 2003; Lakatos et al., 2007; Schroeder et al., 2001).

Auditory cortex—especially the primary auditory cortex—has been traditionally viewed as a pure sensory area acting as a (non)linear filter on signals passing through thalamus to "higher" cortical areas. However, the presence of nonauditory activity in primary auditory cortex strongly argues against such a simple view.

1.4 Rats in auditory research

Auditory research uses a variety of experimental animals to elucidate the function and structure of auditory system. Experimental preparations for studying the auditory system at the level of single neurons include non-human primates (Beaton and Miller, 1975; deCharms et al., 1998; Durif et al., 2003; Hocherman et al., 1976; Miller et al., 1972; Populin, 2006; Ryan and Miller, 1978; Ryan et al., 1984), ferrets (Fritz et al., 2005a, 2003, 2005b), guinea pigs (Edeline et al., 1993; Edeline and Weinberger, 1993), gerbils (Ohl and Scheich, 1997), cats (Las et al., 2005; Oatman, 1971, 1976; Tollin et al., 2005; Ulanovsky et al., 2003), bats (Suga and Ma, 2003), songbirds (Hahnloser et al., 2002; Sen et al., 2001; Theunissen and Doupe, 1998), chinchillas (Langner et al., 2002), mice (Linden et al., 2003), and others.

For experiments described in this thesis, we have used rats as experimental animals. Rats are another well-established established group of experimental animals in auditory research (Doron et al., 2002; Kelly and Sally, 1988; Sally and Kelly, 1988), used successfully in both anesthetized (DeWeese et al., 2003; DeWeese and Zador, 2006; Kilgard and Merzenich, 1998a,b; Wehr and Zador, 2003, 2005, and many others) and awake experiments (Gaese and Ostwald, 2003; Nakamura, 1999; Ono et al., 1985, and others).

Rats (rodents in general) offer several advantages when used as experimental preparations (especially when compared to humans and non-human primates). First, rats are relatively inexpensive to maintain. Second, the choice of rodents allows us to exploit the full armamentarium of modern electrophysiological, molecular and imaging approaches that are currently difficult or impossible to apply in primates, such as whole-cell patch-clamp methods (Wehr and Zador, 2003; Zhang et al., 2003), viral-mediated delivery of genes of interest (Rumpel et al., 2005), or molecules that allow monitoring (Miesenböck, 2004) or even perturbing neuronal activity *in-vivo* (Boyden et al., 2005) at the level of single neurons. Finally, rats can be trained to perform sophisticated behavioral tasks (Feierstein et al., 2006; Uchida and Mainen, 2003).

1.5 Recording neuronal activity

Most of what we know about activity in the brain at the level of single neurons comes from various recordings of electrical neuronal activity. The rest of this thesis offers analysis and description of neuronal activity recorded using patch-clamp technique with glass pipette, in cell-attached and whole-cell configurations (Hamill et al., 1981; Margrie et al., 2002). In this section we briefly summarize the differences between patch-clamp recordings and conventional recording techniques, which are important for the rest of the text.

Since the advent of modern electrophysiology (Adrian, 1926; Adrian and Matthews, 1927; Hartline, 1925), recording individual action potentials in response to sensory stimuli quickly became standard in auditory neurophysiology as well (Adrian et al., 1938; Hallpike et al., 1937). Recordings using tungsten electrodes (Hubel, 1957) almost immediately became the mainstream magic wand of electrophysiology, and we will refer to them as the conventional recording technique.

The conventional recording technique uses metal electrode which acts as an antenna when inserted into neural tissue. The electrode then "picks-up" electrical signal (i.e. action potentials) from nearby neurons. Good recording (good single unit isolation) with conventional electrodes requires a sufficient number of spikes; skilled practitioners typically search for neurons with sufficiently high firing rates and large spikes, which are easily distinguishable from other neurons. Although it is possible for a committed investigator to isolate even neurons with a low spontaneous firing rate, such neurons rarely appear in analyses, possibly because they do not offer enough signal (i.e. too few spikes), or are im-

mediately discarded. Conventional recording techniques thus appear selectively biased toward neurons providing strong signal.

By contrast, recording with a glass patch pipette in cell-attached mode (extracellular mode recording action potentials only) is not explicitly biased toward active and responsive neurons, or neurons with large action potentials, and provides excellent single unit isolation (DeWeese et al., 2003; Margrie et al., 2002). With cell-attached recording, single unit isolation depends on the physical contact between the glass electrode tip and the neuron. The selection bias of cell-attached recording is thus based on the neuron's "patchability," rather than on the firing rate or responsiveness of the target neuron; only to the extent that patchability is correlated with functional characteristics such as firing rate or responsiveness would cell-attached recording (indirectly) bias the sampled population.

Recently, extracellular recording techniques with metal electrodes underwent significant improvements in the form of tetrodes (for example Feierstein et al., 2006; Harris et al., 2000; Henze et al., 2000; Mehta et al., 2002), or silicone probes (for example Barthó et al., 2004; Csicsvari et al., 2003). Experimenters using these techniques can record activity of more neurons in parallel, and identify neurons based on different spike parameters (signatures). Such recordings are then less biased toward active neurons. Nevertheless they still rely on neuronal activity, and neurons with extremely low firing rates can still be missed, simply because collecting a sufficient number of spikes might take considerable amount of time.

Patch-clamp recordings, in whole-cell configuration, also offer an opportunity to record *intracellular* activity of neurons. Most of what is known about intracellular neuronal activity comes from studies in *brain slices* (Edwards et al., 1989; McCormick et al., 1985; Stevens and Zador, 1998, and many others). Recently, several laboratories introduced whole-cell recording techniques in *anesthetized in-vivo* preparations in various sensory systems: auditory (DeWeese and Zador, 2006; Tan et al., 2004; Wehr and Zador, 2003, 2005; Zhang et al., 2003), visual (Anderson et al., 2000; Carandini and Ferster, 2000; Ferster and Jagadeesh, 1992; Jagadeesh et al., 1993, 1997; Pei et al., 1991), somatosensory (Brecht et al., 2003; Brecht and Sakmann, 2002; Bruno and Sakmann, 2006; Bureau et al., 2004; Petersen et al., 2003b; Zhu and Connors, 1999), and olfactory (Margrie and Schaefer, 2003; Schaefer et al., 2006). Whole-cell recordings in-vivo are very challenging, and their application in *awake* animals is technically difficult. Only very recently, several studies have described successful whole-cell recordings

in somatosensory (barrel) cortex of awake (unanesthetized) mice and rats (Brecht et al., 2004; Crochet and Petersen, 2006; Lee et al., 2006; Margrie et al., 2002; Petersen et al., 2003a).

1.6 Thesis outline

This thesis is divided into six chapters (following this Introduction), followed by three appendices containing supplementary information. Each chapter was intended to be self-contained, so they do not have to be read in the order they are presented.

Second chapter (Sec. 2) contains a detailed description of experimental techniques: surgery, recording, and training techniques we used in awake head-fixed rats. We have also included a detailed description of all sets of stimuli we used to probe neurons, analytical methods used to analyze data, and description of computational models used in other parts of the thesis.

Third chapter (Sec. 3) focuses on description of single-neuron responses in primary auditory cortex of awake head-fixed rats. The primary emphasis of this part is on the sparse representation of various auditory stimuli we used to probe neurons, and the heterogeneity of responses of single neurons. To characterize population responses to sound in the auditory cortex we asked the question "What is the typical response to acoustic stimuli?" instead of what is usually asked "What is the stimulus that evokes a response?" We found that the population response was sparse, with many unresponsive neurons. In addition, the responsive neurons showed a great variety of responses. This heterogeneity of neuronal responses ("response zoo," courtesy of Anthony M. Zádor) was, however, surprisingly well characterized by lognormal distribution of firing rates.

The observation that firing rates in awake auditory cortex were lognormally distributed was even more interesting given the observation of lognormal distribution of synaptic weights in the cerebral cortex. The fourth chapter (Sec. 4) focuses on mechanisms which could give rise to lognormal distribution of firing rates, as well as synaptic weights. We proposed specific types of correlations among synaptic connections, and formulated a multiplicative learning rule which led to the observed distributions.

We were also able to characterize intracellular activity of neurons in awake auditory cortex. The fifth chapter (Sec. 5) contains analysis of so-called up and down states in awake auditory cortex. We show that up and down states—the "signature" subthreshold dynamics so often described in various

cortical areas of anesthetized animals—were rare in the primary auditory cortex of awake rats, instead, subthreshold dynamics was consisted of brief, infrequent fluctuations of membrane potential.

The experiments described and analyzed in chapters 2–4 were conducted in naïve awake rats. As behavior or attention can influence neuronal activity even in primary sensory areas, we developed a setup for head-fixed behavior. In the sixth chapter (Sec. 6) we describe the sound discrimination task we have used to study behavior in head-fixed rats. We present a comparison of basic behavioral parameters between restrained and unrestrained rats, as well as evidence of nonauditory modulations of single neuron activity in auditory cortex.

Finally, in the seventh chapter (Sec. 7) we briefly summarize results from chapters 3–6.

The first appendix shows raster plots of all hundred neurons we identified to evaluate response significance. The rasters are presented to supplement the data analysed in Sec. 3. The second appendix shows raster plots of neuronal activity of fifteen neurons recorded in head-fixed rats performing the sound discrimination task. The rasters supplement data presented in Sec. 6. The third appendix presents schematics of surgical implants, and the lickometer we used in experiments with awake head-fixed rats. The schematics supplement experimental procedures described in Sec. 2.

1.7   Disclosures

All experiments were conducted under supervision of Anthony Zador, as my scientific mentor. Michael R. DeWeese performed 31 cell-attached and 5 whole-cell recordings of those analyzed in chapters 3–5. Chapter 4 was prepared in collaboration with Alexei Koulakov, who also developed mathematical proofs of several statements mentioned in the chapter (not shown in this text).

Chapter 2

Experimental procedures

We have developed an experimental preparation that allows us to use patch-clamp recording techniques in awake head-fixed rats. This chapter starts with a description of the surgical procedure we used to prepare animals for head-fixed recordings. Then we continue with a description of the experimental setup we developed to record from both naïve and trained head-fixed rats. We also describe *in-vivo* patch-clamp recording techniques, together with acoustic stimuli used in our experiments, and custom data acquisition (plus stimulus delivery) system. The final parts of this chapter focus on data analysis, a detailed description of simulation experiments, and a description of training procedures for head-fixed behavior. Brief description of the relevant experimental methods is also included in each of the following results chapters.

2.1 Experimental animals

We used rats as experimental animals. For experiments on *naïve* head-fixed animals we used Sprague-Dawley rats, 21–30 days old at the time of surgery. For *behavioral* experiments we used Long-Evans rats, 23–25 days old at the beginning of training, 33–50 days old at the time of surgery.

2.2 Surgery

Rats 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. Part of scalp was removed

11
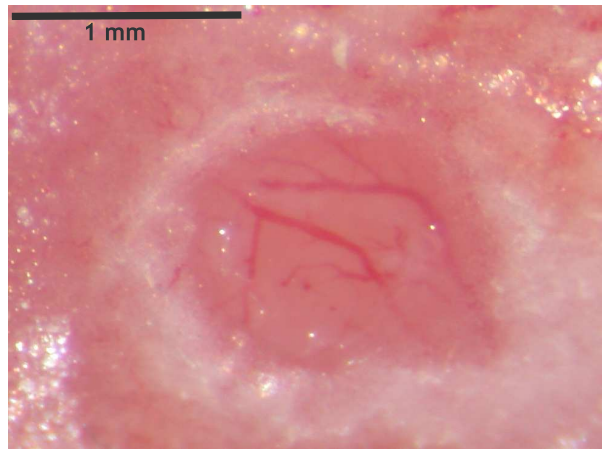
Figure 2.1: Example of a craniotomy performed during preparation for head-fixed recordings.

together with underlying tissue, and the exposed bone was cleaned and dried out. The left temporal muscle was cut and partially removed to enable access to the left temporal bone. The exposed area— except for the craniotomy site (see below)— was dried out and covered with a thin layer of either Krazy Glue (Krazy Glue, Columbus, OH), or VetBond (3M, St. Paul, MN). The layer of dry glue was then covered with a thin layer of dental acrylic (except for the craniotomy site) and left to dry.

A small craniotomy (maximum size of $1.5\,\text{mm} \times 1.5\,\text{mm}$) and durotomy were then performed over the left (primary) auditory cortex (Fig. 2.1). Perfect durotomy with no bleeding was of utmost importance for the success of recordings because it preserved the health of the cortex and prevented contamination of electrode tips. The position of the craniotomy was determined by its distance from bregma ($4.5\,\text{mm}$ posterior, and $4\,\text{mm}$ ventral), and its relationship to other bone sutures. The presence of clear auditory single-unit responses and/or local field potentials were further used as physiological criteria to confirm the location of the auditory cortex. Based on the anatomical landmarks and physiological criteria we expect that the neurons recorded in this thesis were in the primary auditory cortex (Doron et al., 2002).

The brain surface was covered with Kwik-Cast (World Precision Instruments, Sarasota, FL) immediately after performing the durotomy as well as between recording sessions. The whole area was protected by a plastic well with a removable cap (Appendix C). An aluminum headpost was attached to the skull with Relyx Luting Cement (3M ESPE, St. Paul, MN). A silver chloride ground wire was implanted on the back of the animal, with the silver chloride pellet on one end positioned subcutaneously,

and a golden pin on the other end attached to the headpost. After surgery, the animals were allowed at least 24 hours of recovery before the first recording session.

2.3 Recording setup

During the recording (or training) session, the head of the animal was fixed in the headpost holder and the animal was positioned inside a plastic tube, which provided loose restraint of body movements (Fig. 2.2). Rats were not forced to enter the recording tube, and usually just ran inside. Occasionally, if a rat did not voluntarily enter the tube, the session was postponed.

The animals sat quietly, occasionally moved their limbs, groomed, whisked, etc. The behavioral state of each animal was monitored by a closed video circuit. Excessive movement, signs of stress, or discomfort of the animal were used to indicate the end of the session.

During behavioral sessions, a custom made "lickometer" was placed in front of rat's mouth to monitor behavioral responses (licking) and deliver water drops as rewards. The lickometer contains a water delivery port (a modified 18 gauge needle) in the middle, below it there is a vacuum port providing suction to remove excess water, if necessary. Tongue licks were defined and measured as interruptions of infrared beam positioned across the end of water port. The lickometer is a modified version of choice ports, originally developed by Naoshige Uchida and Zachary Mainen (Uchida and Mainen, 2003) for training freely moving rats. Water rewards were delivered via a gravitational system from a 50 ml reservoir (syringe), and water delivery was controlled by solenoid valves (NResearch, West Caldwell, NJ). Air-puffs as behavioral penalties were delivered from an air outlet (via plastic tubes) to a modified 16 gauge syringe needle. The syringe needle was aimed at the animal's snout. Air pressure at the outlet was set to 30–35 psi (pounds-force per square inch; for people actually using SI units, the pressure would correspond to 207–241 kPa).

For recordings, the plastic cap and Kwik-Cast layer were removed and the cortex was covered with physiological buffer (in mM: NaCl, 127; Na_2CO_3, 25; NaH_2PO_4, 1.25; KCl, 2.5; $MgCl_2$, 1; and glucose, 25) mixed with 1.5% agar. Each recording session usually lasted for 1.5–2 hours. We recorded from each animal during several recording sessions (usually 2 or 3 sessions per rat). Because the glass electrodes we used (see below) had a rather wide shank and the cortical area available for recordings
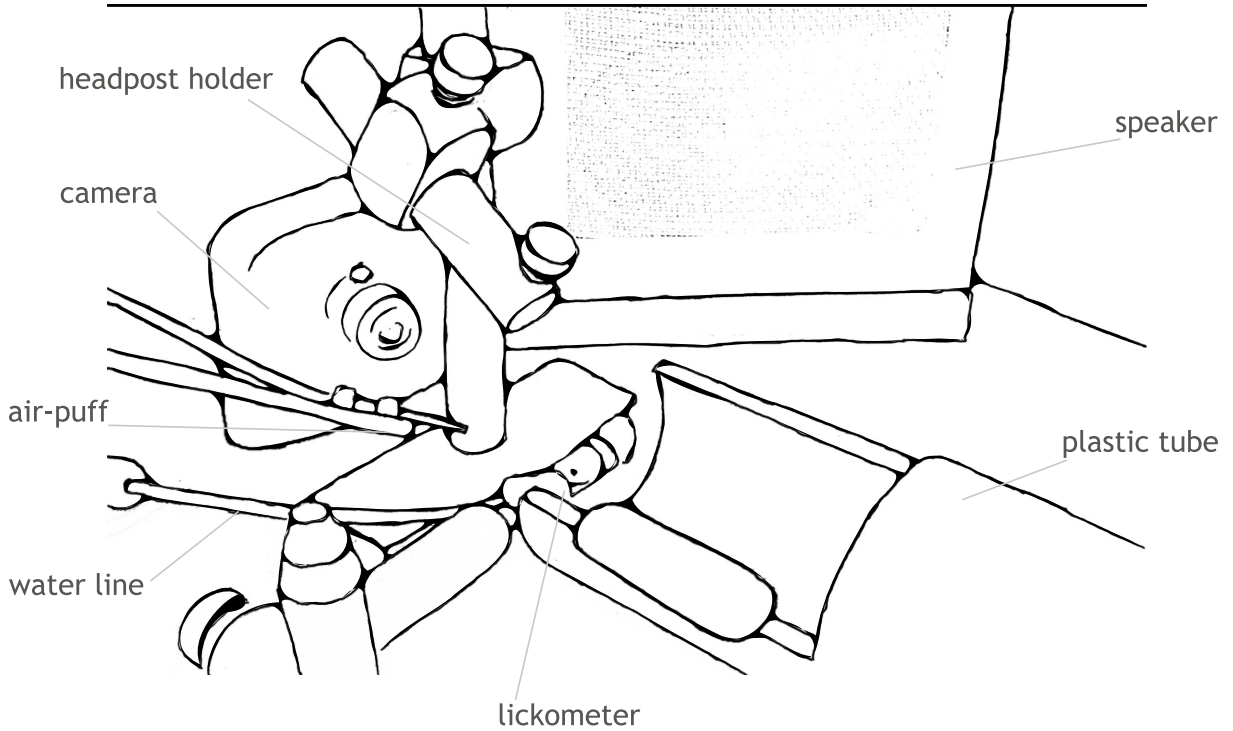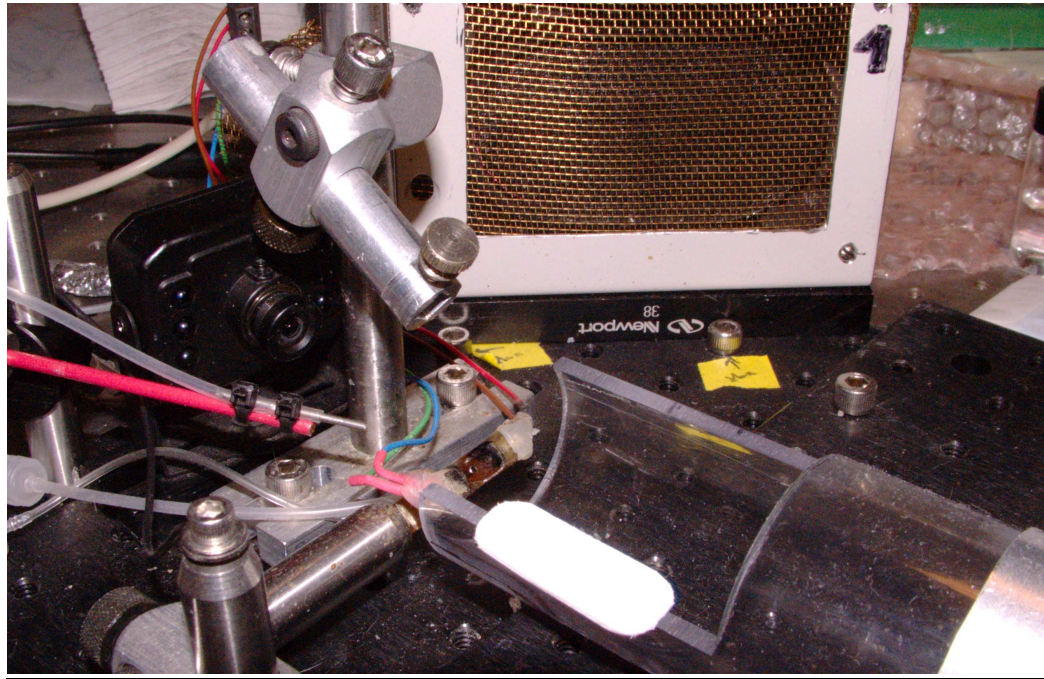
Figure 2.2: Experimental setup for patch-clamp recordings in awake head-fixed rats. Plastic tube, head-post holder were mounted on a recording platform, together with a speaker (facing the right ear of the animal.) Behavioral state and body movements were monitored by video camera mounted directly in front of the animal. The recording platform was mounted on an air-table to prevent vibrations from outside the sound booth. A lickometer was placed in front of the rat's mouth during behavioral sessions to monitor licking and deliver water.

was small, the number of recording sessions was limited by the total number of electrode penetrations. We usually performed up to 15–20 electrode penetrations per animal.

2.4 Electrophysiology

Recordings were obtained using standard blind patch-clamp recording techniques in cell-attached or whole-cell configuration (see also DeWeese et al., 2003; Machens et al., 2004; Wehr and Zador, 2003). Electrodes were pulled from filamented, thin-walled, borosilicate glass (outer diameter, 1.5 mm; inner diameter, 1.17 mm; World Precision Instruments, Sarasota, FL) on a vertical two-stage puller (Narishige, East Meadow, NY). For recording, electrodes were filled with internal solution containing (in mM): KCl, 10; KGluconate, 140; HEPES, 10; $MgCl_2$, 2; $CaCl_2$, 0.05; MgATP, 4; Na_2GTP, 0.4; Na_2Phosphocreatine, 10; BAPTA, 10; and biocytin, 1%, pH 7.25; diluted to 290 mOsm. Resistance to bath was 3.5–5.0 MΩ before seal formation. Recordings were obtained using Axopatch 200B (Axon Instruments, Union City, CA) and a custom data acquisition system written in MATLAB (Mathworks, Natick, MA, see also Sec. 2.9), with a sampling rate of either 4 kHz or 10 kHz.

Before inserting an electrode into cerebral cortex, positive pressure (150–200 mbar) was applied to the electrode (glass pipette) to prevent debris from accumulating in the electrode tip as it would travel through cortex. After an electrode was in the cortex positive pressure was lowered to approximately 100–120 mbar, and upon reaching 150–200 μm depth, the pressure was further lowered to 70–75 mbar. With this positive pressure the pipette was slowly advanced (in 1–2 μm steps) and the resistance was continuously monitored in voltage-clamp mode. Voltage pulses (-10 mV, 30 ms) were delivered at a rate of 7–8 Hz and the evoked changes in current were monitored to assess the pipette's resistance. Increase in the pipette resistance usually signalled the presence of cellular membrane. Once the pipette's resistance increased to 6–8 MΩ, positive pressure was released and a seal between the electrode tip and membrane was established. However, sometimes suction (usually 10–20 mbar, but occasionally up to 100 mbar) had to be applied to establish a stable seal.

Cell-attached recording requires a seal of only 10–20 MΩ, whereas whole-cell recordings require at least 1 GΩ seals. Therefore, in the cell-attached configuration almost every neuron encountered can be patched. Although virtually all neurons encountered could be patched, not all recordings were sufficiently stable and the seals were "lost" before the full set of stimuli was presented. Neurons were

A

5 mV

B

50 mV

0.5 s
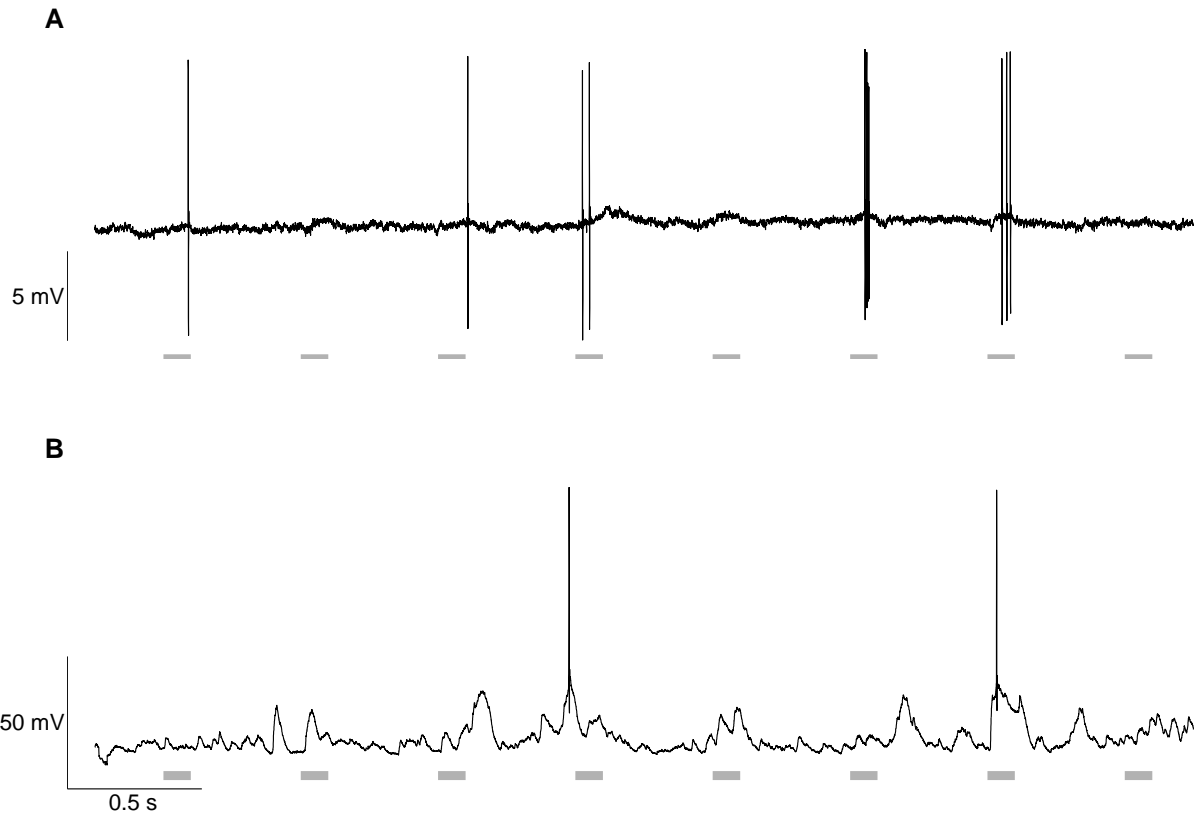
Figure 2.3: Examples of patch-clamp recordings in awake head-fixed rats. **(A)** An example trace of neuronal activity recorded in cell-attached configuration. **(B)** An example trace recorded in whole-cell configuration. Note the ten-fold difference in signal amplitude. Grey rectangles show position of sound stimuli (tones of various frequencies and intensities.)

recorded in all depths, with more neurons encountered in the middle cortical layers, as estimated by the micromanipulator travel distance. At least one spike was recorded for each neuron.

Both cell-attached and whole-cell recordings offered an excellent signal to noise ratio (Fig. 2.3). Spike amplitudes in cell-attached (extracellular) recordings were in the millivolt range, often in tens of millivolts. Signal amplitudes for whole-cell (intracellular) recordings were in tens of millivolts, as the signal varied from resting membrane potential (for example -70 mv) to peaks of action potentials (for example 20 mV).

2.4.1   Technical aspects of whole-cell recordings in awake rats

Most of what is known about intracellular neuronal activity comes from studies in *brain slices* (Blanton et al., 1989; Cauli et al., 1997; Edwards et al., 1989; McCormick et al., 1985; Stevens and Zador, 1998;

Stuart et al., 1993, and many others). Recently, several laboratories introduced whole-cell recording techniques *in-vivo* in various sensory systems: auditory (DeWeese and Zador, 2006; Tan et al., 2004; Wehr and Zador, 2003, 2005; Zhang et al., 2003), visual (Anderson et al., 2000; Carandini and Ferster, 2000; Ferster and Jagadeesh, 1992; Jagadeesh et al., 1993, 1997; Pei et al., 1991), somatosensory (Brecht et al., 2003; Brecht and Sakmann, 2002; Bruno and Sakmann, 2006; Bureau et al., 2004; Petersen et al., 2003b; Zhu and Connors, 1999), and olfactory (Margrie and Schaefer, 2003; Schaefer et al., 2006).

The vast majority of in-vivo studies of intracellular activity has been performed in *anesthetized* preparations. Whole-cell recordings in-vivo are very challenging, and their application in *awake* animals is technically difficult. Only very recently, several studies have described successful whole-cell recordings in somatosensory (barrel) cortex of awake (unanesthetized) mice and rats (Brecht et al., 2004; Crochet and Petersen, 2006; Lee et al., 2006; Margrie et al., 2002; Petersen et al., 2003a).

Obtaining a gigaohm seal between the electrode tip and neuronal membrane is a necessary condition for a successful whole-cell recording. Overall, we were able to obtain gigaohm seals in about one third of our attempts (penetrations). However, not all seals were successfully transformed into a whole-cell configuration, usually due to sudden movement of the animal. Similar success rates have been reported for both anesthetized and awake recordings (Margrie et al., 2002).

Factors that probably influenced the success rate of our recordings included the animal's movement, and possible contamination of electrode tips when penetrating cortical surface. In addition, the actual number of attempts to obtain whole-cell configuration (for this study) was limited. We have interleaved whole-cell and cell-attached recordings in order to maximise amount of data obtained from each animal. Because each attempt to obtain a gigaohm seal requires a clean fresh electrode tip, we were usually able to obtain good seal only on the first neuron—in a given penetration—we encountered. When the attempt failed, we continued searching for cells and used cell-attached recordings to record from more neurons in the same penetration. Each recording session lasted approximately 1.5–2.5 hours, therefore the number of penetrations (especially when we recorded from multiple neurons) was limited to 3–5 per recording session.

We have successfully recorded intracellular activity for several minutes (see Fig. 5.1 on page 92 for an example). The recordings were sufficiently stable and offered an excellent signal-to-noise ratio. Recordings described in this text lasted between 2–13 min (mean 7.4 min, median 7 min, n = 19 neu-

17

rons), with one additional neuron from which we recorded for approximately 2 hours. Good stability of recordings was evident during periods of animal movements, such as whisking, licking or even grooming. Such movements usually did not cause any apparent changes in recording quality. Gross, sudden movements of the whole body (such as adjusting body position), however, usually led to an abrupt end of recording.

2.5   Acoustic stimuli

We have used several types of acoustic stimuli (sounds), all of which were presented free-field in a double-walled sound booth (Industrial Acoustics Company, Bronx, NY). For experiments on naïve head-fixed rats, free-field stimuli were presented at 97.656 kHz using the TDT System 3 (Tucker-Davis Technologies, Gainesville, FL) connected to an amplifier (Stax SRM 313, STAX Ltd, Japan), which drove a calibrated electrostatic speaker (taken from the left side of a pair of Stax SR303 headphones) located 8 cm lateral to, and facing, the contralateral (right) ear. For behavioral experiments on head-fixed rats, free-field stimuli were either presented with the setup just described, or were presented at a 200 kHz sampling rate using a custom built real-time Linux system driving a high-end Lynx L22 audio card (Lynx Studio Technology Inc., Newport Beach, CA) connected to the same amplifier and calibrated speaker as mentioned previously.

2.5.1   Tones

The main sets of stimuli for *cell-attached recordings* consisted of 100 ms long pure-tone pips of 16, 20, or 64 different frequencies logarithmically spaced between 1–40 kHz (81 % of recordings, 134 out of 166) presented at either 20, 50, 80 dB sound pressure level (SPL) (n = 43), or at 0, 30, 60 dB SPL (n = 15), or at 0, 20, 40, 60 dB SPL (n = 76). For the rest of the recordings (19 %, 32 out of 166) the stimulus protocol contained 100 ms long pure-tone pips of 28 frequencies logarithmically spaced between 2–48 kHz presented at 60 dB SPL. All tones were repeatedly presented in a fixed pseudo-random order at a rate of 2 tones per second. A full tuning curve was obtained for each neuron, i.e. each neuron was probed at least once with each tone of a particular stimulus protocol.

We used the same sets of tone stimuli for *whole-cell recordings* (n = 19 neurons out of 20). The neuron shown in Fig. 2.6 on page 27 was probed with 50 ms long tones, logarithmically spaced between 0.1–40 kHz presented at 0, 30, 60 dB SPL.

2.5.2 Frequency modulated sweeps

In 22 neurons (13% of recordings, 22 out of 166) we also presented frequency-modulated sweep stimuli. Sweeps covered the frequency range from 1 to 40 kHz, and both upward (from 1 to 40 kHz) and downward (from 40 to 1 kHz) going sweeps were presented at 6 different rates (25, 50, 75, 100, 125, 150 octaves/s) for each neuron.

2.5.3 Natural sounds

For *cell-attached recordings* natural sound stimuli were presented for 28 neurons, 23 of these neurons were also presented with pure tones (14% of recordings, 23 out of 166), and five neurons were presented only with natural sounds. The natural sound stimuli were taken from a commercially available audio compact disc, *The Diversity of Animal Sounds* (Cornell Laboratory of Ornithology), originally sampled at 44.1 kHz and resampled at 97.656 kHz for stimulus presentation (Machens et al., 2004). The sounds chosen had no special relevance to the rats (unlike e.g. rat pup calls), and therefore are less likely to engage specialized processing mechanisms; to the extent that these sounds are representative of the acoustic environment of humans, they are also representative for rats, which often share the same habitat as humans. Altogether, four natural sound segments were presented for each neuron, with at least 4 repeats of each segment per neuron. The segments included *Jaguar call* (track 3, seconds 2 to 11 for total duration of 10 s), *Bowhead Whale* (track 9, seconds 1 to 10, 10 s duration), *Knudsen's Frog* (track 11, seconds 1 to 10, 10 s duration), and *Bearded manakin* (track 19, seconds 0.1 to 5.1, 5 s duration). The peak amplitude of each segment was normalized to the ± 10 V range of the TDT system, which corresponded to 80 dB SPL.

Six neurons were probed with the same set of natural sounds in *whole-cell recordings.*

2.5.4 Behavioral stimuli

In the sound-discrimination go/no-go task we used pure tones and frequency modulated tones (warbles) as stimuli. Tones were 500 ms long pure-tone pips of 1, 2, 4, 8, 16, 32, and 40 kHz presented at 70 dB SPL.

Warbles (500 ms) were synthesized according to:

$$warble(t) = A(\sin(2\pi f_c t + \frac{ind_m}{f_m}(\cos(2\pi f_m t + \phi_m - \frac{\pi}{2})) + \phi_c), \qquad (2.1)$$

where A stands for the final amplitude of $warble$ (our range was ± 1 V), f_c stands for carrier frequency of the underlying tone, ϕ_c is the phase of the carrier, f_m modulation frequency, ϕ_m modulation phase, and ind_m modulation index. Note, that our modulation index (ind_m) represents frequency deviation from the modulation frequency, whereas the usual definition would define modulation index as $\frac{ind_m}{f_m}$.

2.6 Data analysis

2.6.1 Spike extraction and analysis for cell-attached recordings

Spikes recorded in cell-attached mode were extracted from raw voltage traces by applying a high-pass filter and thresholding (Fig. 2.4). Spike times were then assigned to the peaks of suprathreshold segments, and rounded to the nearest millisecond.

Individual spikes can assume very different shapes even in a single cell (Fig. 2.4). In some cases we observed bursts of spikes, during which spike amplitude sometimes decreased severalfold. For example, for the cell shown in Fig. 2.4, both single spikes and bursts were sometimes evoked approximately 40 ms following tone termination. Such large changes in spike characteristics could result either in a failure of spike detection, or errors in spike sorting in conventional extracellular tungsten recordings.

Spikes were recorded at a sampling rate of 4 kHz for 88 neurons in our main set of cell-attached recordings (n = 166 neurons), and 10 kHz for the remainder of the population. For the analysis of spike shape (Sec. 3.9 on page 65), the spike waveforms recorded at 4 kHz were resampled to 10 kHz (using

20

A

2 mV

500 ms

B

1 mV

500 ms
25 ms

*

*

C

1 mV

10 ms
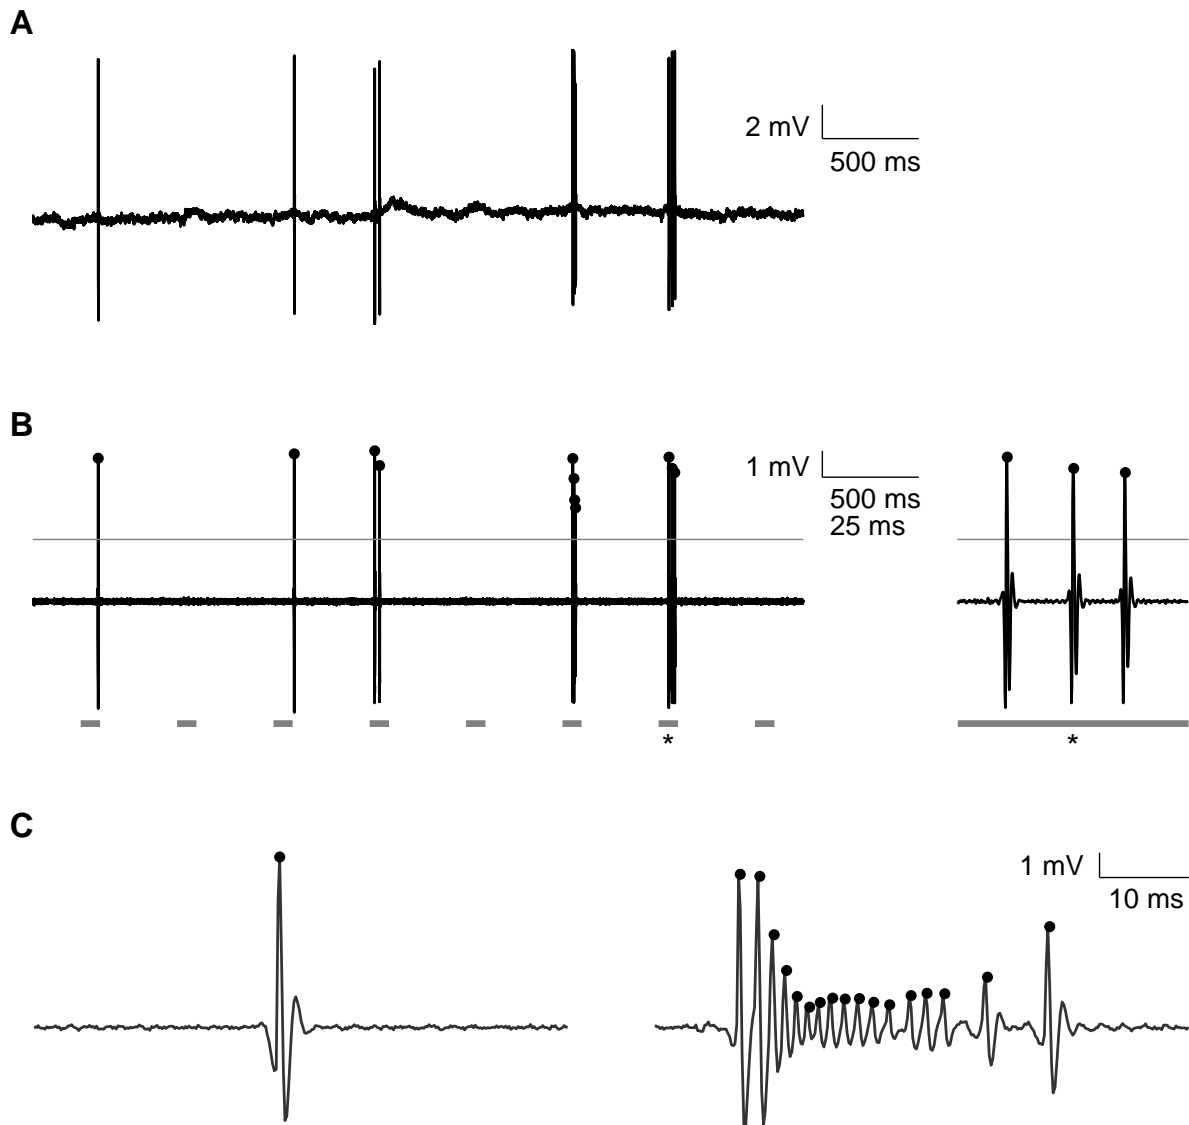
Figure 2.4: **(A)** Cell-attached recordings provide an excellent signal-to-noise ratio. Four second long raw voltage trace recorded in cell-attached mode in the auditory cortex of an unanesthetized rat. **(B)** Cell-attached recording allows for high quality single unit isolation. Voltage trace from A was high-pass filtered, and spikes were easily identified in the trace after thresholding (grey line). Spike times (dots) were assigned to peaks of suprathreshold segments. Grey squares indicate the positions of stimuli (pseudorandom sequence of 100 ms long tones of different frequencies and attenuations). The trace segment on the right shows the last 60 ms of the response to the 7[th] stimulus (asterisk). **(C)** Spike shape can change significantly in a single neuron. Spike times (dots) are assigned to peaks of suprathreshold segments. The burst in the right panel preceded the single spike shown in the left panel by approximately 5 seconds. Both examples were tone-evoked, and occurred about 40 ms after stimulus termination.

Matlab resample function). We then computed the mean spike waveform for each neuron, and defined *spike width* as the time difference between the peak (maximum amplitude) and valley (minimum amplitude following the peak) of the waveform. Because the spike waveforms were (re)sampled at 10 kHz, spike widths are rounded to the nearest tenth of a millisecond. For each cell we also computed the *amplitude index*, the absolute value of peak-to-valley-ratio, of the mean spike waveform.

2.6.2   Evoked response analysis for cell-attached recodings

Responses to stimuli were divided into 50 ms duration time bins. In addition, tone-evoked responses were also binned in frequency space. We use the term *response bin* to refer to subdivisions of a response in general, as defined below for various stimuli. When we explicitly refer to binning in time, or frequency, we use the terms *response epoch,* or *octave bin* respectively.

Tones

Tone-evoked responses were divided into four 50 ms long *response epochs* (see also Fig. 3.3 on page 49). The *spontaneous* epoch was defined as the 50 ms long period preceding stimulus onset. The *early* epoch was defined as the first 50 ms of stimulus duration, the *late* epoch as the last 50 ms of stimulus duration, and the *off* epoch as the first 50 ms after stimulus termination. In frequency space the responses were grouped into one-octave-wide bins, which resulted in 4 or 5 frequency bins (*octave bins*) per cell (depending on the stimulus protocol used, see Sec. 2.5.1 above).

The *spontaneous firing rate* for each cell was computed as a mean of firing rates across all trials in the spontaneous epoch for a given cell. *Evoked firing rates* were computed for each combination of response epoch and octave bin as a mean of firing rates of all trials in the specific octave-epoch combination.

The distribution of firing rates across octave bins for each response epoch was fit with a log-normal distribution (see also Fig. 3.4 on page 51, and Fig. 3.5 on page 52). To fit each distribution, the octave bins with zero firing rate were removed, and the mean and variance of the distribution of log-transformed firing rates were computed. The mean and variance obtained directly from the data were then used as parameters for the normal distribution fit to log-transformed firing rates. The goodness-of-fit for each distribution was assessed using the Kolmogorov-Smirnov test.

The significance of stimulus-evoked changes in firing rates was evaluated with the Wilcoxon signed-rank test, i.e., a non-parametric paired, two-sided test of the hypothesis that the difference in firing rates between the matched trials in two different epochs comes from a distribution whose median is zero. For each octave and *early*, *late*, and *off* response epochs we tested on a trial-by-trial basis whether the stimulus-evoked firing rate increased or decreased significantly compared to the corresponding *spontaneous* epoch. For this test, we only considered cells with at least 20 trials per octave bin (69 %, 100 cells out of 145).

For the analysis of *responsiveness of single neurons* (see Sec. 3.4 on page 53) the evaluation of significance involved 15 comparisons for most of the cells, because the responses of most cells were binned into 15 response bins (5 octave bins times 3 response epochs). Therefore, we used a significance criterion of either $p < 0.0033$ (for 15 comparisons, 0.05/15), or $p < 0.0042$ (for 12 comparisons, 0.05/12) to keep the overall significance criterion for each cell at $p < 0.05$. To be considered tone-responsive, a cell had to show a significant change in firing rate (increase or decrease) in at least one response bin.

For the *population response* analysis (see Sec. 3.4 on page 53, and Sec. 3.5 on page 56) the response bins from all neurons were considered independent and their responsiveness was evaluated with the Wilcoxon signed-rank test using a significance criterion of $p < 0.01$. To evaluate the population response in the *early* response epoch the fraction of bins showing a significant increase, decrease, or no change in the firing rate was computed for each octave bin in the *early* response epoch. The fraction of responsive bins in the *early* response epoch was then defined as the mean of the octave-bin fractions in the epoch. Analogous computations were carried out for the *late* and *off* response epochs. To compute the population response across all epochs, the fraction of responsive bins was computed from all response bins (from all neurons) pooled together.

Careful inspection revealed no clear examples of frequency tuning sharper than about one octave, suggesting that it would be appropriate to pool together responses to tones within an octave. To confirm systematically that our results were robust to this choice we repeated this analysis with half-octave wide (i.e. narrower) frequency bins, two, and four octaves wide (i.e. wider) frequency bins, and 50 ms long response epochs. To control for neurons with more transient or sustained responses we performed the population response analysis with 25, 75, and 100 ms duration response epochs and one-

octave-wide frequency bins. The results of these analyses, however, were the same as for the basic analysis with one-octave-wide frequency bins and 50 ms duration response epochs (not shown).

Frequency-modulated sweeps

Responses to frequency-modulated sweeps (FM-sweeps) were subdivided into 50 ms duration response epochs. Slower sweeps, with 25 or 50 octaves/s, contained four or two 50 ms epochs, respectively, during the stimulus presentation. Faster sweeps (75, 100, 125, 150 octaves/s) contained one 50 ms epoch. For all sweep rates we also added an *off* epoch starting either at the sweep termination (for 25, 50, 75, 100 octaves/s), or immediately after the response epoch (for 125, 150 octaves/s). Each response was thus divided into 32 response bins (including upward and downward moving sweeps). Therefore, for the analysis of significance of sweep-evoked responses in individual neurons we used a significance criterion of $p < 0.0016$ (0.05/32). To compute the population response to FM-sweeps all response bins were considered statistically independent and their responsiveness was computed using a significance criterion of $p < 0.01$.

Natural sounds

Responses to natural sounds were also divided into 50 ms duration response epochs. Ten second long segments thus contained 200 response bins each, and five second long segments contained 100 response bins. Natural sound-evoked responses were used only for the analysis of stimulus-evoked changes in firing rate, because none of the recordings met our criterion for the test of evoked response significance (i.e. at least 20 trials per response bin.)

2.6.3   Analysis of whole-cell recordings

Spikes recorded in the whole-cell configuration were detected in the same way as in cell-attached configuration. High-pass filter was applied to raw voltage traces, and the filtered trace was thresholded. Spike times were then assigned to the peaks of suprathreshold segments and rounded to the nearest millisecond.

Voltage traces were recorded with action potentials superimposed on subthreshold activity. To analyze only the subthreshold activity, action potentials were removed from raw voltage traces by me-
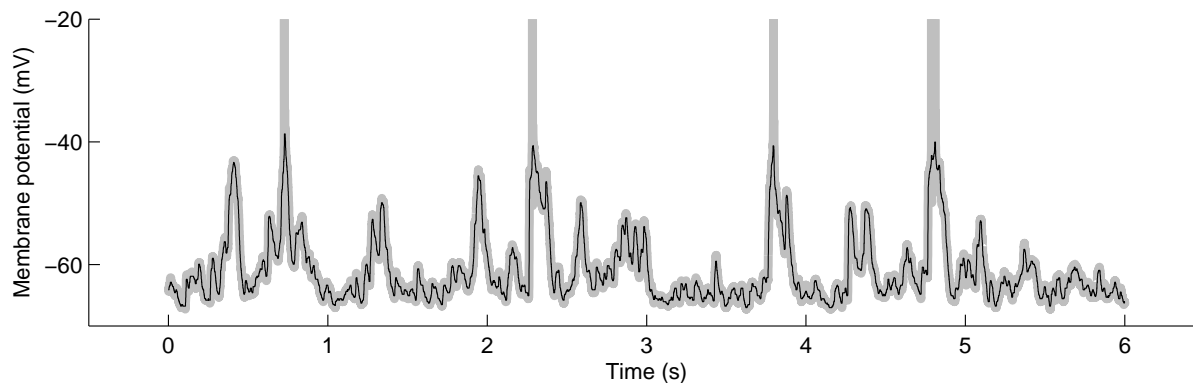
Figure 2.5: Median filter was applied on the raw voltage traces to remove action potentials. The **thick grey line** shows the raw voltage trace recorded in unanesthetized auditory cortex. The action potentials shown were truncated at -20 mV. The black line shows the filtered voltage trace (4 ms median filter). Note that the filtered trace preserves subthreshold voltage fluctuations.

dian filtering (Fig. 2.5). We used 4 ms median filter, where each point in the voltage trace was replaced with a median of itself and 16 or 40 (depending on sampling rate) surrounding points (Jagadeesh et al., 1997). Therefore, large fast fluctuations (action potentials) were removed, while slower fluctuations (up to very fine details) were preserved. All subthreshold traces shown were median filtered, unless otherwise specified.

Traces were selected for further analysis after careful visual inspection. Nonstationary traces, and traces in which the fifth percentile value of the membrane potential across the trace was above 50 mV were excluded from further analysis.

Kurtosis analysis

For each recorded trace we computed a *kurtosis excess* k of membrane voltage distribution (DeWeese and Zador, 2006; Olshausen and Field, 2004):

$$k = \frac{1}{n} \sum_{i=1}^{n} \frac{(v_i - v_{mean})^4}{\sigma^4} - 3, \tag{2.2}$$

where n is the number of data points (samples) in a voltage trace, v_i is the membrane potential at data point i, v_{mean} is the mean membrane potential across the trace, and σ is the standard deviation of membrane potential across the trace. The kurtosis of a normal distribution is equal to zero, whereas for

heavy-tailed distributions (distributions with tall peak and long tails) kurtosis is greater than zero. For this analysis we used 7 s long voltage traces (neurons 2–15), except for neuron 1, for which we recorded data in 4 s long trials. Traces from neurons recorded with continuous data acquisition (neurons 16–20) were divided into 7 s long non-overlapping segments.

We used kurtosis (henceforth we shall use the term kurtosis instead of kurtosis excess) to measure "sparseness" of subthreshold activity. Occasional deviations from mean membrane potential (bumps) lead to long tailed distribution of membrane potential values, and kurtosis higher than zero. Note that, because the recordings were performed in awake animals, recordings were more susceptible to pulmonary and cardiac pulsations, recording instabilities, animal movements (e.g. grooming), etc.These events typically lead to a decrease in kurtosis because they smear out the peak in the mean of the voltage histogram near the resting potential (DeWeese and Zador, 2006).

Bump analysis

Subthreshold membrane potential activity usually appeared as occasional voltage fluctuations above resting membrane potential (*bumps*, see Sec. 5 on page 89 for details). The *resting membrane potential* for each trace was defined as the fifth percentile value of the membrane potential across the trace. We quantified durations of bumps by measuring the amount of time the membrane potential stayed above a given threshold. "Gaps" shorter than 20 ms were included into bumps, as they might have reflected "noisy" fluctuations of membrane potential. We used *absolute thresholds* of 10, and 15 mV above resting potential and computed durations of continuous voltage segments above given threshold. For each trace, *maximum potential amplitude* was defined as the difference between absolute maximum potential (in each trace) and resting membrane potential for the trace. We also used *relative thresholds* defined as 20 %, and 40 % of the maximum potential amplitude and computed durations of continuous voltage segments above threshold. The 20 % threshold corresponded to 7.33 ± 1.18 mV when expressed in absolute units. Note that such a threshold is probably quite low (see for example Petersen et al., 2003a; Stern et al., 1997) and thus may lead to an overestimation of the actual "bump" duration.

Because bump durations can vary from several milliseconds to about 1–1.5 s (especially for the lowest thresholds, see Fig. 5.4 on page 97) we divided the detected bumps into three classes: bumps shorter than 100 ms, bumps between 100–200 ms, and longer than 200 ms (up-states). For each neuron
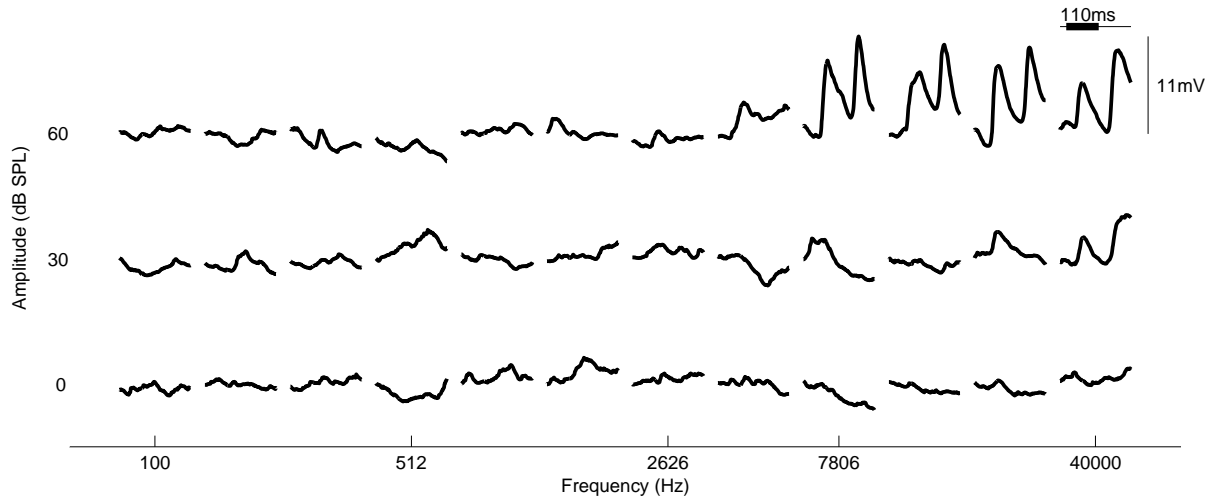
Figure 2.6: Subthreshold tuning curve of a neuron in unanesthetized auditory cortex. Frequency-intensity tuning curve shows robust subthreshold activity for this neuron. Note the prominent on and off responses, i.e. responses to the onset and the offset of some stimuli. The trace shown are average responses (of median-filtered traces) to 50 ms tones. Each trace contains 10 ms of spontaneous activity before the tone (see upper right corner).

we computed the fraction of bumps belonging to each class, as well as the fraction of up-time neurons spent in bumps (up-states) of different durations (Fig. 5.5).

Subthreshold tuning in awake auditory cortex

One of the basic characteristics of neurons in auditory cortex is the tuning curve. Subthreshold frequency-intensity tuning curve of one neuron is shown in Fig. 2.6. Each trace represents 10 ms of spontaneous activity followed by response to 50 ms tones of various frequencies and intensities. This neuron displayed prominent on and off responses (two bumps), and the responses covered about 4-5 octaves of frequency range at 60 dB SPL. The best frequency of this neuron was around 40 kHz, where best frequency was defined as the tone frequency which elicited discernible responses at threshold sound intensity (i.e. the lowest sound intensity applied still capable of eliciting responses). The tuning curve is comparable to subthreshold tuning curves from anesthetized preparations (Machens et al., 2004; Tan et al., 2004; Zhang et al., 2003), although no example [known to the author of this text] in auditory literature shows such prominent subthreshold off response.

27
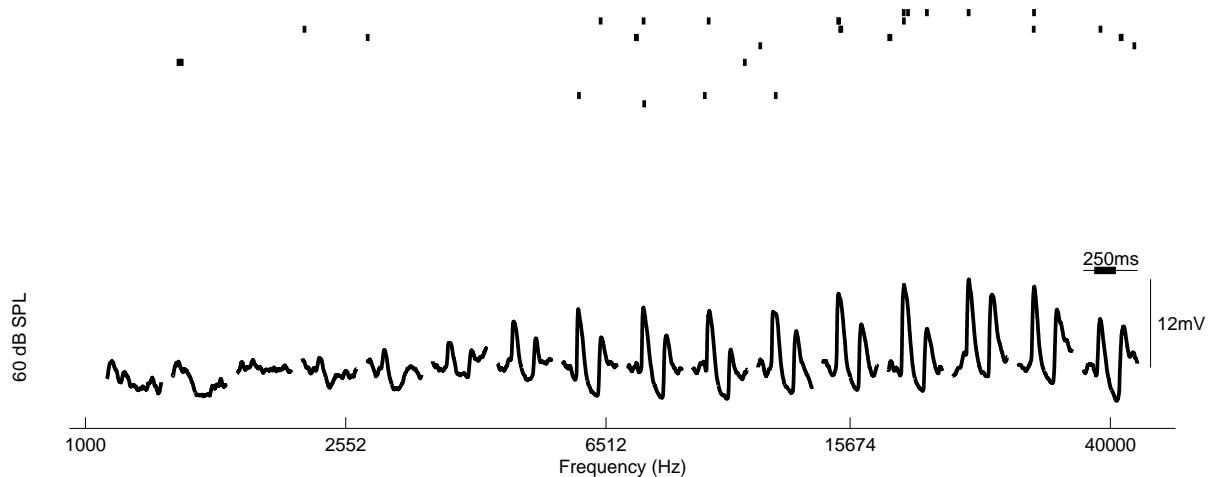
Figure 2.7: Subthreshold tuning can be broader than suprathreshold tuning. *Bottom* traces show average subthreshold responses to 100 ms tones presented at 60 dB SPL. Each trace contains 50 ms of spontaneous activity (see upper right corner). Responses to four nearby frequencies were averaged together for each average trace shown in this plot. *Top* rasters show positions of spikes detected in the raw voltage trace. Notice that detectable subthreshold responses covered wider frequency range than the spiking output. Note also the low spiking activity arising from robust subthreshold dynamics.

While subthreshold bumps represent a combination of synaptic inputs, spiking responses represent neuronal output. In Fig. 2.7 we compared frequency tuning of subthreshold responses and spiking (suprathreshold) responses. The neuron responded with prominent subthreshold on and off responses to 100 ms tones at 60 dB SPL, with responses covering frequency 2.5–40 kHz frequency range (4 octaves). However, the spiking output was rather sparse with only a few spikes elicited. Spike responses covered a frequency range of about 6.5–40 kHz, which was more than 1 octave less than the subthreshold tuning. Thus subthreshold tuning curves can be wider than suprathreshold tuning curves, which suggests that neuronal synaptic inputs cover wider frequency range than can be detected using extracellular recording techniques (DeWeese and Zador, 2000; Ojima and Murakami, 2002).

28

2.6.4 Cell counts

Cell-attached recordings

Results summarized in Sec. 3 on page 44, and Sec. 4 on page 73 are based on analysis of cell-attached recordings in awake head-fixed rats. We recorded from 166 neurons[1] (in 25 animals) out of which we identified 145 neurons (in 24 animals) with at least 8 trials per octave bin.

For the analysis of *stimulus-evoked changes in firing rate* we identified neurons with at least *8 trials* per response bin (*5 trials* for natural sounds). For the analysis of *significance of stimulus-evoked responses* we identified neurons with at least *20 trials* per response bin.

Tones

We recorded from 166 neurons (100 %), while presenting pure-tone pips. For further analysis of firing rates evoked by 50 or 60 dB tones we identified 145 neurons (87 %) with at least 8 trials per response bin. For the analysis of evoked response significance we further identified a subset of 100 neurons (60 %) with at least 20 trials per response bin.

For 91 neurons (55 %) we also presented 30 or 40 dB tones. All of these neurons were used for the firing rate analysis, and 62 neurons (37 %) from this subset—those with at least 20 trials per response bin—were used for the analysis of evoked response significance. Accordingly, out of 43 neurons (26 %) presented with 80 dB tones we selected 22 (13 %) for firing rates analysis, and 6 (4 %) with at least 20 trials per octave bin for the analysis of evoked response significance.

Frequency modulated sweeps

Frequency modulated sweeps were presented for 22 (100 %) neurons, all of which were used for the firing rates analysis. Seventeen neurons (77 %) with at least 20 trials for each sweep rate and direction were further selected for the analysis of significance of sweep-evoked responses.

Natural sounds

Natural sounds were presented for 28 neurons (100 %). Twenty-seven neurons (96 %) with at least 5 trials for each natural sound segment were identified for the analysis of stimulus-evoked firing rates.

---

[1]Thirty-one recordings (19 %) were performed by Michael R. DeWeese.

29

Whole-cell recordings

Results summarized in Sec. 5 on page 89 are based on analysis of data obtained using whole-cell recordings in awake head-fixed rat. Altogether, we report results from 20 intracellular whole-cell recordings[2] in the unanesthetized auditory cortex of 13 Sprague-Dawley rats (24–30 days old). One recorded trace has been used as an example in Fig. 2 in DeWeese and Zador (2006).

2.7 Simulation experiments

2.7.1 Discriminability of firing patterns

To compare stimulus representations in sparse and dense neuronal populations we computed the similarity of pairs of spike patterns that were either both generated by neurons with firing rates drawn from lognormal (sparse) distributions, or both generated by neurons with firing rates drawn from truncated Gaussian (dense) distributions. For the simulated lognormal distribution of firing rates we used parameters given by the distribution of spontaneous firing rates (Sec. 3.3), with a mean of 1.3, and a standard deviation of 1.0, both on a logarithmic scale. To create a corresponding (truncated) Gaussian distribution of firing rates we matched the mean firing rate and entropy of the lognormal distribution, which corresponded to a Gaussian with a mean of 4.2 sp/s and a standard deviation of 5.2 sp/s, on a linear scale, with negative firing rates replaced by rates drawn again from the same distribution until the distribution contained only non-negative firing rates.

To simulate responses of neuronal populations, we first drew two patterns, X and Y, of firing rates—each of these rate patterns was a vector of n = 200 values, representing the firing rate of each neuron in the population. For the sparse patterns each element of each vector was drawn from the sparse distribution; similarly for the dense patterns, each element was drawn from the dense distribution. We then generated 100 individual spike patterns from each rate pattern by treating each element as the rate (in a 10 ms window) of a Poisson process.

---

[2]Five recordings (25 %) were performed by Michael R. DeWeese.

We defined the *discriminability* $q(X, Y)$ between a pair of rate patterns (i.e. between two sets of firing rates X and Y over a population of neurons) as:

$$q(X, Y) = \frac{\langle x \cdot x' \rangle \langle y \cdot y' \rangle}{\langle x \cdot y \rangle^2} \qquad (2.3)$$

where x and x' are patterns of spike counts drawn from X, and y and y' are spike patterns drawn from Y; the brackets denote averages over all pairs of spike patterns (instantiations of the Poisson spike trains). We then used a *discriminability score (Q)*—the average discriminability over pairs of rate patterns—to quantify how different spike patterns drawn from the different underlying distributions were:

$$Q = \langle q(X, Y) \rangle_{rates} \qquad (2.4)$$

We thus could compare Q for patterns X and Y drawn from a sparse distribution to Q for patterns X and Y drawn from a dense distribution.

The higher the discriminability score, the more discriminable are the spike patterns drawn from one pattern of rates compared to another pattern of rates. A discriminability score close to 1 means that spike patterns from one rate pattern are (on average) the same as spike patterns from another rate pattern.

2.7.2 Hebbian learning for sparse and dense firing rate representations

To compare learning of neuronal activity patterns arising in sparse (lognormal) or dense (Gaussian) distribution of firing rates we used a single neuron model with Hebbian synapses. Parameters for sparse and dense distributions were estimated from data, as described in the previous section.

From the firing rate distributions we generated a *training set* of k firing rate patterns, each consisting of n neurons with firing rates chosen randomly from the same firing rate distribution. One neuronal pattern was randomly chosen to be the *target* pattern, and the remaining patterns were labeled as *nontarget* patterns. Every other pattern in the training set was then replaced by the target pattern. From the training set of firing rate patterns we then generated a set of spike patterns representing population outputs during 10 ms windows: each firing rate was replaced by a number of spikes generated by

a Poisson process (λ = firing rate) in a 10 ms window. Each input spike pattern in the training data set represented a 10 ms snapshot of neuronal spiking activity. Note that while the target *firing rate patterns* were identical, the actual target *spike patterns* were not (although they were similar, see also previous section) because they were generated by a stochastic Poisson process.

We simulated learning in a single sigmoidal neuron with n inputs (corresponding to n neurons in the input patterns) and one output. In each trial (a single presentation of input pattern) the sigmoidal neuron computed its output (response) as a weighted sum of its inputs transformed by a sigmoidal function. After the trial, the neuronal response was used to compute a new set of synaptic weights.

Thus, the neuronal response y^t in trial t was computed as:

$$y^t = \sigma(\sum_{i=1}^{n} w_i^{t-1} x_i^t), \tag{2.5}$$

where w_i^{t-1} are current synaptic weights (computed in the previous trial) associated with current inputs x_i^t, and $\sigma = 1/(1 + e^{-100z+6.2})$ is a sigmoid function with z as parameter (i.e., $z = \sum_i w_i x_i$).

The weights w_i were initialized with values from 0 to 1 drawn from a uniform distribution, and then normalized so that $\sum w_i = 1$. After the presentation of each input pattern, weights were adapted according to a Hebbian learning rule:

$$w_i^t = w_i^{t-1} + \eta\ y^t x_i^t \tag{2.6}$$

where $t = 1, \ldots, k$ denotes trials (presentations of individual input patterns), η is the learning rate, y^t is the current neuronal response (i.e. , postsynaptic activity), x_i^t denotes the i^{th} input in the current trial (i.e. , presynaptic activity), w_i^0 is the set of initial synaptic weights. w_i^k is thus the final set of synaptic weights after presenting all k input patterns. We repeated the simulation experiment 1000 times, each time drawing a different set of $n = 100$ neurons, and $k = 100$ trials. The value at each point (trial t) of Fig. 3.15C was computed as $\sigma(\sum_i w_i^t x_i^t)$ for each experiment and then averaged across experiments. The line thickness in panel C represents the standard error of the mean.

2.8 Head-fixed behavior

We have developed sound discrimination go/no-go task to study effects of *vigilance/alertness* on neuronal responses in auditory cortex in head-fixed rats (Fig. 2.8). In this section we describe the structure of the task, and of individuals trials, and training procedures.

2.8.1 Sound discrimination go/no-go task

In the sound discrimination go/no-go task, rats were trained to respond (i.e. lick) to target sounds while listening to a continuous stream of 500 ms sounds (separated by silent intervals). The rats were also trained not to respond to distractor sounds or during inter-trial intervals. We have used pure tones and warbles (frequency modulated tones) as stimuli (see also Sec. 2.5). Some rats were trained with tones as targets and warbles as distractors; for some rats we presented warbles as targets and tones as distractors. For each rat, however, we kept the task contingencies the same. In the following description we assume that warbles were presented as targets, and pure tones were presented as distractors. Distractor tones were of different (eight) frequencies to ensure that rats were not discriminating between concrete sounds, but rather discriminating different *quality* of sounds, i.e. tone vs frequency modulation. Because rats did not initiate trials, we have defined *trial* in this task as inter-stimulus interval, i.e. the period from the beginning of one sound stimulus to the beginning of the next one. Each trial thus contained a single sound stimulus.

Target and distractor trials were randomly interleaved. *Target* trials started with a target sound, followed by the *reward period* (Fig. 2.9.) Responses (i.e. licks) during the target sounds or reward period were rewarded with a drop of water. The reward period was followed by a *penalty period,* duration of which was drawn randomly from an exponential distribution with a mean of 1–2 s, therefore, the duration of the penalty period occasionally exceeded 10 s or more. The penalty period was followed by beginning of the next trial, and thus corresponded to the inter-trial interval in a general sense. Any licks during the penalty period were penalized by an air-puff (500 ms) and a restart of the penalty period (with the same parameters.) *Distractor* trials started with a distractor sound immediately followed by a penalty period. Therefore, any licks during distractor trials were penalized by air-puffs and restarted the penalty period (if the lick happened during the penalty period).
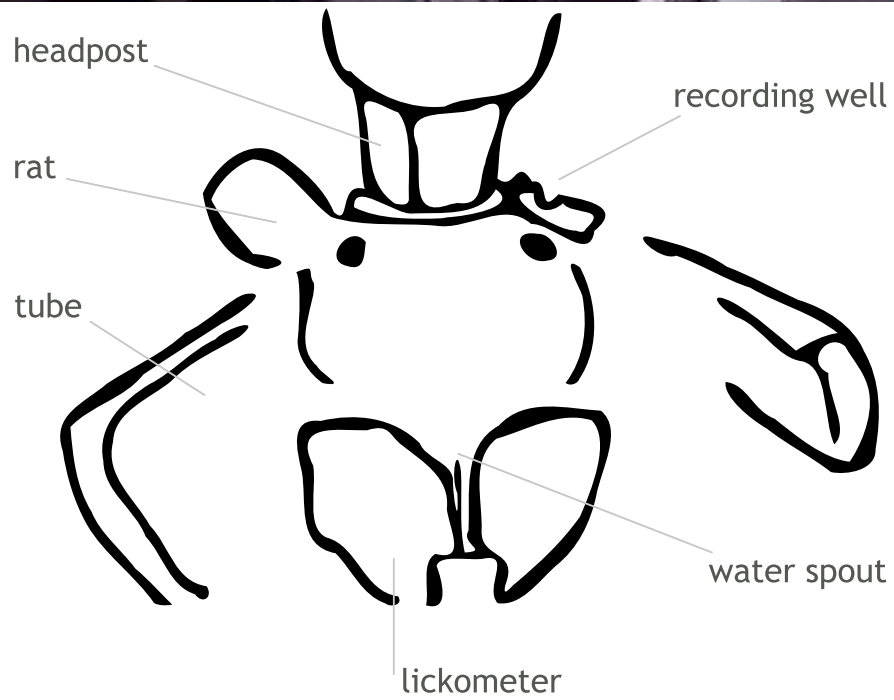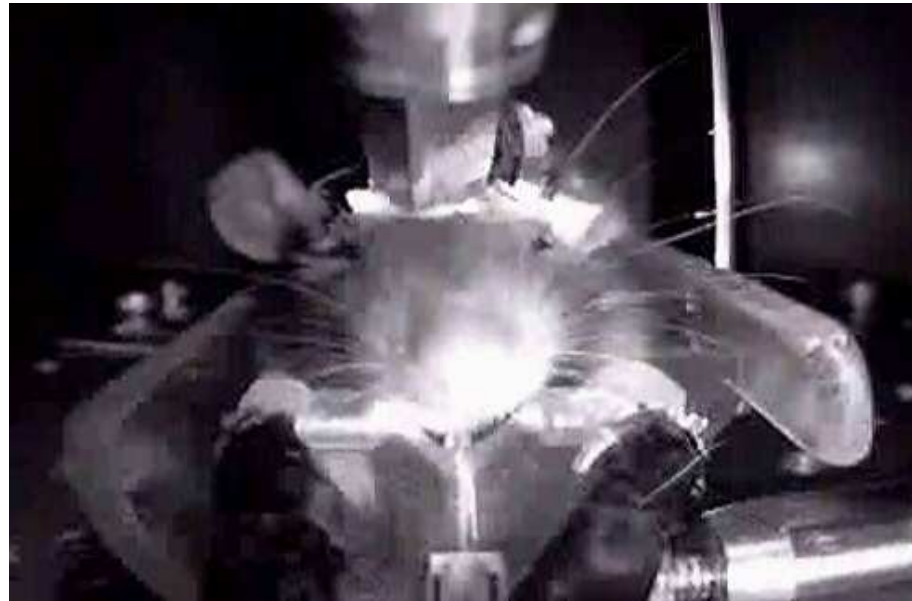
Figure 2.8: Training and recording setup for head-fixed behavior. During (re)training and recording head-fixed sessions, a rat was positioned inside a plastic tube with a lickometer placed in front of his/her/its mouth.
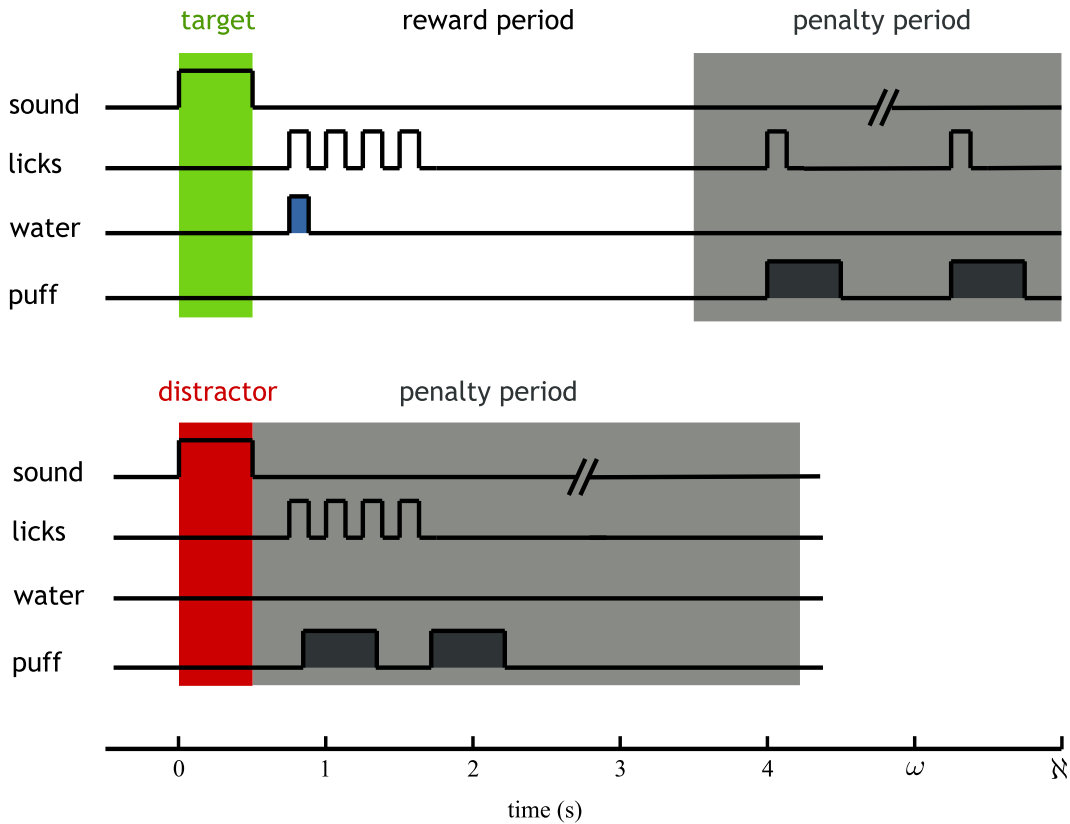
Figure 2.9: Trials of sound discrimination go/no-go task. Target trials (top) started with 500 ms target sound, followed by a 3 s reward period, during which responses (i.e. licks) were rewarded with a drop of water. The reward period was followed by a (random duration) penalty period, during which licks were penalized by air-puffs. Distractor trials (bottom) started with a 500 ms distractor sound, immediately followed by a penalty period. Any licks during distractor trial were penalized by an air-puff. Target and distractor trials were randomly interleaved during a behavioral session. See also Fig. 6.1.

In our task, rats were then trained to *discriminate* different sounds by responding to target sounds and not responding to anything else. In addition, rats were trained to *"understand"* task structure, i.e. wait for a sound and lick only when a (target) sound is presented.

We computed *performance* in the task as a measure combining how well a rat can discriminate sounds, and how well a rat can follow task structure. We defined *responses* (r_i) in individual trials (t_i) as 1 for correct trials, and 0 for error trials. Correct trials included target trials during which the rat licked in response to a target, and distractor trials during which the rat did not lick. Error trials included the rest of the trials.

35

Discrimination performance (d) was computed as a running average of responses r_i in k consecutive trials:

$$d_i = \frac{1}{k} \sum_{j=i-k+1}^{i} r_j, \tag{2.7}$$

where $i \geq k$. We have used $k = 20$ trials for our computations. Therefore, $d_i = 1$ for 20 consecutive correct trials, and $d_i = 0$ for 20 consecutive error trials.

We computed a *penalty period extension* (ε) as a measure of how well rats followed task structure. For each trial we computed ε_i as a ratio of the actual duration of the penalty period ε_i^{act} and the intended duration of the penalty period ε_i^{int} for that trial:

$$\varepsilon_i = \frac{\varepsilon_i^{act}}{\varepsilon_i^{int}}, \tag{2.8}$$

where $\varepsilon_i = 1$ when the rat followed trial structure, and $\varepsilon_i > 1$ when the rat licked outside the reward period. Recall that for each trial the duration of the penalty period was drawn randomly before the trial (*intended* duration). Whenever the rat licked during the penalty period, the period was restarted with the same parameters (in addition to air-puff). Therefore, when the rat *did not* follow the task structure (i.e. licked during penalty period), the *actual* penalty period duration was longer than intended and $\varepsilon > 1$. When the rat did not lick during the penalty period, $\varepsilon_i^{act} = \varepsilon_i^{int}$, and $\varepsilon_i = 1$. Then we computed $\bar{\varepsilon}_i$ as a running average of the penalty period extension ε_i in $k = 20$ consecutive trials: $\bar{\varepsilon}_i = (1/k) \sum_{j=i-k+1}^{i} \varepsilon_i$, where $i \geq k$.

We defined performance p_i in each trial as a ratio of discrimination performance to a running average of the penalty period extension:

$$p_i = \frac{d_i}{\bar{\varepsilon}_i}. \tag{2.9}$$

Performance, as we defined it, was bounded between 0 and 1. The best performance, $p_i = 1$, was achieved whenever the rat performed 20 consecutive correct trials, with no licks during the penalty period in trial t_i. On the other hand, $p_i = 0$ for twenty consecutive error trials, regardless of any (possible) licks during the penalty period.

2.8.2   Analysis of response parameters

We used the terms *lick* or *licking* to describe the "stereotypical oromotor ingestive responses, which include tongue protrusions, lateral tongue and mouth movements" (Whishaw and Kolb, 2005). Licks—as detected by our lickometers—corresponded to brief interruptions of an infrared beam in the lickometer (Sec. 2.3 on page 13). Licking was the only behavioral response, especially in the head-fixed position. We have, therefore, measured and compared several licking parameters (lick duration, lick latency, licking frequency) between freely moving and head-fixed sessions.

For analysis of *lick durations* in freely moving and head-fixed sessions we identified licks shorter than 200 ms. Longer licks were rare and usually corresponded to interruptions of the infrared beam by various body parts, for example tail in freely moving sessions. Licks identified in all *operant* freely moving sessions were pooled together for each animal, and licks identified in all head-fixed sessions were also pooled together. To compare distributions of lick durations between (operant) freely moving and head-fixed sessions we first divided duration distributions into 5 ms wide bins, and then used nonparametric Wilcoxon rank-sum test (Mann-Whitney U test) to compare the resulting discrete probability density functions.

We have defined *response latency* as the latency of the first lick in a response lick train (we only considered response latencies 10–1000 ms for this analysis). To compute response latencies in freely moving session we pooled together first lick latencies from all target trials in the last three operant freely moving sessions. Analogously, we pooled together first lick latencies from all target trials in head-fixed sessions. We then compared distributions of response latencies using the nonparametric Wilcoxon rank-sum test (Mann-Whitney U-test).

Licking frequency in freely moving sessions (head-fixed sessions) was computed by first aligning lick trains from target trials in all operant freely moving sessions (all head-fixed sessions). The aligned lick trains were added together to compute "laster PSTH," i.e. peristimulus time histograms from licking rasters. We then estimated the power spectral density of laster PSTHs using Welch's method of spectral estimation. Main licking frequency was then defined as the frequency of the maximum peak in 3–40 Hz range. The distribution of licking frequencies from freely moving sessions was compared to the distribution of licking frequencies from head-fixed sessions using Wilcoxon rank-sum (Mann-Whitney U-test).

2.8.3 Training protocol

Rats were usually trained in batches of four, housed together in one cage. All rats were picked randomly at postnatal age of 23–24 days. Initially we tried training rats as young as 20 days, but rats younger than 23 days usually sat passively until they reached 23–24 days of age. We did not specifically select for male or female rats.

Rats were water-deprived one day before the first training session, with otherwise free access to food. During training, rats received water during training sessions in the form of rewards, and for a limited time (20–30 min) each day at a randomly chosen time after the last session for that day. After the first day of water deprivation the rats lost about 10 % of their body weight. Eventually, after a few days, their growth resumed and rats gained 1–2 g of body weight a day (in the water-deprivation regime).

Animals were trained on a training platform similar to the recording platform described in Sec. 2.3, with the plastic tube closed on both sides to prevent the rats from escaping. As a *reward* for correct licking responses we used sugar water, 5% sucrose in millipore water, to ensure continuous water delivery without air bubbles. We have not found any difference in training progress or performance when using 5% aspartame solution, or pure water. We used 500 ms air-puffs to *penalize* rats for licking incorrectly.

Training for the sound discrimination task consisted of two main phases, freely moving and head-fixed. Rats were trained freely moving for the full task, and retrained in the head-fixed position after reaching sufficient performance.

The first two sessions in the freely moving phase were introductory. Rats were trained to associate a target sound with water reward in the *classical conditioning* paradigm. Each target sound was followed by a drop of (sugar) water, which was in turn followed by a random-length inter-trial interval. Rats were *not* penalized for licking incorrectly during these classical conditioning sessions.

Starting with the third session, we switched to an *operant conditioning* paradigm, and introduced air-puffs as penalty for incorrect licks. Rewards were no longer delivered automatically after each target, but rats had to initiate rewards by responding (licking) during the reward period. In most cases, rats switched seamlessly to operant conditioning; in rare cases we had to manually deliver several rewards as a "reminder." If a particular rat was licking actively during the third session, we introduced an air-puff as a penalty for incorrect responses. We started with shorter air-puffs (10–100 ms), and in-

creased air-puff duration to 500 ms, usually during the same session. After the first air-puff, most of the rats stopped performing for several minutes, but then gradually resumed licking.

We introduced a single distractor stimulus, randomly interleaved with targets after two, occasionally three operant conditioning sessions with targets only. We started with 30 % of distractors for one session, continued with 50 % of distractors in the following (sometimes the same) session, and rarely—for "very active" rats—used 70 % of distractors and 30 % of targets.

After about two more sessions with a single distractor we introduced another distractor in one session, two more distractors in the following 1–2 sessions, and finally four more distractors to reach 8 different distractors. Thus in about 10–12 sessions rats were exposed to the full task, and their discrimination performance usually reached at least 80 %.

One day after the last freely-moving session, rats underwent surgery (Sec. 2.2). We have performed surgeries both on water-deprived rats and rats which received water for one day before surgery, and found no difference. After surgery, rats were allowed at least 24 hours to recover, after which they were water-deprived again.

Two days (usually) after surgery, we started retraining the implanted water-deprived rats in a head-fixed position. Most rats started licking, and eventually performing the task, during the first 20–30 min in the head-fixed position. After two to three head-fixed retraining sessions the rats were ready for head-fixed recordings.

2.9 Data acquisition system

All data was obtained using a custom made data acquisition and stimulus delivery system written in Matlab. This section describes the design and implementation of a modular system for continuous data acquisition, with flexible on-line stimulus delivery.

Originally, data was obtained using a legacy trial-based system designed and used for recordings in anesthetized animals. However, with more and more recordings in awake head-fixed animals, the original system—although very useful for its original purpose—became unsuitable for three main reasons. First, a considerable portion of the electrophysiological signal—the inter-trial interval—was not recorded in the trial based system. Depending on the actual experiment design, inter-trial inter-

vals corresponded to 15–25 % of data. Because experiments in awake animals tend to be shorter than experiments in anesthetized animals, inter-trial intervals effectively represented lost valuable data. Second, all stimuli had to be prepared before the experiment. Any change in structure of desired stimuli (type, different parameters, different order of stimuli, etc.) led to resynthesizing of the stimulus file. Third, the system was designed to be rather monolithic. All components of the system were dependent on various other components, and communication among different components (change of parameters, change of path, etc.) was explicitly specified. Therefore any change in the system (adding functionality, adding recording channels, adding another sound channel, etc.) required direct changes in various (often unrelated) parts of the code, often introducing inconsistencies, or even errors.

We have designed a modular, event-based system to address limitations of the original data acquisition system. The system's functionality was split into four main subsystems: data acquisition, stimulus delivery, data processing, and behavioral control (Fig. 2.10). Each subsystem contained several modules, each providing a well-defined, separate task. Splitting the system's functionality into separate modules ensured that changes in different subsystems could be made without any changes in the other subsystems. Thus, the data acquisition subsystem contained modules for analog input (*ai*), analog output (*ao*), digital input/output (*dio*), etc. The stimulus delivery subsystem contained the *stimulusczar* module "supervising" stimulus delivery. Stimuli requested by another module (for example *stimulusprotocol*) were prepared (synthesized/loaded) by *stimulusczar* and uploaded to appropriate hardware by separate modules (for example *soundload*). The data processing subsystem provided means for online data processing by extracting relevant portions of data (*dataguru* module) from continuous stream delivered by *ai*. *Boxmaster* was the main module of the behavioral subsystem, providing behavioral control by uploading the task description, and continuously updating behavioral input supplied by the subject performing the task. The module thus served as an interface for other modules (behavioral protocols) providing on-line data analysis of behavior. Two modules provided additional functionality: *control* module provided basic information about current experiment, and was responsible for coordinating communication among other modules; and the *sealtest* module used to search for cells during patch-clamp recordings.

Communication among different modules was implemented as event-based. In this paradigm, each module implemented "responses" to specific, defined "events." When an event occurred, for ex-

control ←——————→ sealtest

daq
ai, ao,
dio,...

stimuli
stimulusczar,
stimulusprotocol,
soundload,
pulse,...

data
dataguru,
datawatch,
scope,...

behavior
boxmaster,
behavioral
protocols

other
timestamp,
cyberamp,
experlog
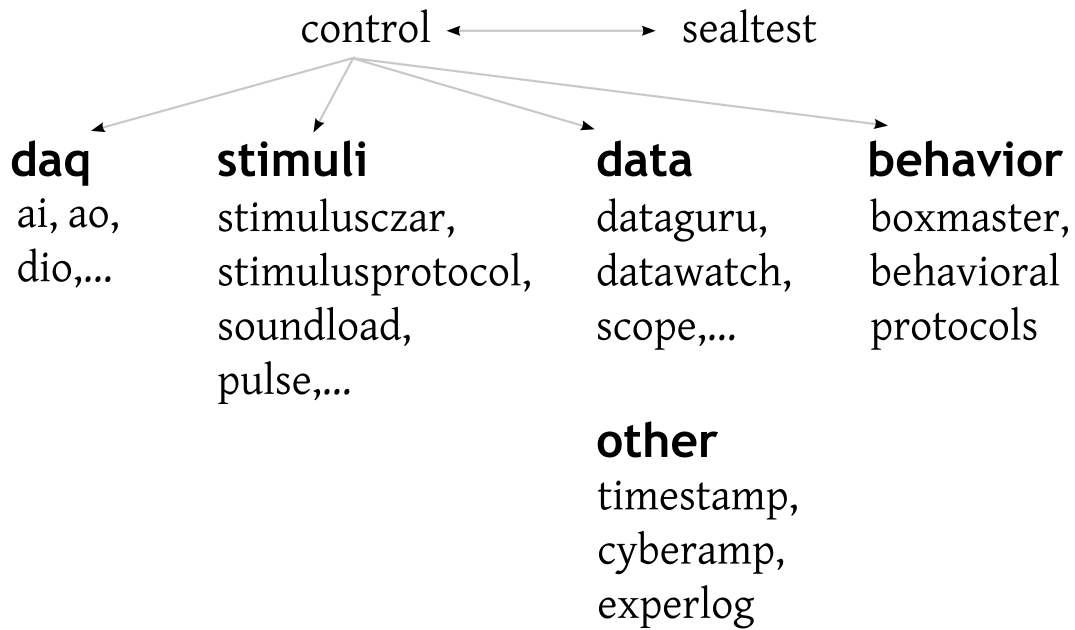
Figure 2.10: Structure of data acquisition system. Core functionality is split into four subsystems: daq (data acquisition), stimuli (stimulus delivery), data (data processing), and behavior (behavior control). Each subsystem contains several modules providing separate tasks. *Control* and *sealtest* are modules providing additional functionality. Other modules provide specialized tasks, not necessarily related to any of the main subsystems.

ample the user pressed a button to start recording, modules were notified that a particular event has occurred, and appropriate portions of the code in "responsive" modules were executed. Thus, events were essentially anonymous, and different modules could be replaced with completely different portions of the code, different (number of) modules, as long as the new code provided the same *functionality* as the original code.

The system was implemented in Matlab (version 6.5.1, Mathworks, Natick, MA), and requires Data Acquisition (DAQ) Toolbox (version 2.2) (Fig. 2.11). We have tried other (later) versions of Matlab and DAQ Toolbox, but due to the apparent Mathworks policy of consistently introducing inconsistencies in new versions of their products, all data was obtained using the original Matlab version. On the hardware side of data acquisition we used PCI-MIO-16E-4 data acquisition card (National Instruments Corporation, Austin, Tx).

We originally used RP2.1 Real-Time Processor (TDT, Alachua, Fl) for three purposes: to deliver sounds, to control behavior, and to create a timestamp signal. Delivering sounds was the main purpose. All sounds were created as Matlab vectors sampled at 97.656 kHz, uploaded to the RP proces-
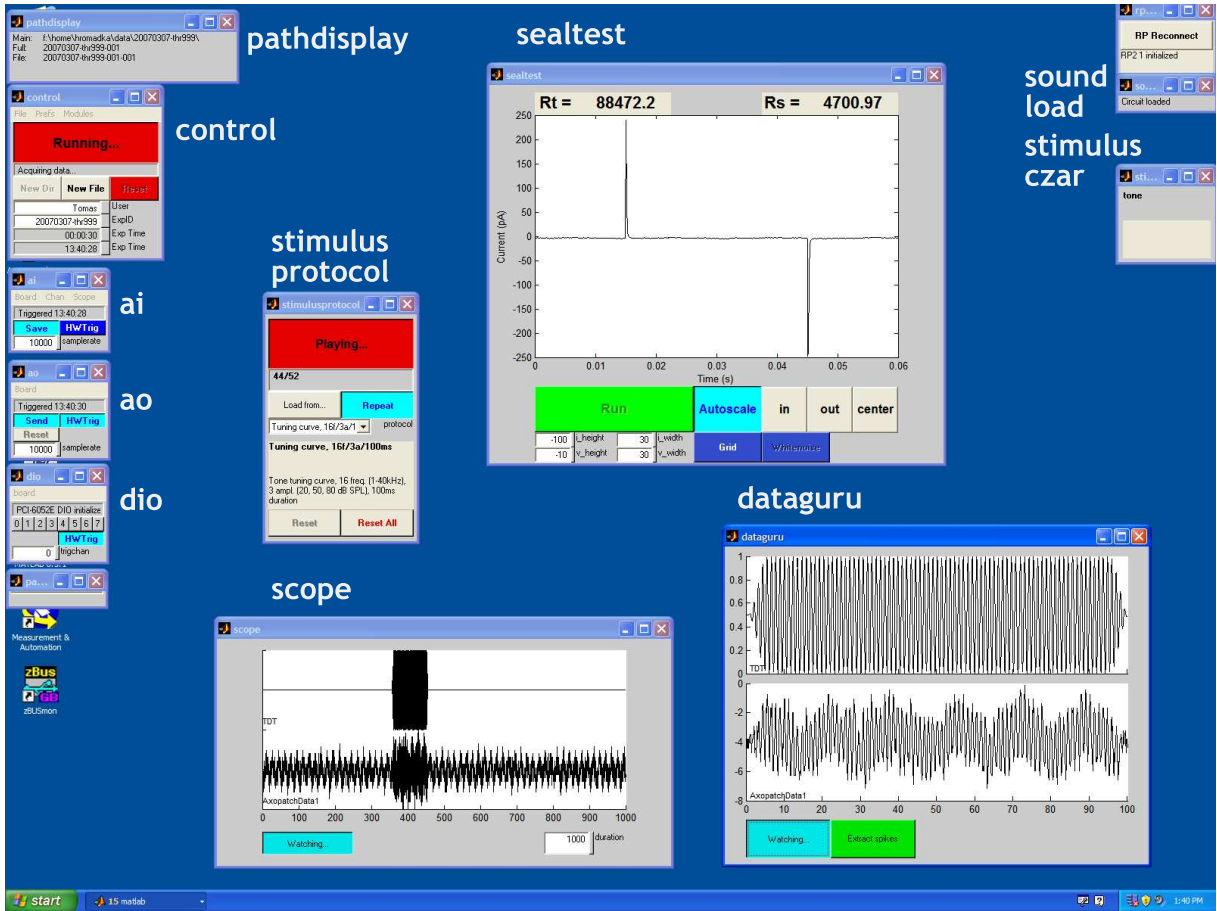
Figure 2.11: An example of the data acquisition session. Screenshot of a "dummy" data acquisition session, showing the actual implementation of the modular structure. Each module is represented by a single window. See text for details.

sor, and played after being triggered by the data acquisition card. Behavioral control was implemented in a separate RP processor, and used a state machine paradigm based on a design originally developed and coded by Lung-Hao Tai. Additionally, in some experiments we delivered a timestamp signal (a continuous stream of Manchester encoded 22-bit integers), used to synchronize the electrophysiological signal with the separately recorded video signal. The timestamp signal was recorded simultaneously in one of the recording channels and one audio channel of the video signal.

We have recently replaced RP processors with a standalone GNU/Linux system with a real-time kernel (version 2.14.29). Real-time linux system runs on a separate computer, and uses a high-end audio card (Lynx L22, Lynx Studio Technology Inc., Newport Beach, CA) to deliver sounds at 200 kHz. The system—behavioral state machine and Lynx sound card driver—was developed by Calin Culianu.

The modular, event-based design of our data acquisition system enabled us to smoothly make changes in a working environment by simply adding or replacing software modules (usually a single one) without making any changes in the system's core. Thus adding new functionality (behavioral subsystem), or replacing existing hardware (real-time GNU/Linux instead of RP processor) was straightforward. Finally, the versatility of the system is best documented by its various configurations as used by several members of our laboratory: patch-clamp recordings, dual patch-clamp recordings, behavioral training, patch-clamp recordings in behaving rats, extracellular tungsten recordings, recordings with rigid tetrodes (in anesthetized, awake naïve, and behaving rats); delivering sounds, current pulses, light via LEDs, etc.

Chapter 3

Sparse representation of sounds in auditory cortex of unanesthetized rats

How does a population of cortical neurons encode a sensory stimulus such as a sound? At one extreme, the neural representation could be dense, engaging a large fraction of neurons, each with a broad receptive field. At the other extreme, the neural representation could be sparse, at any moment in time engaging only a small fraction of neurons, each highly selective with a narrow receptive field. Although a dense code under some conditions makes the most efficient use of the "representational bandwidth" (DeWeese et al., 2005) available in a neuronal population—why should a large fraction of neurons remain silent most of the time?— sparse models have recently gained support on both theoretical (Asari et al., 2006; Olshausen and Field, 2004; Smith and Lewicki, 2006) and experimental (Baddeley et al., 1997; Brecht et al., 2004; Dan et al., 1996; Hahnloser et al., 2002; Margrie et al., 2002; Szyszka et al., 2005; Vinje and Gallant, 2000) grounds. However, it is not at present clear which of these is a better model of sensory representations in the auditory cortex. In order to distinguish between these alternatives experimentally, we must know what fraction of neurons responds to a given stimulus.

The direct experimental approach to measuring the density of a cortical code would begin by simultaneously recording sound-evoked responses of all the neurons in the auditory cortex to an ensemble of stimuli; one could then simply count the number of spikes elicited by each stimulus. Unfortunately, currently available recording techniques do not permit such a direct approach. An alternative approach is to record the activity of a representative subset of neurons serially, and infer the population response

44

from this sample. In this way, the population code could in principle be inferred by sequentially sampling a large population of single unit responses.

We have used cell-attached recording in the primary auditory cortex of unanesthetized rats to sample the population response to brief tones and other stimuli. Because we were interested in the population response, we presented the same stimulus ensemble to each neuron, rather than optimizing the stimulus to drive each neuron to fire maximally (O'Connor et al., 2005; Wang et al., 2005). Thus, we could assess the fraction of neurons that responded to each stimulus we presented. Our data, therefore, address the question: "What is the typical response across the entire neuronal population to a particular stimulus?" rather than: "What is the optimal stimulus for a particular neuron?"

We find that the typical population response in unanesthetized auditory cortex is sparse. Consistent with previous findings in barrel cortex (Margrie et al., 2002), some neurons had very low spontaneous firing rates (≤ 0.01 spikes/s); at the other extreme, some neurons had driven rates in excess of 50 spikes/s. However, a given stimulus typically elicited a high firing rate (≥ 20 spikes/s) in only about 3 % of the population. Note that sparseness as used here refers only to the fraction of neurons active at a given instant; it is quite possible that each neuron might, under the appropriate conditions (e.g. when presented with an optimal stimulus), participate in a representation by firing at a high rate. Our results represent the first quantitative experimental support for the hypothesis that the representation of sounds in the auditory cortex of unanesthetized animals is sparse.

3.1 Results

We recorded responses of neurons in the auditory cortex of head-fixed unanesthetized rats. Because our approach was to construct the population response one neuron at a time, we did not optimize the stimulus ensemble to conform to the response properties of each neuron, but instead probed many neurons with the same ensemble. In this way, we could reconstruct the overall population response.

The primary stimulus ensemble consisted of pure-tone pips of different frequencies and amplitudes (145 neurons in 24 animals, see Sec. 2.5). For some neurons we also tested additional stimulus ensembles, consisting of either the same tones presented at different amplitudes (113 neurons), frequency modulated sweeps (22 neurons), or natural sounds (27 neurons, see Sec. 2.5, and Sec. 2.6.4 for details). Recordings were at all cortical depths (Fig. 3.1).
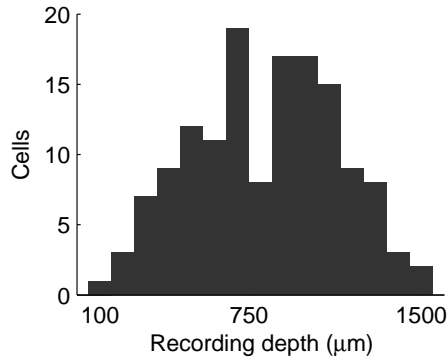
Figure 3.1: Frequency histogram of recording depths shows that all depths were represented in our sample. All depths (n = 141) were given by the micromanipulator reading, and as such are only approximate. For four cells the depth information was not available, and they are not included in the figure.

Our goal was to record the responses to stimuli generated by a representative sampling of neurons in the auditory cortex. We therefore chose to record with a glass patch pipette in cell-attached mode, a method which is not explicitly biased toward active and responsive neurons, or neurons with large action potentials, and which provides excellent single unit isolation (DeWeese et al., 2003; Margrie et al., 2002) (see Sec. 2.6.1 on page 20). With cell-attached recording, single unit isolation depends on the physical contact between the glass electrode tip and the neuron. The selection bias of cell-attached recording is thus based on the neuron's "patchability," rather than on the firing rate or responsiveness of the target neuron; only to the extent that patchability is correlated with functional characteristics such as firing rate or responsiveness would cell-attached recording (indirectly) bias the sampled population. By contrast, good single unit isolation with conventional extracellular (e.g. tungsten; Hubel, 1957) electrodes requires a sufficient number of spikes; skilled practitioners typically search for neurons with sufficiently high firing rates and large spikes. Although it is possible for a committed investigator to isolate even neurons with a low spontaneous firing rate, for the purposes of this study, cell-attached recording seemed a particularly suitable choice.

3.2   Neuronal responses are heterogeneous

Consistent with the earliest studies of unanesthetized auditory cortex (Davies et al., 1956; Evans and Whitfield, 1964; Hubel et al., 1959), tones evoked a wide range of response patterns. Tones could elicit either an increase or a decrease in a neuron's firing rate (compared to the background firing rate), or

46

both; the change could be transient, delayed, or sustained; and the response pattern could be different for different frequencies in a single neuron.

Fig. 3.2 shows some examples of the range of response types we observed. In one neuron (Fig. 3.2A), tones elicited a transient, short latency response of the sort commonly observed in the barbiturate-anesthetized auditory cortex. In a second neuron (Fig. 3.2C), tones elicited a suppression of background activity. In a third neuron (Fig. 3.2B), higher frequency tones (~8-40 kHz) elicited vigorous sustained firing; interestingly, lower-frequency tones elicited transient responses in the same neuron, emphasizing that the distinction between "transient" and "sustained" applies to responses, not neurons. Other more complex response patterns were also observed (Fig. 3.2D-G.) Finally, half of the neurons tested (50 %, see below) showed no change in firing rate for any stimulus presented (Fig. 3.2H). Because a given neuron could show very different response patterns to stimuli of different frequencies (e.g. Fig. 3.2B), we could not find a simple and objective scheme for organizing neurons into a small number of distinct classes, such as "transient," "sustained," "off," etc.The neurons shown in Fig. 3.2 are a subset; the complete set of responses from the entire data set are shown in Appendix A.

3.3 Population response is lognormally distributed

We first analyzed the basic population response elicited by tones, beginning with the response to tones presented at 50 or 60 dB. We divided the tone-evoked response into four 50 ms long "epochs": *spontaneous, early, late, and off* (Fig. 3.3, also see Sec. 2.6.2). To ensure a sufficient number of trials for assessing the statistical significance of putative changes in firing rate compared to the background, we grouped responses across nearby frequencies (one-octave-wide bins; 4 or 5 octave bins for each response epoch). Control analyses using narrower (half-octave) bins gave similar results (see Sec. 2.6.2), as expected from the relatively broad frequency tuning of neurons in the rat primary auditory cortex (Kilgard and Merzenich, 1999; Sally and Kelly, 1988); see also Moshitch et al. (2006).

Both spontaneous and evoked firing rates were typically low (see Tab. 3.1 and Fig. 3.4). The median spontaneous firing rate across the population was 2.8 spikes/s. The mean was somewhat higher (4.9 spikes/s) because it was dominated by a relatively small set of neurons—possibly interneurons (see Sec. 3.8 on page 62)—with high spontaneous rates.

47
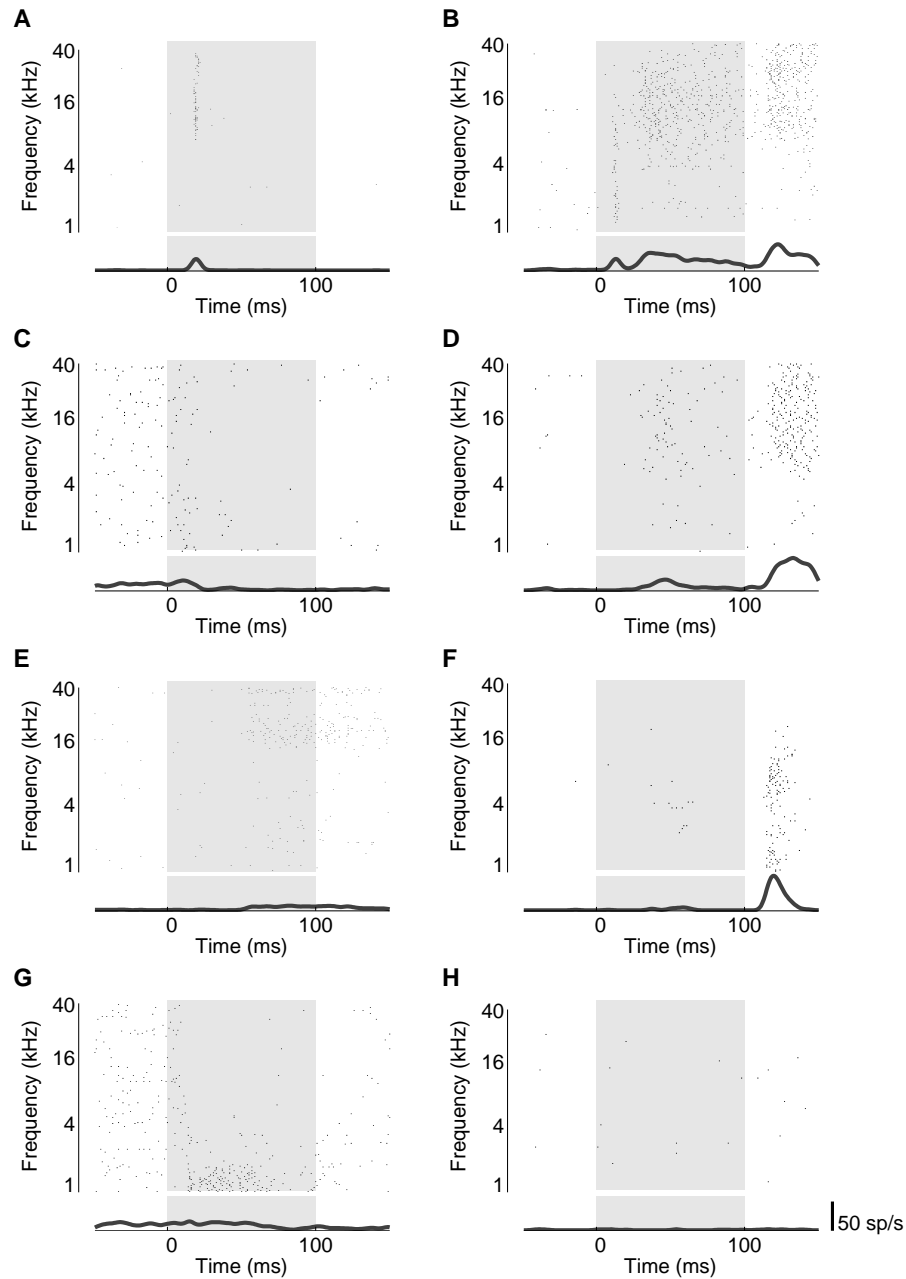
Figure 3.2: Tone-evoked responses in the auditory cortex of unanesthetized rats are heterogeneous. The panels show response dynamics of eight representative neurons to 60 dB tones. Firing rate curves in the bottom of the rasters were computed by first summing the spikes in 1 ms bins over *all* frequencies shown, and then convolving the resulting peristimulus time histogram (PSTH) with a Gaussian ($\sigma = 3$ ms). In each panel, dots represent individual spikes, the grey shaded region indicates the tone duration (100 ms). (**A**) transient onset response (maximum firing rate 22 spikes/s); (**B**) transient onset response followed by sustained excitatory response followed by off response (50 spikes/s); (**C**) suppressive response (19 spikes/s); (**D**) late onset response followed by strong off response (62 spikes/s); (**E**) late onset response (9 spikes/s); (**F**) off response (66 spikes/s); (**G**) sustained response combined with suppressive response (17 spikes/s); (**H**) non-responsive cell (2 spikes/s). See also Appendix A for more examples.
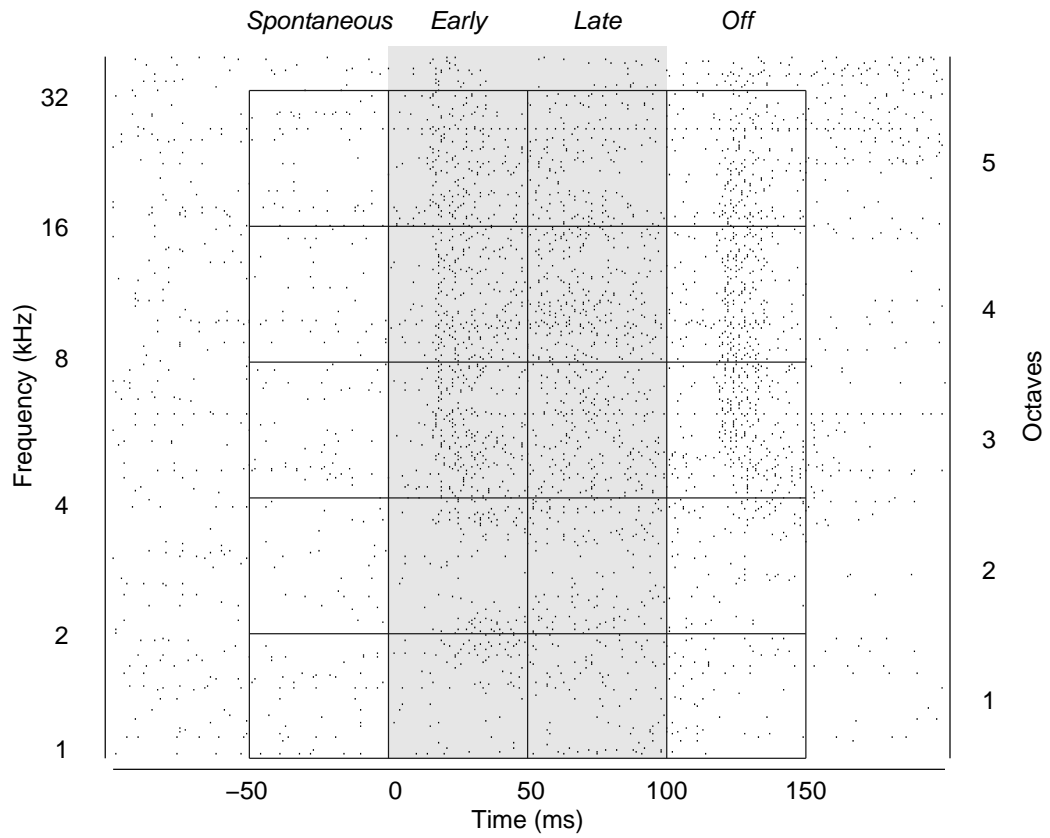
Figure 3.3: Cells were characterized by their activity during each of the response epochs: *spontaneous*, *early*, *late*, and *off*, each 50 ms long. *Spontaneous* epochs cover spontaneous activity before the stimulus, *early* and *late* epochs cover first and second half of the stimulus duration (100 ms) respectively, and *off* epochs cover 50 ms period after stimulus termination. In frequency space, individual trials were grouped into one-octave-wide bins, and averaged to provide a firing rate value for each octave bin. This figure shows a spike raster plot for an example neuron (with a sustained excitatory response), where each row represents a single trial, and each dot marks the occurrence of a spike. Shown are responses to 1–40 kHz tones (60 dB SPL, left ordinate.) Individual trials were grouped into 5 spontaneous, and 15 evoked response bins (right ordinate.) Note that the top quarter of an octave is not included in any of the bins.

Epoch	Firing rate (spikes/second)	
	median	mean±std
Spontaneous	2.8	4.9±6.1
Early	2.7	7.0±11.2
Late	2.0	5.4±9.3
Off	2.3	6.0±10.1
Early + Late[1]	2.4	6.2±10.3
All evoked epochs[2]	2.4	6.2±10.2

Table 3.1: Firing rates for different response epochs were typically low. Each value was computed across all octave bins from all 145 cells for the corresponding epoch(s) (n = 693 octave bins for each epoch). Firing rate for a given octave bin was defined as mean firing rate of all trials grouped inside that octave bin. The table shows firing rates of responses to 50 or 60 dB SPL tones (n = 145 neurons). [1]"Early + Late" values were computed across all octave bins from *early* and *late* response epochs, which cover the entire stimulus duration (100 ms). [2]For each cell, firing rates across all response bins were pooled to give an estimate of evoked firing rate.

Evoked firing rates showed the same pattern: a low median (2.0–2.7 spikes/s) and a somewhat higher mean (5.4–7.0 spikes/s). The higher mean rates reflect the fact that in some neurons, some frequencies evoked vigorous firing (Fig. 3.4, bottom). However, such well-driven responses were the exception rather than the rule; as quantified below, most neurons did not respond vigorously to any of the tones presented. Note that for a neuron to contribute on average at least one spike to the population representation of a sound in a 50 ms window, its evoked firing rate must exceed 20 spikes/s.

To assess whether the low firing rates resulted from some intrinsic defect of the spike generating mechanism, perhaps introduced by the cell-attached recording method, we extracted the shortest inter-spike interval for each neuron. In most neurons, the shortest ISI was less than 10 ms (median shortest ISI = 4 ms, n = 145). Thus the low firing rates do not appear to arise from an intrinsic inability of neurons to fire rapidly, but instead presumably arise from differences in the synaptic drive received by different neurons.

The distribution of spontaneous firing rates across the population was remarkably well fit with a lognormal distribution—that is, the logarithm of the firing rates showed a Gaussian distribution (Fig. 3.4BC). Because the lognormal distribution has a "heavy tail," most spikes were generated by just a few neurons: About 16 % of neurons—the subset of 23 neurons firing at higher than 9.5 spikes/s—accounted for 50 % of all spikes.
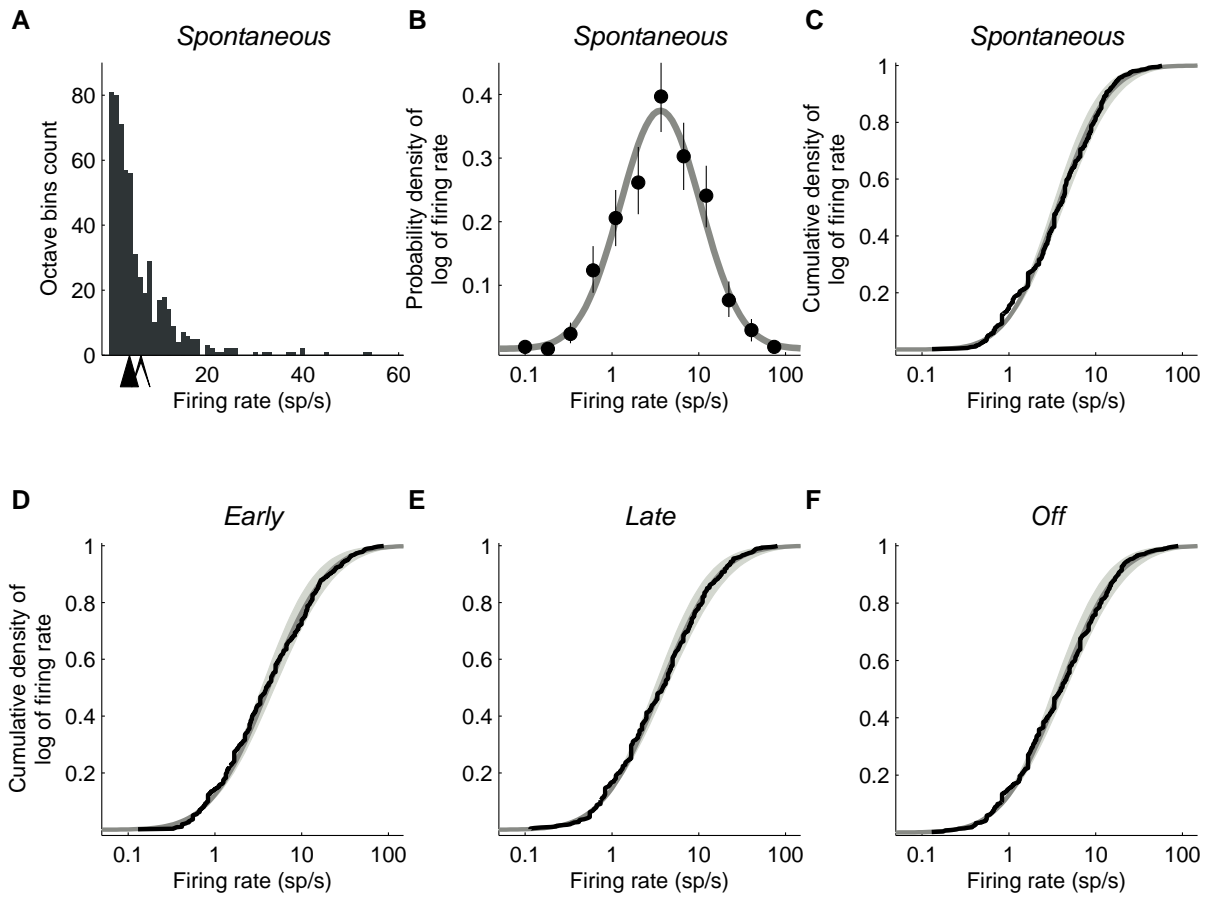
Figure 3.4: Firing rates of most neurons were low and followed a lognormal distribution. **(A)** Frequency histogram of nonzero spontaneous firing rates in individual octave bins (n = 567 octave bins, from 145 neurons). Each neuron contributed a maximum of 4 or 5 data points (because each neuron had 4 or 5 octave bins per epoch). The filled arrow shows the position of the median spontaneous firing rate, and the open arrow shows the position of the mean spontaneous firing rate. **(B)** The distribution of spontaneous firing rates (dots) was fit with a lognormal distribution (grey line), the mean and variance of which were given by the mean and variance of the original firing rate distribution (see Sec. 2.7). The lognormal distribution appears as a normal distribution on a (semi-)logarithmic scale. The error bars show 95 % confidence intervals determined by bootstrapping. **(C)** Same data as in (B) plotted as a cumulative density function of the spontaneous firing rate (dark line) overlayed with the cumulative density function of the lognormal fit (grey line). **(D), (E), (F)** The cumulative density functions of evoked firing rates for *early* (D), *late* (E), and *off* (F) response epochs also followed lognormal distributions (same format as in (C).
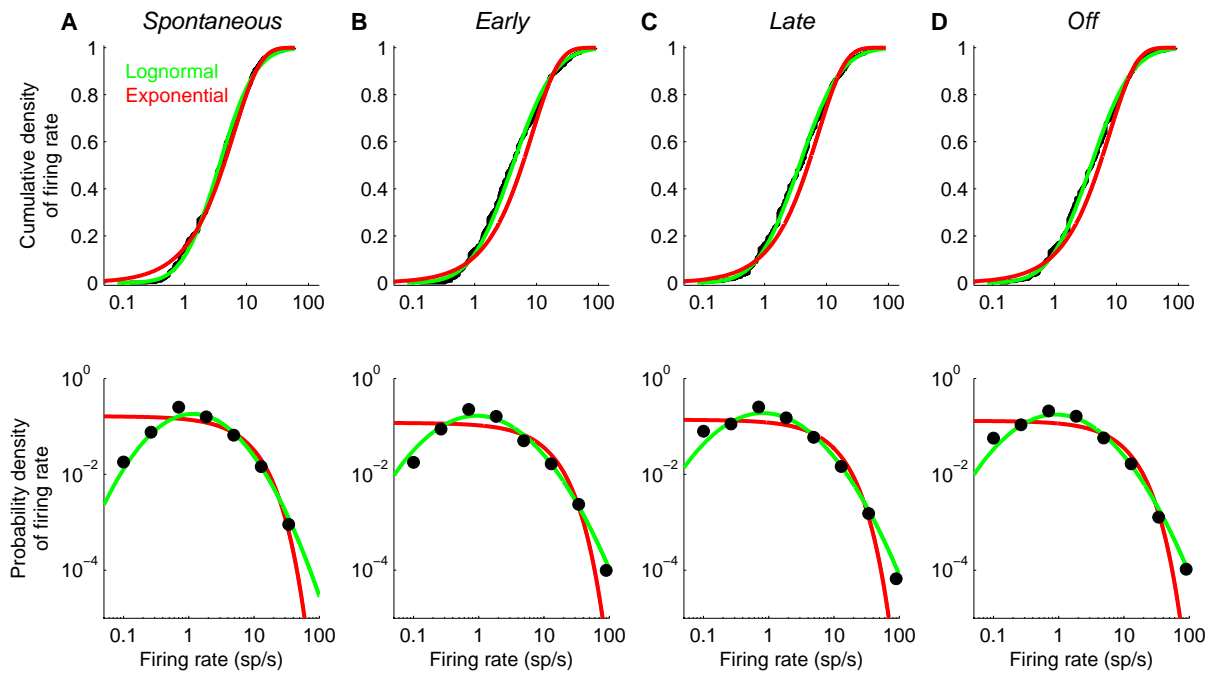
Figure 3.5: A lognormal distribution provides a better fit to the data than an exponential distribution.
(A) The cumulative density (top) and probability density (bottom) functions of the data (black points)
are better fit by a lognormal distribution than an exponential distribution for *spontaneous* firing rates.
The mean and standard deviation of the lognormal fit were given by the mean and standard deviation
of the distribution of (natural) logarithms of the firing rates. The mean of the exponential fit was given
by the mean of the firing rate distribution. **(B)**, **(C)**, **(D)** Same format as in (A) for *early, late* and *off*
epochs respectively.

Since the observed firing rate distributions contained many more low firing rates than high firing rates we compared lognormal and exponential fits to the observed firing rate distributions (Fig. 3.5). Parameters of both fits were given by parameters of the observed firing rate distributions. That is, for each epoch the mean and standard deviation of lognormal fit were given by mean and standard deviation of distribution of (natural) logarithms of firing rates, and the mean of the exponential fit was given by the mean of the observed firing rate distribution. Lognormal distribution provided a better fit to the data than exponential distribution in each response epoch. Both cumulative density and probability density functions of the lognormal fits provided excellent approximations to the data.

Although lognormal distributions have widely been used to describe the interspike interval distributions from a single neuron, population responses are usually reported to be exponentially distributed (Abeles et al., 1990; Baddeley et al., 1997; Gaese and Ostwald, 2003); this is, to our knowledge, the first report that firing rates across a population of neurons are lognormally distributed.

3.4   Population response is sparse

What is the typical response across the entire neuronal population to a particular stimulus? That is, what fraction of the neurons responds to a given stimulus? Fig. 3.6 depicts the evoked population response to a tone at each epoch (*early, late, off;* see Sec. 2.6.2). Each column shows the fraction of the neuronal population responding in a given response bin, and can be thus interpreted as a snapshot of neural activity over the population at a given instant in time.

The typical stimulus-evoked population response was sparse. The majority of neurons showed no discernable response to any tone during any response epoch (Fig. 3.6 top, grey bars); an example of such an unresponsive neuron was shown in Fig. 3.6H. During each 50 ms response epoch only about 10 % of neurons showed any significant stimulus-locked increase in firing rate (Fig. 3.6 top, *Inc*), and a smaller fraction showed a significant stimulus-locked decrease (Fig. 3.6 top, *Dec*). Note that the definition we use here for responsiveness is quite inclusive: even if a tone elicits only a 1 spike/s increase over the baseline firing rate, this response might still be deemed responsive if the spontaneous rate was sufficiently low for us to detect a change.

The form of sparseness we report has sometimes been termed "population sparseness," to distinguish it from "lifetime sparseness" (Olshausen and Field, 2004; Willmore and Tolhurst, 2001). Lifetime
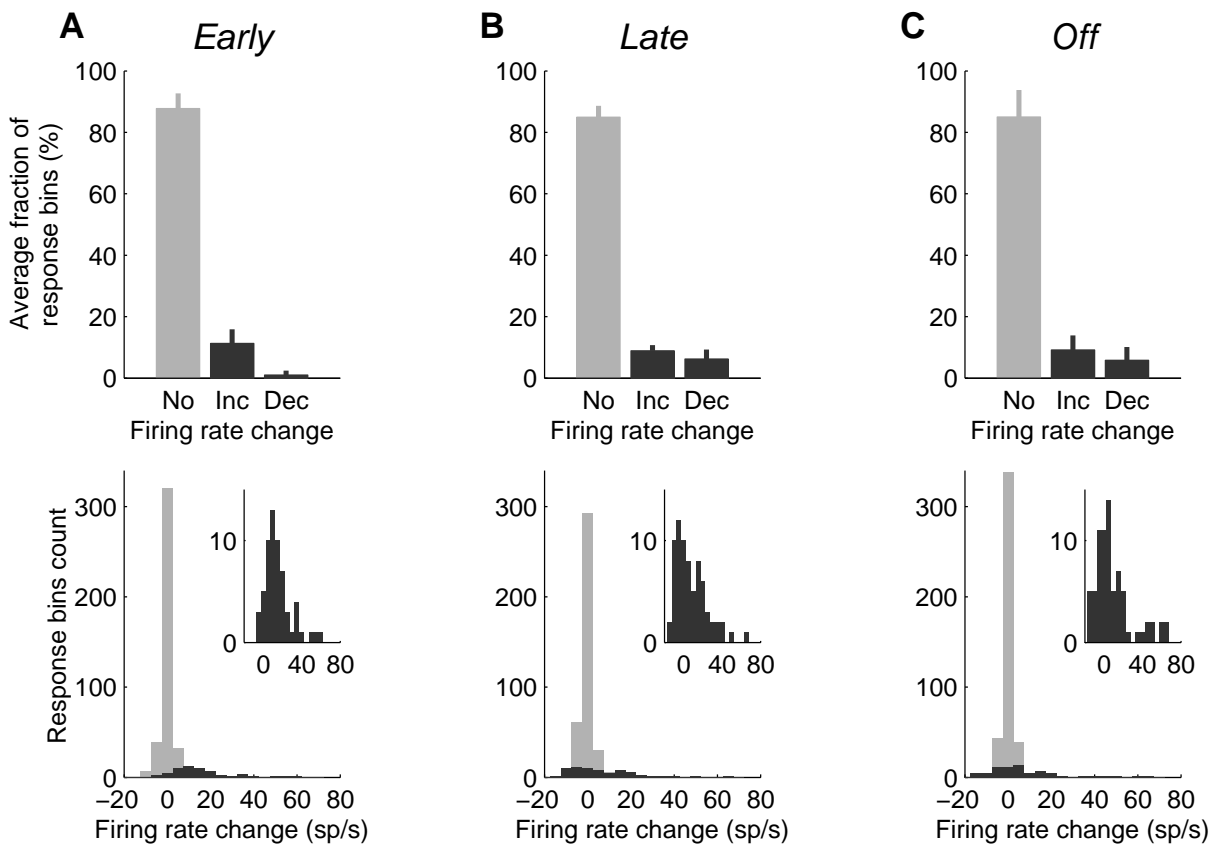
Figure 3.6: Only a small fraction of the neuronal population participated in the response to an acoustic stimulus at any instant (see Sec. 2.6 for details). **(A)** The top panel shows the fractions of response bins per octave displaying no change (No), a significant ($p = 0.01$) increase (Inc), or a significant decrease (Dec) in the stimulus-evoked firing rate during the *early* response epoch. Error bars show standard deviations of the fractions of response bins. The bottom panel shows frequency histograms of firing rate changes in the *early* epoch. The grey histogram shows firing rate changes for response bins from the 'No' bar from the top panel, and the black histogram shows firing rate changes for response bins from the 'Inc' and 'Dec' bars from the top panel. The inset shows the black histogram in more detail. Both panels in (A) include data from 468 response bins (n = 100 neurons). **(B), (C)** Same format as above for *late* (B) and *off* (C) epochs.

sparseness refers to the selectivity of a single neuron probed with different stimuli, and can be assessed for a single neuron during a single unit experiment. Population sparseness refers to the response of the population to a given stimulus. Responses in visual cortex have been reported to show population sparseness (Weliky et al., 2003), but population sparseness has not previously been assessed in auditory cortex.

Even among the minority of neurons that responded to a particular stimulus, the typical response was modest. Fig. 3.6 (bottom) quantifies the distribution of changes elicited by a typical tone across the population. The median stimulus-evoked firing rate was approximately the same as the spontaneous rate (2.8 vs. 2.7 spikes/s, spontaneous vs. early; see Tab. 3.1), and the distribution of evoked responses was barely distinguishable from the spontaneous distribution (cf. Fig. 3.4). Only a very small fraction— about 2–3 %—of the population showed a well-driven ($\geq$20 spikes/s) response during any epoch (Fig. 3.6; note that an increase of 20 spikes/s implies only a single extra spike in the 50 ms response bin we consider). However, the activity in this small fraction raised the mean (as opposed to the median) firing rate by nearly 50 % during the early epoch (4.9 vs. 7.0 spikes/s; changes during later epochs were smaller; see Tab. 3.1). Thus the presentation of a tone caused only a barely discernable change in the activity of most of the population; but an appreciable number of extra spikes was concentrated in a small fraction of neurons.

Half of the cells (50 %) did not show any significant change (increase or decrease) in firing rate during any response epoch, to any tone. At the other extreme, a few broadly tuned cells showed significant changes in firing rate in all (4 or 5) octave bins (i.e. across the whole frequency space tested) for at least one of the response periods.

It might appear that the sparseness we report is incompatible with the broad frequency tuning of rat auditory cortical neurons. However, we found that sparseness was not achieved through narrow frequency tuning. Instead, it arose through a combination of factors. First, 50 % of the neural population failed to respond to any of the simple stimuli we presented. Second, responses were often brief; in many neurons, the change in firing rate was limited to just one of the three response epochs. Thus, sparseness of the response in time contributed to the overall sparseness of the population response. Finally, even when changes occurred they were typically small; the increase in firing rate exceeded 20 spikes/s in only

about a quarter of the statistically significant responses. As a result, only a small fraction of neurons responded vigorously to any tone, even though frequency tuning was broad.

3.5   Other stimuli

In a subset of neurons, we tested stimuli other than the 50 or 60 dB standard tones used to probe most of the population. For some neurons we evaluated responses to quieter (30 or 40 dB SPL; n = 62 neurons) or louder (80 dB SPL; n = 6 neurons) tones. As shown in Fig. 3.7, the quieter tones elicited responses in a smaller fraction (3% increase, 2% decrease in firing rate) of the population than our 50 or 60 dB standard (10% increase, 4% decrease), but the louder tones did not elicit responses in a significantly larger fraction (11% increase, 6% decrease, t-test p = 0.66). Thus the sparseness of responsiveness to pure tones did not appear to result from the amplitude of the stimuli we used.

In order to assess whether such sparse population encoding was restricted to pure tones, in some neurons we also tested other stimuli. In one subset of neurons, we also tested FM-sweeps (Sec. 2.6.2). In a different subset of neurons, we tested the response to an ensemble of complex (natural) sounds (Machens et al., 2004). The population response (Fig. 3.7) elicited by sweeps (n = 17), as well as firing rates evoked by natural sounds (n = 27) and sweeps (n = 22) was similar to that elicited by tones (Fig. 3.8).

The estimated distribution of evoked firing rates might be dependent on sample size and number of trials per each response bin used in analysis. Although the sample size of neurons probed with natural sounds (n = 27) that we used to estimate firing rate changes might seem small compared to the size of the main data set of neurons probed with tones (n = 145), we show that limiting ourselves to a smaller set of neurons probed with tones did not change our estimates (Fig. 3.9). We subsampled the set of neurons probed with tones (n = 145) to match set of neurons probed with natural sounds both in size (n = 27) and number of trials per response bin. For each tone-neuron drawn randomly, the number of trials in each response bin was matched (trials were chosen randomly) to a given neuron probed with natural sounds. Therefore, a randomly generated sample of such neurons matched the set of neurons probed with natural sounds both in size and distribution of number of trials. Distributions of stimulus-evoked changes in firing rates (Fig. 3.9AB; see also Fig. 3.6) remained essentially the same. The estimated fraction of
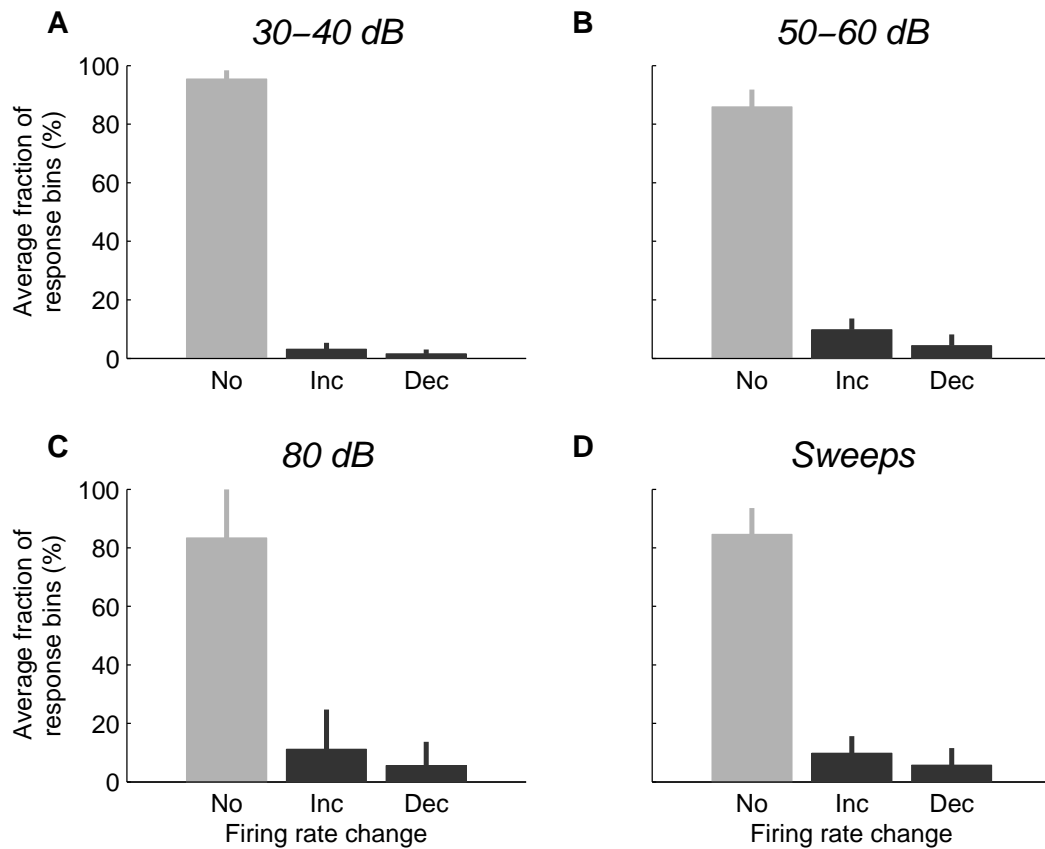
Figure 3.7: Only a small fraction of the neuronal population was engaged in the response to any stimulus. **(A)**, **(B)**, **(C)**, **(D)** Each panel shows fractions of response bins with either no change (No), or a significant ($p < 0.01$) increase (Inc), or decrease (Dec) in stimulus-evoked firing rate for the given stimulus. For this analysis the response bins were pooled together **(A)** Responses to 30 or 40 dB tones ($n = 62$ neurons). **(B)** Responses to 50 or 60 dB tones ($n = 100$ neurons). For responsive fractions in individual response epochs, see Fig. 3.6. **(C)** Responses to 80 dB tones ($n = 6$ neurons). **(D)** Responses to 54 dB sweeps ($n = 17$ neurons). Error bars show standard deviations of the fractions of response bins. Note that only neurons for which we assessed response significance (Sec. 2.6.2) were used for analysis in this figure.
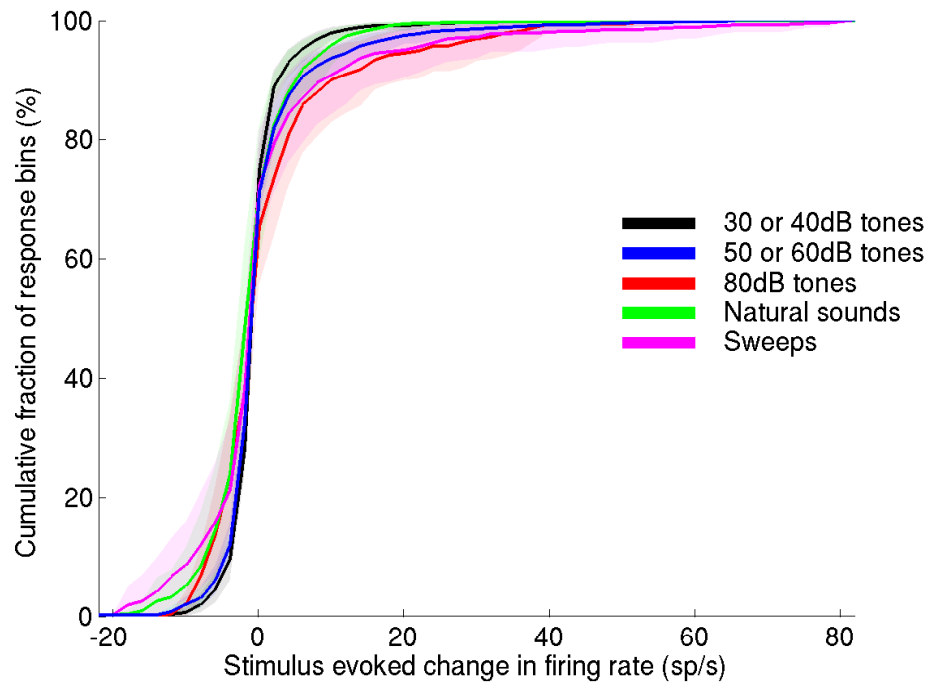
Figure 3.8: Only a small fraction of the population showed a well-driven stimulus-evoked response at any instant, regardless of which stimulus ensemble was used. The figure shows the cumulative fraction of stimulus-evoked changes in firing rate for various stimuli. The overlapping colored patches show 95 % confidence intervals obtained by bootstrap. The large overlap between the responses to the different stimulus ensemble suggests that stimulus-dependent differences in evoked firing rates were rather small. Different colors correspond to different stimuli: black—30 or 40 dB tones (n = 91 neurons, 1365 response bins); blue—50 or 60 dB tones (n = 145 neurons, 2079 response bins); red—80 dB tones (n = 22 neurons, 330 bins); green—natural sounds (n = 27 neurons, 18900 bins); and magenta—sweeps (n = 22 neurons, 704 bins).

A Natural sounds (n=27)

B *Matched* tones (n=27)

C Tones (n=27)

D *Matched* tones (n=27)
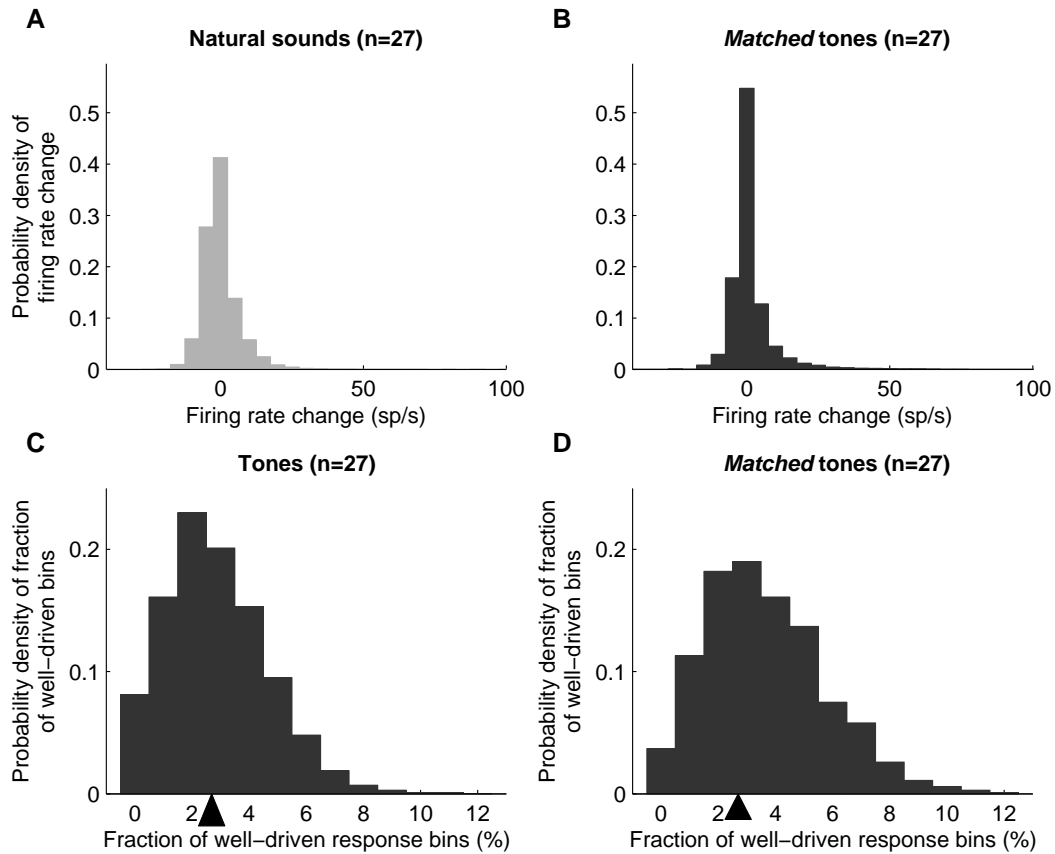
Figure 3.9: Smaller sample size did not influence estimates of response distributions and fraction of well-driven bins. In panels **(A)**, **(B)** we show that smaller sample size did not change response distributions for neurons probed with natural sounds and neurons probed with tones. **(A)** Response distribution for neurons probed with natural sounds (n = 27). Response was defined as difference between evoked firing rate and spontaneous firing rate for each response bin. All response bins were pooled for this figure. **(B)** Response distribution for matched set of neurons probed with 50–60 dB tones. To match dataset from panel A 1000 sets of n = 27 neurons were drawn (without replacement) from the original set of neurons probed with tones (n = 145). For each selected neuron the number of trials in each response bin was matched (trials were chosen randomly) to a given neuron probed with natural sounds. Therefore, all 1000 randomly drawn sets of neurons probed with tones matched set of neurons probed with natural sounds both in size and distribution of number of trials. Panel shows mean frequency histogram of 1000 randomly drawn sets of neurons. In panels **(C)**, **(D)** we show that smaller sample size did not change the estimate of fraction of well-driven bins. Well-driven response bins were defined as response bins for which the (absolute) difference between evoked and spontaneous firing rate was at least 20 sp/s. Panels show frequency histograms of fractions of well-driven response bins estimated by bootstrap from original data set. Arrow shows estimate of fraction of well-driven bins (Sec. 3.4). **(C)** For this panel we randomly drew (with replacement) 1000 sets of n = 27 neurons probed with tones to match number of neurons probed with natural sounds. This panel shows frequency histogram of fractions of well-driven response bins computed from 1000 such randomly drawn sets of neurons. **(D)** Twenty-seven neurons were drawn randomly (with replacement) from set of neurons probed with tones to match set of neurons probed with natural sounds both in size and number of trials per response bin. For each randomly drawn set of neurons we pooled response bins together and computed fraction of well-driven bins. The panel shows frequency histogram of fractions of well-driven response bins computed from 1000 randomly drawn matched sets of neurons.

59

highly responsive cells (Sec. 3.4) did not change by using a smaller sample of neurons (Fig. 3.9CD). Decreasing the sample size of neurons probed with tones did not change our population estimates.

Thus, although we cannot rule out the possibility that some stimuli other than those we tested might elicit responses in a larger fraction of neurons in auditory cortex, the fact that stimuli ranging from the simplest (tones and sweeps) to the most complex (natural sounds) evoke such sparse population responses leads us to conclude that sparse responses are not anomalous.

3.6 Binary responses

In the anesthetized auditory cortex, tones typically elicit transient responses (Heil, 1997), a subset of which show high trial-to-trial reliability which was termed "binary" (DeWeese et al., 2003). These highly reliable neurons were of interest because they provided an existence proof that the cortex is capable of precisely controlling spike count as well as spike timing, and therefore challenged conventional theories that posited that cortical computation is necessarily noisy (reviewed in DeWeese et al., 2005).

However, it has recently been suggested that such binary neurons were an artefact of anesthesia (Wang et al., 2005). This claim was surprising, because apparently binary responses can be found in the published literature on auditory cortex of alert animals (e.g. Fig. 4A from Chimoto et al. (2002); Fig. 4A from Barbour and Wang (2003)). Fig. 3.10 shows further examples of binary responses from the unanesthetized rat preparation used in this study. These responses indicate that binary spiking is not an artefact of anesthesia.

How common are such binary responses? Because the goal of the present experiments was to assess typical responses to a broad range of different stimuli, we did not present each stimulus a sufficient number of times to assess reliability, so we can provide only a lower bound on the prevalence. Moreover, we made no attempt to optimize stimuli to maximize firing reliability; it may be that, with the appropriate choice of stimuli, it is possible to elicit binary responses from many or most neurons in auditory cortex, in the same way that it is possible with the appropriate choice of stimuli to elicit sustained firing. Indeed, we observed sustained firing (>20 spikes/s evoked increase in firing rate over the 100 ms stimulus duration) in only 4/100 neurons, which is comparable to the lower bound of 2/100 neurons in which we were able to confirm binary spiking. Interestingly, one neuron (shown in Fig. 3.2B) fell into both categories. Note that categories such as sustained or transient firing characterize only
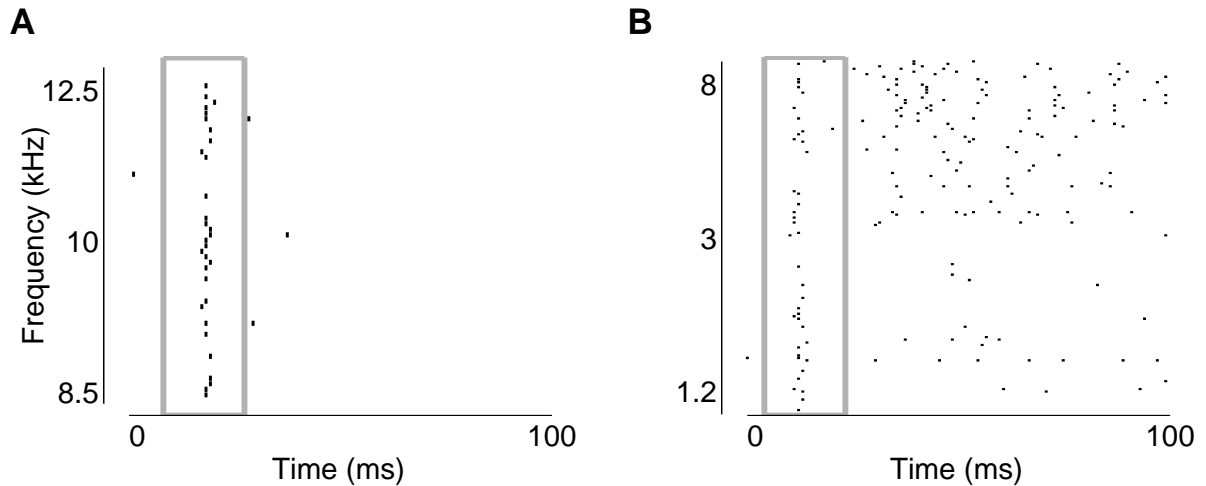
Figure 3.10: Some neurons in the unanesthetized auditory cortex exhibit responses with low trial-to-trial variability in spike count. The spike rasters show details of responses for two neurons from the corresponding panels of Fig. 3.2. For each of these neurons, every trial consisted of either one spike or no spikes, with no multi-spike responses during the onset transient occurring inside the 20 ms window (grey rectangles); the stimulus was presented from 0 to 100 ms. Responses such as these, which were termed "binary" since they consist of ones and zeros, can be shown to have significantly lower trial-to-trial variability than less structured (e.g. Poisson) responses at these firing rates (DeWeese et al., 2003). Examples such as these from the unanesthetized cortex demonstrate that binary responses are not anesthesia artifacts.

the mean of the response as opposed to binary firing, which depends on both the mean and the variance. Nevertheless, it appears that the prevalence of binary and sustained responses to simple tones are comparable.

3.7   Neither spatial nor laminar position predicts response pattern

The heterogeneity of response patterns to simple tones led us to wonder whether neurons with similar properties might be clustered into nearby regions of the cortex; for example, neurons with predominantly transient responses might be found in one region, and sustained neurons might be found in another. In some cases, therefore, we recorded from multiple cells in a single electrode penetration. Because the recording electrodes were aligned approximately perpendicular to the cortical surface, the cells recorded in a single electrode penetration likely belonged to the same or neighboring cortical column.

61

We did not detect any clustering of response patterns; highly responsive cells were often very near unresponsive cells. Fig. 3.11 shows an example with five neurons recorded over two penetrations (three in one penetration, and two more in a penetration approximately 50–100 μm ventro-caudal from the first penetration). In the first penetration, one neuron was unresponsive, one showed suppression over a wide range of frequencies, and the third showed enhanced firing over an even wider range of frequencies. In the second penetration, both neurons were unresponsive. The fact that unresponsive neurons were often mixed closely with responsive neurons indicates that unresponsiveness need not indicate gross cortical damage (cf. Evans and Whitfield, 1964) or recording from a region of cortex that was unresponsive to the stimuli we were presenting, but that instead neurons with different selectivity are comingled.

We also wondered whether firing rate was correlated with cortical depth (Fig. 3.12). We segregated neurons ($n = 141$) recorded at different cortical depths into 6 groups corresponding to the cortical layers (Games and Winer, 1988). We compared the firing rates using multiple comparisons based on Kruskal-Wallis test and found that the spontaneous (Fig. 3.12A) and mean evoked firing rates (Fig. 3.12B) were not significantly different, with the exception of layer II, which displayed firing rates significantly lower ($p < 0.01$) than the other cortical layers (layer I contained only one neuron and was not included in the comparisons).

3.8 Very responsive neurons may be narrow-spiking interneurons

Because we could record from only a relatively small number of neurons in a single penetration, we cannot rule out the possibility that more thorough sampling of all the nearby neurons in a region might reveal subtler forms of spatial or laminar organization that escaped our detection. Alternatively or additionally, responsiveness might be correlated with single neuron properties such as type, morphology, and molecular expression pattern. Although in this study we did not recover neurons for histological analysis and so could not assess whether there was a correlation with morphology or molecular expression pattern, we did attempt to correlate responsiveness with cell type.

Cortical neurons can be grouped into two broad classes: excitatory neurons that release glutamate at their synapses; and inhibitory interneurons, which release GABA. Most cortical neurons are excitatory. GABAergic neurons can have diverse morphological, physiological, or molecular charac-
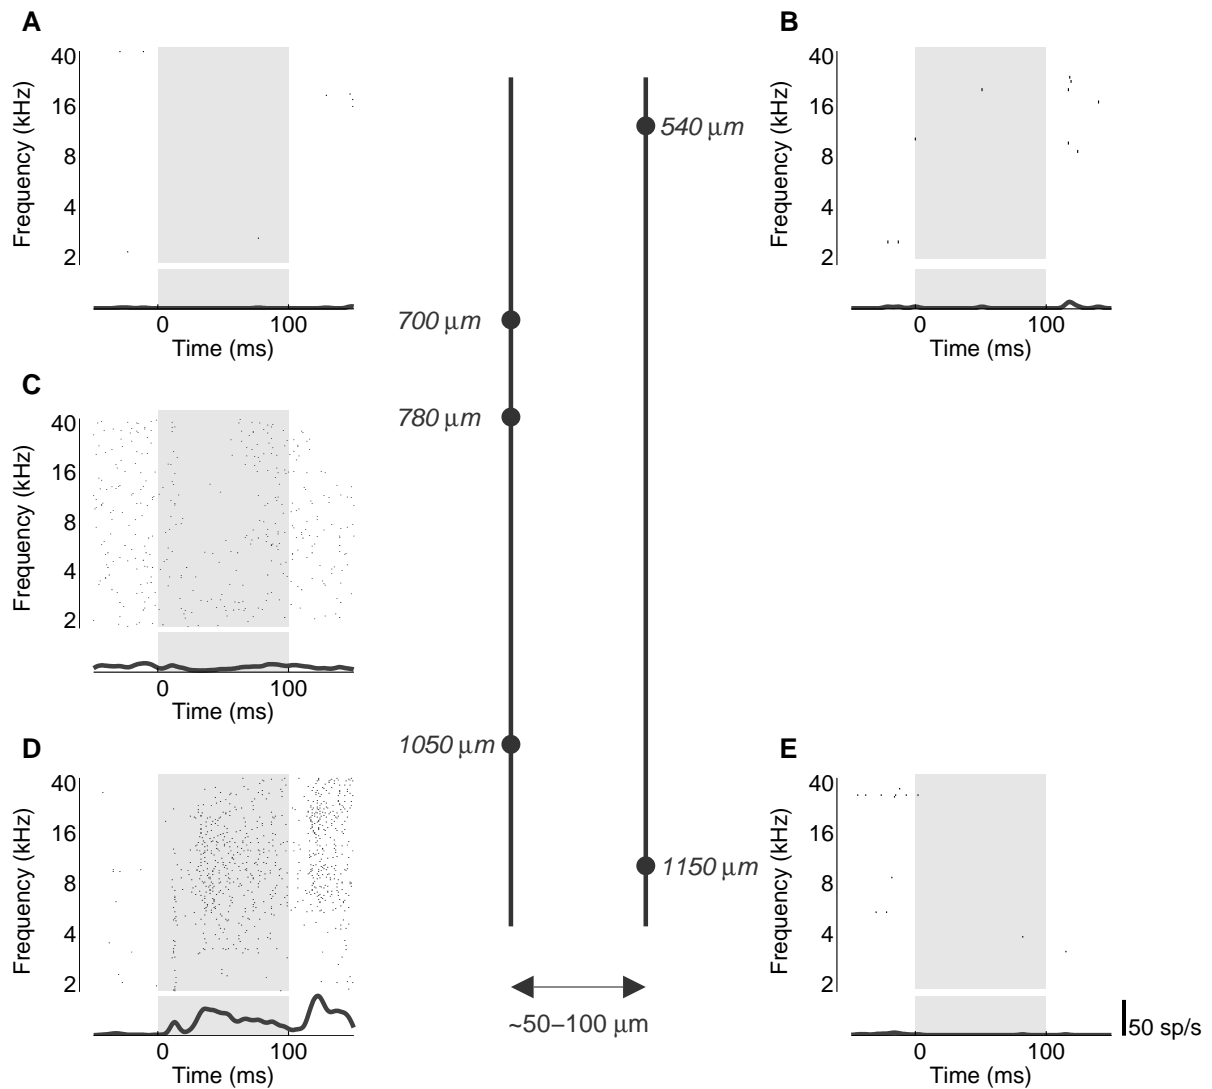
Figure 3.11: Neurons recorded in a single penetration can show very different tone-evoked responses. Response rasters are shown from neurons recorded in two penetrations (left and right columns); the penetration depicted in the right column was about 50-100 μm ventro-caudal from the penetration depicted on the left. Dots represent individual spikes and grey shaded regions indicate the tone duration (60 dB, 100 ms). Firing rate curves at the bottom of the rasters were computed by first summing the spikes in 1 msec bins over *all* frequencies shown, and then convolving the resulting PSTH with a Gaussian ($\sigma = 3$ ms). Depths of recordings were measured perpendicular to the cortical surface, as given by micromanipulator readings. The cell shown in panel (D) is the same as that in Fig. 3.2B
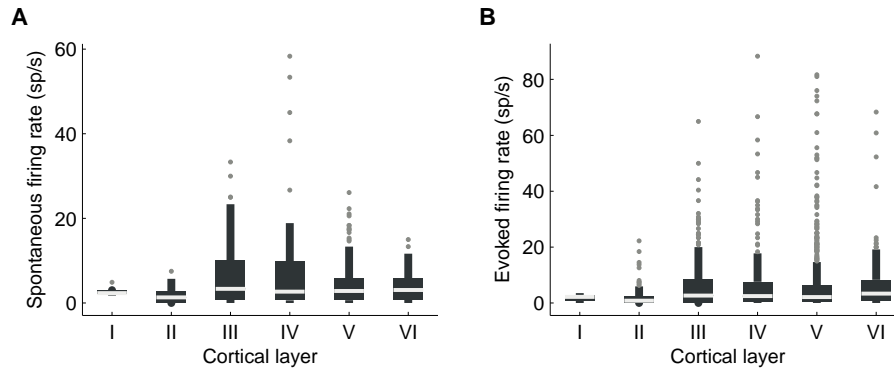
Figure 3.12: Firing rates were similar across cortical layers. Neurons (n = 141) were segregated into 6 groups corresponding to cortical layers (see text for details). **(A)** Spontaneous firing rates were computed for each response bin. **(B)** Evoked firing rates were computed in all evoked response bins. Boxes show the positions of the lower quartile (25^{th} percentile), median (white horizontal line), and upper quartile (75^{th} percentile). Whiskers extend to the most extreme values within 1.5 times the interquartile range from the ends of the box. Grey circles show positions of values beyond the ends of whiskers.

teristics (Markram et al., 2004). Excitatory and inhibitory neurons can also be distinguished based on a variety of physiological parameters (Cauli et al., 1997; McCormick et al., 1985). In particular, the firing rate of some inhibitory interneurons—the so-called fast-spiking subtype—is higher when stimulated by current injection. Spike width and shape has been used in previous studies to assign spikes recorded extracellularly in vivo to putative excitatory and inhibitory neurons in hippocampus (Csicsvari et al., 1999), and cortex (Barthó et al., 2004). We therefore asked whether spike shape might predict response patterns in our sample.

Based on previous studies (for example Cauli et al., 1997; McCormick et al., 1985), we expected that fast-spiking interneurons would likely have narrow and symmetric spikes. For each cell we therefore computed the spike width, and also the "spike amplitude index" as a measure of spike symmetry (Sec. 2.6.1). For our population of cells the spike widths ranged from 0.4 ms to 1.9 ms, with a median value of 0.9 ms. We defined the spike amplitude index as the absolute value of the spike peak-to-valley-ratio. A spike amplitude index of unity indicates a perfectly symmetrical spike, whereas a value greater than unity indicates a tall spike, and a smaller value indicates a spike with a deep valley; a fast-spiking interneuron would be expected to have a low spike amplitude index. Spike amplitude indices ranged from 0.8 to 34.3, with a median value of 2.0.

64
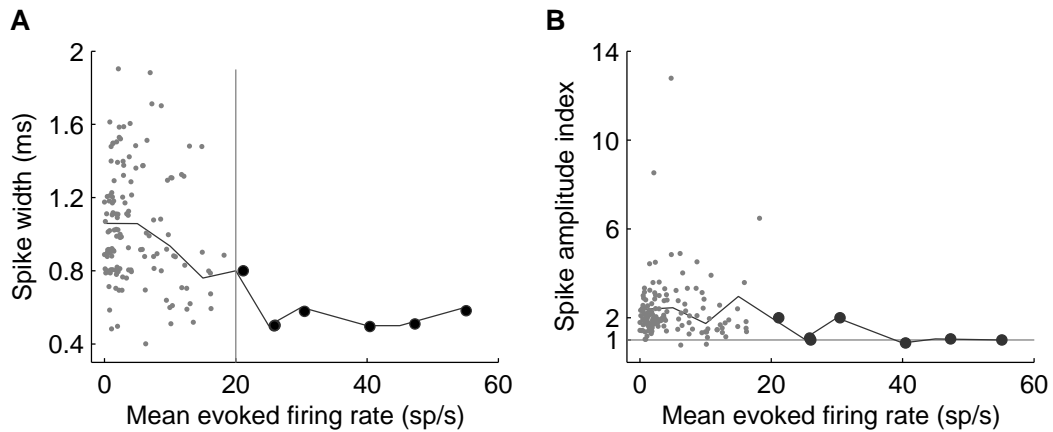
Figure 3.13: Very responsive neurons might be narrow-spiking interneurons. Neurons with higher evoked firing rates tended to have narrow, symmetrical spikes. **(A)** Spike width decreased with increasing evoked firing rate (n = 144 neurons). **(B)** Neurons with higher mean evoked firing rates displayed more symmetrical spikes. Perfectly symmetrical spikes would fall on the grey horizontal line. The large dots in both panels denote cells with mean evoked firing rates > 20 spikes/s. The black lines in both panels indicate average values in 5 spikes/s bins. The mean evoked firing rate for each cell was computed as the mean firing rate across all response bins. In panel (A) all points were jittered slightly so that overlying points could be seen. One cell had a very large spike amplitude index of 34.3 (spike width 1.9 ms, mean evoked firing rate 1.3 spikes/s) and was excluded from both panels.

Neurons with higher evoked firing rates tended to have narrower spikes (Fig. 3.13A), suggesting that interneurons were over-represented among the most responsive neurons. Indeed, the seven most responsive neurons—those with a mean evoked firing rate (computed across all octave bins and response epochs) higher than 20 spikes/s—had narrow spikes with spike widths less than or equal to 0.9 ms. Neurons with high firing rates also tended to have symmetrical spikes (Fig. 3.13B). Although spike width and shape are only at best crude surrogates for cell type, the striking correlation between these quantities and tone responsiveness suggest that a substantial fraction of the most responsive neurons may be interneurons.

3.9   Sparse coding for reliable stimulus representation and learning

Why should the cortical representation of an acoustic stimulus be sparse? Several explanations have been advanced. One set of proposals focuses on the energy used by neural activity; sparse representations involve fewer spikes and thereby minimize the energetic costs associated with a neural representation (Attwell and Laughlin, 2001; Laughlin and Sejnowski, 2003; Levy and Baxter, 1996). Other

proposals focus on the advantages of sparse representations for computation. For example, it has been suggested that the statistics of natural sensory environments are sparse, and that a sparse code provides a natural match to such environments (Lewicki, 2002; Olshausen and Field, 2004, 1997). Recently, it was shown how a sparse overcomplete representation could be used to solve the "cocktail party problem" (i.e. separate a single auditory stream from several mixed together) (Asari et al., 2006).

Fig. 3.14 demonstrates one benefit of sparse representations in the present context. This simple example is not intended as a model of auditory cortex, but merely to illustrate some of the basic intuitions underlying sparse representations. We compared the representation of auditory stimuli by two hypothetical neuronal populations, one dense and the other sparse. In the *sparse* population (Fig. 3.14AB), firing rates were drawn from the lognormal distribution we observed (mean = 1.3, std = 1.0, both on a logarithmic scale), whereas in the *dense* population firing rates are drawn from a hypothetical Gaussian distribution (Fig. 3.14CD), the mean firing rate and entropy (a measure of the representational capacity) of which were matched to the observed distribution (mean = 4.2, std = 5.2, elements from the dense distribution with negative firing rates were discarded and drawn again from the same distribution until the dense distribution contained only non-negative firing rates; see Sec. 2.7 for details).

To examine the ability of each population to represent a pair of distinct stimuli, we drew two patterns of firing rates (X and Y), corresponding to the two stimuli, from the sparse distribution; and similarly drew two patterns from the dense distribution. It seems reasonable to suppose that a good neural representation of a pair of distinct stimuli should allow the stimuli to be easily discriminated. Specifically, since at any instant an organism only has access to a set of spikes rather than to the underlying firing rates, the stimuli should be discriminated on the basis of the pattern of spikes across the population—that is, on the basis of the spikes representing a single instantiation of the firing rates. The question, then, is how well a spike train drawn from X can be discriminated from Y, and how this discriminability depends on whether X and Y are drawn from sparse or dense distributions.

The spike trains drawn from the sparse distribution were dominated by a few outliers—a few neurons with high firing rates—which could be used to reliably discriminate the pattern X from Y even by eye. By contrast, the absence of such outliers in the spike trains drawn from the dense distribution made it difficult to discriminate these patterns. This intuition can be quantified by a discriminability measure Q (Sec. 2.7), which confirmed that the sparse representations were consistently more easily

A *Sparse distribution 1*

B *Sparse distribution 2*

Sparse patterns 1

Sparse patterns 2

C *Dense distribution 1*

D *Dense distribution 2*

Dense patterns 1

Dense patterns 2

35 spikes per
100 trials
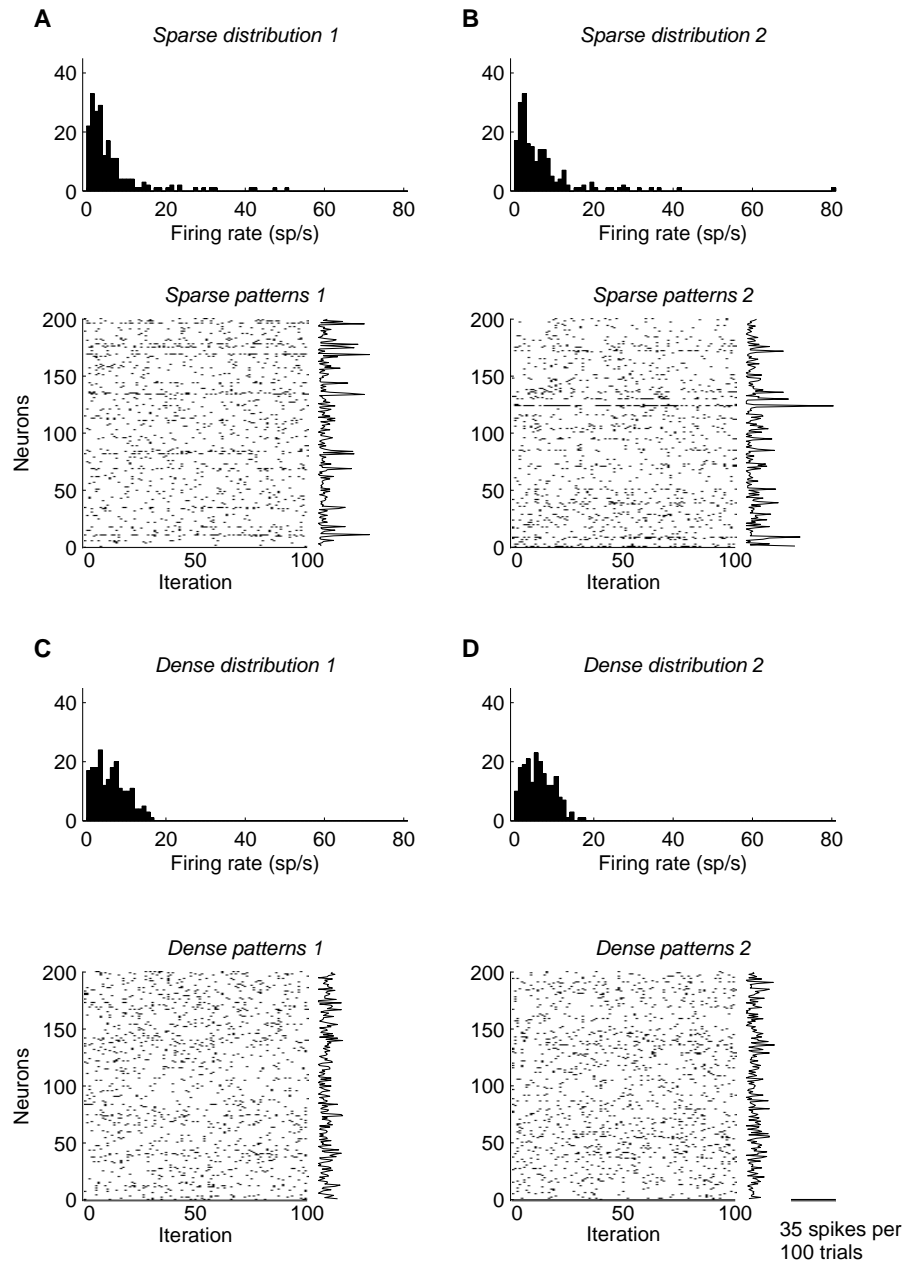
Figure 3.14: Spike patterns generated from sparse distributions of firing rates are dominated by a small number of neurons with higher firing rates and are more distinct than patterns generated from dense distributions. **(A), (B)** Two examples of spike patterns generated from sparse (lognormal) distributions of firing rates. The top panels show two distributions of firing rates for 200 neurons drawn from the same underlying lognormal distribution matching our data. Each of the bottom panels show 100 spiking patterns (columns) generated for each neuron in the population, with summary neuronal activity for the 100 patterns depicted as a histogram on the right side of each panel; note the tall peaks in each histogram that greatly facilitate the discriminability between these two sparse patterns of activity. Each spike pattern represents a snapshot of neuronal activity during 10 ms (Sec. 2.7). Dots represent the occurrence of at least one spike in 10 ms wide windows. **(C), (D)** Two examples of spike patterns generated from dense (normal) distribution of firing rates. Parameters for the normal distribution were chosen to match mean firing rate, total amount of spikes, and entropy of the lognormal distribution. Same format as in (A) and (B).

67

discriminated than the dense ones ($Q = 5.0$ for sparse representations, compared to $Q = 1.9$ for dense representations).

3.10 Hebbian learning of spike patterns from sparse and dense distributions

Similarity of evoked activity patterns might also facilitate their learning. We have studied a single neuron model with Hebbian synapses to compare learning of neuronal activity patterns arising in sparse (lognormal) or dense (Gaussian) distribution of firing rates (Fig. 3.15). Parameters for sparse and dense distributions were estimated from our data, as described in previous section.

From the firing rate distributions we first generated a training set of $k = 100$ *firing rate patterns*, each consisting of $n = 100$ neurons with firing rates chosen randomly from the same firing rate distribution. Every other pattern in the training set was the same *target* pattern. From the training set of firing rate patterns we then generated a set of *spike patterns* representing a 10 ms snapshot of neuronal spiking activity. Note that while the *target firing rate patterns* were identical, the actual *target spike patterns* were not (although they were similar, see previous section), because they were generated by a stochastic Poisson process (see Sec. 2.7.2).

We simulated learning in a single sigmoidal neuron with n inputs corresponding to n neurons in the input patterns, and one output (see Sec. 2.7.2 for details). In each trial (a single presentation of input pattern) the sigmoidal neuron computed its output as a weighted sum of its inputs transformed by a sigmoid function. After the trial, synaptic weights were updated according to the Hebbian learning rule, i.e. weights corresponding to stronger inputs grew stronger.

Fig. 3.15 documents the learning process in detail for both sparse and dense representations. The top panels show the two distributions of firing rates from which we generated input firing rate patterns. Panel B shows examples of input spike patterns generated from firing rate patterns for $n = 100$ neurons and $k = 100$ trials. Red dots show spikes generated by the target pattern (every second trial), black dots show spikes from the nontarget patterns. The histograms on the right side of each of the input patterns show spiking activity for target patterns (red line), and nontarget patterns (black line). Note that the red histogram is easily distinguishable from the black histogram for the sparse distribution.
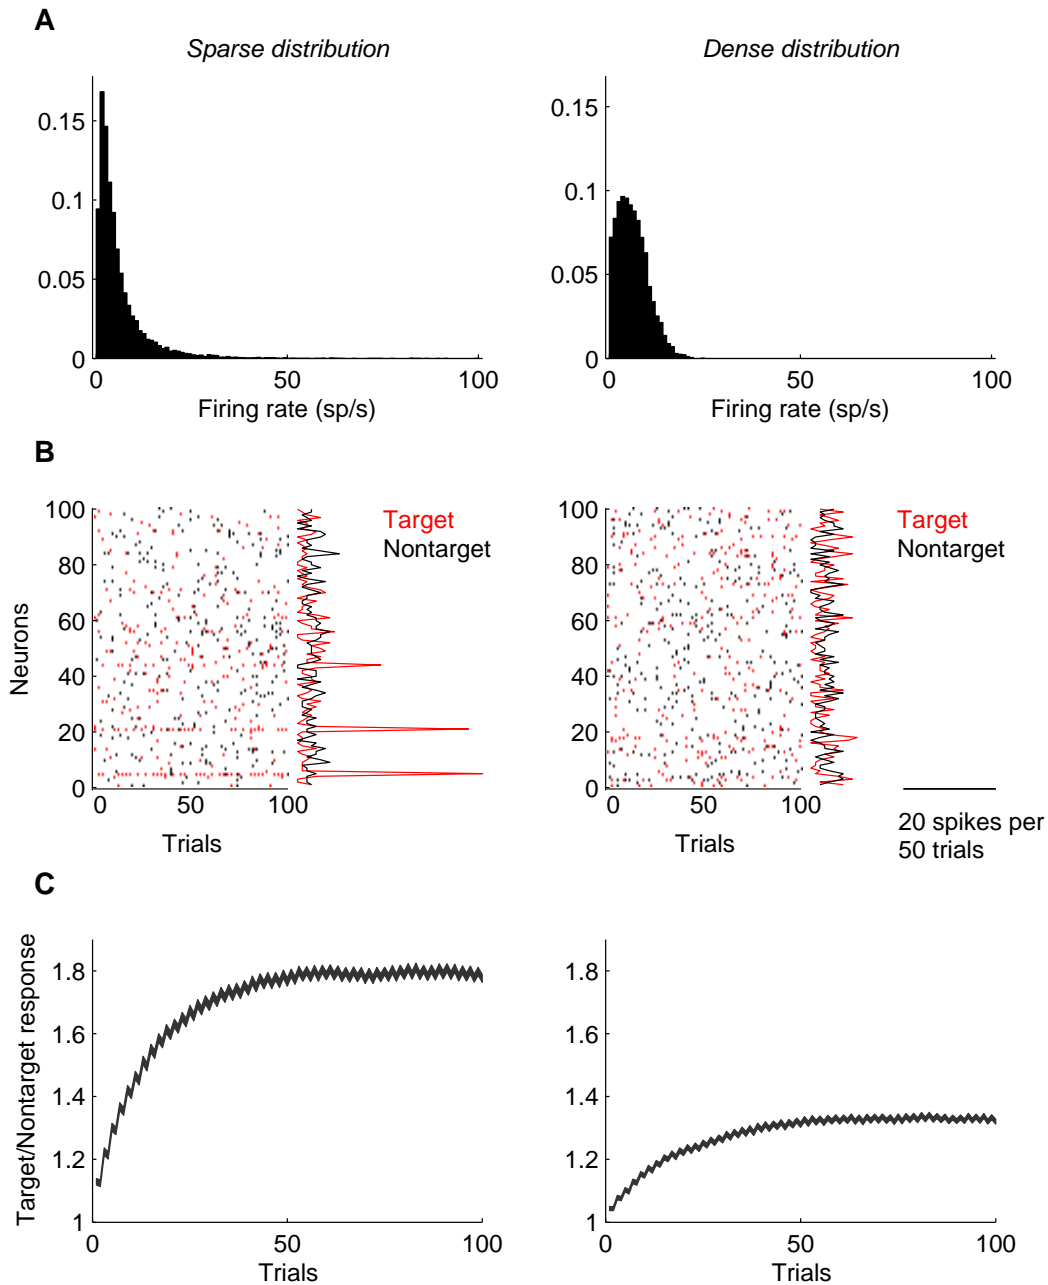
Figure 3.15: Hebbian learning is easier and faster for neuronal patterns derived from sparse distributions of firing rates. **(A)** The distribution of firing rates of dense and sparse populations used in the simulation. The parameters of the sparse (lognormal) distribution were taken from the observed data; the parameters of the dense distribution were matched in mean firing rate and entropy. **(B)** Sample rasters for the target (red) and background (black) patterns, and the associated firing rate histograms. A spike pattern from the target distribution is presented to a model neuron with Hebbian synapses every other trial; non-target spike patterns presented on alternating (background) trials are drawn from a new distribution each trial. The target sparse distribution contains a few outliers that dominate learning, whereas the absence of such outliers in the dense distribution leads to slower and less efficient learning. **(C)** The ratio of responses elicited by target patterns on a model neuron with Hebbian synapses (see Eq. 2.6) to non-target responses is shown. The ratio grows faster and asymptotes to a higher level for the sparse distribution, indicating that learning is faster and more efficient.

We repeated the simulation experiment 1000 times, each time drawing a different set of $n = 100$ neurons, and $k = 100$ trials. Panel C shows the ratio of target to nontarget responses averaged across experiments (line thickness represents the standard error of the mean). For sparse representation, the ratio grew faster and reached higher asymptotic value, indicting that learning was faster and more complete than for dense representations.

Model neurons with Hebbian synapses learned to discriminate sparse patterns more rapidly and completely than dense patterns. Moreover, spike patterns discriminations was easier for sparse than dense patterns (see previous section). Thus the presence of a handful of neurons with high firing rates can facilitate stimulus discrimination and learning in a simple model.

3.11 Discussion

We used cell-attached techniques in the auditory cortex of unanesthetized rats to measure the responses of individual neurons to simple acoustic stimuli; we then used this data set to infer the stimulus-evoked activity across the population. The distribution of firing rates across the population was lognormal rather than exponential, and stimuli typically elicited a high firing rate in only about 3 % of the population. Such sparse representations may offer computational advantages, including faster and more complete learning of auditory patterns.

3.11.1 Sparse representations in cortex

Experimental evidence for sparse coding has been found in a range of experimental preparations, including the visual (Baddeley et al., 1997; Vinje and Gallant, 2000), motor (Brecht et al., 2004), barrel (Margrie et al., 2002) and olfactory systems (Perez-Orive et al., 2002; Rinberg et al., 2006; Szyszka et al., 2005), the zebra finch auditory system (Hahnloser et al., 2002), and cat LGN (Dan et al., 1996). However, the sparseness of representations in the auditory cortex has not been explicitly addressed in previous work. Our results constitute the first direct evidence that the representation of sounds in the auditory cortex of unanesthetized animals is sparse.

Our data support the "efficient coding hypothesis," (Barlow, 1961) according to which the goal of sensory processing is to construct an efficient representation of the sensory environment. Sparse

codes can provide efficient representations for natural scenes (Olshausen and Field, 2004, 1997). Sparse representations may also offer energy efficient coding, where fewer spikes are required compared to dense representations (Attwell and Laughlin, 2001; Laughlin and Sejnowski, 2003; Levy and Baxter, 1996).

A growing body of theoretical work on sparse representations suggest they may be useful for computation (Asari et al., 2006; Guyonneau et al., 2004; Lewicki, 2002; Olshausen and Field, 2004, 1997; Smith and Lewicki, 2006). Sparse representations have become increasingly important in statistical machine learning (Girosi, 1998). One intuition underlying this approach is that it can be easier to recognize a sparse pattern in a high-dimensional space than a dense pattern in a low dimensional space. This was illustrated in Fig. 3.14, where spike trains drawn from a sparse distribution could more easily be discriminated than those drawn from a dense distribution. This discriminability in turn can make the patterns easier to learn rapidly (Fig. 3.15). Thus an advantage of sparse cortical representations may be to facilitate rapid learning of arbitrary auditory patterns.

3.11.2 Sparse representations and optimal stimuli

The sparse and heterogeneous responses we report are consistent with many previous single unit studies of auditory cortex in unanesthetized animals, including the classical studies (Evans and Whitfield, 1964); (see also Abeles and Goldstein, 1972; Brugge and Merzenich, 1973; Gerstein and Kiang, 1964; Miller et al., 1974; Volkov and Galaziuk, 1989). In many anesthestized preparations (e.g. barbiturate and ketamine), sound-evoked responses are typically transient (DeWeese et al., 2003; Doron et al., 2002; Heil, 1997; Phillips and Irvine, 1981). With the resurgence of work in the awake preparation in the last decade, many researchers have emphasized the much richer repertoire of responses in awake animals, including especially sustained responses (Barbour and Wang, 2003; Chimoto et al., 2002; Gaese and Ostwald, 2003; Recanzone, 2000).

Our study complements recent work aimed at identifying "optimal" stimuli—stimuli that elicit high sustained firing rates (deCharms et al., 1998; O'Connor et al., 2005; Wang et al., 2005). The fact that a stimulus can be optimized to drive a particular neuron well, tells us little about how this stimulus is represented across the population. Our data suggest that only a minority of neurons are engaged in the representation of many stimuli; indeed, the fact that most stimuli drive most neurons only weakly

explains why finding the optimal stimulus for any given neuron can be such a challenge. Thus although there may be an optimal stimulus for any given neuron, most stimuli are not optimal for most neurons, and so are represented sparsely across the population.

Chapter 4

The emergence of lognormal distributions in cortical networks

Spontaneous neuronal activity in cerebral cortex depends on the underlying cortical circuitry. Cortical neurons receive most of the synaptic inputs from other cortical neurons via intracortical synapses (Ahmed et al., 1994; Benshalom and White, 1986; McGuire et al., 1984; Peters and Payne, 1993). Cortical networks, however, appear to be only sparsely connected, with most of the neurons not connected to other neurons (Holmgren et al., 2003; Markram, 1997; Markram et al., 1997; Thomson et al., 2002). Synaptic strengths themselves can be variable, and it has been shown recently that the distribution of synaptic strengths in cortex follows a lognormal distribution (Song et al., 2005), i.e. the logarithm of synaptic strengths is normally distributed (Fig. 4.1). This means that the majority of existing connections are rather weak with occasionally (very) strong connection.

Synaptic connectivity itself depends on neuronal activity. Common mechanisms thought to underlie this dependence include activity-dependent synaptic changes (Bi and Poo, 2001; Malenka and Nicoll, 1999), and activity-dependent axon and dendrite dynamics (Engert and Bonhoeffer, 1999; Lang et al., 2004; Maletic-Savatic et al., 1999; Matsuzaki et al., 2004; Ruthazer et al., 2003). We have shown (Sec. 3.3) that the firing rates (both spontaneous and stimulus evoked) of neurons in awake auditory cortex follow a lognormal distribution (Fig. 4.1). Therefore, the majority of neurons in the auditory cortex appear to be quite silent, with few very active neurons.

73
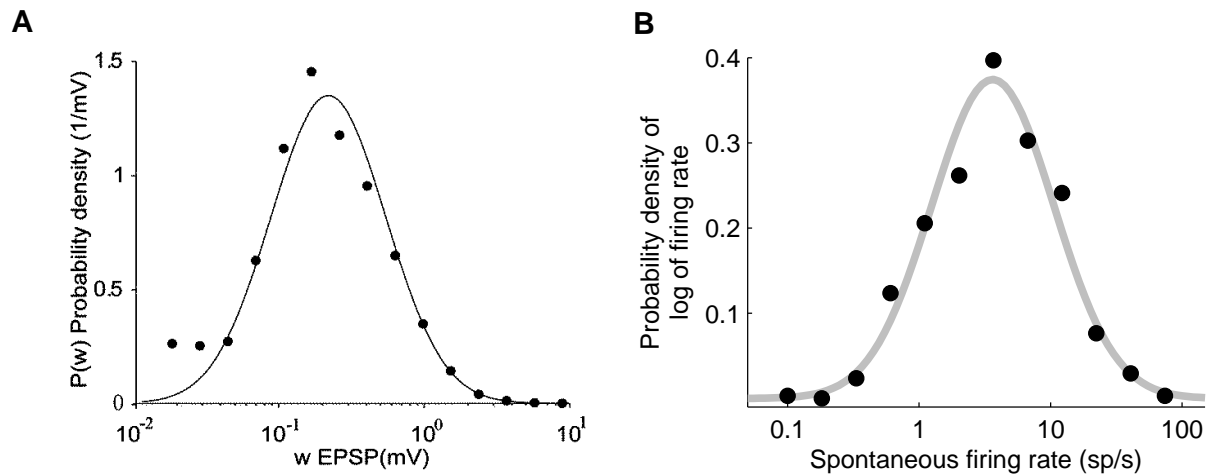
Figure 4.1: Lognormal distributions in the cerebral cortex. **(A)** Distribution of intracortical synaptic strengths is lognormal, i.e. appears normal on a (semi)logarithmic scale. Reprinted from Song et al. (2005). **(B)** Distribution of spontaneous firing rates in auditory cortex of awake rats follows a lognormal distribution.

Lognormal distributions of both synaptic strengths and firing rates provide experimentally verified constraints for population activity and wiring in cerebral cortex. However, it is not immediately clear how a lognormally distributed population of firing rates could arise from an underlying network with lognormally distributed connection strengths. Intuitively (see also below), if neuronal activity is dependent on an underlying set of (lognormally distributed) *independent* synapses/connections, then the distribution of neuronal activity is expected to be normal (Gaussian). It is also unclear, how the underlying lognormal distribution of synaptic strengths arises in connection to the related neuronal activity.

In this chapter we first formulate a linear model of spontaneous network dynamics. Based on the model, we propose that synaptic input correlations are sufficient for the neuronal network to display lognormal firing rate distribution. We then show how both lognormal distribution of synaptic strengths and lognormal distribution of firing rates can arise via a multiplicative Hebbian learning rule; then we briefly discuss the experimental predictions based on our model.

4.1 Spontaneous activity in a network with recurrent connections

In the following we describe neuronal activity in a network of n neurons, each of which is characterized by its firing rate $f_i, i = 1, 2, \cdots, n$ (Fig. 4.2). Connection strengths among (pairs of) neurons are given
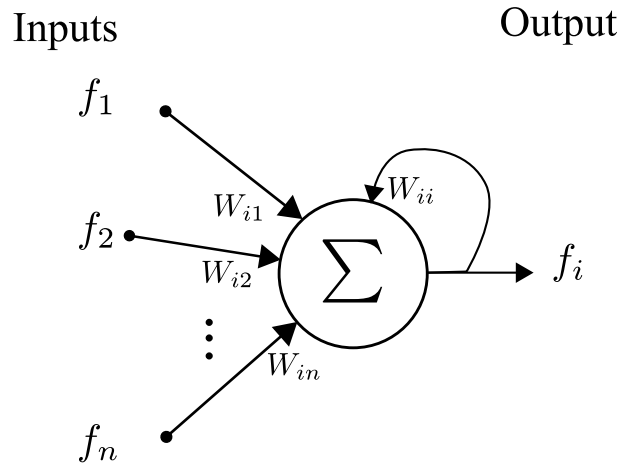
Figure 4.2: Simulations used linear neurons as basic units of computation. Output of i^{th} neuron (f_i) was computed as a weighted sum of its inputs ($f_1, f_2, \cdots, f_n$). Synaptic connection from neuron j to neuron i had an associated weight (synaptic strength) W_{ij}.

by the square connection matrix $\boldsymbol{W}$, with W_{ij} equal to strength of connection from neuron j to neuron i (Fig. 4.3). To describe spontaneous activity in this network we assume that spontaneous firing rates are primarily determined by recurrent intracortical connections (Douglas et al., 1995). Therefore, the network activity can be described with the following linear model:

$$f_i = \sum_{j=1}^{n} W_{ij} f_j, \tag{4.1}$$

where activity (firing rate) f_i of each neuron is computed as a weighted sum of activities of all neurons in the network. The i^{th} row of matrix $\boldsymbol{W}$ then contains *input* synaptic strengths for i^{th} neuron, and j^{th} column contains *output* synaptic strengths for j^{th} neuron.

When written in a compact matrix form, Eq. 4.1 becomes:

$$\boldsymbol{f} = \boldsymbol{W}\boldsymbol{f}, \tag{4.2}$$

which is an eigenvector problem for matrix $\boldsymbol{W}$. Thus the vector of firing rates $\boldsymbol{f}$ is an eigenvector of synaptic connection matrix $\boldsymbol{W}$ with a corresponding unit eigenvalue. Such a description is consistent with the notion of $\boldsymbol{f}$ representing spontaneous activity, i.e. not varying (much) in the absence of external inputs. In the following we assume that the synaptic connectivity matrix has a unit princi-
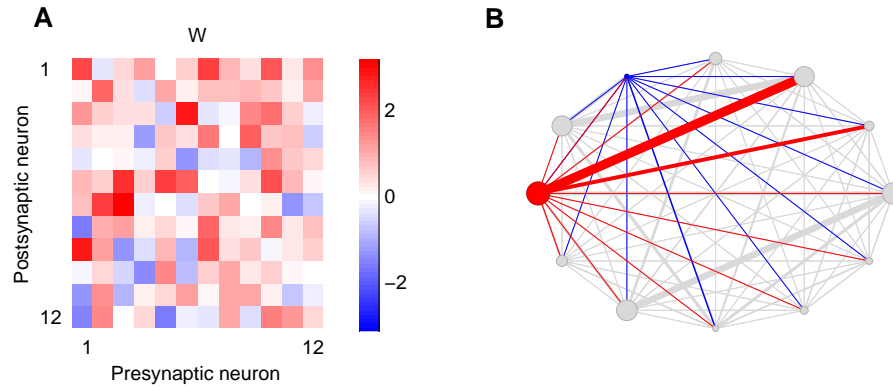
Figure 4.3: Connectivity matrix W described connection strengths, and its eigenvector f provided spontaneous firing rates. **(A)** Connectivity matrix for network in (B). Connection strengths were randomly drawn from lognormal distribution. Strengths of connections are given in arbitrary units. Note that, in general, synaptic strengths can be zero, i.e. some connections might be "missing." Note also that connections are unidirectional, so that, in general $W_{ij} \neq W_{ji}$, i.e. the connectivity matrix W does not have to be symmetric. **(B)** Cortical activity was modelled as a recurrent network. Here we show an example network of $n = 12$ neurons (circles), where each neuron is connected (lines) to all other neurons. Circle diameters are plotted proportional to neuron's firing rate (f_i), and line widths are proportional to synaptic strengths of *incoming* connections for each neuron. Red, and blue neurons are neurons with maximum, and minimum firing rate, respectively.

pal eigenvalue, with all other eigenvalues much lower than one. Then, the firing rates are given by the principal eigenvector f of W.

To model the *dynamics* of the network given by Eq. 4.2 we rewrite the equation in terms of temporal evolution of firing rates, $f(t + \Delta t) = W(t)f(t)$. Such interpretation leads to an important constraint on the connectivity matrix values. If the synaptic weights are large, firing rates in subsequent iterations will grow without bounds. Analogously, small weight values will cause firing rates to decay over time. To ensure stable firing rates, the connectivity matrix must be normalized in each step. Such normalization could be achieved by many general mechanisms thought to stabilize neuronal activity, such as changes in synapse number (Engert and Bonhoeffer, 1999), changes in neuronal excitability (Desai et al., 1999), or activity dependent synaptic scaling (Desai et al., 2002; Turrigiano et al., 1998).

4.2 Independent connections among neurons

In a straightforward implementation of Eq. 4.1, we first generated a synaptic connectivity matrix with lognormally distributed uncorrelated elements (Fig. 4.4). The parameters of the lognormal distribution

were mean = 0.5 (on a log scale), and unit standard deviation, which were similar to experimentally observed values (Song et al., 2005).

However, spontaneous firing rates (given by the principal eigenvector) were not lognormally distributed. In fact, all firing rates were very similar, as shown in Fig. 4.4. The distribution of firing rates was indeed very different from the experimentally observed distribution (Fig. 4.1), in which firing rates varied across several orders of magnitude.

The failure of the simple model with *uncorrelated* lognormally distributed synaptic weights is summarized in a cartoon in Fig. 4.4. In the figure, neurons are depicted by circles where the diameter is proportional to the firing rate. Connecting lines correspond to synaptic connections, and line width is proportional to connection strength. Because the connection matrix was uncorrelated, the strengths of incoming connections for each neuron came from a distribution with the same mean. Therefore, if the number of inputs to each neuron was large enough, each neuron received the same average input, which led to very similar neuronal firing rates. The example in Fig. 4.4 shows 12 neurons (and their connections) drawn randomly from the principal eigenvector f. We identified two neurons (out of the 12) with maximum (red) and minimum (blue) firing rate. Note that both firing rates (diameters) and incoming weights (line widths) of the colored neurons were very similar.

Model dynamics based on a connectivity matrix with uncorrelated synaptic weights failed to replicate experimental data. Therefore, we explored different types of correlations among synaptic weights.

4.3   Output weight correlations do not lead to lognormal distribution

We first introduced correlations among *output* synaptic weights for each neuron. These correlations would correspond to correlations among synapses on the axon of each neuron. Recall that output synaptic weights for a particular neuron j were stored in the j^{th} column of the connectivity matrix W, therefore, the connectivity matrix with output weight correlations appeared as a column matrix (Fig. 4.5). Correlations were introduced by first generating a connectivity matrix with lognormally distributed uncorrelated elements (mean = 0.5, std = 1), and then multiplying each column with a value randomly drawn from a lognormal distribution with the same parameters. The resulting distribution of synaptic weights was then normalized to have unit standard deviation. It can be shown that the distribution of the
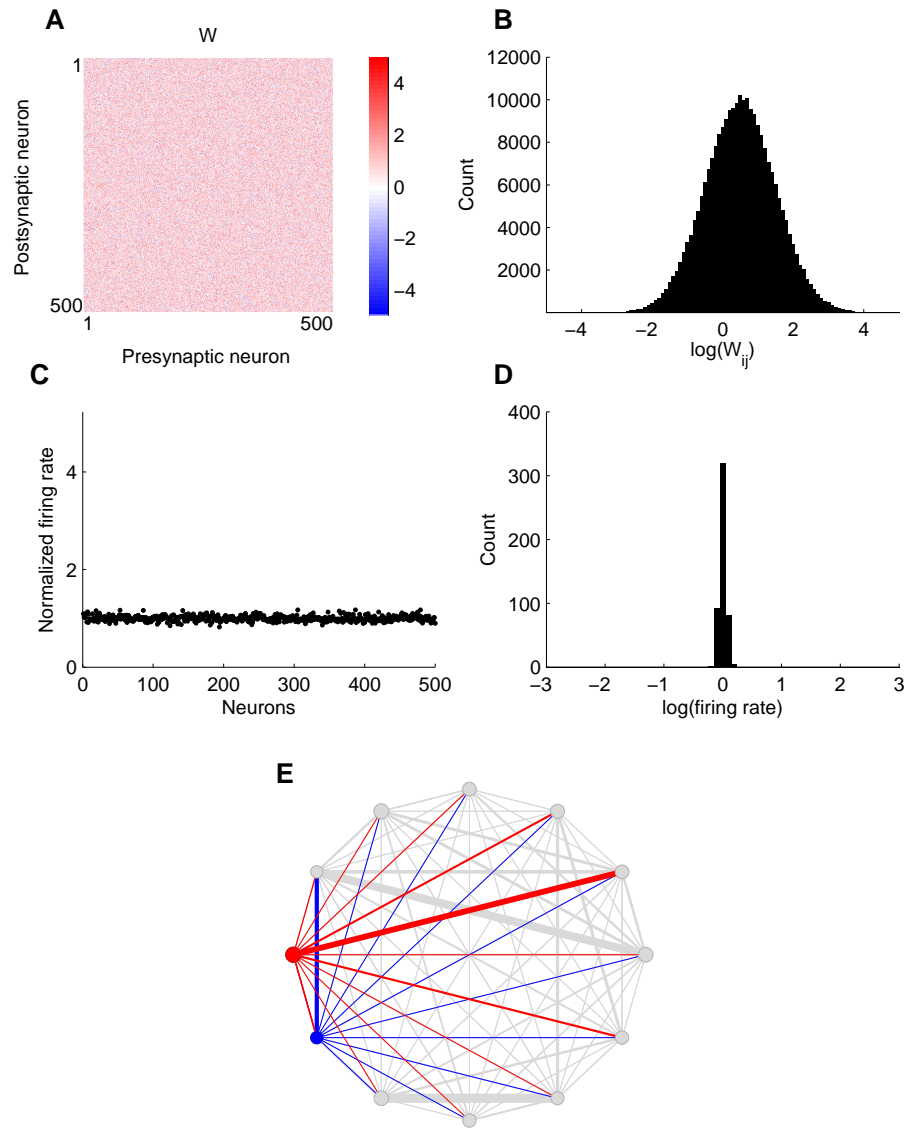
Figure 4.4: Random uncorrelated connection weights do *not* lead to lognormal distribution of spontaneous firing rates. **(A)** Synaptic connectivity matrix (W) for 500 neurons. Synaptic strengths spanned several orders of magnitude. We therefore plotted logarithms of synaptic strengths (color scale) to stress their normal distribution in a log-space. Because synaptic strengths were uncorrelated, the weight matrix looked like a "white-noise" matrix. **(B)** Distribution of synaptic weights in W was lognormal, i.e. appeared normal when plotted on a logarithmic scale. **(C)** Spontaneous firing rates arising from neuronal network with synaptic weights given by W were similar. The set of firing rates was given by principal the eigenvector of W (see Eq. 4.1.) The rates were normalized to have a unit average. **(D)** Distribution of firing rates (from C) was not lognormal, all firing rates were approximately the same. **(E)** Synaptic weights and firing rates of 12 randomly chosen neurons tended to be similar. Each circle corresponds to one neuron, with diameter proportional to its firing rate. Thickness of connecting lines corresponds to strengths (synaptic weights) of *incoming* connections for each neuron. Neurons (and connections) shown in red and blue show two neurons with maximum, resp. minimum firing rates from the sample shown. Note that firing rates (diameters) and synaptic weights (line widths) tend to be similar.

products of two numbers drawn from lognormal distributions is itself lognormal. Briefly, let's assume that variables η and ν follow normal distributions. Then variables e^η and e^ν follow lognormal distributions, and their product $e^\eta e^\nu = e^{\eta+\nu}$ also follows lognormal distribution, because it is an exponential of a sum of two normally distributed variables.

However, output (column) correlations in the connectivity matrix W did not lead to lognormal distribution of firing rates (Fig. 4.5). Although synaptic weights were lognormally distributed, firing rates (principal eigenvector) were not. The cartoon in Fig. 4.5 demonstrates that all neurons received connections with, on average, the same synaptic strengths. Note that although red and blue lines connected to grey neurons (output connections) tended to have the same thickness for each grey neuron, the lines connected to red and blue neurons (input connections) had, on average, the same width.

4.4   Input weight correlations lead to lognormal distribution

We then introduced correlations among *input* synaptic weights for each neuron. These correlations would correspond to correlations among synapses on a dendritic arbor of each neuron. The connectivity matrix with input weight correlations appeared as a row matrix (Fig. 4.6), because input synaptic weights for a given neuron i were stored in the i^{th} row of matrix W. Correlations in the connectivity matrix W were introduced by first generating lognormally distributed uncorrelated values (mean = 0.5, std = 1) and then multiplying each row with value drawn randomly from lognormal distribution with the same parameters. The resulting distribution was then normalized to have unit variance.

Spontaneous firing rates generated by our model displayed a wide range of values, and were lognormally distributed. It can also be shown that principal eigenvector values (firing rates) of a (substantially large) row connectivity matrix with lognormally distributed values are also lognormal. The reason why row-correlated connectivity matrix yielded lognormally distributed spontaneous firing rates is shown in Fig. 4.6. Note the incoming connections of two colored neurons: a red one with maximum firing rate, and a blue one with minimum firing rate. The red neuron received connections with much higher average synaptic strength (line width) than the blue neuron. Note also that the red and blue lines leaving individual grey neurons had very different widths, corresponding to very different *output* synaptic strengths.
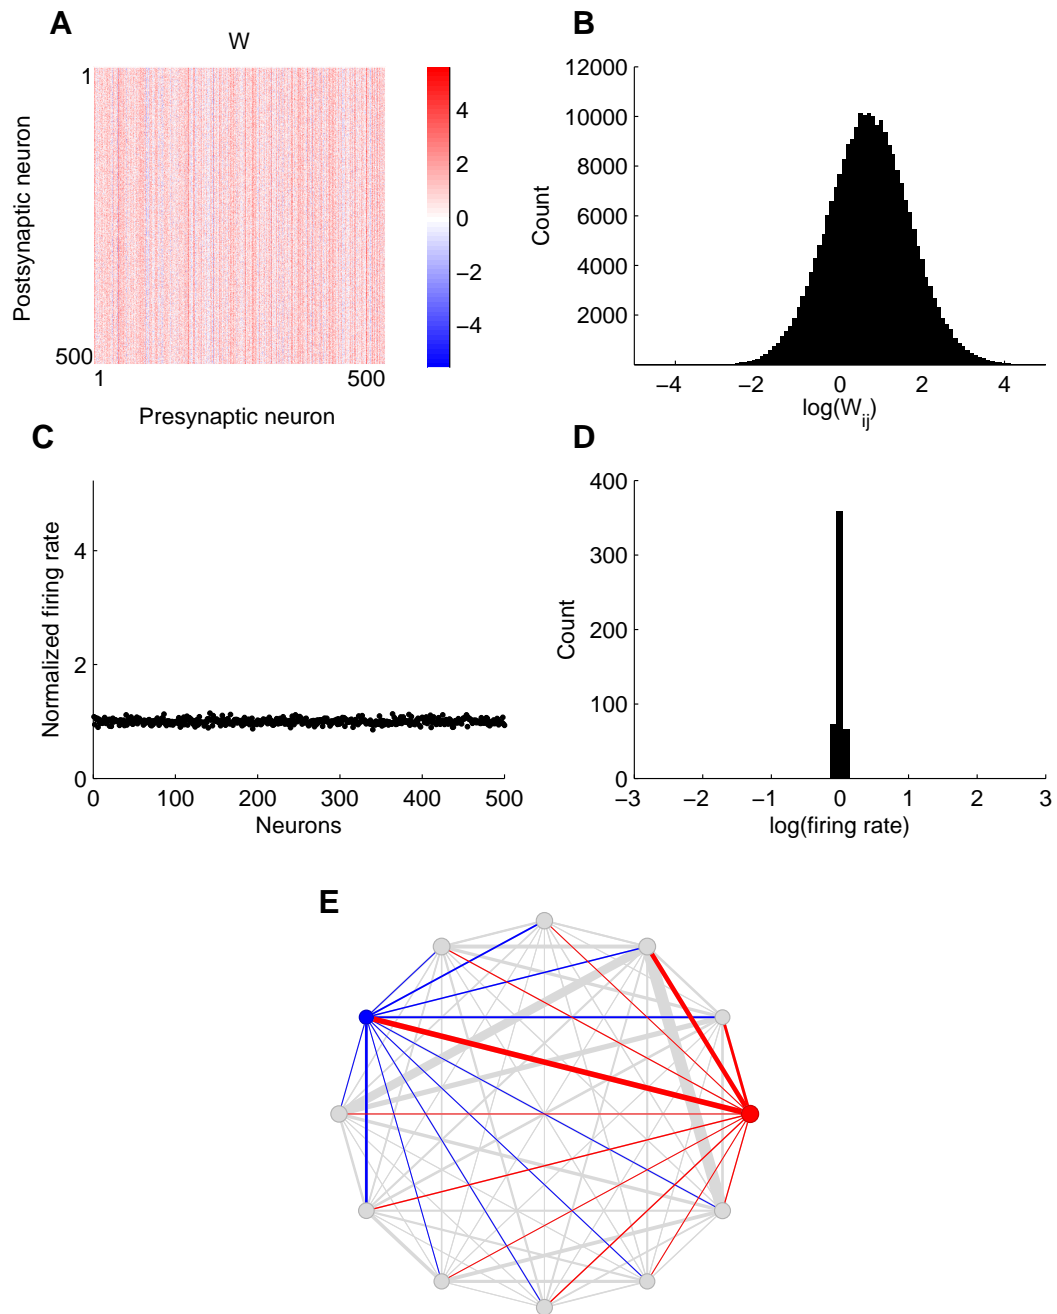
Figure 4.5: Correlations among synaptic weights on the same axon (output correlations) do *not* lead to lognormal distribution of spontaneous firing rates. Same format as in Fig. 4.4. **(A)** Synaptic connectivity matrix (W) for 500 neurons. Note the vertical "stripes" showing column correlations. **(B)** Distribution of synaptic weights was lognormal. **(C)**, **(D)** Spontaneous firing rates (principal eigenvector of W) were very similar. Firing rates were normalized to have unit average. **(E)** Firing rates and synaptic weights tended to be similar, as shown on an example of 12 randomly chosen neurons. Red and blue circles show neurons with maximum, resp. minimum firing rates (out of he sample shown), with their corresponding incoming connections.
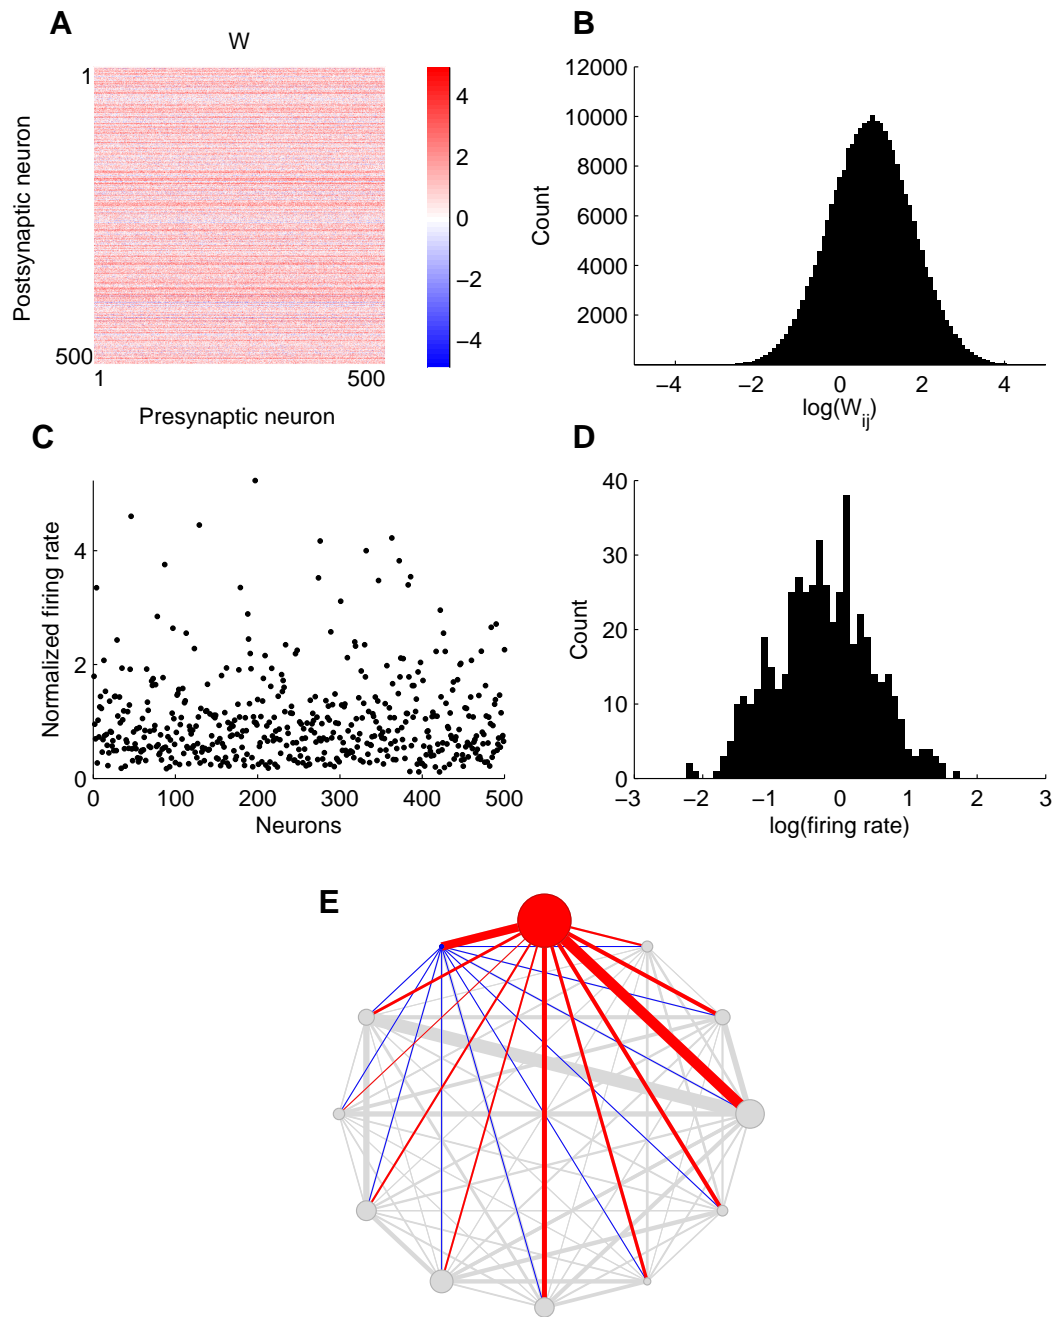
Figure 4.6: Correlations among synaptic weights on the same dendrite (input correlations) lead to log-normal distribution of spontaneous firing rates. Same format as in Fig. 4.4. **(A)** Synaptic connectivity matrix (**W**) for 500 neurons. Note the horizontal "stripes" showing input correlations. **(B)** Distribution of synaptic weights was lognormal. **(C)** Spontaneous firing rates varied with occasional high values. Firing rate were normalized to have unit average. **(D)** Distribution of spontaneous firing rates was log-normal. **(E)** Set of 12 randomly chosen neurons shows few neurons with high firing rates (circles), and higher corresponding synaptic weights (line widths). Red and blue circles show neurons with maximum and minimum firing rates (and their corresponding incoming connections).

These results suggest that row-correlated connectivity matrix with lognormally distributed synaptic weights is sufficient to replicate experimental observations of lognormally distributed synaptic weights and lognormally distributed spontaneous firing rates.

4.5   Required connectivity can arise by a multiplicative learning rule

In previous sections we assumed that synaptic weights given by connectivity matrix W were already lognormally distributed. As we mentioned above, synaptic weights can be influenced by neuronal activity and vice versa. What are then the mechanisms which could lead to lognormal distribution of synaptic weights?

The interplay between neuronal activity and strength of neuronal connections is often described in terms of Hebbian plasticity rules. In general, Hebbian rules describe changes in synaptic strengths depending on correlations between pre- and postsynaptic firing (Abbott and Nelson, 2000), where the "often-used" (effective) synapses are strengthened and ineffective ones weakened. In this section we propose a *multiplicative Hebbian learning rule* capable of producing the required structure of connectivity matrix with lognormal distribution of synaptic weights.

The learning rule for synaptic weights defines how the weights change in subsequent time steps. In the Hebbian learning rule, the change is dependent on presynaptic and postsynaptic activity (firing rates f_i and f_j, see also Eq. 4.1). In contrast to the conventional Hebbian learning rule we propose that the synaptic weight changes are depedent on the weights as well:

$$\Delta W_{ij}(t) = \varepsilon_1 f_i^{\alpha}(t) W_{ij}^{\beta}(t) f_j^{\gamma}(t) - \varepsilon_2 W_{ij}(t), \tag{4.3}$$

where ΔW_{ij} stands for synaptic weight change in time t, W_{ij} is a synaptic weight of a connection from neuron j to neuron i, and f_i, f_j stand for firing rates (activity) of neurons i and j respectively; ε_1, ε_2, α, β, and γ are parameters. The two terms in Eq. 4.3 describe Hebbian learning ($f_i^{\alpha} W_{ij}^{\beta} f_j^{\gamma}$) and (passive) synaptic decay ($-W_{ij}$). Parameters ε_1 and ε_2 determine contributions of Hebbian learning and synaptic decay.

In each time step (iteration) Δt, synaptic weights are updated according to:

$$W_{ij}(t + \Delta t) = W_{ij}(t) + \Delta W_{ij}(t). \qquad (4.4)$$

Power parameters α, β, γ are essential parameters of the learning rule. The sum of parameters $\alpha + \beta$ must be below 1. If $\alpha + \beta > 1$, synaptic weights would grow in each time step, given that the synaptic decay is weaker than Hebbian potentiation. It can be shown that if $\alpha + \beta \to 1^-$ then the distribution of synaptic weights (given by $\boldsymbol{W}$) becomes lognormal.

We demonstrate this result in Fig. 4.7. We simulated the proposed multiplicative Hebbian rule in a population of $n = 500$ neurons. For simulation parameters we used $\varepsilon_1 = \varepsilon_2 = 0.1$, and $\alpha = \beta = \gamma = 0.48$. We started the simulation by generating a binary connectivity matrix $\boldsymbol{W}(0)$ with randomly assigned values $W_{ij}(0)$ of 0 or 1 determining whether neuron j was connected to neuron i ($W_{ij}(0) = 1$), or not ($W_{ij}(0) = 0$). The probability of $W_{ij} = 1$ was set to 0.15, similar to probabilities reported in literature (Thomson and Bannister, 2003; Thomson et al., 2002). The initial connectivity matrix $\boldsymbol{W}(0)$ was thus equivalent to *adjacency matrix* describing connectivity regardless of connection strength. Note that $\boldsymbol{W}(0)$ was not necessarily a symmetric matrix, because all connections $W_{ij}(0)$ were generated randomly. Matrix $\boldsymbol{W}(0)$ then described a directed graph with $n = 500$ vertices.

After 100 iterations, the connectivity matrix $\boldsymbol{W}$ developed a "plaid" structure exhibiting both row and column correlations (Fig. 4.7). The proposed learning rule preserved the initial sparse circuit connectivity given by $\boldsymbol{W}(0)$, i.e. new connections were not formed and existing connections did not disappear. Both synaptic connection weights and spontaneous firing rates followed lognormal distributions (Fig. 4.7). The example network in Fig. 4.7 (compare to Fig. 4.6) shows that neuronal firing rates (circle diameters) varied, and neurons with higher activity (red) received stronger connections (wider connecting lines) than neurons with lower activity (blue). The proposed multiplicative learning rule was thus capable of reproducing both experimental observations: lognormal distribution of synaptic weights and lognormal distribution of (spontaneous) firing rates.

A W

B Adjacency matrix

C

D

E
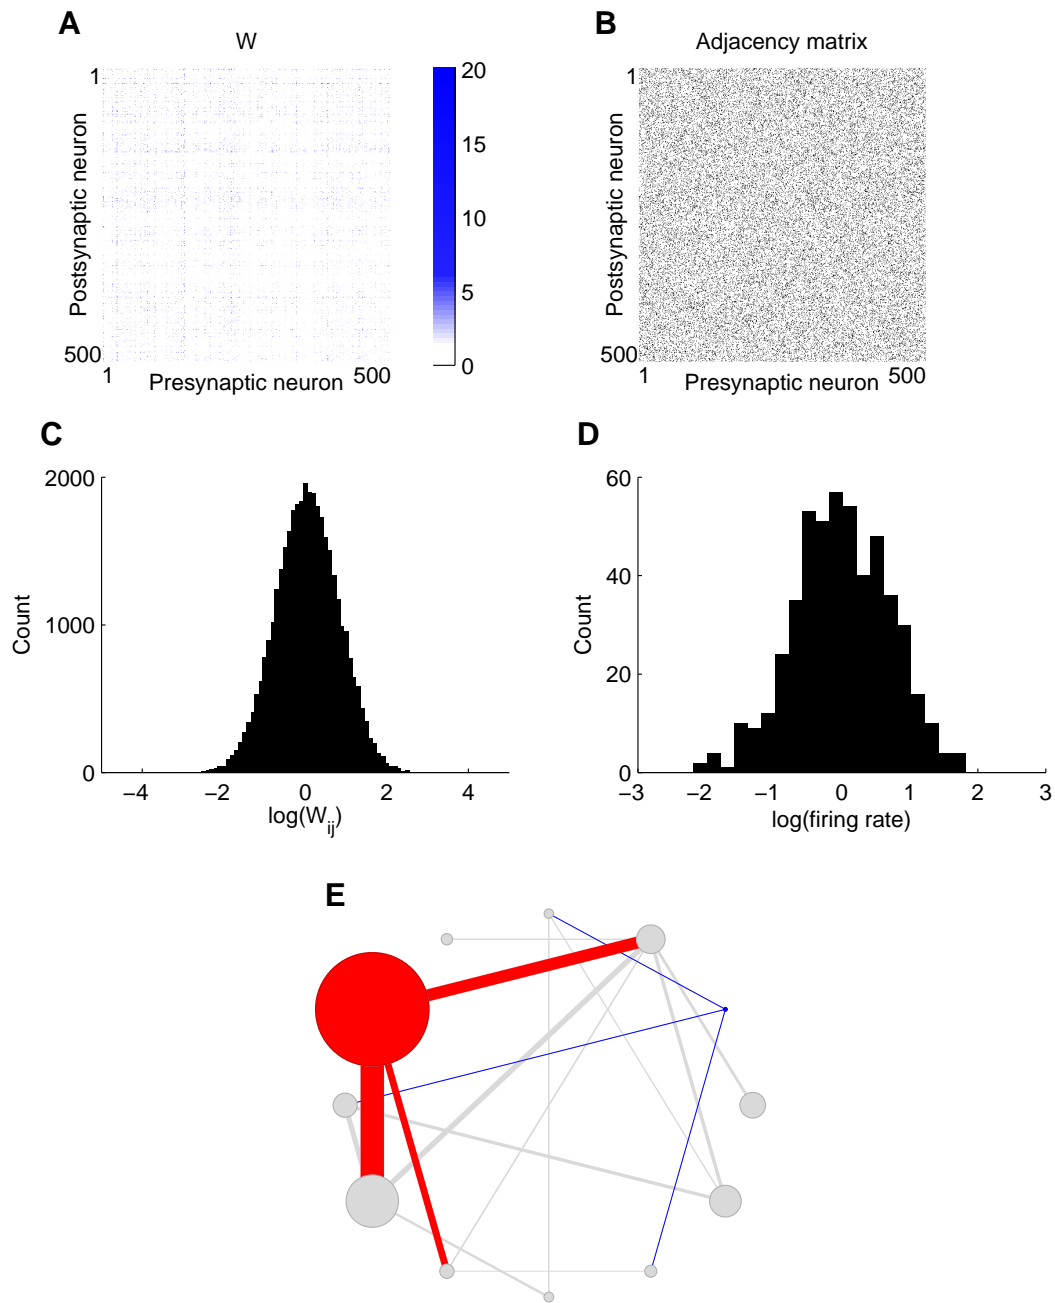
Figure 4.7: **(A)** Synaptic connectivity matrix (500 neurons) with "plaid" structure (horizontal and vertical "stripes".) This matrix arose after 100 iterations of multiplicative Hebbian learning rule (see text for details.) Note that here we do *not* plot $\log(W)$ as in the previous figures. **(B)** Adjacency matrix for W was not symmetric, i.e. synaptic connections formed a directed graph. **(C)**, **(D)** Distributions of synaptic weights (C) and firing rates (D) were lognormal, i.e. appeared as normal on logarithmic axis. **(E)** Set of randomly selected 12 neurons shows occasional neurons with high firing rates (circles), with stronger corresponding synaptic weights (line widths). Red and blue circles show neurons with maximum and minimum firing rates and their corresponding incoming connections. See also Fig. 4.6 for comparison.

4.6 Model predicts synaptic correlations

We have described a theoretical model of cortical dynamics, which is able to capture two experimental observations about synaptic weights and neuronal firing rates. Our model implies that row-correlations in the connectivity matrix are sufficient for lognormal distribution of firing rates in the given circuit. This means that synapses on a given dendrite (dendritic arbor) would have similar strengths. However, mean synaptic strength across the dendrites of one neuron should be different, on average, from mean synaptic strength across the dendrites of another neuron. Furthermore, if the two lognormal distributions arise according to the multiplicative Hebbian rule, one would expect to find correlations among output synapses (on the same axon) as well, thanks to the "plaid" structure of the connectivity matrix. Two signature properties should then be present: on the same dendrite synaptic strength are correlated with different values from dendrite to dendrite; and synaptic strengths are correlated on axons of individual cells.

We suggest that two signatures of the proposed model could be detected experimentally. Two-photon imaging techniques can be used to measure properties of individual synapses localized on different neurons. Postsynaptic indicators of synaptic strength would allow the measurement of *input* correlations, i.e. detecting the row structure in a connectivity matrix. Such postsynaptic indicators would include AMPA receptors, NMDA receptors, PSD-95 density in individual synapses (Gray et al., 2006; Kopec et al., 2006; Shi et al., 1999). Detecting changes in synaptophysin, synaptobrevin, or VAMP fluorescence (Ahmari et al., 2000; Bozza et al., 2004; Miesenböck et al., 1998; Park et al., 2006) could be used to estimate *output* synaptic weights, i.e. the column structure in the synaptic weight matrix. The presence of both predicted signatures would then suggest that synaptic weights and firing rates are "interconnected" via a multiplicative Hebbian learning rule.

4.7 Discussion

We have studied properties of cortical networks which could give rise to lognormal distributions of synaptic weights and firing rates. We propose that correlations among neurons' synaptic inputs lead to a lognormal distribution of firing rates in a neuronal population, provided that synaptic strengths themselves are lognormally distributed. In addition we propose a multiplicative Hebbian learning rule which

can give rise to specific correlations in the synaptic weight matrix and, consequently, to lognormally distributed synaptic weights and firing rates.

We have used a simple linear model based on the observed distributions of firing rates and synaptic weights. The model does not incorporate any possible temporal relationship among neuronal inputs, such as spiking input correlations. Additional factors—which might describe various properties of neurons—are also not included in the model. Suppose, for example, that the excitability of each neuron was described by a "neuronal excitability factor" α_i. Then, the circuit dynamics described by $f_i = \alpha_i \sum W_{ij} f_j$, with lognormally distributed α_i and independent W_{ij}, would lead to lognormally distributed f_i. Such "excitability," however, would not have any associated learning rule, like the one proposed above (Sec. 4.5). Thus, although our proposed linear model is *sufficient* to explain the experimental observations and make predictions, it is not *necessary* to explain the observed data. There is a large class of other possible models—nonlinear, utilizing temporal input correlations, or utilizing additional factors, such as "excitability"—which might also account for the observations.

4.7.1 Lognormal distribution of firing rates

We used cell-attached patch clamp recordings to record neuronal activity. Cell-attached recording differs from conventional extracellular recording methods—especially from tungsten recordings with single electrode—in its selection bias (see also Margrie et al., 2002). In conventional recording, single-unit isolation requires neural activity: the experimenter typically searches for neurons by isolating either stimulus-evoked or spontaneous spikes. In either case, neurons with low firing rates—spontaneous or evoked—tend to be undersampled. During cell-attached recording, by contrast, a glass patch pipette is advanced until it is in physical contact with the neuron, so neurons are selected solely on the basis of the experimenter's ability to form an electrical seal between the pipette tip and the neuronal membrane. Seal formation does not require neuronal activity, so even neurons with very low firing rates are as likely to be included in the sample as those with high firing rates.

Although lognormal distributions have been used widely to describe the interspike interval distributions from a single neuron, population responses are usually reported to be exponentially distributed (Abeles et al., 1990; Baddeley et al., 1997; Gaese and Ostwald, 2003). We propose that the different sampling bias of cell-attached recordings enabled us to sample the low tail of firing rate distribution.

It would be of interest to see whether lognormal distribution of firing rates is present in neuronal data sets obtained using different recording techniques with similar recording biases, such as optical (Ohki et al., 2005), tetrodes (Feierstein et al., 2006), or silicone probes (Barthó et al., 2004).

4.7.2   Signatures of multiplicative Hebbian plasticity

We propose that lognormal distribution of synaptic weights (Song et al., 2005) could arise in a recurrent network as a result of a multiplicative Hebbian learning rule. The presence of this plasticity rule could be inferred from the "plaid" structure of the synaptic connectivity matrix. Correlations among input synaptic weights ("row" correlations) and correlations among output synaptic weights ("column" correlations) could be detected experimentally by measuring pre or postsynaptic synaptic strengths respectively. Some correlations have already been observed (Song et al., 2005, Supplementary material).

Interestingly, the distribution of input synaptic weights has been *indirectly* measured previously (Konur et al., 2003). The authors measured the distribution of *diameters* of synaptic spines on dendritic arbors of pyramidal neurons in cortex and hippocampus. Since spine volume is proportional to spine diameter cubed, and spine volume seems to be linearly correlated with synaptic strength (as given by the size of the postsynaptic density and the number of presynaptic docked vesicles), spine diameter could serve as an indirect measure of (input) synaptic strength. The coefficient of variation of spine diameters in the same neuron was around 0.3, whereas the coefficient of variation for the means of spine diameters across cells was 0.14. For the row matrix in Fig. 4.6 the coefficient of variation of cube roots of input synaptic weights (rows) was around 0.23, and the coefficient of variation across cells/rows was 0.24. Moreover, distributions of cube roots of synaptic weights in each row resembled slightly skewed distribution of spine diameters in Konur et al. (2003).

We propose that one could test for the presence of correlations by first measuring the spread of the distribution of (logarithm of) mean synaptic strengths within cells (row distribution for example), and then measuring the spread of an analogous distribution obtained from a reshuffled (randomized) set of synaptic weights. If the randomization procedure destroys any existing correlations, or if there were no correlations present in the original sets, then the mean synaptic strengths *after randomization* will be similar and their distribution will be narrow. Thus the distribution of (logarithms of) mean synaptic

weights significantly wider than an analogous distribution of mean synaptic weights after reshuffling would indicate the presence of correlations among synaptic weights.

4.7.3   Lognormal distributions

The normal (Gaussian) distribution is probably the most often assumed (and consequently used) distribution of observations across scientific disciplines. However, almost all physical and chemical laws in nature are ruled by multiplication, not by addition (Arita, 2005). And given that many independent multiplicative processes will give rise, in their limits, to lognormal distribution, it should come as no surprise that many processes in nature were identified as being lognormally distributed. For example, such different sets of observations as survival times after cancer diagnosis, size of crystals in ice cream, size of oil drops in mayonnaise, length of spoken words in phone conversations, and the age of marriage of women in Denmark (Limpert et al., 2001, and references therein) all follow lognormal distributions. It remains to be seen whether any of the mechanisms we propose (or similar) could be responsible for generation of lognormal distributions in other systems.

Chapter 5

Up-states are rare in awake auditory cortex

Dynamics of activity in cortical networks is dependent on complex interactions between neuronal activity, synaptic strengths, and intrinsic neuronal properties (Getting, 1989; Harris-Warrick and Marder, 1991). An intriguing example of an intrinsic neuronal property is bistability of the neuronal membrane potential, often described in terms of "up" and "down" states. During the two states, membrane potential transits between two main values, resting *down* state, and depolarized *up* state, staying in both states often for several seconds, while spending very little time in intermediate values (Cowan and Wilson, 1994; Steriade et al., 1993a).

Up and down states have attracted a considerable attention, as they have been suggested to underlie persistent activity in cortical networks (McCormick et al., 2003; Wang, 2001). The persistent activity, in turn, is thought to underlie such interesting processes as short-term (working) memory (Camperi and Wang, 1998; Goldman-Rakic, 1995; Marder et al., 1996; Wang, 2001), or attention (Chance et al., 2002; McCormick et al., 2003). Large, persistent changes in membrane potential can also occur in a response to behavioral state changes (Steriade et al., 2001), or sensory stimulation (Anderson et al., 2000; Carandini and Ferster, 1997). Up and down states have been observed in frontal cortical areas (Lewis and O'Donnell, 2000; Léger et al., 2005), somatosensory (Sachdev et al., 2004; Steriade et al., 2001), visual (Lampl et al., 1999), olfactory areas (Luo and Katz, 2001; Margrie and Schaefer, 2003), striatum (Stern et al., 1997; Wilson and Kawaguchi, 1996), and many others.

The precise origin and function of up and down states, however, remain unknown. Various studies suggested the involvement of intracortical mechanisms (Steriade et al., 1993b), intrinsic network properties (Lampl et al., 1999; McCormick et al., 2003), pairwise correlations among neuronal inputs (Benucci et al., 2004), or activity-dependent synaptic depression among excitatory neurons (Holcman and Tsodyks, 2006).

Answering the question of the significance of up-states is complicated by the fact that there is no universally agreed upon definition of what constitutes up and down states. Up-states are usually presented as stereotyped plateaus at which the membrane potential remains for a prolonged period (seconds). However, some studies have *failed* to observe up and down states, for example in somatosensory cortex (Bruno and Sakmann, 2006; Wilent and Contreras, 2005), or visual cortex (Anderson et al., 2000) where many neurons did not display typical up and down states. Up and down states were also not previously observed in studies performed in auditory cortex (DeWeese and Zador, 2006). Importantly, the vast majority of studies describing the canonical long up and down states have been conducted *in-vitro* or *in-vivo* in anesthetized preparations, and it is unclear whether the presence of up and down states represents a specific feature of anesthetized, or sleeping brain (Mahon et al., 2003). Few studies conducted in unanesthetized preparations suggested disruption of up-states depending on behavioral state (Crochet and Petersen, 2006; Steriade et al., 2001).

Here we used whole-cell patch-clamp recording techniques to record intracellular activity of neurons in auditory cortex of awake head-fixed rats. We have generalized the definition of up-states, and defined several membrane potential thresholds. We found that long up-states were rare in awake auditory cortex, with only 4 % of up-states longer than 200 ms. Most neurons displayed only brief up-states (bumps) and spent most of the time near their resting potential, typically spending less than 4 % of recording time in up-states longer than 200 ms.

5.1 Whole-cell recordings in awake rats are stable

We have used the whole-cell patch-clamp technique to record intracellular activity of neurons in auditory cortex of awake head-fixed rats. Altogether we recorded from 20 neurons in 13 animals (Sprague-Dawley rats, postnatal day 24–30).

We have successfully recorded intracellular activity for several minutes (Fig. 5.1A). The recordings were sufficiently stable and offered an excellent signal-to-noise ratio. Some recordings displayed small and very slow fluctuations in the resting membrane potential. Since we have analysed 7 s long traces (Sec. 2.6.3 on page 24), any small fluctuations in the resting membrane potential on a scale of minutes did not affect the analyses described below. Indeed, the coefficient of variation of the resting membrane potential (for each trace) across all neurons in our sample was 3.95±2.26 %, which would correspond to a mean difference of 2.5 mV from trace to trace for a neuron with -60 mV resting potential. Recordings described in this text lasted between 2–13 min (mean 7.4 min, median 7 min, n = 19 neurons), with one additional neuron from which we recorded for approximately 2 hours.

5.2 Membrane potential dynamics is typically "bumpy"

The basic features of subthreshold activity in auditory cortex of awake rats are demonstrated by an example trace shown in Fig. 5.1B. Most of the time, the membrane potential stayed near the resting level, and occasionally, large rather brief excursions of the membrane potential were observed. Sometimes the membrane potential excursions were high enough to reach threshold for action potential(s). Such excursions will be henceforth called bumps ("brief ones") or up-states ("longer ones"). Bumps (up-states) could be either spontaneous (i.e. with no apparent relationship to the auditory stimulus presented), or stimulus evoked (i.e. following the presentation of a sound stimulus after a brief delay). Both types of bumps are shown in Fig. 5.1B. The very first brief bump was probably spontaneous, and most of the other bumps were stimulus evoked (note the temporal relationship between stimuli and voltage trace in the figure). Also note that there were no apparent large fluctuations in resting membrane potential. All voltage traces shown in the remainder of this chapter will be median-filtered (i.e. action potentials have been removed; see Sec. 2.6.3 and Fig. 2.5 for details).

The basic impression of subthreshold activity as mostly staying near the resting level and occasionally displaying brief excursions is shown in more detail in Fig. 5.2. The three neurons shown displayed both spontaneous and tone-evoked bumps, but most of the time their activity stayed close to the resting membrane potential. The bottom trace in Fig. 5.2B shows a 1 minute of recording, documenting that the basic characteristics did not change even on longer time scales.
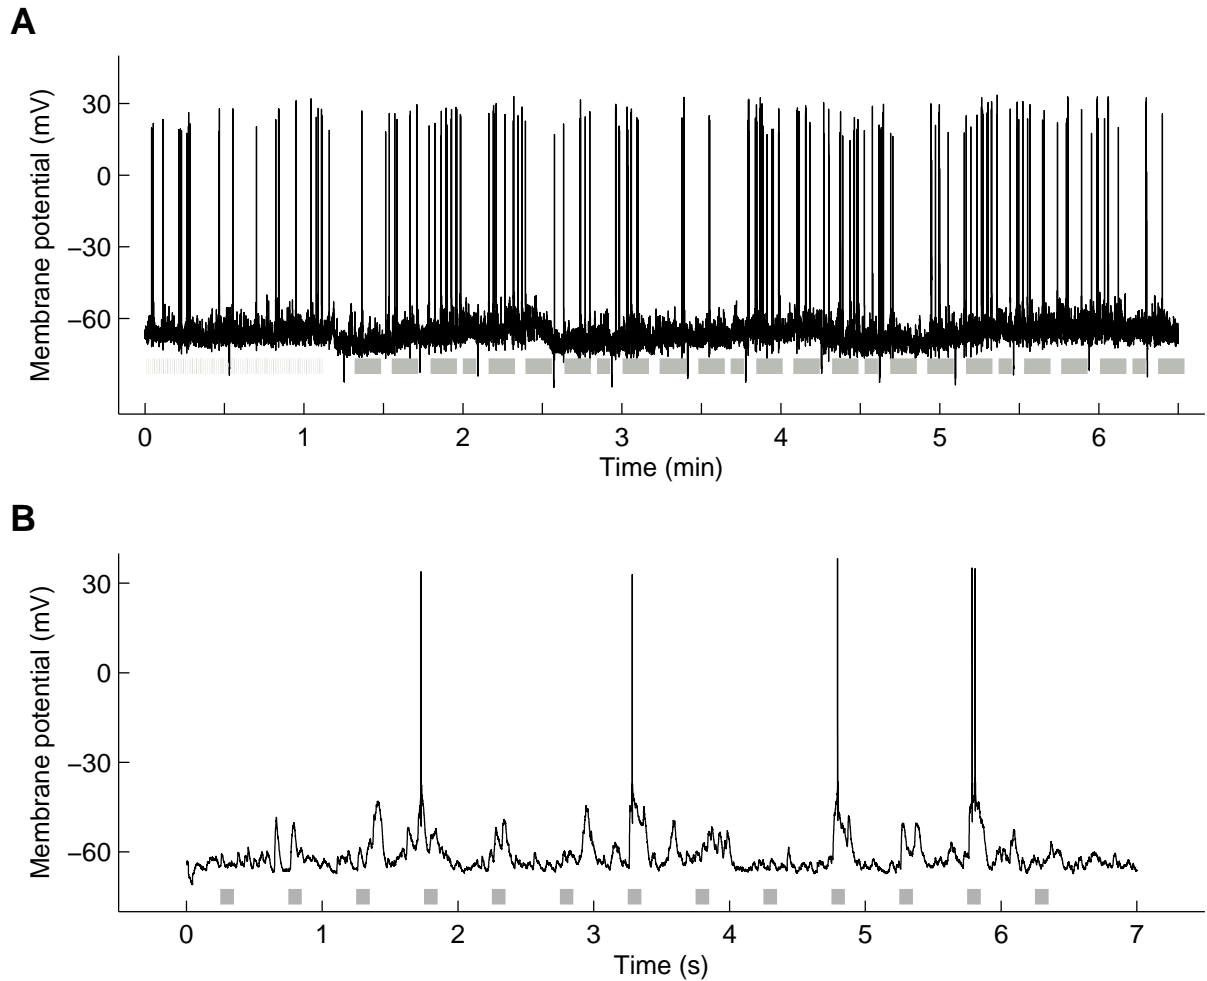
Figure 5.1: **(A)** Whole-cell recordings in unanesthetized auditory cortex can be stable for several minutes. Notice the small and slow fluctuations in resting membrane potential ("baseline") which could be caused by animal movements or electrode seal instability. Grey rectangles show positions of sound stimuli. Brief rectangles (up to approximately 1.1 min) show positions of 100 ms tones, long rectangles (rest of the trace) show positions of natural sounds played. Note the long time scale compared to other figures in this text. **(B)** Seven seconds long example of raw voltage trace recorded in awake head-fixed rat. Subthreshold activity usually consisted of brief excursions of membrane potential, both spontaneous and sound-evoked. Grey rectangles show positions of 100 ms long tones of different frequencies and intensities. Note the large signal-to-noise ratio of raw signal.

A

40 mV

0.5 s

B

5 s
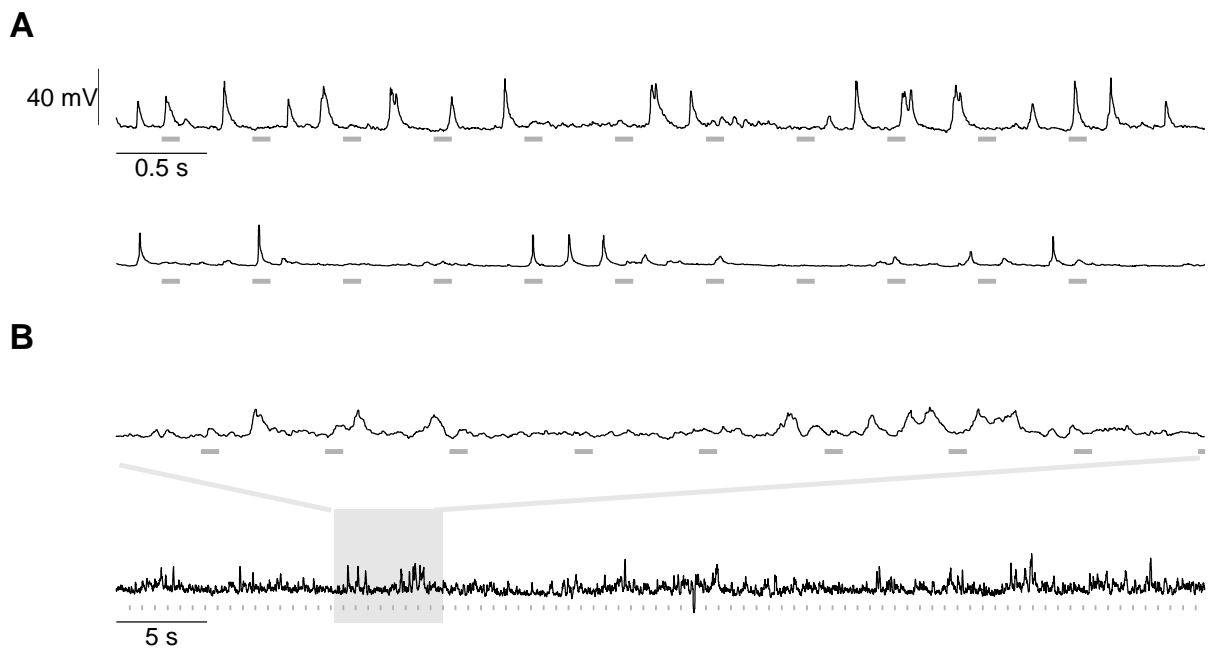
Figure 5.2: Subthreshold activity in unanesthetized auditory cortex was usually "bumpy." Examples of recordings from three different neurons demonstrate the main observation, that subthreshold dynamics in awake auditory cortex consisted of brief, sporadic fluctuations in membrane potential (bumps). **(A)** Each trace shows 6 s of subthreshold activity from different neurons. Grey rectangles show positions of 100 ms tones of different frequencies and intensities. **(B)** Top trace shows 6 s of subthreshold activity out of 1 min of activity shown in the bottom trace (grey square).
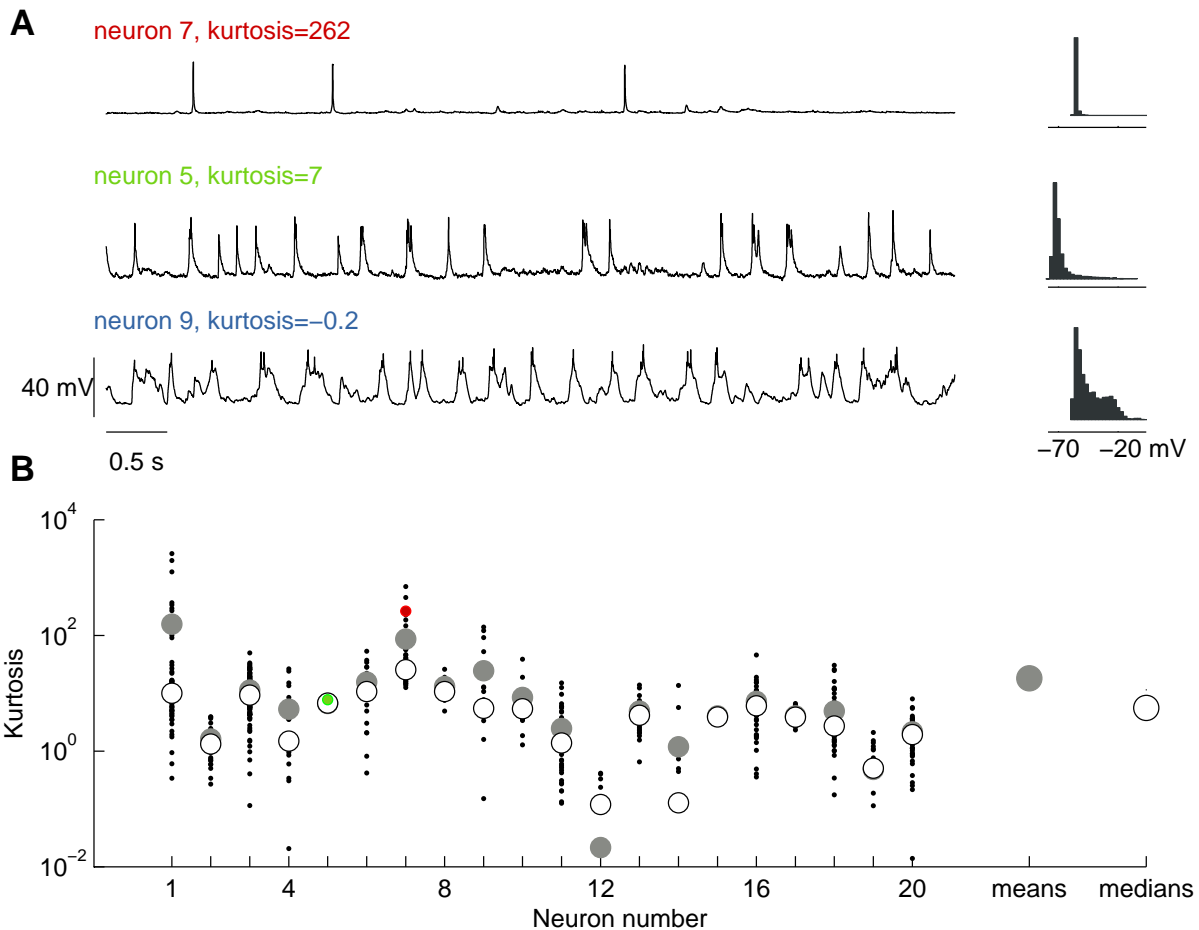
Figure 5.3: Kurtoses of recorded traces are usually large and positive, reflecting brief fluctuations of membrane potential. **(A)** Example traces from three different neurons with very different kurtoses. Histograms on the right show distributions of membrane potential values for each trace. The first trace (from neuron 7, red dot in top panel) displayed very high kurtosis thanks to very brief infrequent bumps. The second trace (from neuron 5, green dot in top panel) displayed still considerable kurtosis, thanks to brief bumps. The third trace (from neuron 9) displayed negative kurtosis close to zero and therefore was not plotted in the top panel. Nevertheless, the trace still contained brief bumps, which were, however, frequent, thus smearing the strong peak in the voltage histogram and driving kurtosis close to zero. Note that the voltage histogram for neuron 9 resembled bimodal distribution often used to describe up and down-states. **(B)** Recorded traces displayed, on average, large and positive kurtoses. Shown are distributions of trace kurtoses for all 20 neurons. **Full grey large circles** show means, and empty large circles show medians of kurtoses for each neuron. Note the logarithmic scale on y-axis. The large circles on the right show *mean of means* and *mean of medians* of kurtoses across neurons, their respective standard deviations were 37.6 and 5.9, both smaller then the circle diameter. Occasionally, some traces displayed kurtoses very close to zero, or even small negative values (see also the bottom trace). These are not displayed, because of logarithmic axis, but their presence can be inferred from asymmetrically positioned medians.

We have quantified the basic impression of "bumpy" activity by computing kurtosis of the membrane potential histogram for each recorded trace (DeWeese and Zador, 2006), see Sec. 2.6.3 on page 24 for details. Kurtosis measures "peakedness" of membrane potential distribution, and as such is large (greater than zero) for distributions with strong peaks and long tails, and zero for normal (Gaussian) distribution (Olshausen and Field, 2004). Therefore, traces with most of the activity near the resting potential and only occasional voltage excursions would have kurtosis greater than zero, because their membrane potential distribution would have a strong peak around the resting potential value, and a long tail corresponding to peaks of (occasional) bumps.

The examples in Fig. 5.3A document kurtosis values for three different neurons. A "canonical" trace from neuron 7 showed very brief infrequent bumps, and thus had a very high kurtosis ($k = 262$), which corresponded to very tall membrane potential histogram with a very long tail (shown on the right). The trace from neuron 5 showed considerably more activity. While the bumps were still brief, the neuron spent more time away from the resting potential, thus smearing the strong peak in the membrane potential histogram, which led to a decrease in kurtosis ($k = 7$). Finally, the third trace in Fig. 5.3 shows an example from neuron 9 displaying even higher activity manifested by a higher number of (wider) bumps. In this case the neuron spent even more time away from the resting potential and the trace kurtosis was $k = -0.2$. The membrane potential histogram of this trace (bottom right) resembled the bimodal histogram often studied in connection with canonical up and down states (for example in Stern et al., 1997).

Overall, kurtoses of subthreshold responses recorded in awake auditory cortex were typically large and positive (Fig. 5.3 top), reflecting occasional brief fluctuations in—otherwise resting—membrane potential. Individual neurons could display a wide range of kurtosis values, but both mean and median kurtoses for each neuron in our sample were positive. The mean and median kurtoses *across all traces* we analyzed were equal to 18 and 5.6 respectively. Note, however, that kurtosis values for traces like neuron 9 above are not plotted in Fig. 5.3 because of the logarithmic y-axis. Nevertheless, the presence of such low kurtoses can be inferred from asymmetrically positioned medians (empty circles) for some neurons.

5.3 Most up-states are brief

Although kurtosis values were typically large, corresponding to bumpy subthreshold dynamics, for some of the recorded traces we have observed low kurtoses, suggesting wide distributions of membrane potential values. Such low kurtosis values could arise in different ways. Membrane potential could fluctuate randomly around its mean value, giving rise to normal (unimodal) distribution of membrane potential values, or could display prominent up and down states, giving rise to bimodal distribution of membrane potential values. After a careful visual inspection, which verified the presence of bumps in recorded traces, we have concluded that random fluctuations were not the main cause of low kurtosis values in our data. Bimodal distribution of membrane potential could arise when neurons display *canonical long up and down states* (in the order of seconds) (Cowan and Wilson, 1994; Steriade et al., 1993a,b); or when neurons display *rapidly changing brief up and down states* (Petersen et al., 2003a) (see also the bottom trace in Fig. 5.3B). These two scenarios (canonical up and down states vs. brief, frequent bumps) would display very similar membrane potential distributions, and therefore, very similar kurtoses. To differentiate between these two scenarios we computed up-state durations in our data set.

We computed *up-state durations* as lengths of continuous membrane potential segments above a given threshold (Fig. 5.4A). The usual methods of computing up-state durations assume bistability of membrane potential and estimate up-state thresholds based on the resulting *bimodality* of membrane potential histograms (Lewis and O'Donnell, 2000; Stern et al., 1997). Alternatively, in the case of traces containing only *well isolated* bumps, one can measure durations at half maximum for each bump (DeWeese and Zador, 2006). However, we have observed a variety of subthreshold dynamics with most traces displaying *unimodal voltage histograms*, with the rest of the traces showing a *variety of bump shapes* (see Figs. 5.2–5.3 for examples). Therefore, we could not rely on the traditional methods and used instead *absolute* thresholds: 10, and 15 mV above resting membrane potential; and *relative* thresholds: 20 %, and 40 % of maximum potential amplitude (see Sec. 2.6.3 on page 24 for details). Gaps shorter than 20 ms were considered spurious fluctuations in membrane potential and were included in bumps (see an example in Fig. 5.4A). The lowest relative threshold (20 % of maximum amplitude) was always lower than 10 mV, and corresponded to 7.33±1.18 mV across all neurons, with occasional values as low as 5.24 mV above resting membrane potential. Both 10 mV and 20 % thresholds appeared to be *lower* than thresholds used to compute up-state durations in other *in-vivo* studies (Léger et al.,
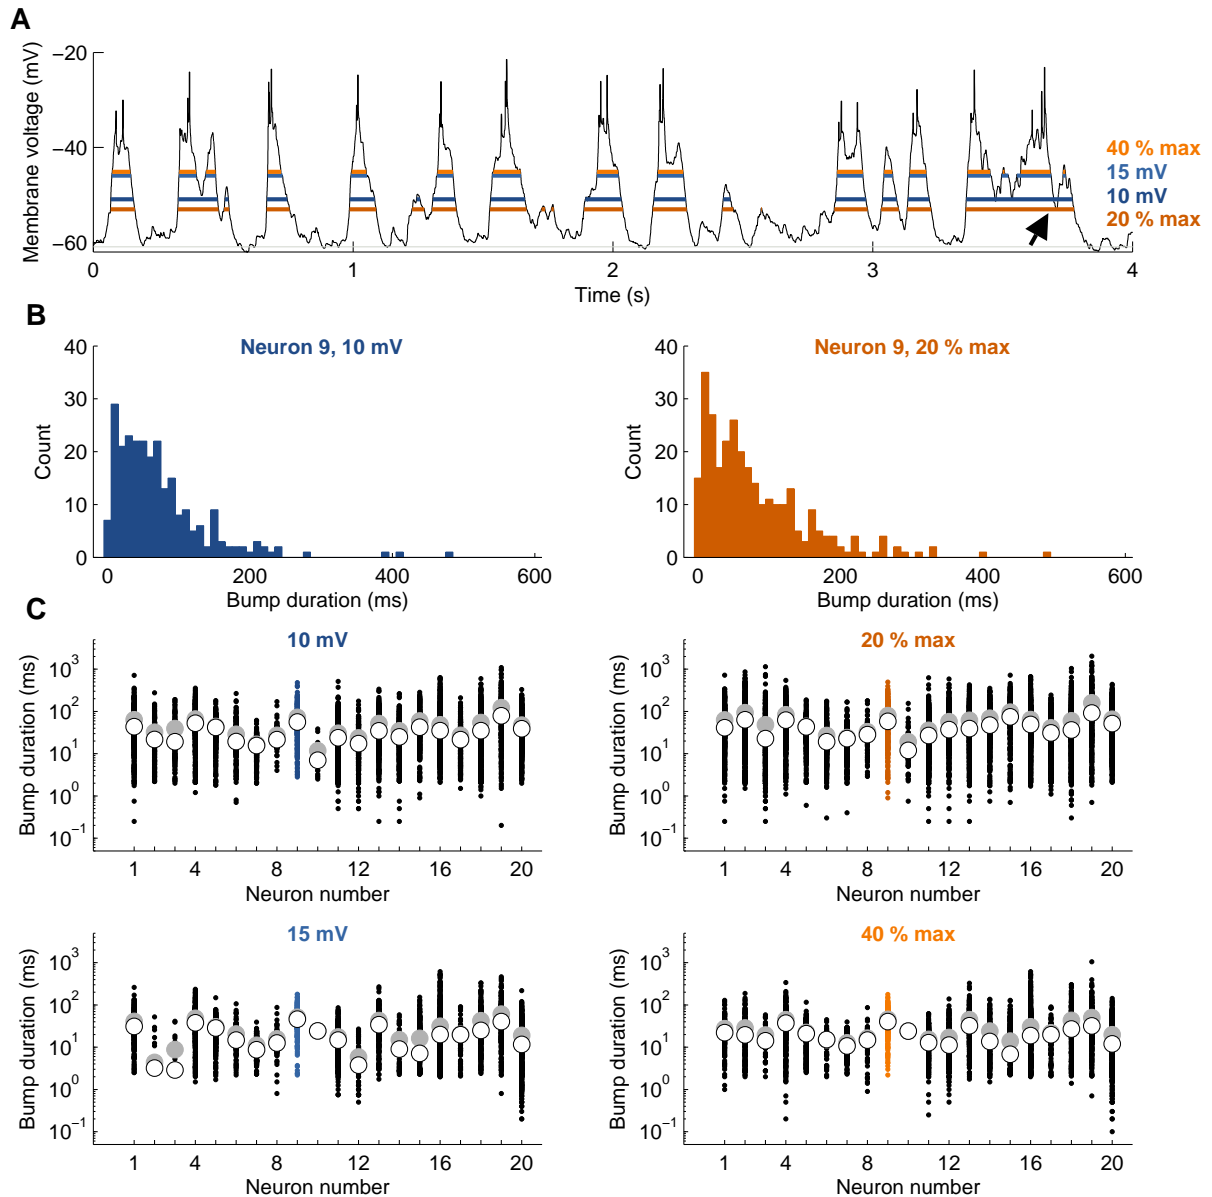
Figure 5.4: Up-states in unanesthetized auditory cortex were usually brief. **(A)** We have quantified duration of up-states by computing durations of continuous segments above a given voltage threshold. An example trace from neuron 9 shows up-state detection for four thresholds: 10 mV, 15 mV, 20 % max, and 40 % max. Thin horizontal grey line shows the value of the resting membrane potential for this trace. Gaps shorter than 20 ms were included in up-states (arrow). **(B)** Distribution of up-state durations detected in all traces from neuron 9 for two thresholds. Most up-states were brief (bumps), with occasional bumps longer than 200 ms. **(C)** Each panel shows distributions of up-state durations (black dots) for all 20 neurons. Large full grey circles show mean bump durations, and large empty circles show median bump durations for each neuron. Note the logarithmic scale on y-axis. *Left* column panels show up-state (bump) durations for absolute thresholds: 10, and 15 mV above resting membrane potential. *Right* column panels show up-state (bump) durations for relative thresholds: 20 %, and 40 % of maximum potential amplitude (see Sec. 2.6.3 on page 24 for details).

2005; Stern et al., 1997). For example, using the method from Stern et al. (1997), up-state threshold for the bottom trace in Fig. 5.3B would be 12 mV above resting potential. Thus up-state (bump) durations computed here are likely to be an *overestimate* compared to other studies.

The membrane potential spent most of the time near its resting level (i.e. below bump threshold). Bumps (up-states) covered only about 12.5±10 % of the recorded traces across all neurons, and appeared with an average "bump frequency" of 2.6±1.4 bump/s (10 mV threshold; 21.5±14 %, and 3.4±1.6 bump/s for 20 % threshold).

Up-states in awake auditory cortex were usually brief (Fig. 5.4B), despite the fact that their durations might have been overestimated. For each neuron the typical bump duration was in the order of tens of milliseconds, with isolated bumps being several hundred milliseconds long. Even for the lowest thresholds the mean bump duration across neurons was less than 60 ms (46±56 ms, median = 30 ms for 10 mV threshold, and 60±77 ms, median = 37.5 ms for 20 % threshold; Fig. 5.4C top). Not surprisingly, bump durations decreased and fewer bumps were detected for the higher thresholds (29±39 ms, median = 18 ms for 15 mV threshold, and 28±39 ms, median = 18 ms for 40 % threshold; Fig. 5.4C bottom).

5.4   Long up-states are rare

For some neurons, bump durations were considerable and reached hundreds of milliseconds, around 1–1.5 s in a few cases, resembling the "canonical" up-states observed in other cortical areas. However, these long up-states were rare, with only 2 % of all bumps longer than 200 ms (188 out of 9245 bumps, 10 mV threshold; 4 % for 20 % threshold) (Fig. 5.5A). Individual neurons displayed a range of fractions of short, intermediate bumps, and long up-states (Fig. 5.5C), with most neurons containing more than 70 % of short bumps (< 100 ms). The low percentage of long up-states could be artificially low due to the high number of very short bumps, and long up-states could still occupy a large portion of the depolarized membrane potential state simply thanks to being longer. We therefore computed the proportion of up-state time (i.e. away from resting membrane potential) neurons spent in bumps of different durations: short bumps (< 100 ms), intermediate up-states (100–200 ms), and long up-states (> 200 ms).

Neurons typically spent only 3.7 % of *total recording time* in up-states longer than 200 ms (20 % threshold, 0.6 % for 10 mV threshold). Most of *up-state time* was spent in brief bumps (Fig. 5.5B),

despite long up-states usually lasting several-fold longer than the brief bumps. Long up-states ($>200\,\text{ms}$) covered, on average, around 9 % or 17 % of up-state time, even for the lowest thresholds. Individual neurons displayed a "continuum" of up-state times spent in bumps of different durations (Fig. 5.5D). This spectrum of up-state proportions was evident especially for the lowest threshold (20 % max). Neurons ranged from "canonical" bumpy neurons with up-state time devoted solely to brief bumps, to neurons spending more of their depolarizations in long up-states, which still, however, formed the minority of all bumps (compare corresponding panels in Fig. 5.5C and D).

In summary, subthreshold dynamics of most recorded traces was consistent with bumpy activity (Fig. 5.3), with few traces displaying behavior consistent with longer up-states. Neurons spent most of the time in down-states (i.e. close to the resting potential), with mostly brief bumps, and occasionally bumps were longer than 200 ms (Figs. 5.4–5.5). Even neurons spending more up-state time in long up-states still contained the majority of brief and intermediate bumps, consistent with rapidly changing patterns of membrane potential fluctuations, rather than the canonical long up-states lasting up to several seconds.

5.5 Discussion

We have recorded intracellular activity of neurons in auditory cortex of unanesthetized head-fixed rats. The recordings were sufficiently stable and permitted us to record subthreshold activity for minutes. Neurons spent most of their time near the resting membrane potential, and their subthreshold dynamics consisted of usually brief, occasionally long up-states. Most neurons displayed brief infrequent bumps of subthreshold activity, and occasionally neurons displayed rapid fluctuations of membrane potential with a few long up-states.

5.5.1 Recording stability

The apparent stability of whole-cell patch-clamp recordings even in awake animals is probably determined by mechanical stability of the gigaohm seal (Hamill et al., 1981; Margrie et al., 2002). We have not observed any major changes in the recorded signal due to brain, cardiac or pulmonary pulsations. Even relatively major movements (licking, grooming, whisking, etc.) did not cause any apparent dis-
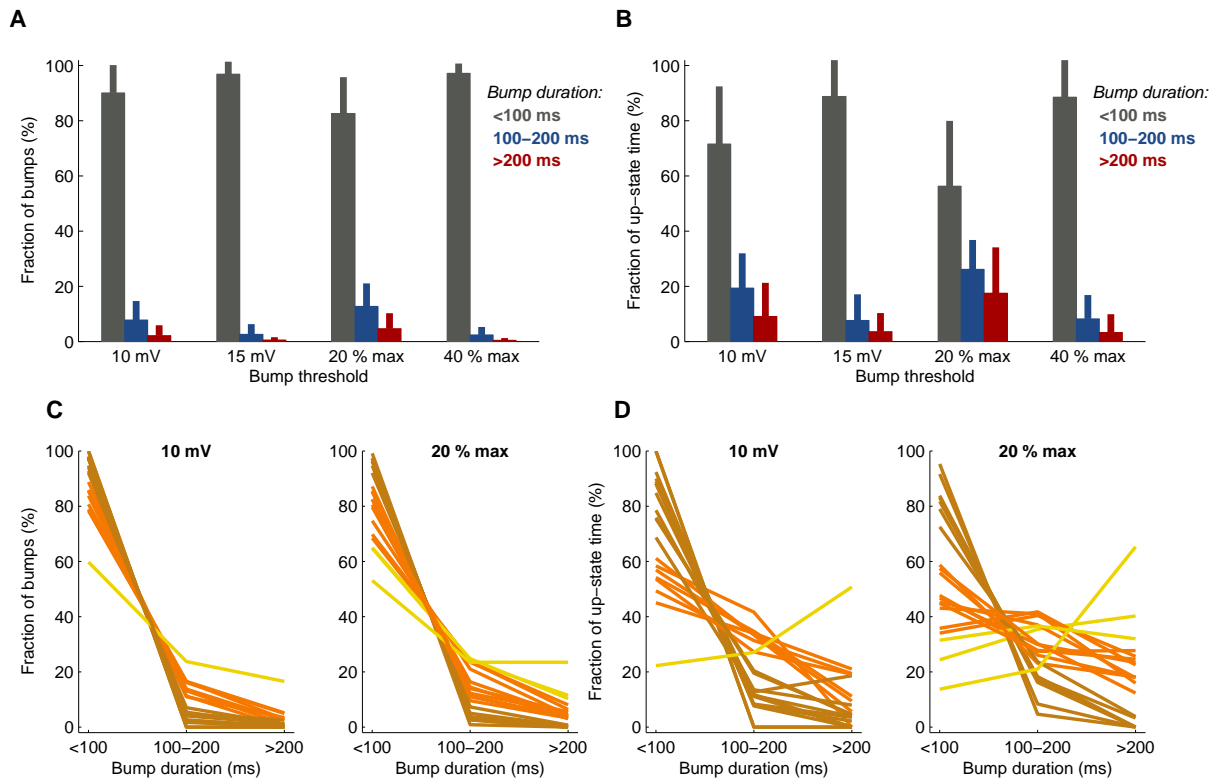
Figure 5.5: Long up-states were rare. Because bump durations can vary depending on the threshold (see Fig. 5.4), bumps were divided into three groups based on bump duration: <100 ms, $100–200$ ms, and >200 ms. **(A)** Fraction of long bumps (>200 mV) was very low for each threshold. Even for the lowest thresholds the fraction of long bumps was around 2 % (10 mV threshold), or 4 % (20 % max). **(B)** The average fraction of time spent in up-state was low for each threshold. Even for the lowest thresholds (10 mV, 20 % max) long bumps (>200 mV) covered only around 9 % (17 %) of up-state time, on average, despite being several fold longer than most of the short bumps. Note that neurons spent most of their time in down-state (see text for details), which is not included in the figure. Note also that our method of bump detection likely overestimated bump durations (see text for details). Error bars show standard deviations. **(C)** Fractions of short, intermediate bumps, and long up-states varied among neurons (left 10 mV, right 20 % max threshold). **(D)** Fractions of up-state time varied among neurons (left 10 mV, right 20 % max threshold).

In panels (C), and (D) neurons were colored based on what fraction of up-state time they spent in *short bumps* (<100 ms): **> 66 %**, **33–66 %**, and **< 33 %**.

ruptions in the recordings. The recording times we observed agreed with recording times mentioned in other studies of intracellular activity in awake animals (Covey et al., 1996; Margrie et al., 2002; Steriade et al., 2001). Whole-cell recordings in awake animals could then be used to further elucidate synaptic mechanisms underlying sensory perception (excitation vs. inhibition) in awake animals, and even to record intracellular activity in animals performing well-defined behavior.

5.5.2 Intracellular activity in auditory cortex

Several studies have focused on intracellular activity of neurons in auditory cortex so far. Studies conducted in auditory cortex of *anesthetized* animals (DeWeese and Zador, 2006; Las et al., 2005; Ojima and Murakami, 2002; Tan et al., 2004; Wehr and Zador, 2003, 2005; Zhang et al., 2003) all support the general impression that subthreshold activity in auditory cortex consists mainly of brief and infrequent voltage excursions, as either stated explicitly in the text, or inferred from the data presented.

The intracellular activity in auditory cortex of *awake* animals has been studied only rarely (Ribaupierre et al., 1972; Serkov and Volkov, 1984; Volkov and Galaziuk, 1985, 1989). These studies were conducted in awake *immobilized* cats using sharp electrodes. Although acutely immobilized animals represent very non-physiological condition, results from these studies also demonstrate brief bumps in subthreshold dynamics with occasional longer "up-state" (see for example Fig. 3 in Volkov and Galaziuk, 1989).

5.5.3 Up and down states in the cerebral cortex

Subthreshold dynamics in various parts of the brain has often been described in terms of a bistable membrane potential, spontaneously fluctuating between two subthreshold values, so-called *up* and *down* states. During depolarized up-states neurons are more active and can fire action potentials, whereas during hyperpolarized down-states neurons sit quietly at their resting potential (but see Petersen et al., 2003a). Neurons toggle their membrane potential between up and down states, with very little time spent in between. Time scales of canonical up and down states are considerably longer than the duration of, for example, action potentials with up-states lasting several seconds. Up and down states have been described and characterized for several decades (Cowan and Wilson, 1994; Steriade et al., 1993a,b; Wilson and Groves, 1981). More recently, the bistability of the membrane potential has been described

in-vitro (Cossart et al., 2003; Sanchez-Vives and McCormick, 2000), as well as in-vivo in frontal cortical areas (Lewis and O'Donnell, 2000; Léger et al., 2005), somatosensory (Petersen et al., 2003a; Sachdev et al., 2004; Steriade et al., 2001; Timofeev et al., 2001), visual (Anderson et al., 2000; Lampl et al., 1999), olfactory (Luo and Katz, 2001; Margrie and Schaefer, 2003), other areas (Paré et al., 1998), and striatum (Kasanetz et al., 2002; Mahon et al., 2003; Stern et al., 1997; Wilson and Kawaguchi, 1996).

Up and down states have not been previously described in auditory cortex (DeWeese and Zador, 2006). Our data presented above are consistent with a predominance of brief voltage excursions, with only occasional long up-states. The mean up-state duration we detected was 60 ms using a low detection threshold. Average up-state durations reported in studies using rigorous criteria for state detection are in the 300–500 ms range (Lewis and O'Donnell, 2000; Stern et al., 1997). It is unclear why auditory cortex would be different in terms of subthreshold dynamics, although it is an intriguing possibility that there is an inherent difference between auditory and other sensory areas. Recording differences might contribute to the lack of up and down states in auditory cortex; indeed even in somatosensory cortex the two states are not always apparent (Bruno and Sakmann, 2006; Wilent and Contreras, 2005).

A more plausible explanation is that the lack of up-down states in auditory cortex is simply caused by differences in nomenclature, and bumps observed in auditory cortex are simply brief up-states (DeWeese and Zador, 2006). We have presented evidence supporting a wide range of up-state durations, with occasional long up-states (> 200 ms). Neurons spent various times in up-states of different durations, depending on the voltage threshold used for up-state detection. Data presented in several studies (Anderson et al., 2000; Jagadeesh et al., 1992) support the view of a wide range of up-state durations, with only about 13–60 % of cells identified as two-state based on bimodal voltage histograms in (Anderson et al., 2000).

Another intriguing possibility is that up and down states might represent very specific activity occurring during sleep or anesthesia, and as such might be dependent on the state of vigilance (Mahon et al., 2003). Experiments in anesthetized animals designed to resemble waking state, such as stimulation of subcortical inputs (Lewis and O'Donnell, 2000), cortical desynchronization (Kasanetz et al., 2002); and direct experiments in awake animals: transition from sleeping to waking state (Steriade et al., 2001), transition from quiet wakefulness to whisking (Crochet and Petersen, 2006), all caused disruption of up-down states with cortical activity usually switching to up-state. We have not explicitly

controlled the behavioral state of animals in our experiments, and although we are confident that our rats were not sleeping during recordings (based on video monitoring of animal activity), it is possible that the longest up-states we observed might have occurred during different behavioral/attentional states. Further experiments recording intracellular activity in awake animals during well controlled behavior would be needed to elucidate the origin and function of up-down states.

Chapter 6

Sound discrimination in freely moving and head-fixed rats

The transformation of an external sensory stimulus into an internal percept is thought to require several (but probably many) computational steps, most of which are arguably unknown at present. One of the hard computational problems in auditory neuroscience is that of auditory scene analysis (cocktail-party problem), in which we are able to extract one auditory stream from a (noisy) background and perceive it as a distinct auditory object. A fundamental feature of auditory scene analysis is auditory stream segregation, a phenomenon in which interleaved sequences of sounds are decomposed into separate perceptual streams (Bregman, 1990). The ability of a listener to segregate auditory streams is affected by, and possibly dependent on, attention (Cusack et al., 2004).

The primary auditory cortex has been traditionally viewed as a primary sensory area, i.e. a unimodal area primarily involved in the processing of sounds. However, since the early observations (Hubel et al., 1959), multiple studies (Beaton and Miller, 1975; Brosch et al., 2005; Fu et al., 2003; Hocherman et al., 1976; Miller et al., 1972) have suggested that sound-evoked responses in the (primary) auditory cortex are influenced by the behavioral context in which the sounds are presented. This view is further supported by the reported heterogeneity and sparseness of responses to various sound stimuli in auditory cortex of *naïve* awake animals (Chimoto et al., 2002; Recanzone, 2000, see also Sec. 3).

The attentional and behavioral modulations of single neuron activity indicate the involvement of primary auditory cortex in rather elaborate computations, meaning that the neurons in the primary

auditory cortex might not be just simple "feature detectors" , but instead might represent the auditory world in terms of auditory "objects" (Nelken et al., 2003). The primary auditory cortex thus seems to be a good candidate for at least the first step in auditory stream segregation; see (Fishman et al., 2001; Micheyl et al., 2005) for some experimental evidence.

We have developed an experimental setup for studying behavior and neural correlates of behavior in head-fixed rats, with the ultimate goal of studying auditory stream segregation in single neurons. We focused on auditory attention as one particular aspect of auditory stream segregation and—as a first step—we have developed a sound discrimination go/no-go task, in which head-fixed rats listened to a continuous auditory stream and reported which sound they heard. Here we present our *full* account on training rats for the sound discriminations task, and provide comparison of various response parameters in restrained and unrestrained rats. We also present evidence documenting nonauditory modulation of responses of single neurons in rats performing the task.

6.1 Sound discrimination go/no-go task

We have trained 33 Long-Evans rats in a sound discrimination go/no-go task (Fig. 6.1, see also Sec. 2.8 on page 33). Water-deprived rats listened to a continuous stream of (randomly interleaved) target and distractor 500 ms sounds, separated by silent random length inter-trial intervals. The rats were trained to lick in response to target sounds, and all correct responses were rewarded with a drop of water. All other licks, i.e. licks in response to distractor sounds, or licks during inter-trial interval were penalized by air-puff(s), and lengthening of the penalty period (inter-trial interval). We have used pure tones and FM-tones (warbles) as sound stimuli. Individual rats were either trained with tones as targets (warbles as distractors), or warbles as targets (tones as distractors), see Sec. 2.5.4 on page 20 for details. We have not observed any differences between these two types of stimuli, and in the following we will only refer to target and distractor stimuli. Initially we have used 3 s "reward period" following the stimulus in distractor trials (Fig. 6.1), during which licks were penalized. In a later variant of the task, distractor stimuli were followed directly by a penalty period, i.e. distractor trials were, on average, 3 s shorter than target trials (compare Fig. 6.1 with Fig. 2.9 on page 35). We have not observed any difference between these two types of tasks, and in the following we refer to both of these variants simply as sound
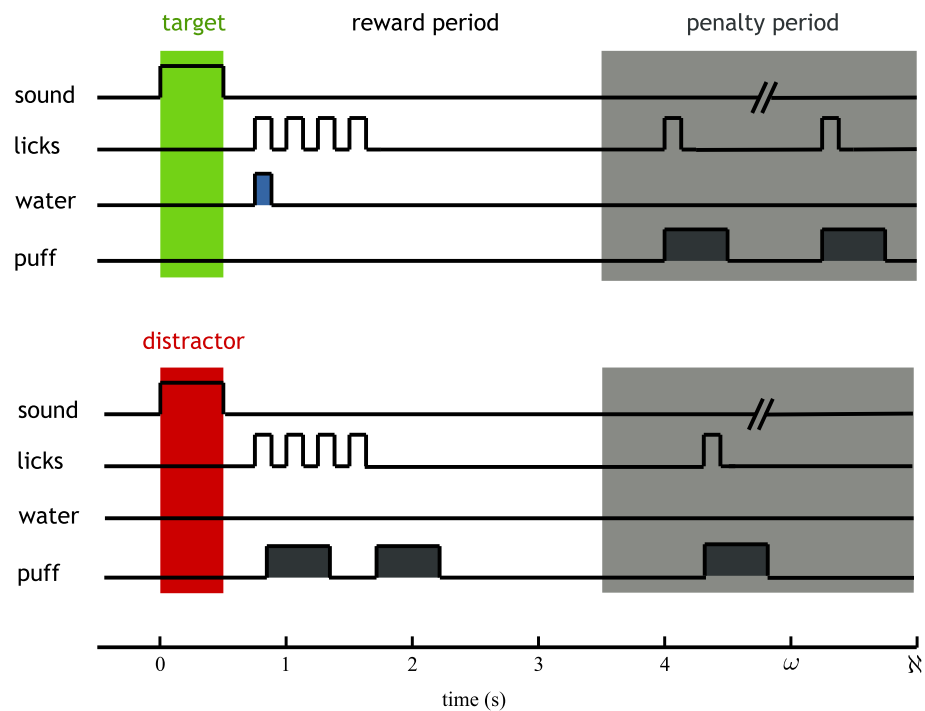
target   reward period   penalty period

sound

licks

water

puff

distractor

sound

licks

water

puff

0   1   2   3   4   ω   ℵ

time (s)

Figure 6.1: Trial structures of soud-discrimination go/no-go task. Target trials (top) started with 500 ms target sound, followed by a 3 s reward period, during which responses (i.e. licks) were rewarded with a drop of water. The reward period was followed by a (random duration) penalty period, during which licks were penalized by air-puffs. Distractor trials (bottom) started with a 500 ms long distractor sound, followed by a "reward period," followed by a penalty period. Any lick during the distractor trial was penalized by an air-puff. We have also used a variant of this task in which distractor sounds were immediately followed by a penalty period, see Fig. 2.9 on page 35. Target and distractor trials were randomly interleaved during a behavioral session.

discrimination go/no-go task. Rats trained for the initial variant of the task were used as examples for figures presented in this text.

Rats were initially trained freely moving for the full task, and retrained in head-fixed position afterwards (see Sec. 2.8.3 on page 38 for details). During the freely moving phase, rats were trained inside a plastic tube (closed on both ends) which provided only a very loose body restraint. Training started with a few *classical conditioning* sessions, during which only target sounds were presented and rats were trained to associate the target sounds with water rewards. Classical conditioning sessions were followed by several *operant conditioning* sessions, i.e. rats had to initiate water delivery by licking. During the operant sessions we gradually introduced distractor stimuli. Finally, rats were implanted with a headpost and recording well, and—after at least 24 hours of recovery—were retrained in the head-fixed position. Twenty-one rats (64 % out of 33 rats) were retrained for at least one head-fixed

training session. We have also recorded the activity of 53 neurons in 10 head-fixed trained animals (30 %, see below).

6.2 Rats perform well when freely moving as well as when head-fixed

A well-trained rat was required to *discriminate* target and distractor stimuli, i.e. to lick in response to target stimuli and not distractor stimuli; as well as to *follow the task structure*, i.e. not to lick during inter-trial intervals.

We have defined *discrimination performance* as a measure of the rat's ability to discriminate between target and distractor stimuli (Fig. 6.2A, see also Eq. 2.7 on page 36, and Sec. 2.8 for details). Correct responses—licks in target trials and no licks in distractor trials—were assigned 1, and incorrect responses—no licks in target and licks in distractor trials—were assigned 0. Discrimination performance was then computed as a running average of correct/incorrect responses in 20 consecutive trials.

The rat's ability to "understand" task structure, i.e. not to lick during penalty period, was quantified using a *penalty period extension* (Fig. 6.2B, Eq. 2.8 on page 36). Recall that whenever a rat licked during the penalty period, the period was restarted with the same parameters again (in addition to an air-puff penalty). Thus, rats that did not follow the task structure would have the actual penalty periods longer than intended, whereas penalty period duration for rats following the task structure would equal the intended duration. Penalty period extensions was defined as ratio of actual to intended period duration, and was computed as a running average in 20 consecutive trials.

Performance in the sound discrimination task was defined as the ratio of discrimination performance to penalty period extension (Fig. 6.2C, Eq. 2.9 on page 36). This measure then included both the ability to discriminate sound stimuli, and the ability to follow task structure.

For most operant freely moving training sessions we used randomly interleaved target and distractor stimuli, which appeared with equal probability (Sec. 2.8.3 on page 38). Chance *discrimination performance* would then be equal to 0.5. Analogously, *discrimination performance* for a rat licking in response to each stimulus, as well as for a rat not responding at all, would be equal to 0.5. The final *performance*, however, could be lower than 0.5, even for rats with very good *discrimination performance*,
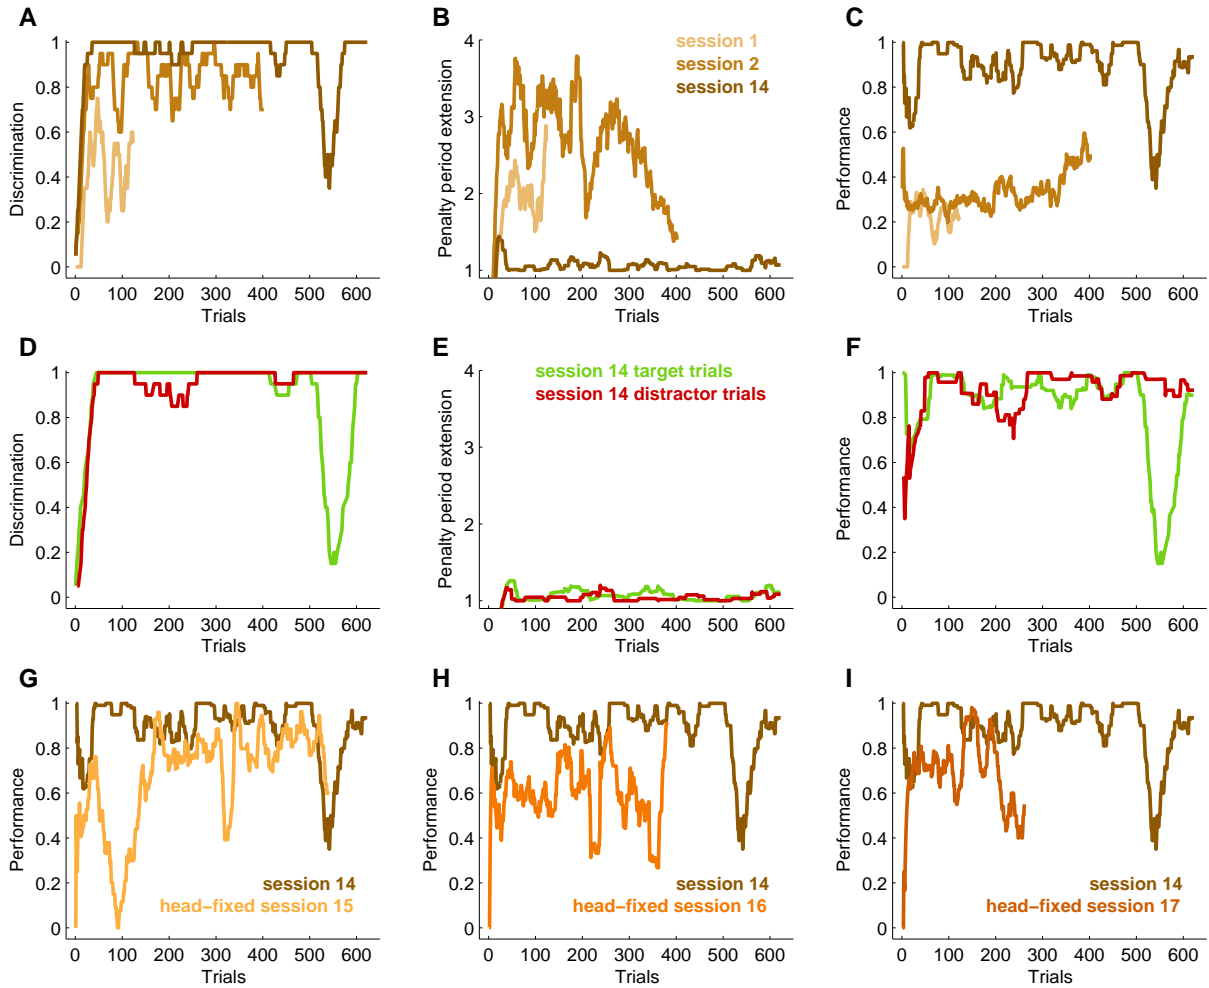
Figure 6.2: Rats performed well when freely moving as well as when head-fixed. This figure shows performance parameters during several behavioral sessions for one rat. The rat performed well during freely moving sessions (top row), both in target and distractor trials (middle row), and also during head-fixed sessions (bottom row). **(A)** Discrimination performance in session 1, session 2, and session 14. Higher value means better discrimination performance. **(B)** Penalty period extension during the same sessions as in panel (A). Lower value (i.e. penalty period closer to its intended duration) means better "understanding" of task structure. **(C)** Performance in sound discrimination task was defined as ratio of discrimination performance to penalty period extension. **(D)** Discrimination performance in target trials and distractor trials during session 14. Target and distractor trials appeared with equal probability and were randomly interleaved. **(E)** Penalty period extension, and **(F)** performance (ratio of discrimination performance to penalty period extension) for target trials and distractor trials during session 14. **(G)** Discrimination performance in (freely moving) session 14 compared with performance in (following) head-fixed session 15. **(H)**, **(I)** Same format as in (G) for head-fixed session 16, and session 17.

All traces in this figure are plotted as running averages of 20 consecutive trials.

because performance also included *penalty period extension*. Inability of a rat to follow task structure well would then lead to higher *penalty period extension* and lower *performance*.

The performance of freely moving rats improved during training (Fig. 6.2 top row). During the initial classical conditioning sessions (session 1, session 2 in the example) rats usually licked actively as documented by the relatively high discrimination performance. Note however, that during the initial sessions rats did not follow the task structure well as shown in Fig. 6.2B. Thus despite the very good discrimination performance in session 2, the actual performance was much lower than in session 14 (Fig. 6.2C), which was the last (operant) freely-moving session before implantation. The top row in Fig. 6.2 documents the improvement in performance in subsequent freely moving training sessions (the discrimination performance is higher and the penalty period is lower).

Rats performed well in both target and distractor trials (Fig. 6.2 middle row). In sessions in which target and distractor trials were randomly interleaved, rats showed good discrimination performance (Fig. 6.2D), with low penalty period extension (Fig. 6.2E), which led to very good overall performance for both types of trials (Fig. 6.2F).

The performance of head-fixed rats was comparable with their previous performance in freely moving sessions (Fig. 6.2 bottom row). Note that the rat started performing quickly in the very first head-fixed retraining session (session 15), and the performance reached the level achieved during the freely moving training (session 14).

Overall, the *performance* of freely moving rats was similar to the performance of head-fixed rats (Fig. 6.3A, $n = 21$ rats, $p = 0.05$, Wilcoxon rank-sum test). Note that the performance measure could be lower than 0.5 because it also included the penalty period extension (Sec. 2.8), and rats that did not follow the task structure reached lower performance values. Performance values were computed as the 75th percentile across operant freely moving, or head-fixed sessions. *Discrimination performance*, however, was better in freely moving animals (Fig. 6.3B, $n = 21$ rats, $p < 0.01$, Wilcoxon rank-sum test). For eighteen rats, however, the 75th percentile of discrimination was > 60 % in the head-fixed session.
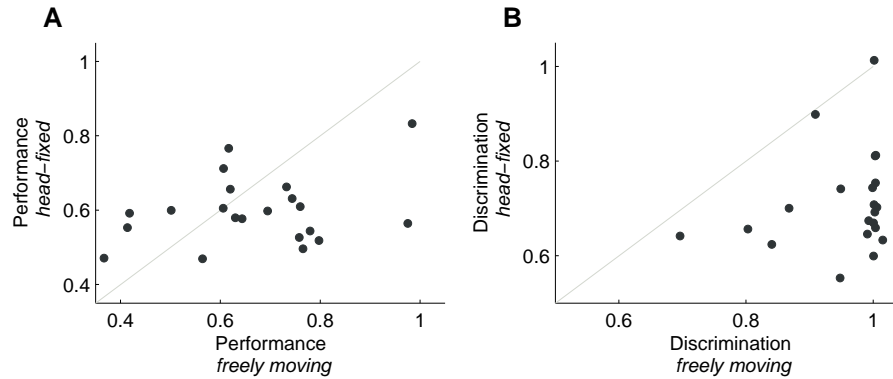
Figure 6.3: Comparison of performance of freely moving and head-fixed rats. **(A)** *Performance* of freely moving and head-fixed rats was similar. Note that the performance measure also includes penalty period extension. **(B)** *Discrimination performance* tended to be better in freely moving rats. All values were computed as the 75th percentile across operant freely moving sessions, or head-fixed sessions. All points in this figure were jittered slightly so that overlying points could be seen.

6.3 Response parameters are similar in freely moving and head-fixed rats

Rats licked actively during training sessions (Fig. 6.4). Licking was usually less prominent during the first training sessions, but as the training progressed, rats would lick more during target trials, and less during distractor trials (Fig. 6.4A). Well-trained rats also licked only occasionally during penalty periods (Fig. 6.4B).

The ability of rats to learn to lick in response to targets is further documented in the bottom of Fig. 6.4. Licking rasters ("lasters") show positions of individual licks in target and distractor trials during three different training sessions for one rat. The rat improved its licking from session 1 to session 14 with many more correct trials, and stimulus locked lick trains. Furthermore, in head-fixed session 17, the lick trains became more stereotypical, with improved response latency. Notice the oscillations in the laster PSTH shown at the bottom of Fig. 6.4E.

Because licking was the only behavioral response of rats in the sound discrimination task (especially in the head-fixed position), we wondered whether head fixation led to any major changes in basic response parameters: durations of individual licks, response latency, and frequency of licking (Fig. 6.5).

Lick durations tended to be shorter in head-fixed rats (Fig. 6.5 top). A typical distribution of lick durations pooled across all operant freely moving and head-fixed sessions for one rat is shown in Fig. 6.5A. Note that we only considered licks shorter than 200 ms for this analysis, as longer licks were
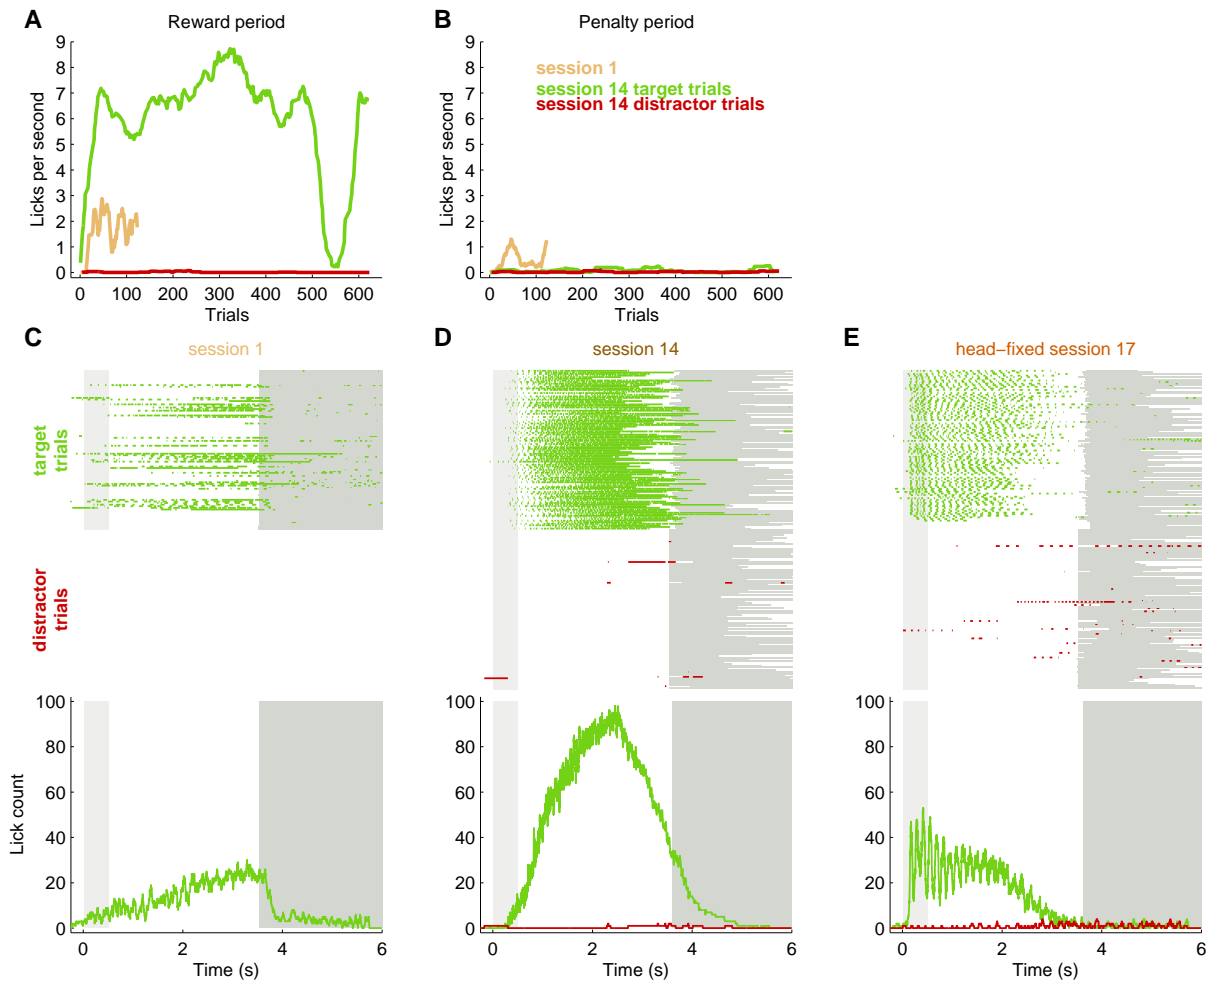
Figure 6.4: Well-trained rats licked more during reward period in target trials, and less during penalty period (top). Licking responses were more stereotypical for well trained rats (bottom). Here we show basic licking characteristics for the same rat as in Fig. 6.2. (**A**) Number of licks per second during reward period. Only target sounds were presented during session 1. Both target and distractor sounds were presented during session 14. (**B**) Same format as in (A) for penalty period. All traces are plotted as running averages of 20 consecutive trials. (**C**) Lick rasters ("lasters") of responses during several training sessions. Top: Licking responses during target trials in session 1 (same rat as in Fig. 6.2). Each colored tick denotes a single lick. Bottom: Time histogram of the rasters above. Note that only target trials were presented in session 1. (**D**) Same format as in (C) for session 14. Target trials and distractor trials were interleaved with equal probability. (**E**) Same format as in (C) for head-fixed sesion 17 (third head-fixed session). Only the first 100 trials from each session (and stimulus type) are plotted in panels (C), (D), and (E).

rare, and usually corresponded to interruptions of an infrared beam by various body parts of a rat. Typical lick durations in operant freely moving sessions—as estimated by medians of lick distributions—were around 40–50 ms (n = 33 rats, Fig. 6.5B). A comparison of lick durations detected in freely moving and head-fixed sessions (n = 21 rats) showed that, on average, head-fixed rats tended to have slightly shorter licks. The distributions of lick durations were narrower then distributions of their counterparts in freely moving sessions (notice the generally smaller standard deviations in Fig. 6.5C). For most rats (n = 15, black circles in Fig. 6.5C) the distributions of lick durations were significantly different (p < 0.01, Wilcoxon rank-sum test) between freely moving and head-fixed sessions.

Response latency, as estimated by first lick latency, tended to improve in the head-fixed position (Fig. 6.5 middle). To estimate response latency we computed first lick latencies during the last three operant freely moving sessions, and all head-fixed sessions. We have only considered first lick latencies longer than 10 ms and shorter than 1 s for this analysis. Typical first lick latencies were around 500 ms in freely moving (n = 33 rats, Fig. 6.5D), as well as head-fixed sessions (n = 21 rats, Fig. 6.5E). A comparison of response latencies computed in freely moving and head-fixed sessions showed that, on average, response latencies were similar, and tended to be shorter in head-fixed sessions. Distributions of first lick latencies were *not* significantly different (empty circles in Fig. 6.5F) for most rats (n = 13 rats, p < 0.01, Wilcoxon rank-sum test).

Rats licked regularly, in theta frequency, both in freely moving and head-fixed sessions (Fig. 6.5 bottom). The oscillations in laster PSTH in Fig. 6.4E suggested an apparent rhythmicity of licking behavior. This rhythmicity was evident in the power spectra of laster PSTHs from different training sessions (Fig. 6.5G). To compute the power spectra we first aligned all lick trains for a given session with respect to the first lick, added the lick trains together to obtain aligned laster PSTH, and then computed the power spectrum density of the PSTH. The peak of the power spectrum between 3–40 Hz then defined the main licking frequency. Typical main licking frequency in freely moving sessions was around 7 Hz (n = 33 rats, Fig. 6.5H), inside the theta frequency band (4–12 Hz). Main licking frequencies in this analysis were estimated as the peaks of power spectra computed from aligned laster PSTHs pooled across all operant freely moving sessions, and all head-fixed sessions. Comparison of licking frequencies between freely moving and head-fixed sessions (n = 21 rats, Fig. 6.4I) showed that head-fixed rats tended to lick at slightly lower frequencies. Licking frequencies in freely moving and head-
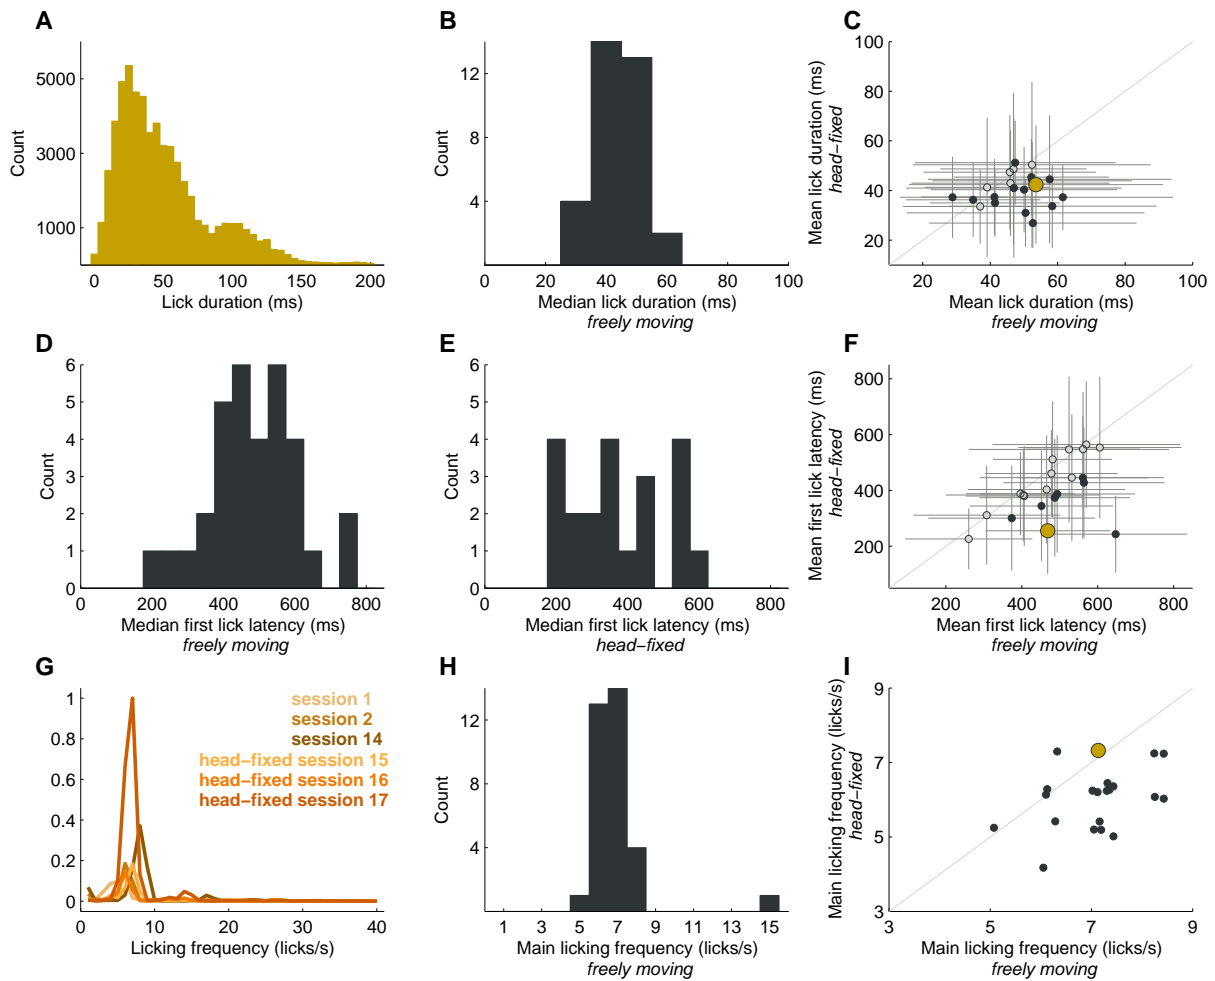
Figure 6.5: Lick parameters were comparable for freely moving and head-fixed rats. Durations of individual licks (top row), first lick latency (middle row), and licking frequency (bottom row) were similar in freely moving and head-fixed sessions. **(A)** Distribution of lick durations for the rat from Figs. 6.2–6.4. Licks from all *operant* behavioral sessions (8 freely moving, 3 head-fixed) were pooled together. **(B)** Distribution of *median* lick durations for all rats (n = 33). **(C)** *Mean* durations of licks identified during freely moving and head-fixed sessions were similar (n = 21 rats). **(D)** Distribution of *median* first lick latencies in *freely moving* sessions (n = 33 rats). **(E)** Distribution of *median* first lick latencies in *head-fixed* sessions (n = 21 rats). **(F)** *Mean* first lick latencies in freely moving and head-fixed sessions were similar (n = 21 rats). **(G)** Frequency spectra of aligned laster PSTHs (see text) for rat from Figs. 6.2–6.4. The plotted spectra were normalized with respect to the peak of session 17. **(H)** Distribution of main licking frequencies in freely moving sessions (n = 33 rats, see text). **(I)** Rats licked at theta frequency both in freely moving and head-fixed sessions (n = 21 rats). All points in this panel were jittered slightly so that overlying points could be seen.

Lick parameters for *freely moving sessions*, and *head-fixed sessions* were computed from licks pooled together from all *operant* freely moving, and head-fixed sessions, respectively. The larger circles in panels (C), (F), (I) show values computed for rat shown in panel (A). Full circles in (C), (F) show animals with significantly different values (see text). Error bars show standard deviations.

fixed sessions were different (p < 0.001, Wilcoxon rank-sum test), with head-fixed licking frequencies slightly lower, but main licking frequencies remained in theta range (4–12 Hz).

Head-fixed rats tended to lick with shorter licks, and licked at slightly lower frequencies. However, head-fixed rats tended to have shorter first lick latencies, i.e. shorter reaction time, when compared to freely moving sessions.

6.4 Lick-locked (nonauditory) responses in auditory cortex

We recorded from 53 neurons (in 10 animals) during the sound discrimination task. During 18 (34 %) of these recordings, however, the animal failed to perform the task. From the remaining 35 recordings we identified 15 (28 %) recordings (in 7 animals) with at least 40 total recorded trials (median = 52 trials per recording).

Neurons displayed a range of responses during the sound discrimination task (Fig. 6.6, see also Appendix B for a complete set of neurons). In accordance with our findings in naïve head-fixed rats (Sec. 3), some neurons showed suppression of their activity (Fig. 6.6A), some showed an increase in their activity (Fig. 6.6B).

The neuron shown in Fig. 6.6A displayed marked suppression of its activity. The suppression had very short latency and relatively long duration (~2 s), and was restricted to trials during which the rat licked in response to the stimulus, i.e. correct target trials and error distractor trials. The remarkable similarity between responses in trials containing different acoustic stimuli (tones vs. warbles), but similar behavioral responses is further documented in Fig. 6.6C. In the figure, spiking responses are plotted relative to lick trains aligned with respect to the first lick. In both cases neuronal activity gradually returned to its previous level about 2 s after the rat *started* licking. This suppression of neuronal activity was not limited to lick trains which occurred as responses, as shown by similar suppression of activity during lick trains which started > 4 s after a stimulus (other). Note, that although there was no sound stimulus presented in the case of other lick trains, there was 500 ms air-puff which started simultaneously with first licks in other and error distractor trials. Because air-puffs sounded like quiet white-noise stimuli we cannot conclude that the suppression of activity was *solely* due to licking. However, the fact that neuronal activity was still suppressed at least 500 ms after any possible acoustic stimulus was over, and that the suppression of neuronal activity appeared to be locked to the start of lick trains regardless
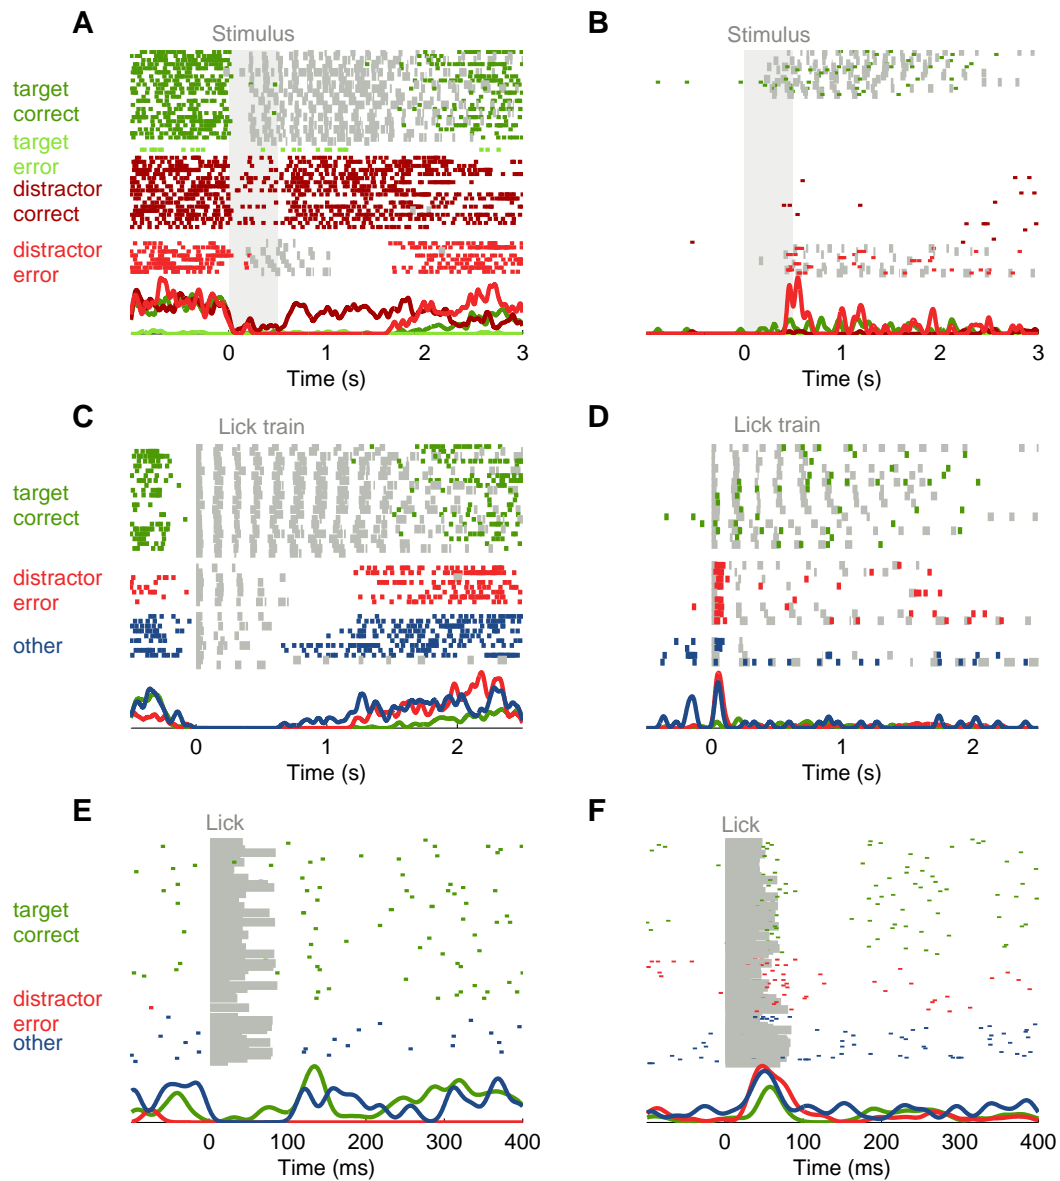
Figure 6.6: Neuronal responses in correct and error trials during a sound discrimination go/no-go task. **(A)** Neurons displayed a range of responses in sound discrimination task. The raster plot shows spiking responses of a single neuron. Colored dots represent individual spikes, light grey shaded region indicates stimulus duration, grey ticks indicate individual licks. Trials are sorted into four groups depending on the stimulus and outcome of each trial. **(C)** Spiking responses during lick trains. Lick trains were aligned with respect to the first lick. Other responses include lick trains which occurred during the penalty period (>4 s after the start of the trial). **(E)** Some neurons displayed lick-locked spiking responses. All licks surrounded by spikes were aligned with respect to the start of each lick. Firing rate curves beneath the rasters were computed by summing the spikes in 1 ms bins, convolving the resulting PSTH with a Gaussian ($\sigma = 20$ ms in A, C, $\sigma = 10$ ms in E), and then normalizing to the maximum peak of all firing rate curves. **(B)**, **(D)**, and **(F)** Same format as in panels (A), (C), and (E) for another neuron. See Appendix B for a complete set of neurons.

of their occurrence, suggested that licking might play a role in modifying neuronal activity in primary auditory cortex. Suppression of spiking activity for this neuron appeared to be lick-locked (Fig. 6.6E). Note, that in case of individual licks there was no acoustic stimulus present.

We have observed similar responses in another neuron (Fig. 6.6B). In this case, licking activity seemed to "evoke" spiking activity, as spikes were almost exclusively observed in correct target and error distractor trials. Spiking activity seemed to be locked to the beginning of lick trains (Fig. 6.6D), but note the spike trains locked to first licks in other and error distractor trials, which were evoked by air-puffs. However, spiking responses aligned to individual licks (Fig. 6.6F) showed that spikes were lick-evoked even in correct target trials (see the intimate relationship between spikes and licks in Fig. 6.6D).

Because correct and error target trials contained exactly the same stimulus, but differed in their behavioral outcomes, changes in neuronal activity during these two types of trials could be considered correlates of changes in the behavioral state of the animal. We have evaluated the difference in spiking responses in correct and error target trials by computing choice probability between spike counts during the first 100 ms of target stimulus, and then used bootstraping to estimate the significance of choice probability values. We have not observed any significant choice probability between responses in correct and error target trials with choice probabilities of 0.51 ± 0.05 (n = 14 neurons). We were only able to compute choice probabilities during the first 100 ms of stimulus, because all later neuronal activity was likely affected ("contaminated") by licking (see above).

6.5 Discussion

We have developed a behavioral setup for studying behavior in head-fixed rats. Using a sound discrimination go/no-go task, we have showed and compared response parameters of unrestrained and head-fixed rats performing the task. Head-fixed rats' performance and response parameters were similar to the parameters measured in the unrestrained condition. Head-fixed rats tended to respond with shorter licks, and licked at lower frequencies than in preceding freely moving sessions. The response latency, however, tended to improve in the head-fixed position compared to freely moving sessions. We have also documented examples of sensorimotor (licking) activity modulating responses of single neurons in the primary auditory cortex of behaving head-fixed rats.

6.5.1   Head-fixed behavior

Experiments using head-fixed (head-restrained) *primates* have a long tradition in auditory research (Beaton and Miller, 1975; Durif et al., 2003; Hocherman et al., 1976; Miller et al., 1972; Populin, 2006). To avoid some disadvantages of primate studies (namely long training time, monkey "resistance" to auditory tasks, and high cost), many groups have chosen other experimental animals and studied neuronal activity in head-fixed animals during different auditory tasks, and in different cortical areas. Tone detection and discrimination have been studied in the primary auditory cortex of ferrets (Fritz et al., 2005a, 2003, 2005b), guinea pigs (Edeline et al., 1993; Edeline and Weinberger, 1993), gerbils (Ohl and Scheich, 1997), and in the lateral hypothalamus (Nakamura and Ono, 1986; Nakamura et al., 1987; Ono and Nakamura, 1985; Ono et al., 1986, 1985), and posterior parietal cortex of rats (Nakamura, 1999); sound localization has been studied in cats (Tollin et al., 2005).

Licking responses are one of the commonly used physiological outputs used to measure the behavioral state of head-restrained animals (see for example Fritz et al., 2003, 2005b; Nakamura, 1999; Ono and Nakamura, 1985, and others), in addition to autonomic responses (Edeline and Weinberger, 1993; Ohl and Scheich, 1997), or eye saccades (Tollin et al., 2005). Licking is an example of a rhythmic orofacial behavior, and shares modal frequencies (theta 4–12 Hz) with other rhythmic behaviors, e.g. sniffing, chewing, sucking, lapping, and whisking (Berg and Kleinfeld, 2003; Gao et al., 2001; Harvey et al., 2001; Kepecs et al., 2006; Nakamura and Katakura, 1995; O'Connor et al., 2002; Sachdev et al., 2001; Vajnerová et al., 2003; Weijnen, 1998). The licking frequencies (and other parameters) that we detected agree with these observations, as well as with studies explicitly providing information about licking behavior in head-fixed rats performing a whisker detection task (Stüttgen et al., 2006; Wiest and Nicolelis, 2003).

Restraining animals can, however, change parameters of rhythmic behaviors (required to perform tasks), such as whisking (Gao et al., 2003; Sellien et al., 2005). The relative inaccessibility of the drinking spout in head-fixed position probably contributed to the decrease in licking frequency in our task (Davis and Smith, 1992; Weijnen, 1998). Such a decrease in licking frequency (or licking in general), together with a more demanding task (discrimination with longer inter-trial intervals) likely contributed to lower performance values in head-fixed position. The overall slight decrease in licking parameters in head-restrained animals could be important when interpreting results from studies relying

on licking as main behavioral output. Especially when animals perform a conditional avoidance task (for example Fritz et al., 2003) during which he/she must stop licking in order to avoid punishment.

6.5.2 Nonauditory modulation of activity in auditory cortex

Primary auditory cortex has been traditionally viewed as a *primary* sensory area, i.e. exclusively specialized in processing of sounds. However, even some of the early studies suggested that neural activity in early stages of auditory processing might be influenced by nonauditory factors; such as "attention" influencing "electric activity" in the cochlear nucleus (Hernandez-Peon et al., 1956), and "attention units" in auditory cortex (Hubel et al., 1959). Since then, multiple studies supported the idea of multi-sensory, or behavioral interactions in auditory cortex, whether they were using evoked potentials (Giard and Peronnet, 1999; Oatman, 1971, 1976; Picton et al., 1971), field potentials (Ghazanfar et al., 2005), magnetoencefalography (Gobbelé et al., 2003; Lütkenhöner et al., 2002), or fMRI (Calvert et al., 1997; Foxe et al., 2002; Johnson and Zatorre, 2005; Kayser et al., 2005; Petkov et al., 2004)

Even at the level of single neurons in the primary auditory area, neuronal responses can be influenced by behavioral contingencies (Beaton and Miller, 1975; Miller et al., 1972), selective attention (Hocherman et al., 1976; Miller et al., 1980), eye position (Fu et al., 2004; Werner-Reiss et al., 2003), or somatosensory stimulation (Brosch et al., 2005; Fu et al., 2003; Lakatos et al., 2007; Schroeder et al., 2001). The "top-down" (attention, behavioral contingencies), and "bottom-up" (somatosensory, eye movements) influences can be seen as enhancing auditory responses, and auditory processing (Lakatos et al., 2007; Schroeder and Foxe, 2005). Auditory cortex—even the primary auditory cortex—should not be viewed as a pure sensory area acting as a (non)linear filter on signals passing through thalamus to "higher" cortical areas. The presence of nonauditory activity in the primary auditory cortex also suggests that better control of behavioral state *and* monitoring of other sensory modalities and motor activity seems necessary when studying neuronal responses in awake animals.

Chapter 7

Conclusions and perspectives

We have developed a setup for studying neuronal activity in head-fixed rats. We have used the setup in combination with patch-clamp recordings to record the activity of single neurons in the auditory cortex of awake naïve and behaving head-fixed rats. We have progressed from characterizing various experimental and theoretical aspects of both extracellular and intracellular neuronal activity in naïve animals toward studying neuronal activity in behaving animals.

Representation of sounds in awake auditory cortex was sparse (Sec. 3), with stimuli typically eliciting a high firing rate in only about 3 % of the population. We have used patch-clamp recordings with glass pipette and sequentially sampled single neurons to obtain an estimate of population activity. Different recording bias of patch-clamp recordings—based on physical contact with neuron, rather than neuronal activity—enabled us to sample neurons with very high, as well as very low firing rates. Our data suggests that only a minority of neurons are engaged in the representation of many stimuli; indeed, the fact that most stimuli drive most neurons only weakly explains why finding the optimal stimulus for any given neuron can be such a challenge. Thus, although there may be an optimal stimulus for any given neuron, most stimuli are not optimal for most neurons, and so are represented sparsely across the population.

Although individual neurons displayed a variety of evoked responses, population activity was well-characterized by a lognormal distribution (Sec. 4). We have studied properties of cortical networks which could give rise to lognormal distributions of firing rates and synaptic weights (Song et al., 2005). We proposed that correlations among neurons' synaptic inputs led to lognormal distribution of

(spontaneous) firing rates in neuronal populations, provided that synaptic strengths themselves were lognormally distributed. In addition, we proposed a multiplicative Hebbian learning rule which can give rise to specific correlations in the synaptic weight matrix and, consequently, to lognormally distributed synaptic weights *and* firing rates.

Intracellular activity in awake auditory cortex was consistent with the idea that neurons receive barrages of synchronized inputs (Sec. 5). Neurons typically spent less than 4 % of their time in up-states (fluctuations of membrane potential) longer than 200 ms, with most neurons displaying only brief bumps in their subthreshold dynamics.

Primary auditory cortex has usually been studied without regard to the behavioral or attentional context of the sounds presented. Having studied sound representation in naïve animals, we sought to study neuronal activity under tight behavioral control (Sec. 6). In a sound-discrimination task, we compared performance and various behavioral parameters in restrained and unrestrained rats, and demonstrated modulation of single-neuron activity in the primary auditory cortex by nonauditory (licking) influences.

The head-fixed experimental setup offers excellent stability, which can be exploited by *in-vivo* imaging techniques in awake animals, as already demonstrated in anesthetized preparations (Ohki et al., 2005; Yaksi and Friedrich, 2006). It would be of interest to see whether techniques such as imaging of population activity, or different recording techniques with recording biases similar to patch-clamp recordings (such as tetrodes (Feierstein et al., 2006), or silicone probes (Barthó et al., 2004)), would yield similar results: sparse representation of sounds, lognormal distribution of firing rates, etc..

We have suggested that neurons with high firing rates might be interneurons, based on extracellular electrophysiological signatures. We have not observed any dependence between neuronal activity and depth of recording, based on an estimate of recording depth. Blind *in-vivo* recordings can provide only estimates of such cell-type specific, or cortical layer-specific effects. Using new molecular approaches, specific populations of neurons (for example, neurons in a specific layer, neurons projecting to or from a specific area, etc.) can be "tagged" and the presence of the tag during blind recording could serve as a *signature* of the tagged populations. Such a tag could be, for example, based on channelrhodopsin-2 (Boyden et al., 2005); a light sensitive cation-selective channel, which can couple a light flash to a spike with millisecond precision. During a recording it would be therefore possible to

distinguish a tagged neuron based on whether the neuron generates a short latency spike in response to light flash.

Perhaps even more interestingly, channelrhodopsin-2 tags could be used to *study behavior* (in head-fixed configuration as well). One could record the activity of a selected subpopulation of neurons in an animal performing a behavioral task, or even perturb activity of a selected population of neurons. Although such tags could be used in freely moving preparations as well, head-fixed animals present a more convenient approach, as one would be allowed more freedom searching for the tagged population.

In summary, the head-fixed experimental setup with awake animals offers excellent control of stimulus parameters, such as stimulus location in auditory experiments, or whisker stimulation in somatosensory experiments. Restrained animals provide great possibility for imaging structure and function in cerebral cortex of awake behaving animals. New molecular approaches could be used for labeling specific subpopulations of neurons in auditory cortex. One could identify labeled subpopulations during *in-vivo* blind recordings, or even perturb activity of specified neuronal subpopulations during behavior.

Appendix A

Tone-evoked responses in awake auditory cortex

Tones evoked a wide range of response patterns in the auditory cortex of unanesthetized rats. To supplement results summarized in Sec. 3 we show raster plots of responses of all neurons for which we assessed response significance (n = 100) in the following five figures .

Each panel in Figs. A.1–A.5 shows response dynamics of one neuron probed with 50–60 dB tones. Firing rate curves beneath each of the rasters were computed by first summing the spikes in 1 ms bins over *all* frequencies shown, and then convolving the resulting peristimulus time histogram (PSTH) with a Gaussian ($\sigma = 3$ ms). In each panel, dots represent individual spikes, and the grey shaded region indicates tone duration (100 ms).
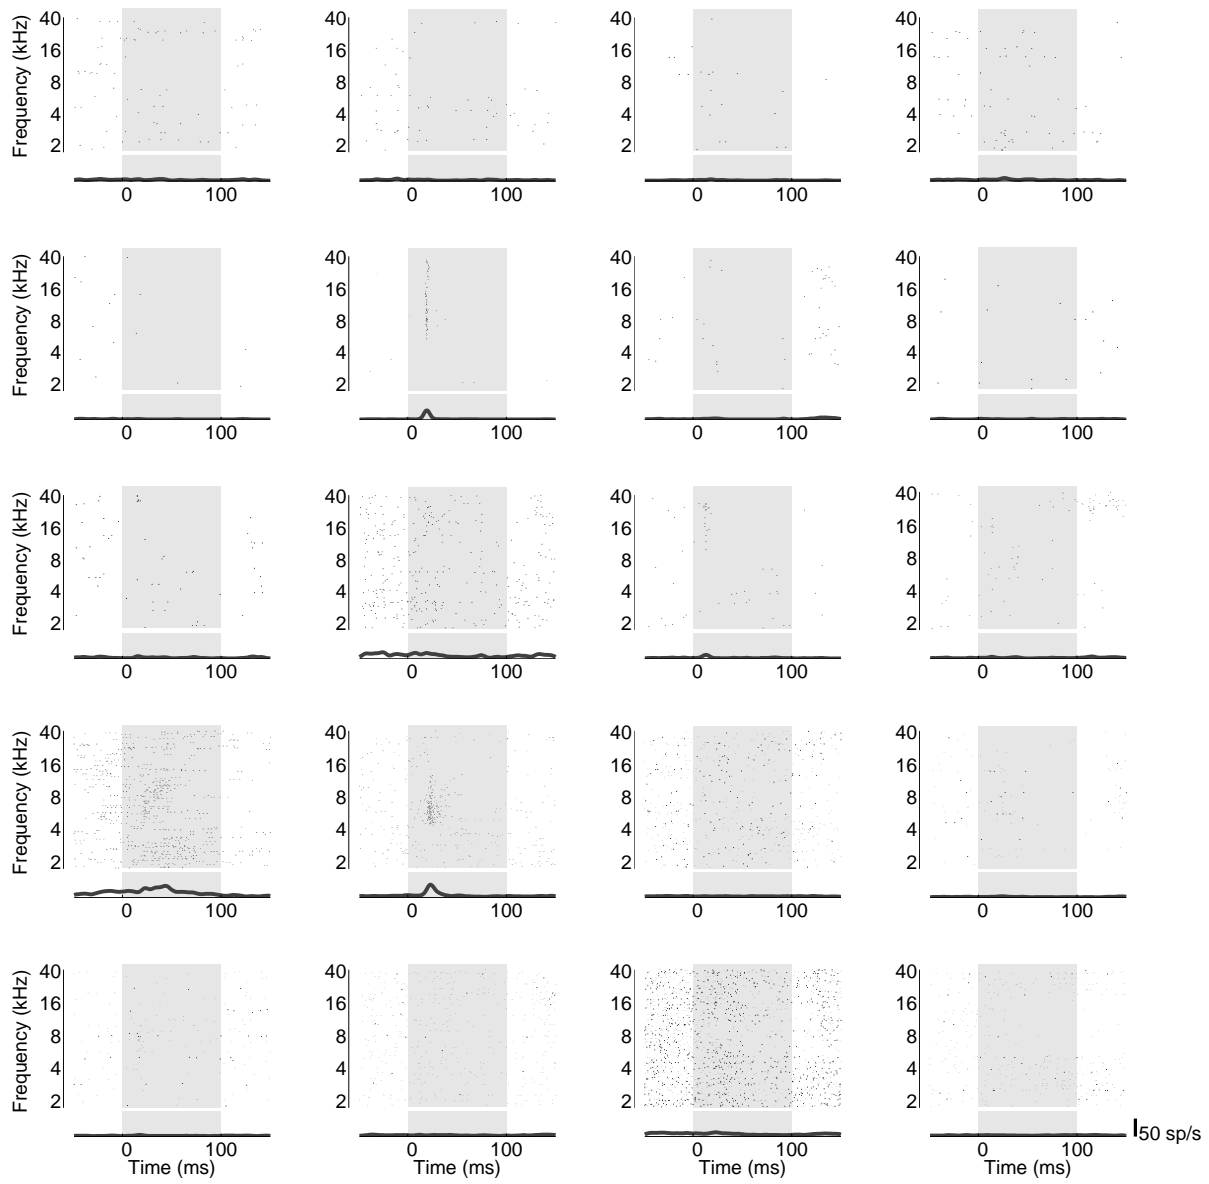
Figure A.1: Tone-evoked responses in awake auditory cortex 1
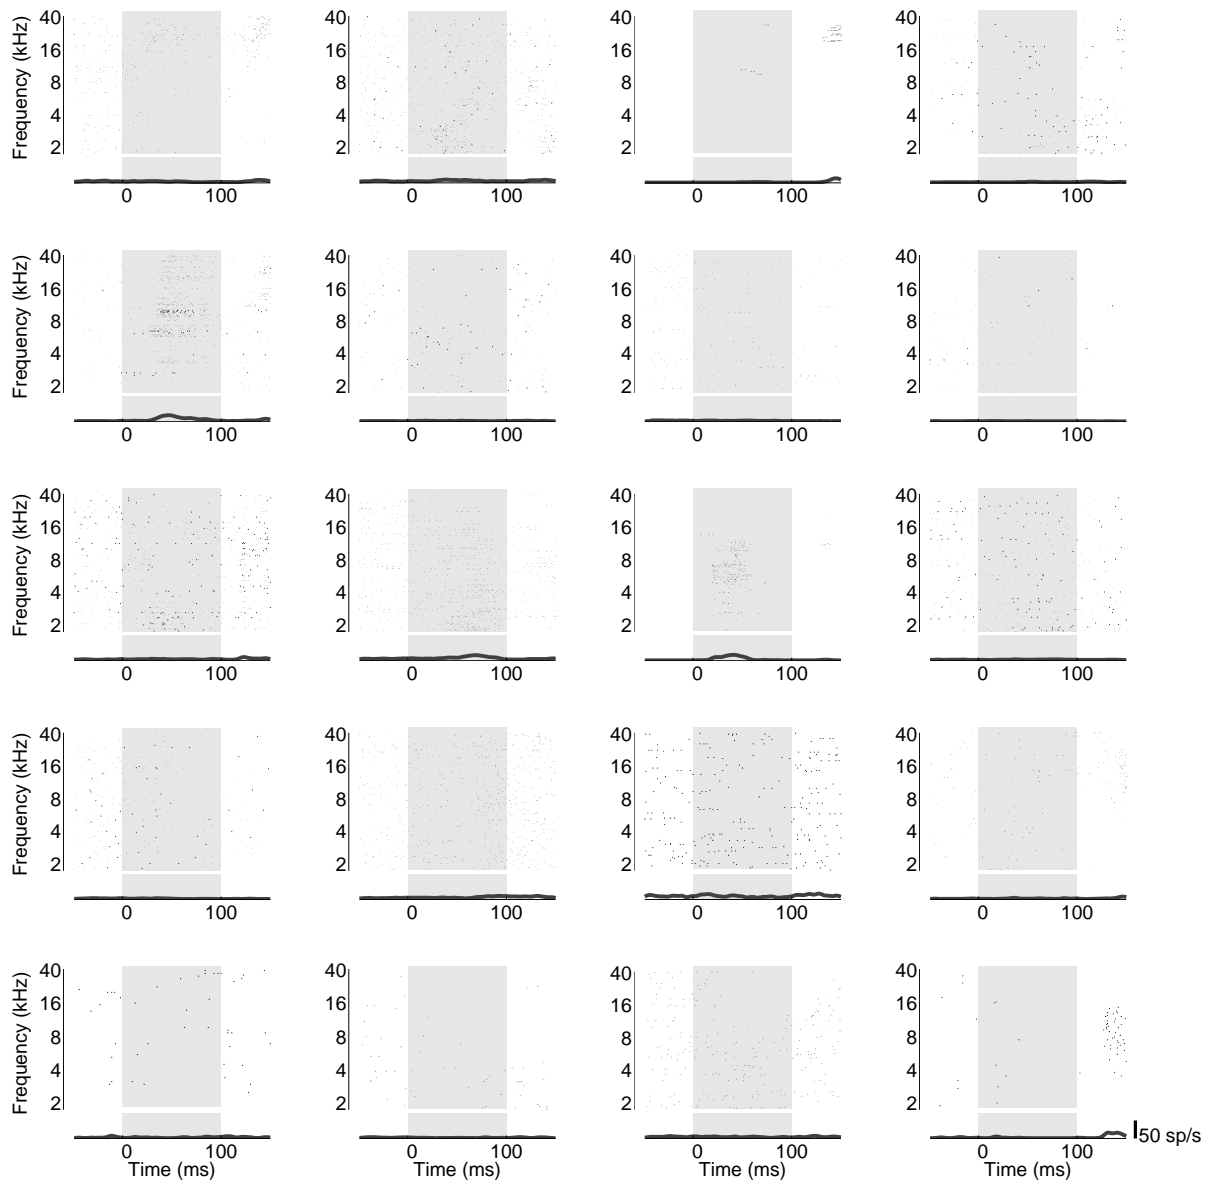
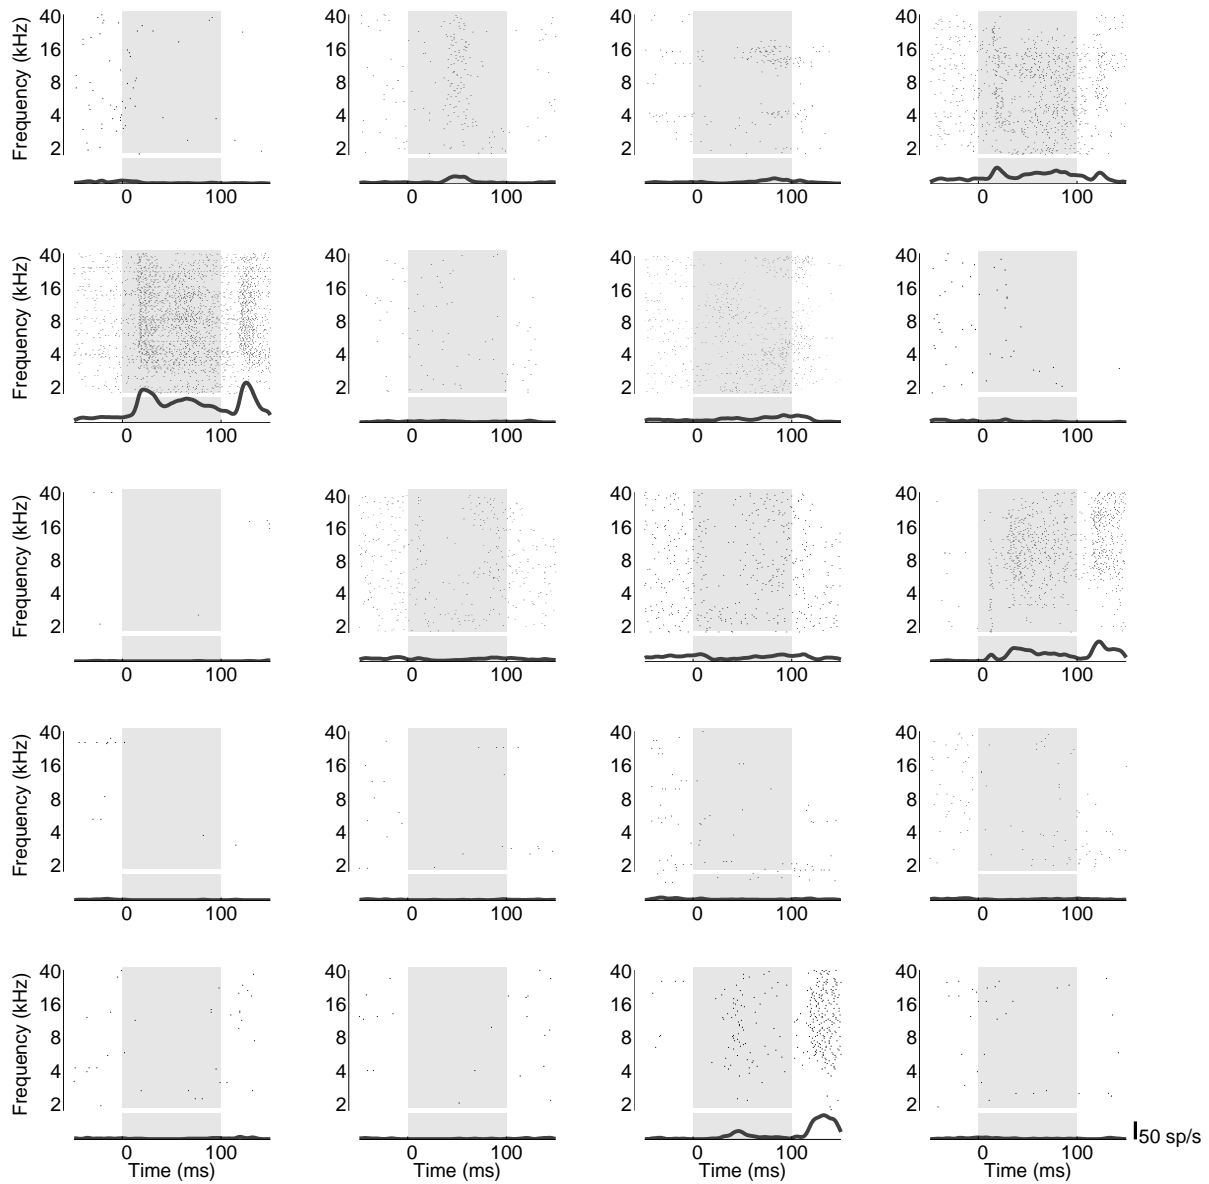Figure A.2: Tone-evoked responses in awake auditory cortex 2

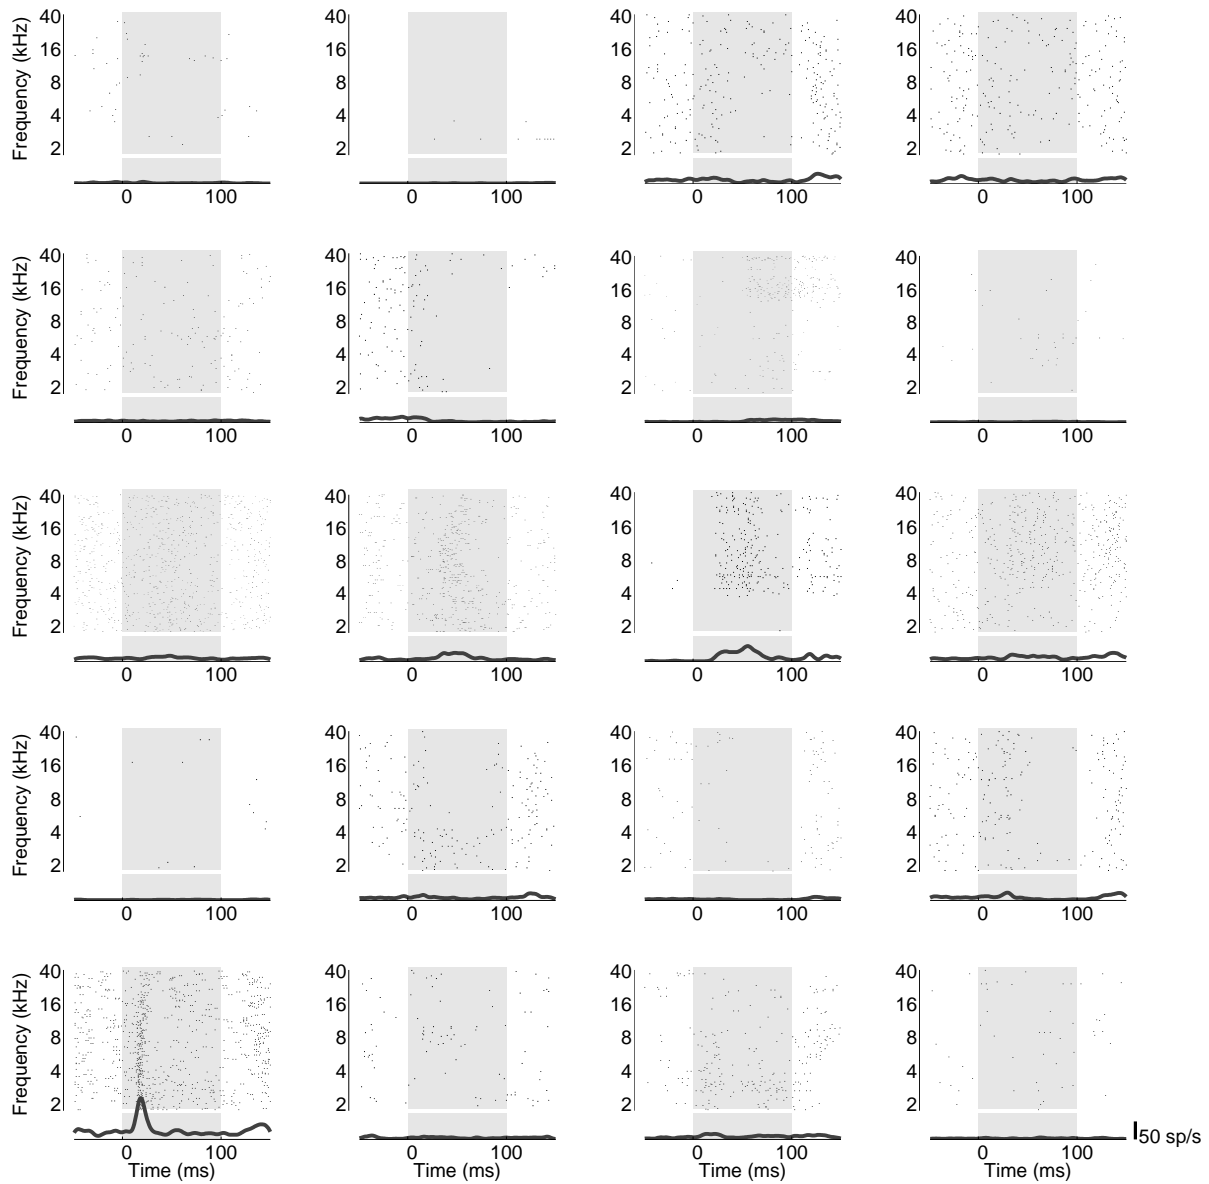Figure A.3: Tone-evoked responses in awake auditory cortex 3

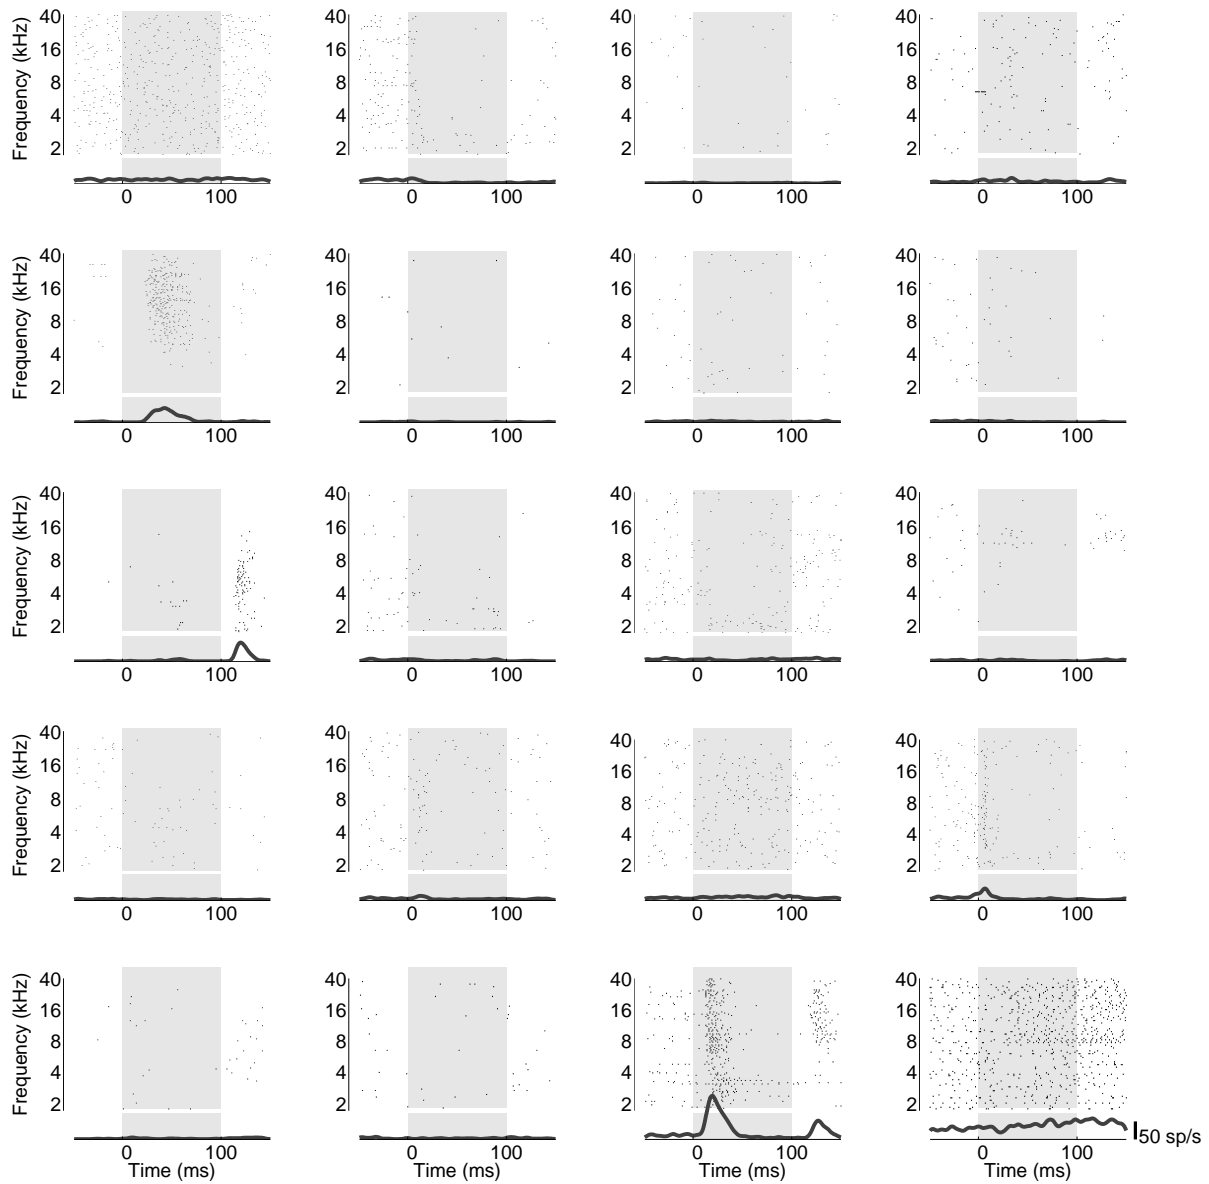Figure A.4: Tone-evoked responses in awake auditory cortex 4

Figure A.5: Tone-evoked responses in awake auditory cortex 5

Appendix B

Neuronal activity in sound discrimination go/no-go task

Neurons displayed a range of responses during a sound discrimination go/no-go task. To suppplement examples presented in Sec. 6.4 we show raster plots for 15 neurons in Figs. B.1–B.8. In each figure, the left column (panels A, C, E) shows a summary for one neuron, and the right column (panels B, D, F) shows summary for different neuron.

Panel **(A)** shows a raster plot of spiking responses for a single neuron. Colored dots represent individual spikes, light grey shaded region indicates stimulus duration, grey ticks indicate individual licks. Trials are sorted into four groups depending on the stimulus and the outcome of each trial. Panel **(C)** shows spiking responses during lick trains, which were aligned to the first lick. Other responses include lick trains which occurred during the penalty period (> 4 s after the start of the trial). Panel **(E)** shows spiking responses to individual licks. All licks surrounded by spikes were aligned to the start of each lick. In all panels, the firing rate curves beneath the rasters were computed by adding the spikes in 1 ms bins, convolving the resulting PSTH with a Gaussian ($\sigma = 20$ ms in A, B, C, D $\sigma = 10$ ms in E, F), and then normalizing to the maximum peak of all firing rate curves. Panels **(B)**, **(D)**, and **(F)** are presented in the same format as panels (A), (C), and (E) for different neuron.

Figure B.1: Neuronal responses during a sound discrimination task 1
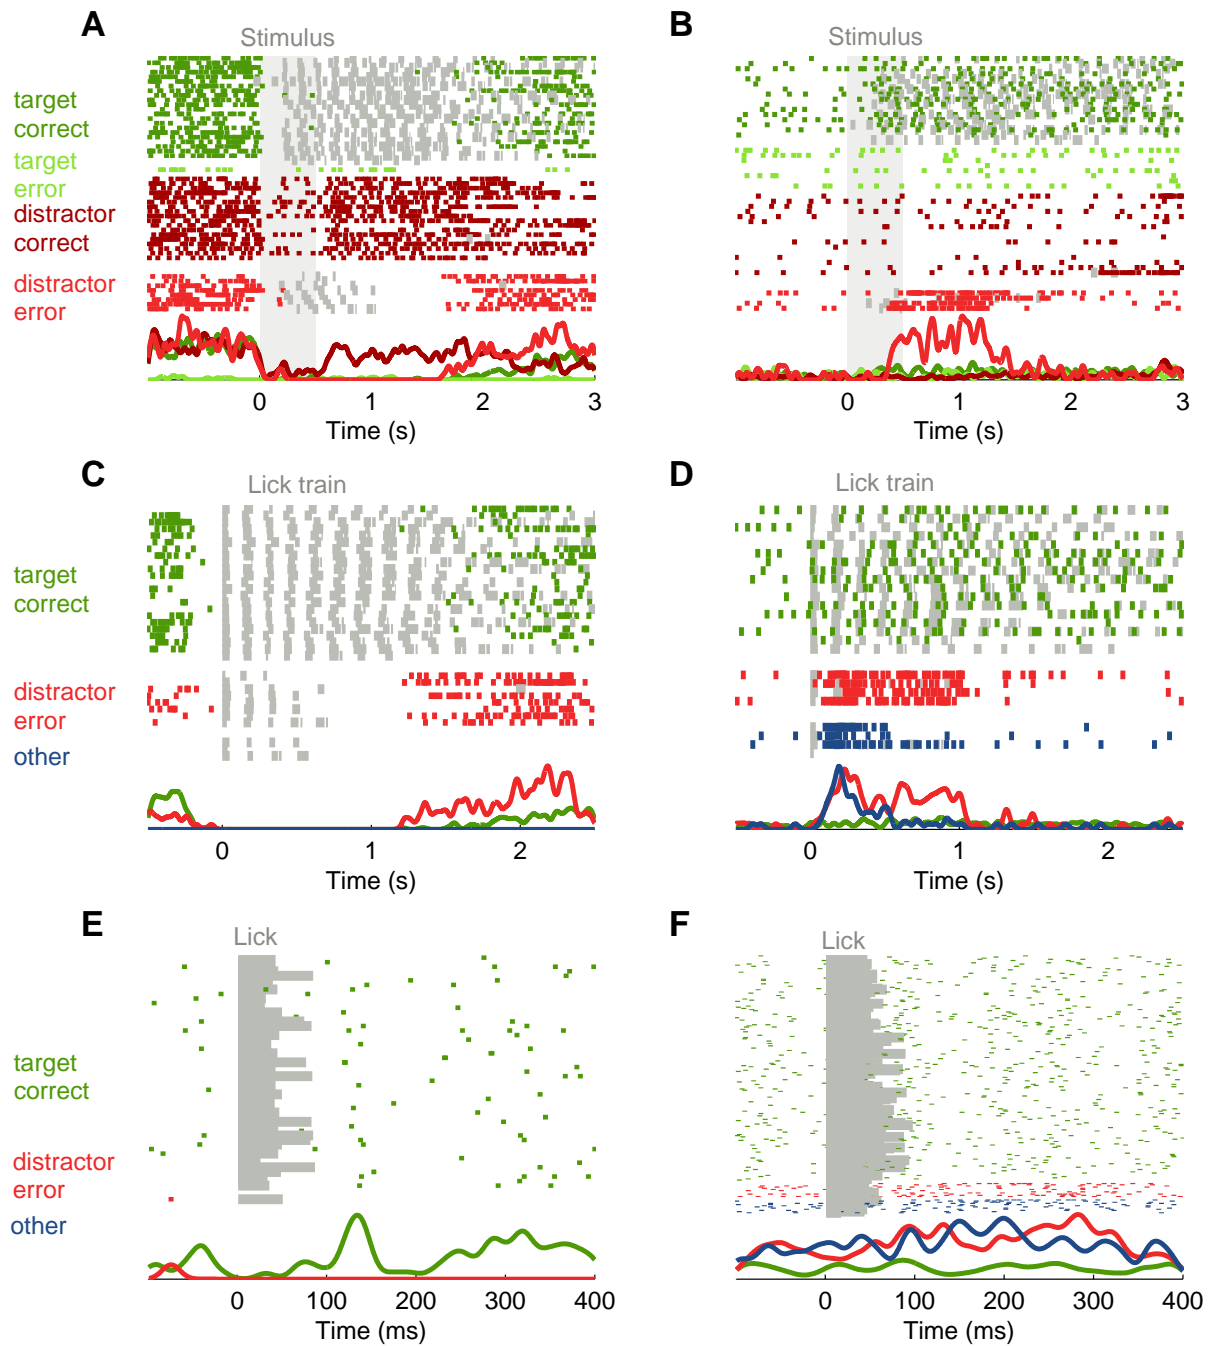
Figure B.2: Neuronal responses during a sound discrimination task 2

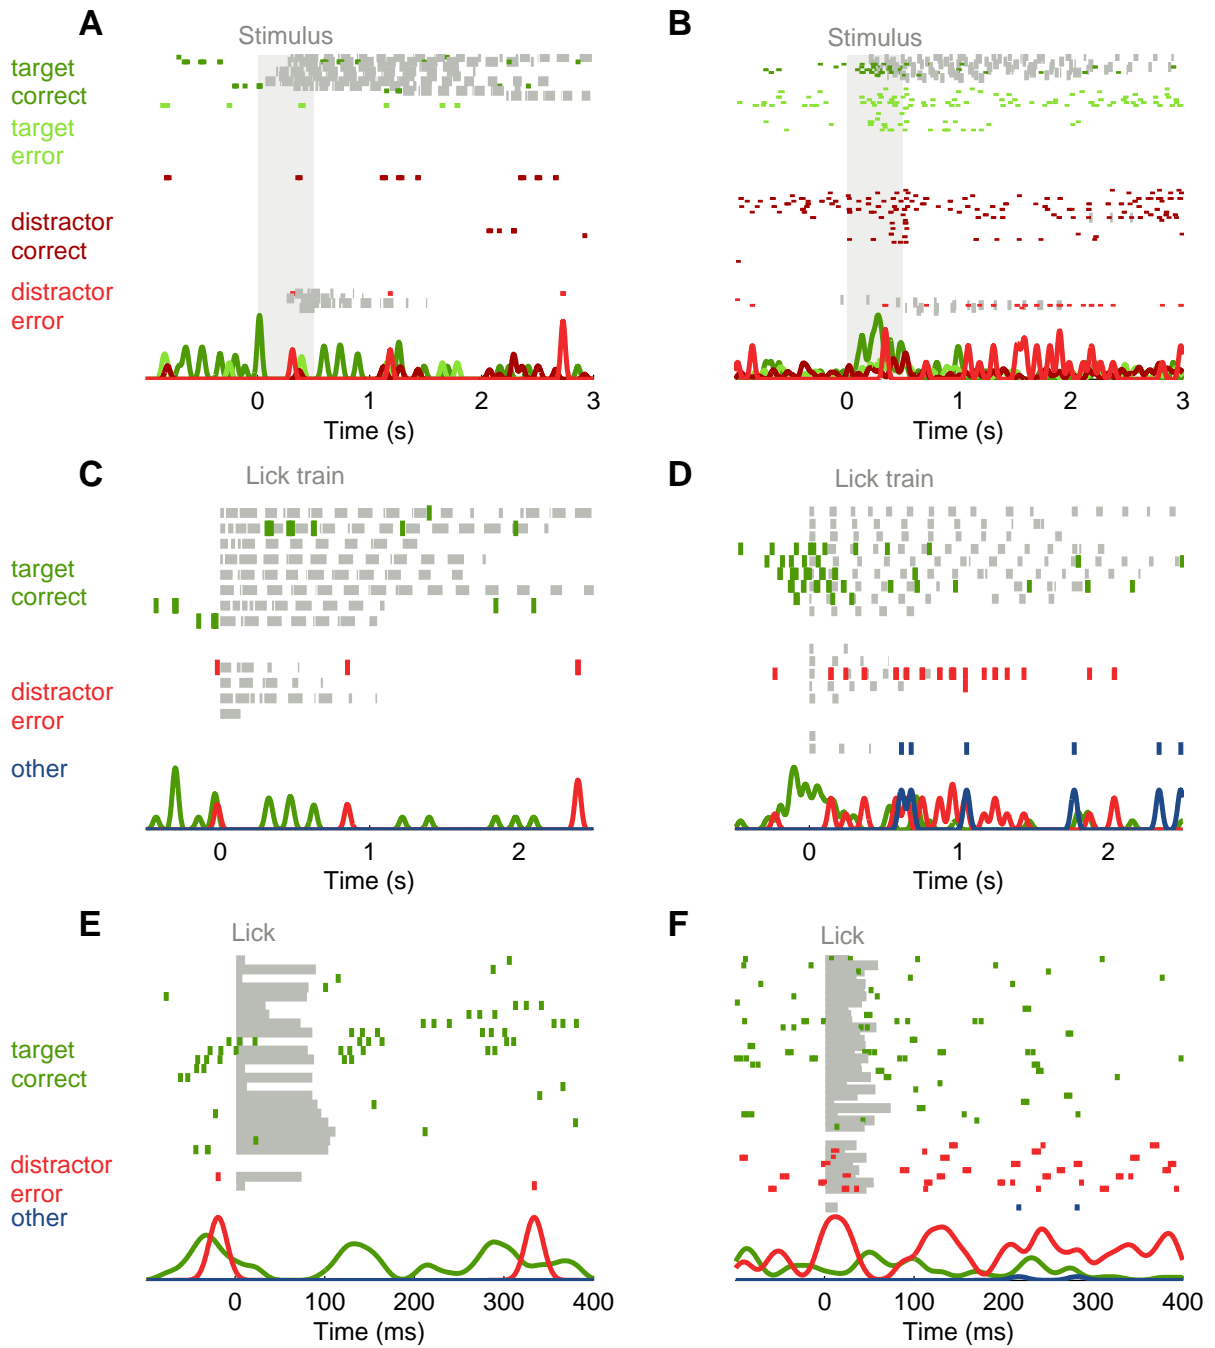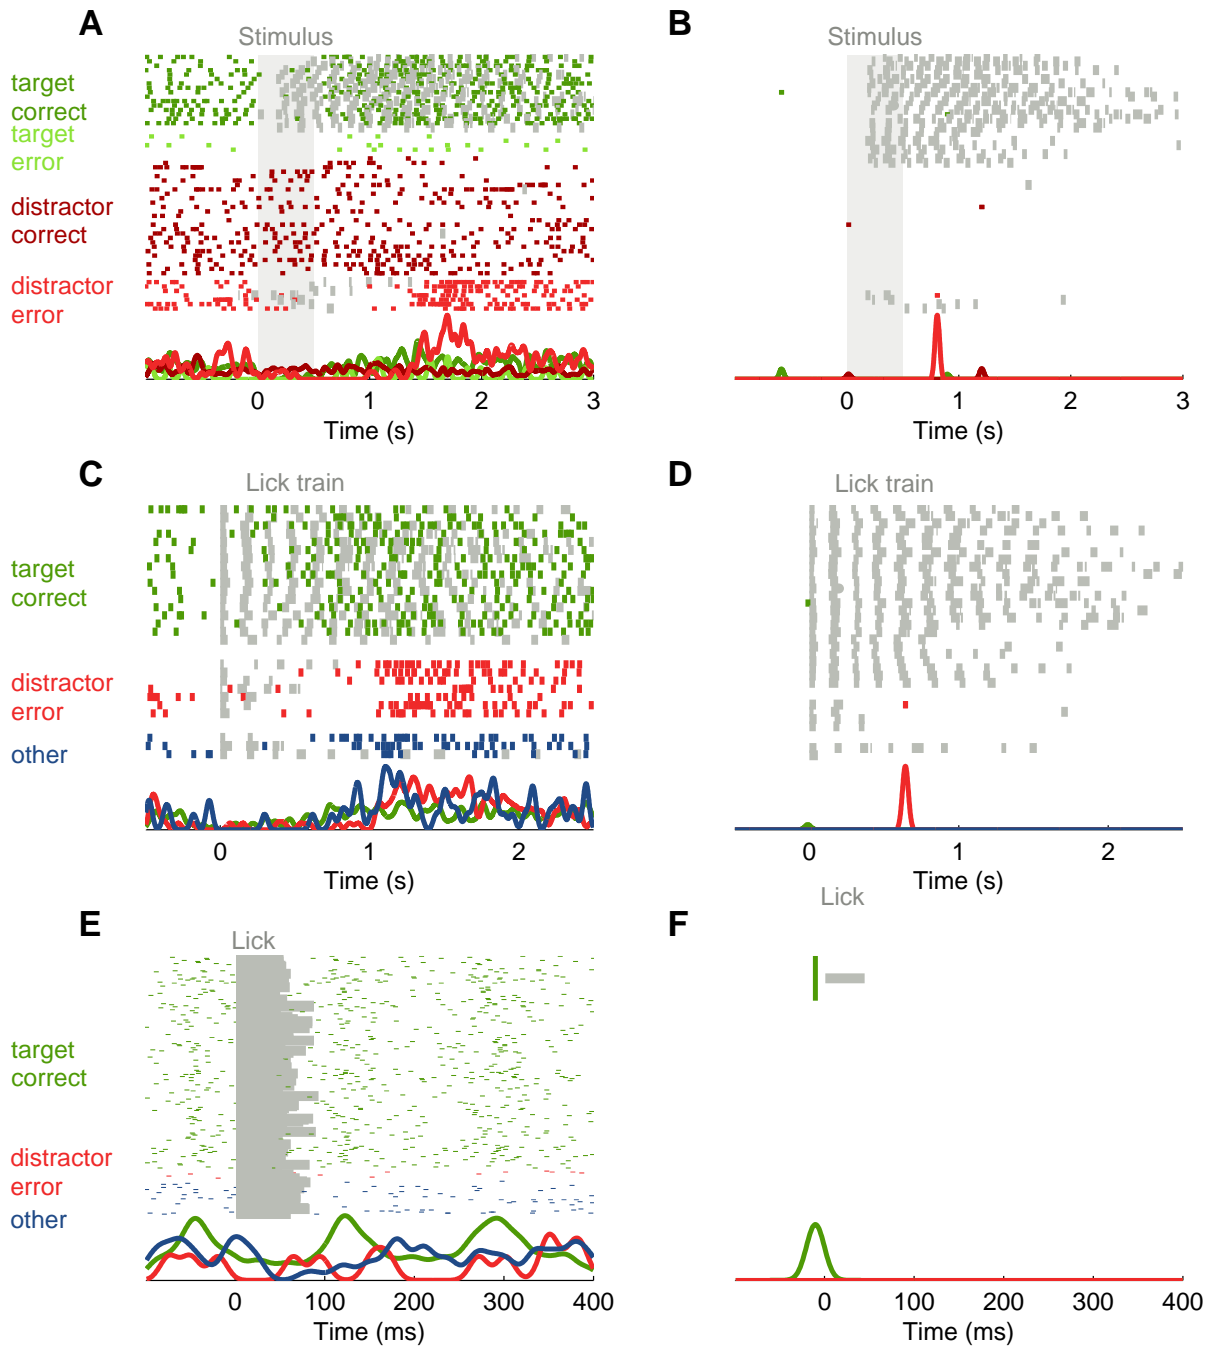Figure B.3: Neuronal responses during a sound discrimination task 3
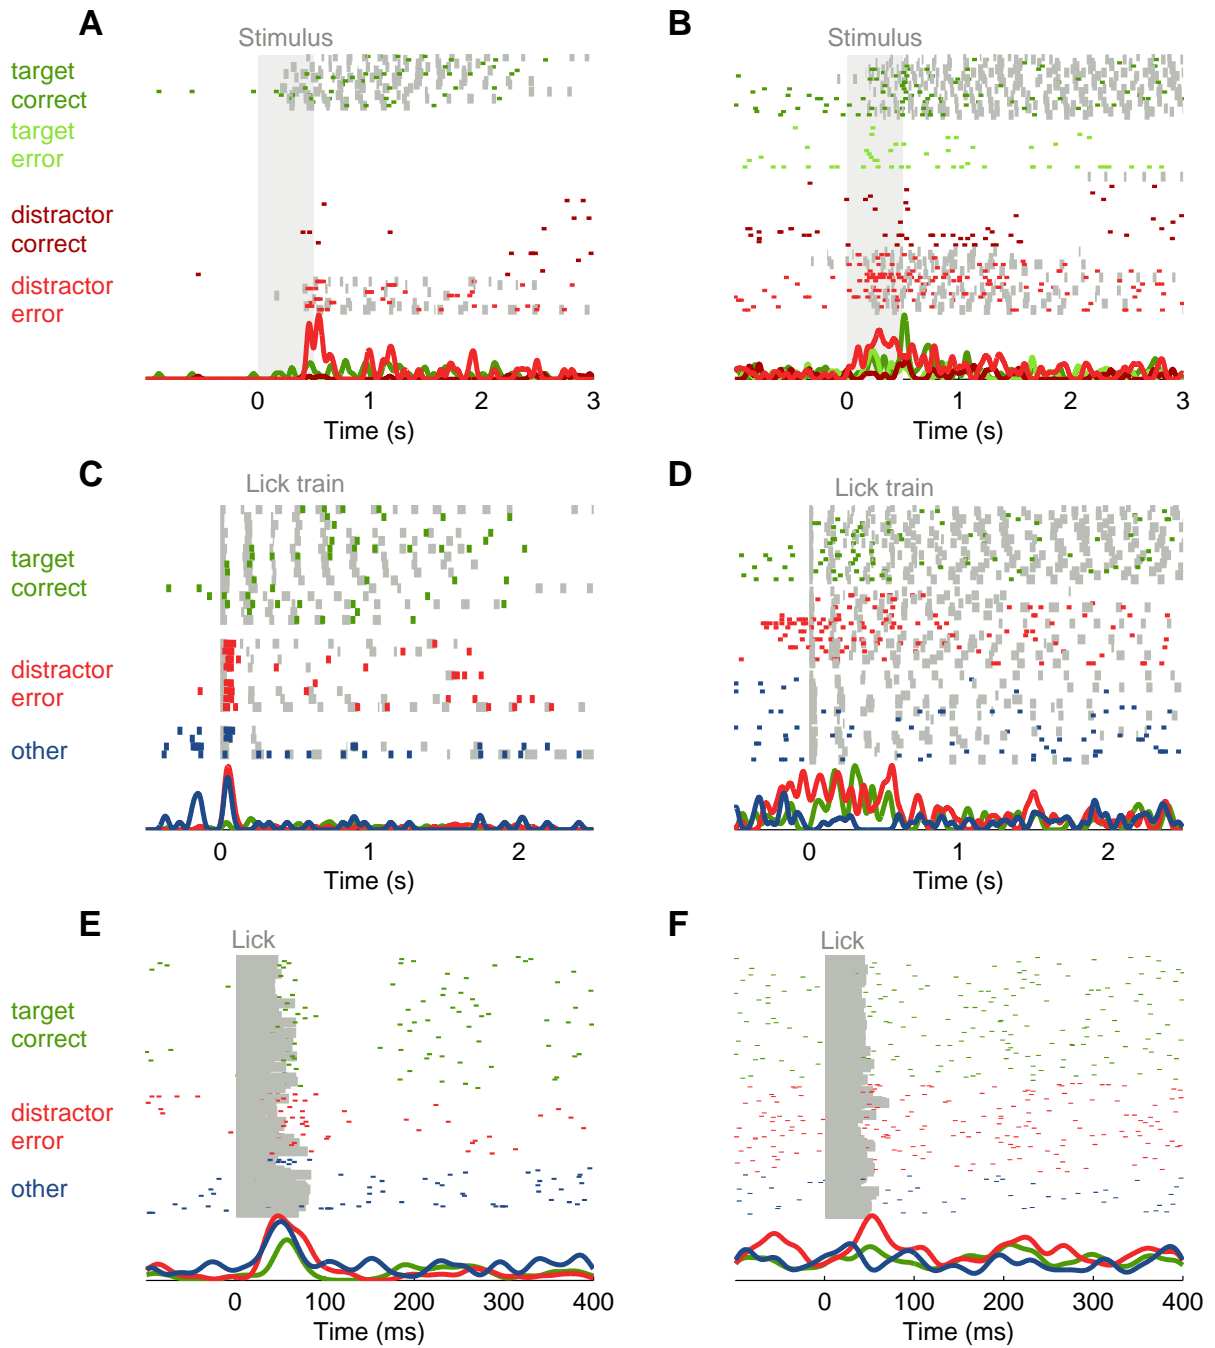
Figure B.4: Neuronal responses during a sound discrimination task 4
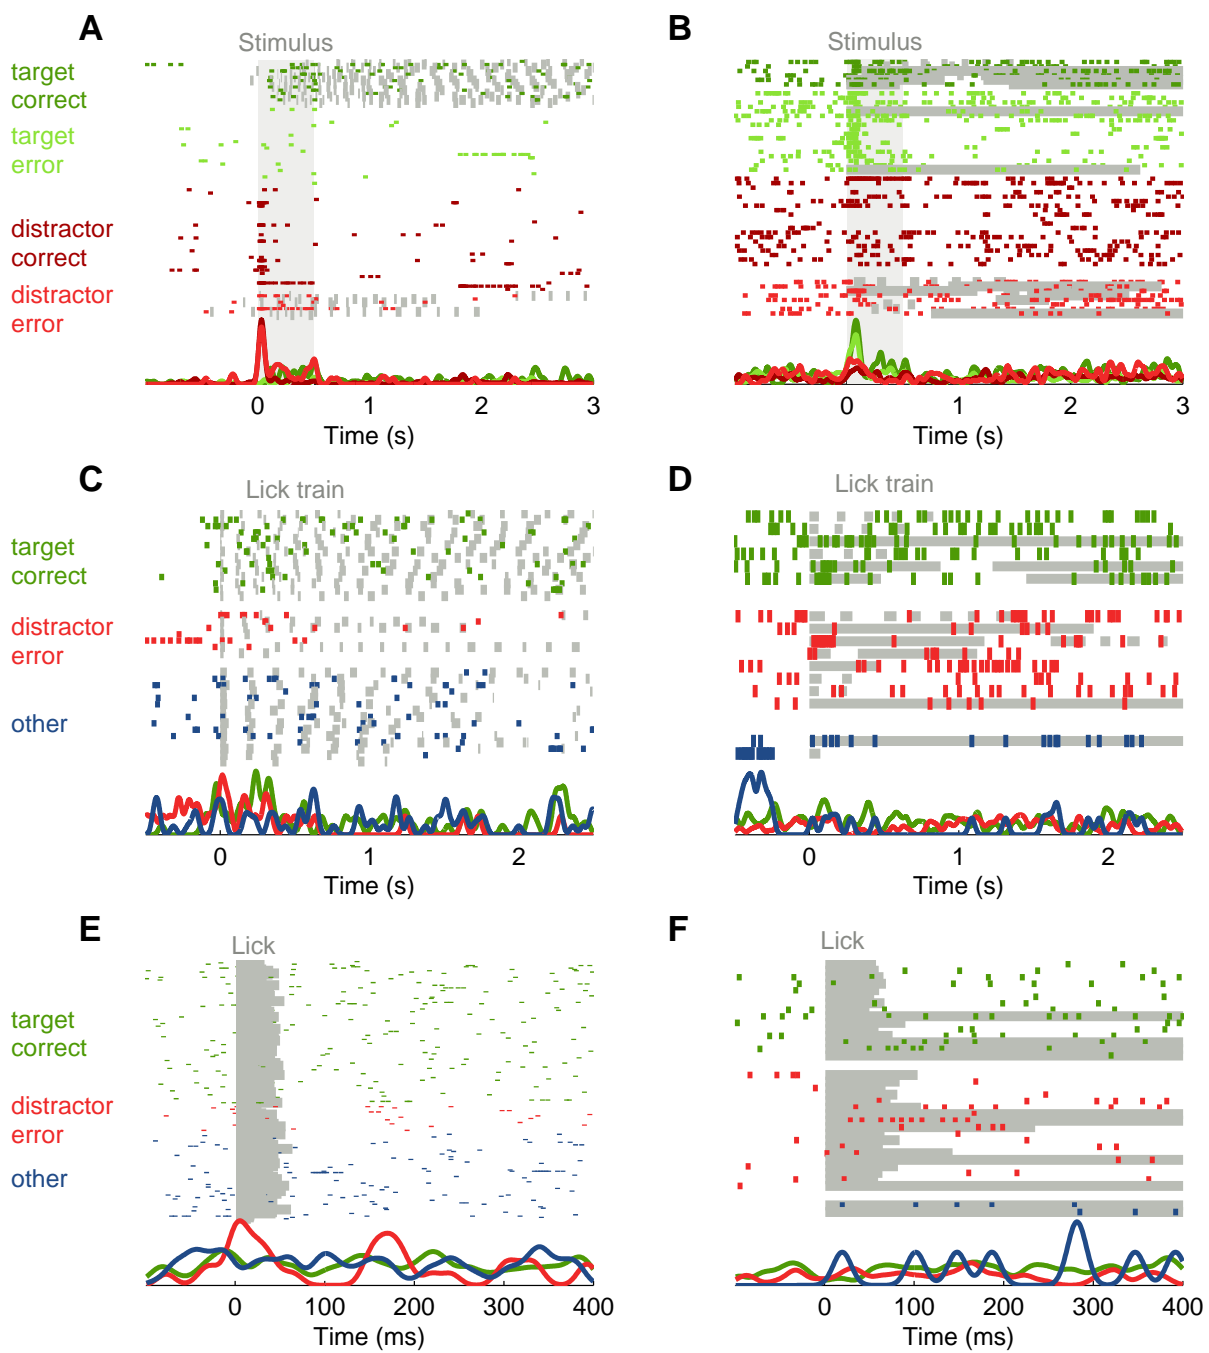
Figure B.5: Neuronal responses during a sound discrimination task 5
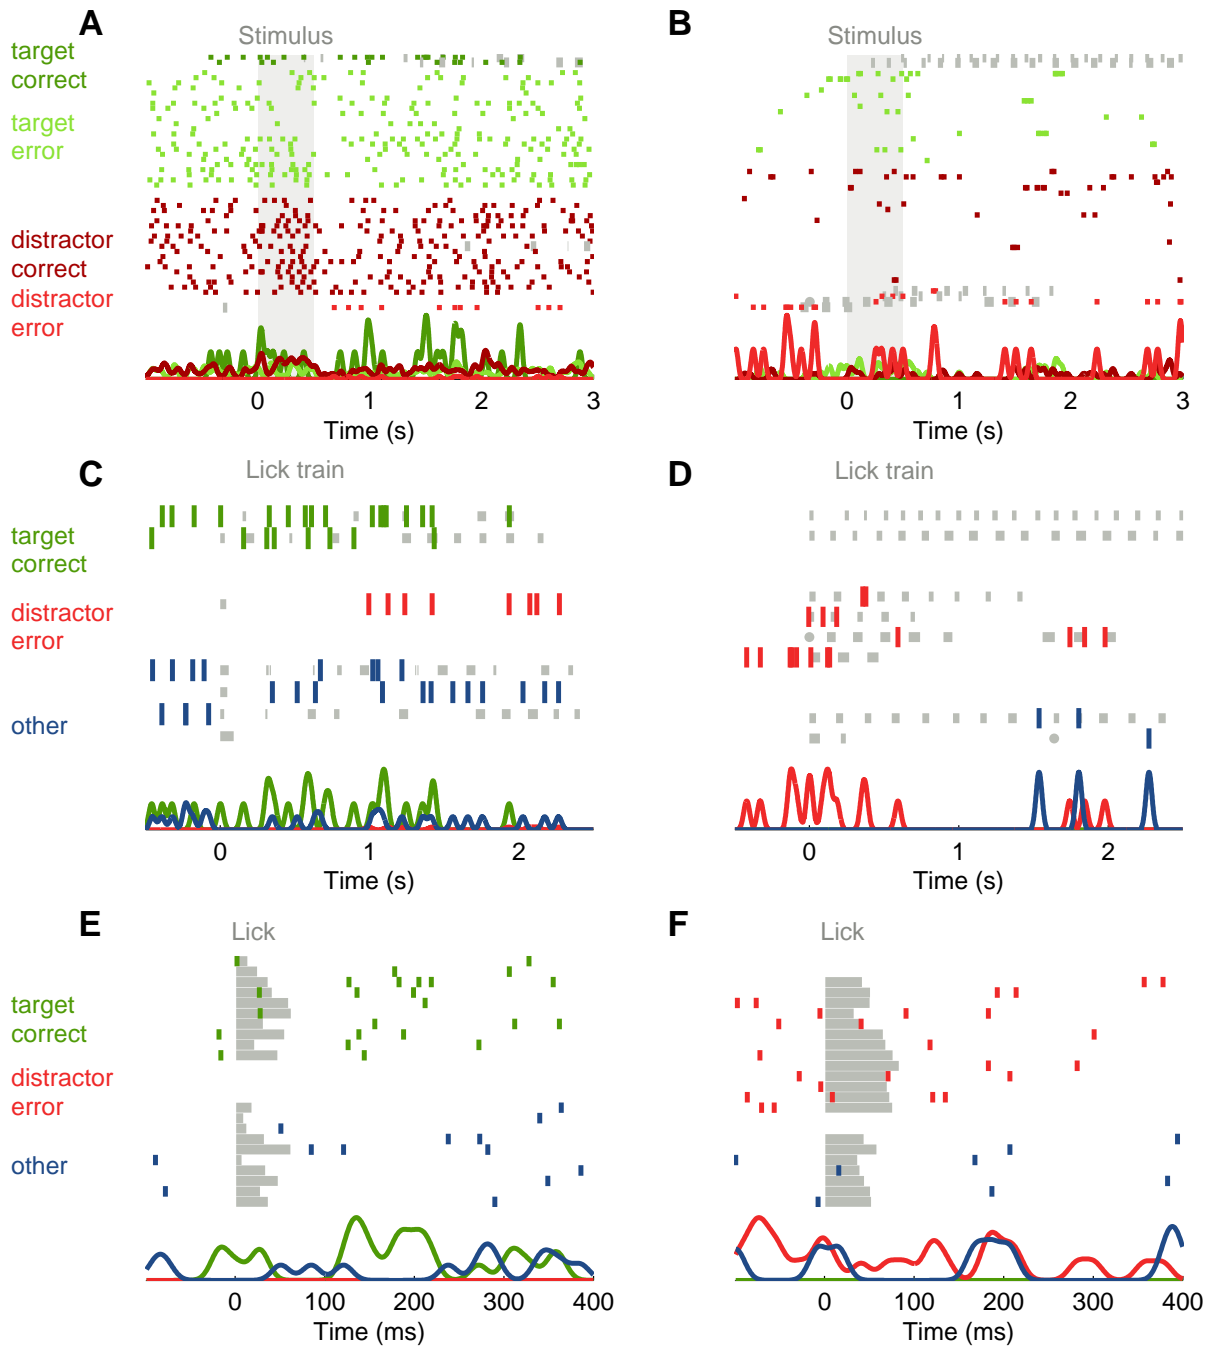
Figure B.6: Neuronal responses during a sound discrimination task 6
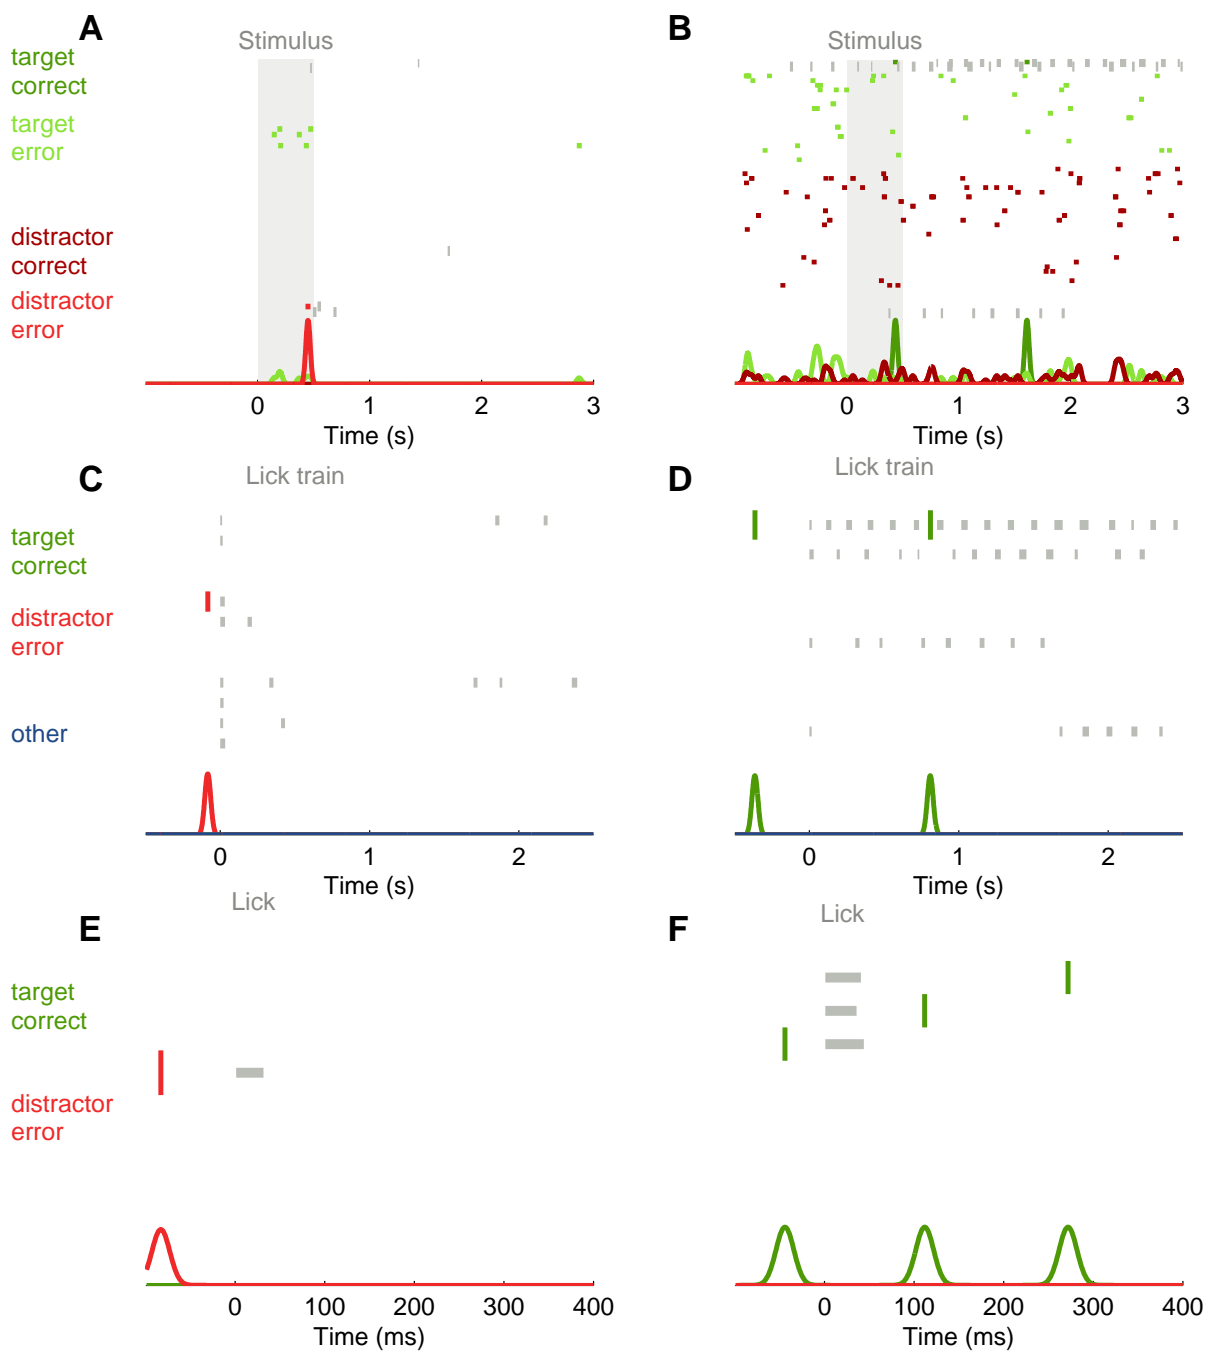
Figure B.7: Neuronal responses during a sound discrimination task 7
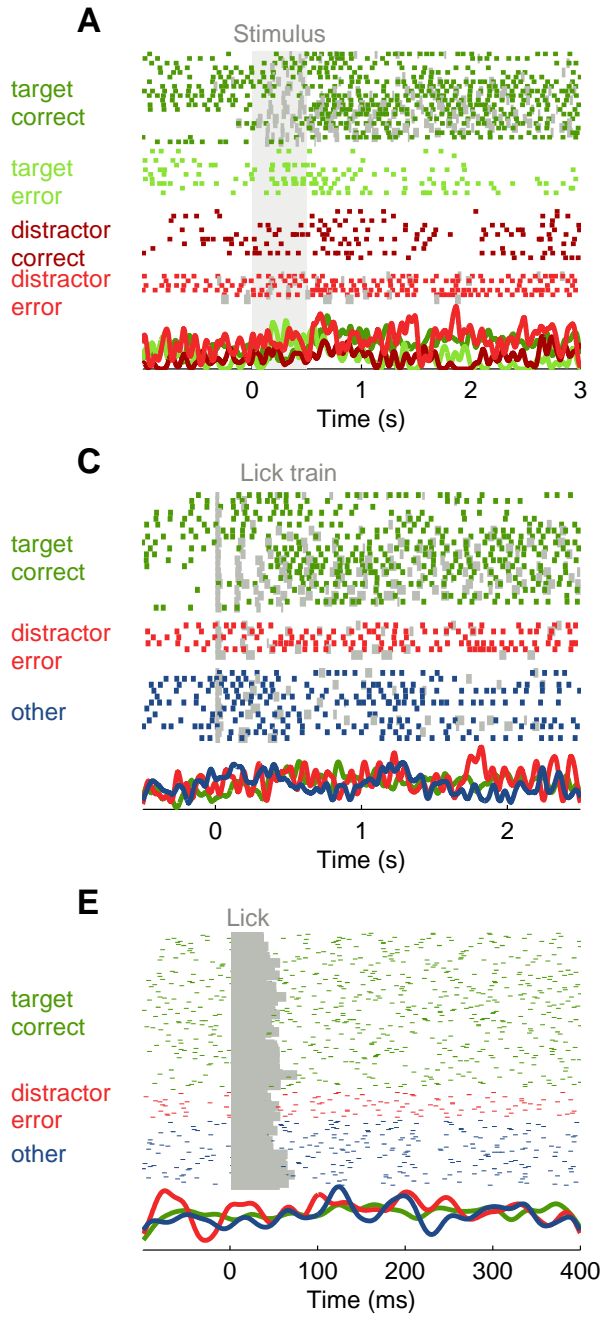
Figure B.8: Neuronal responses during a sound discrimination task 8

Appendix C

Schematics of surgical implants and lickometer

The following five figures show schematics of the surgical implants we used for head-fixed recordings, and a schematic of the lickometer we used to detect licking.

An aluminum headpost (Fig. C.1) was glued to a rat's skull. During experiments, the headpost was held fixed by a headpost holder (Fig. C.2). A plastic capped well (Figs. C.3–C.4) was implanted on top of the craniotomy. The lickometer (Fig. C.5) was placed in front of the rat's mouth, and was used to both deliver water reward and measure the animal's licking.

The implants and the lickometer were made by Robert Eifert, who also drew the schematics.

Figure C.1: Headpost

Figure C.2: Headpost holder

Figure C.3: Recording well
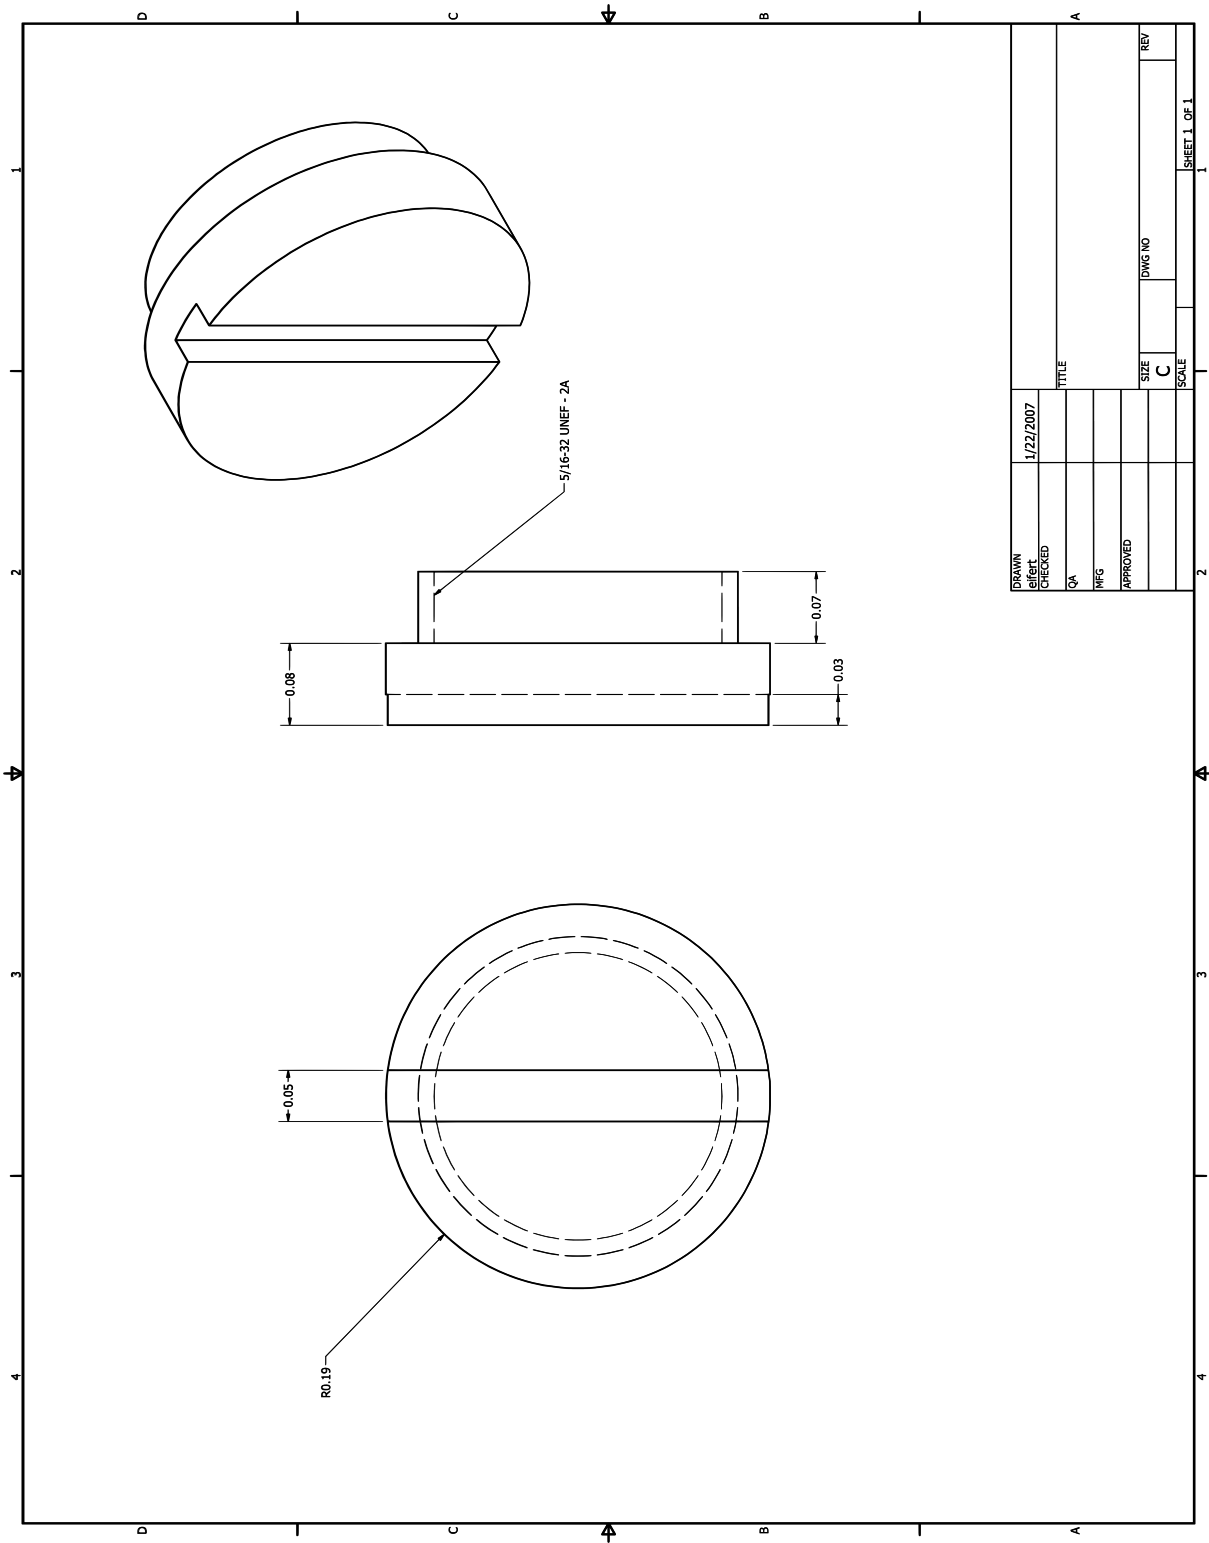
Figure C.4: Recording well cap
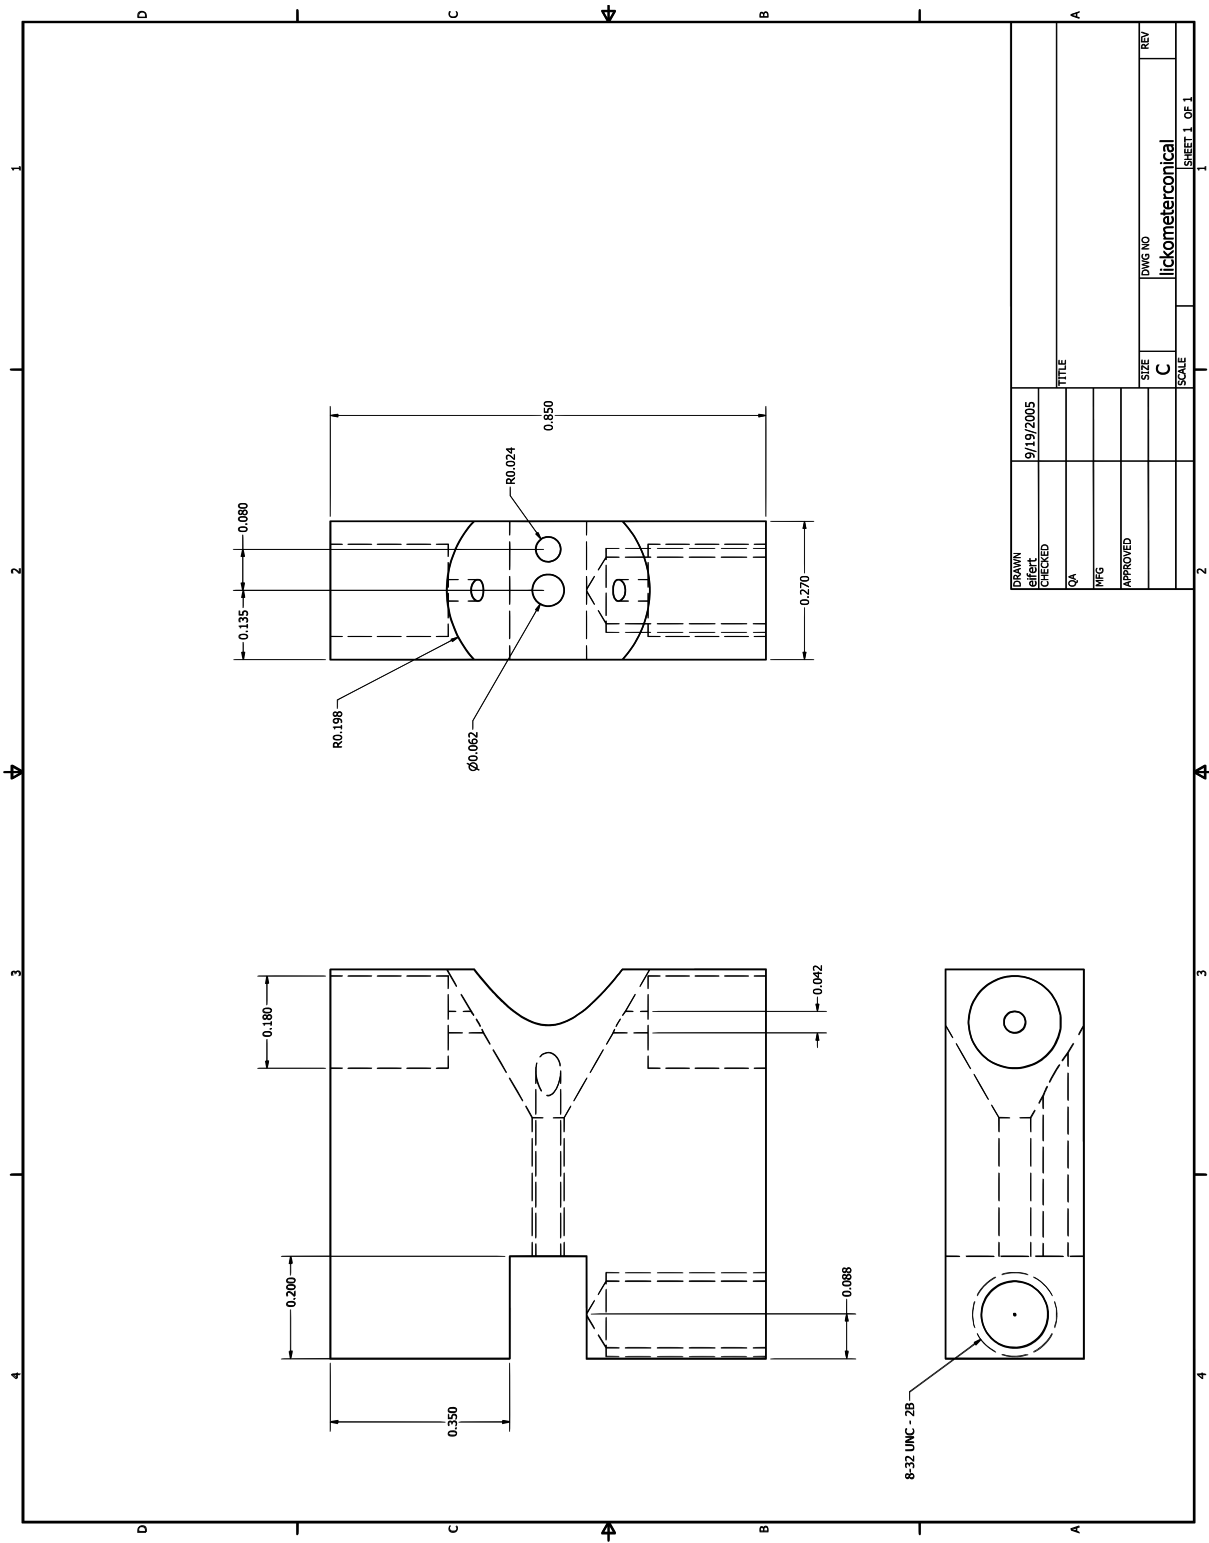
Figure C.5: Lickometer

142

Bibliography

Abbott, L. F. and Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nat Neurosci* **3 Suppl**:1178–1183.

Abeles, M. and Goldstein, M. (1972). Responses of single units in the primary auditory cortex of the cat to tones and to tone pairs. *Brain Res* **42**:337–352.

Abeles, M., Vaadia, E., and Bergman, H. (1990). Firing patterns of single units in the prefrontal cortex and neural network models. *Network* **1**:13–25.

Adrian, E. (1926). The impulses produced by sensory nerve-endings: Part 4. Impulses from pain receptors. *J Physiol* **62**:33–51.

Adrian, E. and Matthews, R. (1927). The action of light on the eye: Part I. The discharge of impulses in the optic nerve and its relation to the electric changes in the retina. *J Physiol* **63**:378–414.

Adrian, E. D., Craik, K. J. W., and Sturdy, R. S. (1938). The electrical response of the auditory mechanism in cold-blooded vertebrates. *Proc R Soc Lond B Biol Sci* **125**:435–455.

Ahmari, S. E., Buchanan, J., and Smith, S. J. (2000). Assembly of presynaptic active zones from cytoplasmic transport packets. *Nat Neurosci* **3**:445–451.

Ahmed, B., Anderson, J. C., Douglas, R. J., Martin, K. A., and Nelson, J. C. (1994). Polyneuronal innervation of spiny stellate neurons in cat visual cortex. *J Comp Neurol* **341**:39–49.

Anderson, J., Lampl, I., Reichova, I., Carandini, M., and Ferster, D. (2000). Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex. *Nat Neurosci* **3**:617–621.

Arita, M. (2005). Scale-freeness and biological networks. *J Biochem (Tokyo)* **138**:1–4.

Armstrong-James, M. and George, M. J. (1988). Influence of anesthesia on spontaneous activity and receptive field size of single units in rat Sm1 neocortex. *Exp Neurol* **99**:369–387.

Asari, H., Pearlmutter, B. A., and Zador, A. M. (2006). Sparse representations for the cocktail party problem. *J Neurosci* **26**:7477–7490.

Attwell, D. and Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *J Cereb Blood Flow Metab* **21**:1133–1145.

Baddeley, R., Abbott, L. F., Booth, M. C., Sengpiel, F., Freeman, T., Wakeman, E. A., and Rolls, E. T. (1997). Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc Biol Sci* **264**:1775–1783.

Bar-Yosef, O., Rotman, Y., and Nelken, I. (2002). Responses of neurons in cat primary auditory cortex to bird chirps: effects of temporal and spectral context. *J Neurosci* **22**:8619–8632.

Barbour, D. and Wang, X. (2003). Auditory cortical responses elicited in awake primates by random spectrum stimuli. *J Neurosci* **23**:7194–7206.

Barlow, H. B. (1961). Possible principles underlying the transformations of sensory messages. In Rosenblith, W. A., editor, *Sensory Communication* pages 217–234. MIT Press.

Barthó, P., Hirase, H., Monconduit, L., Zugaro, M., Harris, K. D., and Buzsáki, G. (2004). Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J Neurophysiol* **92**:600–608.

Beaton, R. and Miller, J. (1975). Single cell activity in the auditory cortex of the unanesthetized, behaving monkey: correlation with stimulus controlled behavior. *Brain Res* **100**:543–562.

Benshalom, G. and White, E. L. (1986). Quantification of thalamocortical synapses with spiny stellate neurons in layer IV of mouse somatosensory cortex. *J Comp Neurol* **253**:303–314.

Benucci, A., Verschure, P. F. M. J., and König, P. (2004). Two-state membrane potential fluctuations driven by weak pairwise correlations. *Neural Comput* **16**:2351–2378.

Berg, R. W. and Kleinfeld, D. (2003). Rhythmic whisking by rat: retraction as well as protraction of the vibrissae is under active muscular control. *J Neurophysiol* **89**:104–117.

Bi, G. and Poo, M. (2001). Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu Rev Neurosci* **24**:139–166.

Blanton, M. G., Turco, J. J. L., and Kriegstein, A. R. (1989). Whole cell recording from neurons in slices of reptilian and mammalian cerebral cortex. *J Neurosci Methods* **30**:203–210.

Boyden, E. S., Zhang, F., Bamberg, E., Nagel, G., and Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. *Nat Neurosci* **8**:1263–1268.

Bozza, T., McGann, J. P., Mombaerts, P., and Wachowiak, M. (2004). In vivo imaging of neuronal activity by targeted expression of a genetically encoded probe in the mouse. *Neuron* **42**:9–21.

Brecht, M., Roth, A., and Sakmann, B. (2003). Dynamic receptive fields of reconstructed pyramidal cells in layers 3 and 2 of rat somatosensory barrel cortex. *J Physiol* **553**:243–265.

Brecht, M. and Sakmann, B. (2002). Dynamic representation of whisker deflection by synaptic potentials in spiny stellate and pyramidal cells in the barrels and septa of layer 4 rat somatosensory cortex. *J Physiol* **543**:49–70.

Brecht, M., Schneider, M., Sakmann, B., and Margrie, T. (2004). Whisker movements evoked by stimulation of single pyramidal cells in rat motor cortex. *Nature* **427**:704–710.

Bregman, A. (1990). *Auditory Scene Analysis*. The MIT Press Cambridge, MA.

Bremer, F. and Bonnet, V. (1949). An analysis of the sensory responses of the acoustic cortex. *Electroencephalogr Clin Neurophysiol* **1**:447–449.

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.

Brugge, J. and Merzenich, M. (1973). Responses of neurons in auditory cortex of the macaque monkey to monaural and binaural stimulation. *J Neurophysiol* **36**:1138–1158.

Bruno, R. M. and Sakmann, B. (2006). Cortex is driven by weak but synchronously active thalamocortical synapses. *Science* **312**:1622–1627.

Bureau, I., Shepherd, G. M. G., and Svoboda, K. (2004). Precise development of functional and anatomical columns in the neocortex. *Neuron* **42**:789–801.

Callaway, E. M. (1998). Local circuits in primary visual cortex of the macaque monkey. *Annu Rev Neurosci* **21**:47–74.

Calvert, G. A., Bullmore, E. T., Brammer, M. J., Campbell, R., Williams, S. C., McGuire, P. K., Woodruff, P. W., Iversen, S. D., and David, A. S. (1997). Activation of auditory cortex during silent lipreading. *Science* **276**:593–596.

Camperi, M. and Wang, X. J. (1998). A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability. *J Comput Neurosci* **5**:383–405.

Carandini, M. and Ferster, D. (1997). A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science* **276**:949–952.

Carandini, M. and Ferster, D. (2000). Membrane potential and firing rate in cat primary visual cortex. *J Neurosci* **20**:470–484.

Cauli, B., Audinat, E., Lambolez, B., Angulo, M., Ropert, N., Tsuzuki, K., Hestrin, S., and Rossier, J. (1997). Molecular and physiological diversity of cortical nonpyramidal cells. *J Neurosci* **17**:3894–3906.

Chance, F. S., Abbott, L. F., and Reyes, A. D. (2002). Gain modulation from background synaptic input. *Neuron* **35**:773–782.

Cheung, S. W., Nagarajan, S. S., Bedenbaugh, P. H., Schreiner, C. E., Wang, X., and Wong, A. (2001). Auditory cortical neuron response differences under isoflurane versus pentobarbital anesthesia. *Hear Res* **156**:115–127.

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.

Cooper, A. (1801). Farther Observations on the Effects Which Take Place from the Destruction of the Membrana Tympani of the Ear; With an Account of an Operation for the Removal of a Particular Species of Deafness. *Phil Trans R Soc* **91**:435–450.

Cooper, A. and Home, E. (1800). Observations on the Effects Which Take Place from the Destruction of the Membrana Tympani of the Ear. By Mr. Astley Cooper. In a Letter to Everard Home, Esq. F. R. S. by Whom Some Remarks are Added. *Phil Trans R Soc* **90**:151–160.

Cossart, R., Aronov, D., and Yuste, R. (2003). Attractor dynamics of network up states in the neocortex. *Nature* **423**:283–288.

Covey, E., Kauer, J. A., and Casseday, J. H. (1996). Whole-cell patch-clamp recording reveals subthreshold sound-evoked postsynaptic currents in the inferior colliculus of awake bats. *J Neurosci* **16**:3009–3018.

Cowan, R. L. and Wilson, C. J. (1994). Spontaneous firing patterns and axonal projections of single corticostriatal neurons in the rat medial agranular cortex. *J Neurophysiol* **71**:17–32.

Crochet, S. and Petersen, C. C. H. (2006). Correlating whisker behavior with membrane potential in barrel cortex of awake mice. *Nat Neurosci* **9**:608–610.

Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999). Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *J Neurosci* **19**:274–287.

Csicsvari, J., Jamieson, B., Wise, K. D., and Buzsáki, G. (2003). Mechanisms of gamma oscillations in the hippocampus of the behaving rat. *Neuron* **37**:311–322.

Cusack, R., Deeks, J., Aikman, G., and Carlyon, R. P. (2004). Effects of location, frequency region, and time course of selective attention on auditory scene analysis. *J Exp Psychol Hum Percept Perform* **30**:643–656.

Dan, Y., Atick, J. J., and Reid, R. C. (1996). Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J Neurosci* **16**:3351–3362.

Davies, P., Erulkar, S., and Rose, J. (1956). Single unit activity in the auditory cortex of the cat. *Bull Johns Hopkins Hosp* **99**:55–86.

Davis, J. D. and Smith, G. P. (1992). Analysis of the microstructure of the rhythmic tongue movements of rats ingesting maltose and sucrose solutions. *Behav Neurosci* **106**:217–228.

deCharms, R., Blake, D., and Merzenich, M. (1998). Optimizing sound features for cortical neurons. *Science* **280**:1439–1443.

Desai, N. S., Cudmore, R. H., Nelson, S. B., and Turrigiano, G. G. (2002). Critical periods for experience-dependent synaptic scaling in visual cortex. *Nat Neurosci* **5**:783–789.

Desai, N. S., Rutherford, L. C., and Turrigiano, G. G. (1999). Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nat Neurosci* **2**:515–520.

DeWeese, M., Wehr, M., and Zador, A. (2003). Binary spiking in auditory cortex. *J Neurosci* **23**:7940–7949.

DeWeese, M. R., Hromádka, T., and Zador, A. M. (2005). Reliability and representational bandwidth in the auditory cortex. *Neuron* **48**:479–488.

DeWeese, M. R. and Zador, A. M. (2000). In vivo whole-cell recordings of synaptic responses to acoustic stimuli in rat auditory cortex. In *Soc Neurosci Abstr* volume 26 page 63714.

DeWeese, M. R. and Zador, A. M. (2006). Non-gaussian membrane potential dynamics imply sparse, synchronous activity in auditory cortex. *J Neurosci* **26**:12206–12218.

Doron, N., Ledoux, J., and Semple, M. (2002). Redefining the tonotopic core of rat auditory cortex: physiological evidence for a posterior field. *J Comp Neurol* **453**:345–360.

Douglas, R. J., Koch, C., Mahowald, M., Martin, K. A., and Suarez, H. H. (1995). Recurrent excitation in neocortical circuits. *Science* **269**:981–985.

Douglas, R. J. and Martin, K. A. C. (2004). Neuronal circuits of the neocortex. *Annu Rev Neurosci* **27**:419–451.

Durif, C., Jouffrais, C., and Rouiller, E. M. (2003). Single-unit responses in the auditory cortex of monkeys performing a conditional acousticomotor task. *Exp Brain Res* **153**:614–627.

Edeline, J. M., Pham, P., and Weinberger, N. M. (1993). Rapid development of learning-induced receptive field plasticity in the auditory cortex. *Behav Neurosci* **107**:539–551.

Edeline, J. M. and Weinberger, N. M. (1993). Receptive field plasticity in the auditory cortex during frequency discrimination training: selective retuning independent of task difficulty. *Behav Neurosci* **107**:82–103.

Edwards, F. A., Konnerth, A., Sakmann, B., and Takahashi, T. (1989). A thin slice preparation for patch clamp recordings from neurones of the mammalian central nervous system. *Pflügers Arch* **414**:600–612.

Ehret, G. (1997). The auditory cortex. *J Comp Physiol [A]* **181**:547–557.

Elhilali, M., Fritz, J. B., Klein, D. J., Simon, J. Z., and Shamma, S. A. (2004). Dynamics of precise spike timing in primary auditory cortex. *J Neurosci* **24**:1159–1172.

Engert, F. and Bonhoeffer, T. (1999). Dendritic spine changes associated with hippocampal long-term synaptic plasticity. *Nature* **399**:66–70.

Evans, E. and Whitfield, I. (1964). Classification of unit responses in the auditory cortex of the unanaesthetized and unrestrained cat. *J Physiol* **171**:476–493.

Feierstein, C. E., Quirk, M. C., Uchida, N., Sosulski, D. L., and Mainen, Z. F. (2006). Representation of spatial goals in rat orbitofrontal cortex. *Neuron* **51**:495–507.

Felleman, D. J. and Essen, D. C. V. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex* **1**:1–47.

Ferster, D. and Jagadeesh, B. (1992). EPSP-IPSP interactions in cat visual cortex studied with in vivo whole-cell patch recording. *J Neurosci* **12**:1262–1274.

Fishman, Y. I., Reser, D. H., Arezzo, J. C., and Steinschneider, M. (2001). Neural correlates of auditory stream segregation in primary auditory cortex of the awake monkey. *Hear Res* **151**:167–187.

Foxe, J. J., Wylie, G. R., Martinez, A., Schroeder, C. E., Javitt, D. C., Guilfoyle, D., Ritter, W., and Murray, M. M. (2002). Auditory-somatosensory multisensory processing in auditory association cortex: an fMRI study. *J Neurophysiol* **88**:540–543.

Fritz, J., Elhilali, M., and Shamma, S. (2005a). Active listening: task-dependent plasticity of spectrotemporal receptive fields in primary auditory cortex. *Hear Res* **206**:159–176.

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.

Fritz, J. B., Elhilali, M., and Shamma, S. A. (2005b). Differential dynamic plasticity of A1 receptive fields during multiple spectral tasks. *J Neurosci* **25**:7623–7635.

Fu, K.-M. G., Johnston, T. A., Shah, A. S., Arnold, L., Smiley, J., Hackett, T. A., Garraghty, P. E., and Schroeder, C. E. (2003). Auditory cortical neurons respond to somatosensory stimulation. *J Neurosci* **23**:7510–7515.

Fu, K.-M. G., Shah, A. S., O'Connell, M. N., McGinnis, T., Eckholdt, H., Lakatos, P., Smiley, J., and Schroeder, C. E. (2004). Timing and laminar profile of eye-position effects on auditory responses in primate auditory cortex. *J Neurophysiol* **92**:3522–3531.

Gaese, B. and Ostwald, J. (2001). Anesthesia changes frequency tuning of neurons in the rat primary auditory cortex. *J Neurophysiol* **86**:1062–1066.

Gaese, B. and Ostwald, J. (2003). Complexity and temporal dynamics of frequency coding in the awake rat auditory cortex. *Eur J Neurosci* **18**:2638–2652.

Galambos, R. (1954). Neural mechanisms of audition. *Physiol Rev* **34**:497–528.

Galambos, R. and Davis, H. (1943). The response of single auditory-nerve fibers to acoustic stimulation. *J Neurophysiol* **6**:39–57.

Games, K. D. and Winer, J. A. (1988). Layer V in rat auditory cortex: projections to the inferior colliculus and contralateral cortex. *Hear Res* **34**:1–25.

Gao, P., Bermejo, R., and Zeigler, H. P. (2001). Whisker deafferentation and rodent whisking patterns: behavioral evidence for a central pattern generator. *J Neurosci* **21**:5374–5380.

Gao, P., Ploog, B. O., and Zeigler, H. P. (2003). Whisking as a "voluntary" response: operant control of whisking parameters and effects of whisker denervation. *Somatosens Mot Res* **20**:179–189.

Gerstein, G. and Kiang, N. (1964). Responses of single units in the auditory cortex. *Exp Neurol* **10**:1–18.

Getting, P. A. (1989). Emerging principles governing the operation of neural networks. *Annu Rev Neurosci* **12**:185–204.

Ghazanfar, A. A., Maier, J. X., Hoffman, K. L., and Logothetis, N. K. (2005). Multisensory integration of dynamic faces and voices in rhesus monkey auditory cortex. *J Neurosci* **25**:5004–5012.

Giard, M. H. and Peronnet, F. (1999). Auditory-visual integration during multimodal object recognition in humans: a behavioral and electrophysiological study. *J Cogn Neurosci* **11**:473–490.

Girosi (1998). An equivalence between sparse approximation and support vector machines. *Neural Comput* **10**:1455–1480.

Gobbelé, R., Schürmann, M., Forss, N., Juottonen, K., Buchner, H., and Hari, R. (2003). Activation of the human posterior parietal and temporoparietal cortices during audiotactile interaction. *Neuroimage* **20**:503–511.

Goldman-Rakic, P. S. (1995). Cellular basis of working memory. *Neuron* **14**:477–485.

Goldstein, M. H., Hall, J. L., and Butterfield, B. O. (1968). Single-unit activity in the primary auditory cortex of unanesthetized cats. *J Acoust Soc Am* **43**:444–455.

Gray, N. W., Weimer, R. M., Bureau, I., and Svoboda, K. (2006). Rapid redistribution of synaptic PSD-95 in the neocortex in vivo. *PLoS Biol* **4**:e370.

Guyonneau, R., Vanrullen, R., and Thorpe, S. J. (2004). Temporal codes and sparse representations: a key to understanding rapid processing in the visual system. *J Physiol Paris* **98**:487–497.

Hahnloser, R. H. R., Kozhevnikov, A. A., and Fee, M. S. (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* **419**:65–70.

Hallpike, C., Hartridge, H., and Rawdon-Smith, A. (1937). On the electrical responses of the cochlea and the auditory tract of the cat to a phase reversal produced in a continuous musical tone. *Proc R Soc Lond B Biol Sci* **122**:175–185.

Hamill, O. P., Marty, A., Neher, E., Sakmann, B., and Sigworth, F. J. (1981). Improved patch-clamp techniques for high-resolution current recording from cells and cell-free membrane patches. *Pflügers Arch* **391**:85–100.

Harlow, J. M. (1848). Passage of an iron rod through the head. *Boston Medical and Surgical Journal* **39**:389–393. Reprinted in *J Neuropsychiatry Clin Neurosci* (1999) **11**(2):281–283.

Harris, K. D., Henze, D. A., Csicsvari, J., Hirase, H., and Buzsáki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol* **84**:401–414.

Harris-Warrick, R. M. and Marder, E. (1991). Modulation of neural networks for behavior. *Annu Rev Neurosci* **14**:39–57.

Hartline, H. K. (1925). The electrical response to illumination of the eye in intact animals, including the human subject; and in decerebrate preparations. *Am J Physiol* **73**:600–612.

Harvey, M. A., Bermejo, R., and Zeigler, H. P. (2001). Discriminative whisking in the head-fixed rat: optoelectronic monitoring during tactile detection and discrimination tasks. *Somatosens Mot Res* **18**:211–222.

Heil, P. (1997). Auditory cortical onset responses revisited. II. Response strength. *J Neurophysiol* **77**:2642–2660.

Heil, P., Rajan, R., and Irvine, D. R. (1992a). Sensitivity of neurons in cat primary auditory cortex to tones and frequency-modulated stimuli. I: Effects of variation of stimulus parameters. *Hear Res* **63**:108–134.

Heil, P., Rajan, R., and Irvine, D. R. (1992b). Sensitivity of neurons in cat primary auditory cortex to tones and frequency-modulated stimuli. II: Organization of response properties along the 'isofrequency' dimension. *Hear Res* **63**:135–156.

Heil, P., Rajan, R., and Irvine, D. R. (1994). Topographic representation of tone intensity along the isofrequency axis of cat primary auditory cortex. *Hear Res* **76**:188–202.

Henze, D. A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K. D., and Buzsáki, G. (2000). Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *J Neurophysiol* **84**:390–400.

Hernandez-Peon, R., Scherrer, H., and Jouvet, M. (1956). Modification of electric activity in cochlear nucleus during attention in unanesthetized cats. *Science* **123**:331–332.

Hocherman, S., Benson, D. A., Goldstein, M. H., 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.

Holcman, D. and Tsodyks, M. (2006). The Emergence of Up and Down States in Cortical Networks. *PLoS Comput Biol* **2**:e23.

Holder, W. (1668). An Account of an Experiment, Concerning Deafness, Communicated to the R. Society, by That Worthy and Learned Divine Dr. William Holder, as Followeth. *Phil Trans R Soc* **3**:665–668.

Holmgren, C., Harkany, T., Svennenfors, B., and Zilberter, Y. (2003). Pyramidal cell communication within local networks in layer 2/3 of rat neocortex. *J Physiol* **551**:139–153.

Hubel, D., Henson, C., Rupert, A., and Galambos, R. (1959). Attention units in the auditory cortex. *Science* **129**:1279–1280.

Hubel, D. H. (1957). Tungsten microelectrode for recording from single units. *Science* **125**:549–550.

Hubel, D. H. and Wiesel, T. N. (1977). Ferrier lecture. Functional architecture of macaque monkey visual cortex. *Proc R Soc Lond B Biol Sci* **198**:1–59.

Imig, T. J. and Adrián, H. O. (1977). Binaural columns in the primary field (A1) of cat auditory cortex. *Brain Res* **138**:241–257.

Jagadeesh, B., Gray, C. M., and Ferster, D. (1992). Visually evoked oscillations of membrane potential in cells of cat visual cortex. *Science* **257**:552–554.

Jagadeesh, B., Wheat, H. S., and Ferster, D. (1993). Linearity of summation of synaptic potentials underlying direction selectivity in simple cells of the cat visual cortex. *Science* **262**:1901–1904.

Jagadeesh, B., Wheat, H. S., Kontsevich, L. L., Tyler, C. W., and Ferster, D. (1997). Direction selectivity of synaptic potentials in simple cells of the cat visual cortex. *J Neurophysiol* **78**:2772–2789.

Johnson, J. A. and Zatorre, R. J. (2005). Attention to simultaneous unrelated auditory and visual events: behavioral and neural correlates. *Cereb Cortex* **15**:1609–1620.

Kasanetz, F., Riquelme, L. A., and Murer, M. G. (2002). Disruption of the two-state membrane potential of striatal neurones during cortical desynchronisation in anaesthetised rats. *J Physiol* **543**:577–589.

Katsuki, Y., Suga, N., and Kanno, Y. (1962). Neural mechanism of the peripheral and central auditory system in monkeys. *J Acoust Som Am* **34**:1396–1410.

Katsuki, Y., Sumi, T., Uchiyama, H., and Watanabe, T. (1958). Electric responses of auditory neurons in cat to sound stimulation. *J Neurophysiol* **21**:569–588.

Kayser, C., Petkov, C. I., Augath, M., and Logothetis, N. K. (2005). Integration of touch and sound in auditory cortex. *Neuron* **48**:373–384.

Kelly, J. B. and Sally, S. L. (1988). Organization of auditory cortex in the albino rat: binaural response properties. *J Neurophysiol* **59**:1756–1769.

Kepecs, A., Uchida, N., and Mainen, Z. F. (2006). The sniff as a unit of olfactory processing. *Chem Senses* **31**:167–179.

Kilgard, M. P. and Merzenich, M. M. (1998a). Cortical map reorganization enabled by nucleus basalis activity. *Science* **279**:1714–1718.

Kilgard, M. P. and Merzenich, M. M. (1998b). Plasticity of temporal information processing in the primary auditory cortex. *Nat Neurosci* **1**:727–731.

Kilgard, M. P. and Merzenich, M. M. (1999). Distributed representation of spectral and temporal information in rat primary auditory cortex. *Hear Res* **134**:16–28.

Konur, S., Rabinowitz, D., Fenstermaker, V. L., and Yuste, R. (2003). Systematic regulation of spine sizes and densities in pyramidal neurons. *J Neurobiol* **56**:95–112.

Kopec, C. D., Li, B., Wei, W., Boehm, J., and Malinow, R. (2006). Glutamate receptor exocytosis and spine enlargement during chemically induced long-term potentiation. *J Neurosci* **26**:2000–2009.

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.

Lakatos, P., Chen, C.-M., O'connell, M. N., Mills, A., and Schroeder, C. E. (2007). Neuronal oscillations and multisensory interaction in primary auditory cortex. *Neuron* **53**:279–292.

Lampl, I., Reichova, I., and Ferster, D. (1999). Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron* **22**:361–374.

Lang, C., Barco, A., Zablow, L., Kandel, E. R., Siegelbaum, S. A., and Zakharenko, S. S. (2004). Transient expansion of synaptically connected dendritic spines upon induction of hippocampal long-term potentiation. *Proc Natl Acad Sci U S A* **101**:16665–16670.

Langner, G., Albert, M., and Briede, T. (2002). Temporal and spatial coding of periodicity information in the inferior colliculus of awake chinchilla (Chinchilla laniger). *Hear Res* **168**:110–130.

Las, L., Stern, E. A., and Nelken, I. (2005). Representation of tone in fluctuating maskers in the ascending auditory system. *J Neurosci* **25**:1503–1513.

Laughlin, S. B. and Sejnowski, T. J. (2003). Communication in neuronal networks. *Science* **301**:1870–1874.

Lee, A. K., Manns, I. D., Sakmann, B., and Brecht, M. (2006). Whole-cell recordings in freely moving rats. *Neuron* **51**:399–407.

Levy, W. B. and Baxter, R. A. (1996). Energy efficient neural codes. *Neural Comput* **8**:531–543.

Lewicki, M. S. (2002). Efficient coding of natural sounds. *Nat Neurosci* **5**:356–363.

Lewis, B. L. and O'Donnell, P. (2000). Ventral tegmental area afferents to the prefrontal cortex maintain membrane potential 'up' states in pyramidal neurons via D(1) dopamine receptors. *Cereb Cortex* **10**:1168–1175.

Liang, L., Lu, T., and Wang, X. (2002). Neural representations of sinusoidal amplitude and frequency modulations in the primary auditory cortex of awake primates. *J Neurophysiol* **87**:2237–2261.

Limpert, E., Stahel, W. A., and Abbt, M. (2001). Log-normal distributions across the sciences: Keys and clues. *BioScience* **51**:341–352.

Linden, J. F., Liu, R. C., Sahani, M., Schreiner, C. E., and Merzenich, M. M. (2003). Spectrotemporal structure of receptive fields in areas AI and AAF of mouse auditory cortex. *J Neurophysiol* **90**:2660–2675.

Luo, M. and Katz, L. C. (2001). Response correlation maps of neurons in the mammalian olfactory bulb. *Neuron* **32**:1165–1179.

Léger, J.-F., Stern, E. A., Aertsen, A., and Heck, D. (2005). Synaptic integration in rat frontal cortex shaped by network activity. *J Neurophysiol* **93**:281–293.

Lütkenhöner, B., Lammertmann, C., Simões, C., and Hari, R. (2002). Magnetoencephalographic correlates of audiotactile interaction. *Neuroimage* **15**:509–522.

Machens, C., Wehr, M., and Zador, A. (2004). Linearity of cortical receptive fields measured with natural sounds. *J Neurosci* **24**:1089–1100.

Mahon, S., Deniau, J.-M., and Charpier, S. (2003). Various synaptic activities and firing patterns in cortico-striatal and striatal neurons in vivo. *J Physiol Paris* **97**:557–566.

Malenka, R. C. and Nicoll, R. A. (1999). Long-term potentiation–a decade of progress? *Science* **285**:1870–1874.

Maletic-Savatic, M., Malinow, R., and Svoboda, K. (1999). Rapid dendritic morphogenesis in CA1 hippocampal dendrites induced by synaptic activity. *Science* **283**:1923–1927.

Marder, E., Abbott, L. F., Turrigiano, G. G., Liu, Z., and Golowasch, J. (1996). Memory from the dynamics of intrinsic membrane currents. *Proc Natl Acad Sci U S A* **93**:13481–13486.

Margrie, T., Brecht, M., and Sakmann, B. (2002). In vivo, low-resistance, whole-cell recordings from neurons in the anaesthetized and awake mammalian brain. *Pflügers Arch* **444**:491–498.

Margrie, T. W. and Schaefer, A. T. (2003). Theta oscillation coupled spike latencies yield computational vigour in a mammalian sensory system. *J Physiol* **546**:363–374.

Markram, H. (1997). A network of tufted layer 5 pyramidal neurons. *Cereb Cortex* **7**:523–533.

Markram, H., Lübke, J., Frotscher, M., Roth, A., and Sakmann, B. (1997). Physiology and anatomy of synaptic connections between thick tufted pyramidal neurones in the developing rat neocortex. *J Physiol* **500 (Pt 2)**:409–440.

Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nat Rev Neurosci* **5**:793–807.

Matsuzaki, M., Honkura, N., Ellis-Davies, G. C. R., and Kasai, H. (2004). Structural basis of long-term potentiation in single dendritic spines. *Nature* **429**:761–766.

McCormick, D., Connors, B., Lighthall, J., and Prince, D. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. *J Neurophysiol* **54**:782–806.

McCormick, D. A., Shu, Y., Hasenstaub, A., Sanchez-Vives, M., Badoual, M., and Bal, T. (2003). Persistent cortical activity: mechanisms of generation and effects on neuronal excitability. *Cereb Cortex* **13**:1219–1231.

McGuire, B. A., Hornung, J. P., Gilbert, C. D., and Wiesel, T. N. (1984). Patterns of synaptic input to layer 4 of cat striate cortex. *J Neurosci* **4**:3021–3033.

Mehta, M. R., Lee, A. K., and Wilson, M. A. (2002). Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* **417**:741–746.

Mendelson, J. R., Schreiner, C. E., and Sutter, M. L. (1997). Functional topography of cat primary auditory cortex: response latencies. *J Comp Physiol [A]* **181**:615–633.

Mendelson, J. R., Schreiner, C. E., Sutter, M. L., and Grasse, K. L. (1993). Functional topography of cat primary auditory cortex: responses to frequency-modulated sweeps. *Exp Brain Res* **94**:65–87.

Merzenich, M. M., Knight, P. L., and Roth, G. L. (1973). Cochleotopic organization of primary auditory cortex in the cat. *Brain Res* **63**:343–346.

Merzenich, M. M., Knight, P. L., and Roth, G. L. (1975). Representation of cochlea within primary auditory cortex in the cat. *J Neurophysiol* **38**:231–249.

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.

Miesenböck, G. (2004). Genetic methods for illuminating the function of neural circuits. *Curr Opin Neurobiol* **14**:395–402.

Miesenböck, G., Angelis, D. A. D., and Rothman, J. E. (1998). Visualizing secretion and synaptic transmission with ph-sensitive green fluorescent proteins. *Nature* **394**:192–195.

Miller, J., Beaton, R., O'Connor, T., and Pfingst, B. (1974). Response pattern complexity of auditory cells in the cortex of unanesthetized monkeys. *Brain Res* **69**:101–113.

Miller, J., Sutton, D., Pfingst, B., Ryan, A., Beaton, R., and Gourevitch, G. (1972). Single cell activity in the auditory cortex of Rhesus monkeys: behavioral dependency. *Science* **177**:449–451.

Miller, J. M., Dobie, R. A., Pfingst, B. E., and Hienz, R. D. (1980). Electrophysiologic studies of the auditory cortex in the awake monkey. *Am J Otolaryngol* **1**:119–130.

Moshitch, D., Las, L., Ulanovsky, N., Bar-Yosef, O., and Nelken, I. (2006). Responses of neurons in primary auditory cortex (A1) to pure tones in the halothane-anesthetized cat. *J Neurophysiol* **95**:3756–3769.

Nakamura, K. (1999). Auditory spatial discriminatory and mnemonic neurons in rat posterior parietal cortex. *J Neurophysiol* **82**:2503–2517.

Nakamura, K. and Ono, T. (1986). Lateral hypothalamus neuron involvement in integration of natural and artificial rewards and cue signals. *J Neurophysiol* **55**:163–181.

Nakamura, K., Ono, T., and Tamura, R. (1987). Central sites involved in lateral hypothalamus conditioned neural responses to acoustic cues in the rat. *J Neurophysiol* **58**:1123–1148.

Nakamura, Y. and Katakura, N. (1995). Generation of masticatory rhythm in the brainstem. *Neurosci Res* **23**:1–19.

Nelken, I. (2004). Processing of complex stimuli and natural scenes in the auditory cortex. *Curr Opin Neurobiol* **14**:474–480.

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.

Nicholls, J. G., Wallace, B. G., Martin, A. R., and Fuchs, P. A., editors (2001). *From Neuron to Brain*. Sinauer Associates, Inc. Sunderland, Massachusetts, U.S.A. 4th edition.

Oatman, L. C. (1971). Role of visual attention on auditory evoked potentials in unanesthetized cats. *Exp Neurol* **32**:341–356.

Oatman, L. C. (1976). Effects of visual attention on the intensity of auditory evoked potentials. *Exp Neurol* **51**:41–53.

O'Connor, K. N., Petkov, C. I., and Sutter, M. L. (2005). Adaptive stimulus optimization for auditory cortical neurons. *J Neurophysiol* **94**:4051–4067.

O'Connor, S. M., Berg, R. W., and Kleinfeld, D. (2002). Coherent electrical activity between vibrissa sensory areas of cerebellum and neocortex is enhanced during free whisking. *J Neurophysiol* **87**:2137–2148.

Ohki, K., Chung, S., Ch'ng, Y. H., Kara, P., and Reid, R. C. (2005). Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature* **433**:597–603.

Ohl, F. W. and Scheich, H. (1997). Learning-induced dynamic receptive field changes in primary auditory cortex of the unanaesthetized mongolian gerbil. *J Comp Physiol [A]* **181**:685–696.

Ojima, H. and Murakami, K. (2002). Intracellular characterization of suppressive responses in supragranular pyramidal neurons of cat primary auditory cortex in vivo. *Cereb Cortex* **12**:1079–1091.

Olshausen, B. and Field, D. (2004). Sparse coding of sensory inputs. *Curr Opin Neurobiol* **14**:481–487.

Olshausen, B. A. and Field, D. J. (1997). Sparse coding with an overcomplete basis set: a strategy employed by V1? *Vision Res* **37**:3311–3325.

Ono, T. and Nakamura, K. (1985). Learning and integration of rewarding and aversive stimuli in the rat lateral hypothalamus. *Brain Res* **346**:368–373.

Ono, T., Nakamura, K., Nishijo, H., and Fukuda, M. (1986). Hypothalamic neuron involvement in integration of reward, aversion, and cue signals. *J Neurophysiol* **56**:63–79.

Ono, T., Sasaki, K., Nakamura, K., and Norgren, R. (1985). Integrated lateral hypothalamic neural responses to natural and artificial rewards and cue signals in the rat. *Brain Res* **327**:303–306.

Orduña, I., Mercado, E., Gluck, M. A., and Merzenich, M. M. (2001). Spectrotemporal sensitivities in rat auditory cortical neurons. *Hear Res* **160**:47–57.

Park, M., Salgado, J. M., Ostroff, L., Helton, T. D., Robinson, C. G., Harris, K. M., and Ehlers, M. D. (2006). Plasticity-induced growth of dendritic spines by exocytic trafficking from recycling endosomes. *Neuron* **52**:817–830.

Paré, D., Shink, E., Gaudreau, H., Destexhe, A., and Lang, E. J. (1998). Impact of spontaneous synaptic activity on the resting properties of cat neocortical pyramidal neurons in vivo. *J Neurophysiol* **79**:1450–1460.

Pei, X., Volgushev, M., Vidyasagar, T. R., and Creutzfeldt, O. D. (1991). Whole cell recording and conductance measurements in cat visual cortex in-vivo. *Neuroreport* **2**:485–488.

Perez-Orive, J., Mazor, O., Turner, G. C., Cassenaer, S., Wilson, R. I., and Laurent, G. (2002). Oscillations and sparsening of odor representations in the mushroom body. *Science* **297**:359–365.

Peters, A. and Payne, B. R. (1993). Numerical relationships between geniculocortical afferents and pyramidal cell modules in cat primary visual cortex. *Cereb Cortex* **3**:69–78.

Petersen, C., Hahn, T., Mehta, M., Grinvald, A., and Sakmann, B. (2003a). Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. *Proc Natl Acad Sci U S A* **100**:13638–13643.

Petersen, C. C. H., Grinvald, A., and Sakmann, B. (2003b). Spatiotemporal dynamics of sensory responses in layer 2/3 of rat barrel cortex measured in vivo by voltage-sensitive dye imaging combined with whole-cell voltage recordings and neuron reconstructions. *J Neurosci* **23**:1298–1309.

Petkov, C. I., Kang, X., Alho, K., Bertrand, O., Yund, E. W., and Woods, D. L. (2004). Attentional modulation of human auditory cortex. *Nat Neurosci* **7**:658–663.

Phillips, D. and Irvine, D. (1981). Responses of single neurons in physiologically defined primary auditory cortex (AI) of the cat: frequency tuning and responses to intensity. *J Neurophysiol* **45**:48–58.

Phillips, D. P., Hall, S. E., and Boehnke, S. E. (2002). Central auditory onset responses, and temporal asymmetries in auditory perception. *Hear Res* **167**:192–205.

Phillips, D. P., Semple, M. N., Calford, M. B., and Kitzes, L. M. (1994). Level-dependent representation of stimulus frequency in cat primary auditory cortex. *Exp Brain Res* **102**:210–226.

Picton, T. W., Hillyard, S. A., Galambos, R., and Schiff, M. (1971). Human auditory attention: a central or peripheral process? *Science* **173**:351–353.

Populin, L. C. (2005). Anesthetics change the excitation/inhibition balance that governs sensory processing in the cat superior colliculus. *J Neurosci* **25**:5903–5914.

Populin, L. C. (2006). Monkey sound localization: head-restrained versus head-unrestrained orienting. *J Neurosci* **26**:9820–9832.

Read, H. L., Winer, J. A., and Schreiner, C. E. (2002). Functional architecture of auditory cortex. *Curr Opin Neurobiol* **12**:433–440.

Recanzone, G. (2000). Response profiles of auditory cortical neurons to tones and noise in behaving macaque monkeys. *Hear Res* **150**:104–118.

Ribaupierre, F. D., Goldstein, M. H., and Yeni-Komshian, G. (1972). Intracellular study of the cat's primary auditory cortex. *Brain Res* **48**:185–204.

Rinberg, D., Koulakov, A., and Gelperin, A. (2006). Sparse odor coding in awake behaving mice. *J Neurosci* **26**:8857–8865.

Rumpel, S., LeDoux, J., Zador, A., and Malinow, R. (2005). Postsynaptic receptor trafficking underlying a form of associative learning. *Science* **308**:83–88.

Ruthazer, E. S., Akerman, C. J., and Cline, H. T. (2003). Control of axon branch dynamics by correlated activity in vivo. *Science* **301**:66–70.

Ryan, A. and Miller, J. (1978). Single unit responses in the inferior colliculus of the awake and performing rhesus monkey. *Exp Brain Res* **32**:389–407.

Ryan, A. F., Miller, J. M., Pfingst, B. E., and Martin, G. K. (1984). Effects of reaction time performance on single-unit activity in the central auditory pathway of the rhesus macaque. *J Neurosci* **4**:298–308.

Sachdev, R. N., Sellien, H., and Ebner, F. (2001). Temporal organization of multi-whisker contact in rats. *Somatosens Mot Res* **18**:91–100.

Sachdev, R. N. S., Ebner, F. F., and Wilson, C. J. (2004). Effect of subthreshold up and down states on the whisker-evoked response in somatosensory cortex. *J Neurophysiol* **92**:3511–3521.

Sally, S. L. and Kelly, J. B. (1988). Organization of auditory cortex in the albino rat: sound frequency. *J Neurophysiol* **59**:1627–1638.

Sanchez-Vives, M. V. and McCormick, D. A. (2000). Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nat Neurosci* **3**:1027–1034.

Schaefer, A. T., Angelo, K., Spors, H., and Margrie, T. W. (2006). Neuronal oscillations enhance stimulus discrimination by ensuring action potential precision. *PLoS Biol* **4**:e163.

Schreiner, C. E. (1995). Order and disorder in auditory cortical maps. *Curr Opin Neurobiol* **5**:489–496.

Schreiner, C. E. and Mendelson, J. R. (1990). Functional topography of cat primary auditory cortex: distribution of integrated excitation. *J Neurophysiol* **64**:1442–1459.

Schreiner, C. E., Read, H. L., and Sutter, M. L. (2000). Modular organization of frequency integration in primary auditory cortex. *Annu Rev Neurosci* **23**:501–529.

Schroeder, C. E. and Foxe, J. (2005). Multisensory contributions to low-level, 'unisensory' processing. *Curr Opin Neurobiol* **15**:454–458.

Schroeder, C. E., Lindsley, R. W., Specht, C., Marcovici, A., Smiley, J. F., and Javitt, D. C. (2001). Somatosensory input to auditory association cortex in the macaque monkey. *J Neurophysiol* **85**:1322–1327.

Scoville, W. B. and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *J Neurol Neurosurg Psychiatry* **20**:11–21. Reprinted in *J Neuropsychiatry Clin Neurosci* (2000) **12**(1):103–113.

Sellien, H., Eshenroder, D. S., and Ebner, F. F. (2005). Comparison of bilateral whisker movement in freely exploring and head-fixed adult rats. *Somatosens Mot Res* **22**:97–114.

Sen, K., Theunissen, F. E., and Doupe, A. J. (2001). Feature analysis of natural sounds in the songbird auditory forebrain. *J Neurophysiol* **86**:1445–1458.

Serkov, F. and Volkov, I. (1984). Intracellular reactions of neurons of the primary auditory area of the cerebral cortex in the cat to tones of different frequency and electric stimulation of nerve fibers of the spiral ganglion. *Neirofiziologiia* **16**:123–131.

Shi, S. H., Hayashi, Y., Petralia, R. S., Zaman, S. H., Wenthold, R. J., Svoboda, K., and Malinow, R. (1999). Rapid spine delivery and redistribution of AMPA receptors after synaptic NMDA receptor activation. *Science* **284**:1811–1816.

Smith, E. C. and Lewicki, M. S. (2006). Efficient auditory coding. *Nature* **439**:978–982.

Song, S., Sjöström, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol* **3**:e68.

Steriade, M., Nuñez, A., and Amzica, F. (1993a). A novel slow (<1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *J Neurosci* **13**:3252–3265.

Steriade, M., Nuñez, A., and Amzica, F. (1993b). Intracellular analysis of relations between the slow (<1 Hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. *J Neurosci* **13**:3266–3283.

Steriade, M., Timofeev, I., and Grenier, F. (2001). Natural waking and sleep states: a view from inside neocortical neurons. *J Neurophysiol* **85**:1969–1985.

Stern, E. A., Kincaid, A. E., and Wilson, C. J. (1997). Spontaneous subthreshold membrane potential fluctuations and action potential variability of rat corticostriatal and striatal neurons in vivo. *J Neurophysiol* **77**:1697–1715.

Stevens, C. F. and Zador, A. M. (1998). Input synchrony and the irregular firing of cortical neurons. *Nat Neurosci* **1**:210–217.

Stuart, G. J., Dodt, H. U., and Sakmann, B. (1993). Patch-clamp recordings from the soma and dendrites of neurons in brain slices using infrared video microscopy. *Pflügers Arch* **423**:511–518.

Stüttgen, M. C., Rüter, J., and Schwarz, C. (2006). Two psychophysical channels of whisker deflection in rats align with two neuronal classes of primary afferents. *J Neurosci* **26**:7933–7941.

Suga, N. and Ma, X. (2003). Multiparametric corticofugal modulation and plasticity in the auditory system. *Nat Rev Neurosci* **4**:783–794.

Sutter, M. L. and Schreiner, C. E. (1995). Topography of intensity tuning in cat primary auditory cortex: single-neuron versus multiple-neuron recordings. *J Neurophysiol* **73**:190–204.

Szyszka, P., Ditzen, M., Galkin, A., Galizia, C. G., and Menzel, R. (2005). Sparsening and temporal sharpening of olfactory representations in the honeybee mushroom bodies. *J Neurophysiol* **94**:3303–3313.

Tan, A. Y. Y., Zhang, L. I., Merzenich, M. M., and Schreiner, C. E. (2004). Tone-evoked excitatory and inhibitory synaptic conductances of primary auditory cortex neurons. *J Neurophysiol* **92**:630–643.

Theunissen, F. E. and Doupe, A. J. (1998). Temporal and spectral sensitivity of complex auditory neurons in the nucleus HVc of male zebra finches. *J Neurosci* **18**:3786–3802.

Thomson, A. M. and Bannister, A. P. (2003). Interlaminar connections in the neocortex. *Cereb Cortex* **13**:5–14.

Thomson, A. M., West, D. C., Wang, Y., and Bannister, A. P. (2002). Synaptic connections and small circuits involving excitatory and inhibitory neurons in layers 2-5 of adult rat and cat neocortex: triple intracellular recordings and biocytin labelling in vitro. *Cereb Cortex* **12**:936–953.

Timofeev, I., Grenier, F., and Steriade, M. (2001). Disfacilitation and active inhibition in the neocortex during the natural sleep-wake cycle: an intracellular study. *Proc Natl Acad Sci U S A* **98**:1924–1929.

Tollin, D. J., Populin, L. C., Moore, J. M., Ruhland, J. L., and Yin, T. C. T. (2005). Sound-localization performance in the cat: the effect of restraining the head. *J Neurophysiol* **93**:1223–1234.

Tunturi, A. R. (1944). Audio frequency localization in the acoustic cortex of the dog. *Am J Physiol* **141**:397–403.

Turrigiano, G. G., Leslie, K. R., Desai, N. S., Rutherford, L. C., and Nelson, S. B. (1998). Activity-dependent scaling of quantal amplitude in neocortical neurons. *Nature* **391**:892–896.

Uchida, N. and Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nat Neurosci* **6**:1224–1229.

Ulanovsky, N., Las, L., and Nelken, I. (2003). Processing of low-probability sounds by cortical neurons. *Nat Neurosci* **6**:391–398.

Vajnerová, O., Bielavská, E., Jiruška, P., and Brožek, G. (2003). Level of vigilance influences licking frequency in rats. *Physiol Res* **52**:243–249.

Vinje, W. E. and Gallant, J. L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**:1273–1276.

Volkov, I. and Galaziuk, A. (1985). Reactions of neurons of the auditory cortex of unanesthetized cats to tones of a characteristic frequency. *Neirofiziologiia* **17**:500–508.

Volkov, I. and Galaziuk, A. (1989). Reactions of tonic-type neurons in the cat auditory cortex to tones of various frequency and intensity. *Neirofiziologiia* **21**:498–506.

Wang, X., Lu, T., Snider, R. K., and Liang, L. (2005). Sustained firing in auditory cortex evoked by preferred stimuli. *Nature* **435**:341–346.

Wang, X. J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends Neurosci* **24**:455–463.

Wehr, M. and Zador, A. (2003). Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature* **426**:442–446.

Wehr, M. and Zador, A. M. (2005). Synaptic mechanisms of forward suppression in rat auditory cortex. *Neuron* **47**:437–445.

Weijnen, J. A. (1998). Licking behavior in the rat: measurement and situational control of licking frequency. *Neurosci Biobehav Rev* **22**:751–760.

Weliky, M., Fiser, J., Hunt, R. H., and Wagner, D. N. (2003). Coding of natural scenes in primary visual cortex. *Neuron* **37**:703–718.

Werner-Reiss, U., Kelly, K. A., Trause, A. S., Underhill, A. M., and Groh, J. M. (2003). Eye position affects activity in primary auditory cortex of primates. *Curr Biol* **13**:554–562.

Whishaw, I. Q. and Kolb, B., editors (2005). *The behavior of the laboratory rat: a handbook with tests*. Oxford University Press New York.

Whitfield, I. C. (1957). The electrical responses of the unanaesthetised auditory cortex in the intact cat. *Electroencephalogr Clin Neurophysiol Suppl* **9**:35–42.

Wiest, M. C. and Nicolelis, M. A. L. (2003). Behavioral detection of tactile stimuli during 7–12 Hz cortical oscillations in awake rats. *Nat Neurosci* **6**:913–914.

Wilent, W. B. and Contreras, D. (2005). Dynamics of excitation and inhibition underlying stimulus selectivity in rat somatosensory cortex. *Nat Neurosci* **8**:1364–1370.

Willmore, B. and Tolhurst, D. J. (2001). Characterizing the sparseness of neural codes. *Network* **12**:255–270.

Wilson, C. J. and Groves, P. M. (1981). Spontaneous firing patterns of identified spiny neurons in the rat neostriatum. *Brain Res* **220**:67–80.

Wilson, C. J. and Kawaguchi, Y. (1996). The origins of two-state spontaneous membrane potential fluctuations of neostriatal spiny neurons. *J Neurosci* **16**:2397–2410.

Winer, J. A., Miller, L. M., Lee, C. C., and Schreiner, C. E. (2005). Auditory thalamocortical transformation: structure and function. *Trends Neurosci* **28**:255–263.

Woolsey, C. and Walzl, E. (1942). Topical projection of nerve fibers from local regions of the cochlea to the cerebral cortex of the cat. *Bull Johns Hopkins Hosp* **71**:315–344.

Yaksi, E. and Friedrich, R. W. (2006). Reconstruction of firing rate changes across neuronal populations by temporally deconvolved Ca2+ imaging. *Nat Methods* **3**:377–383.

Zhang, L. I., Tan, A. Y. Y., Schreiner, C. E., and Merzenich, M. M. (2003). Topography and synaptic shaping of direction selectivity in primary auditory cortex. *Nature* **424**:201–205.

Zhu, J. J. and Connors, B. W. (1999). Intrinsic firing patterns and whisker-evoked synaptic responses of neurons in the rat barrel cortex. *J Neurophysiol* **81**:1171–1183.