

Shy Shoham · Daniel H. O'Connor · Ronen Segev

How silent is the brain: is there a “dark matter” problem in neuroscience?

Received: 21 November 2005 / Revised: 8 February 2006 / Accepted: 16 February 2006 / Published online: 21 March 2006
© Springer-Verlag 2006

Abstract Evidence from a variety of recording methods suggests that many areas of the brain are far more sparsely active than commonly thought. Here, we review experimental findings pointing to the existence of neurons which fire action potentials rarely or only to very specific stimuli. Because such neurons would be difficult to detect with the most common method of monitoring neural activity in vivo—extracellular electrode recording—they could be referred to as “dark neurons,” in analogy to the astrophysical observation that much of the matter in the universe is undetectable, or dark. In addition to discussing the evidence for largely silent neurons, we review technical advances that will ultimately answer the question: how silent is the brain?

Introduction

Sequences of neuronal action potentials are thought to be the basic symbols used to represent and transmit information in the nervous system. Properties of the ongoing pattern of action potentials, such as frequency, timing and possibly synchrony, underlie our ability to perceive, think about, and act upon the world. Despite their central importance, there are many questions about action potential activity in the brain that we cannot yet answer. One such question is basic: what is the mean firing rate of neurons in the brain? Or, put another way: how silent is the brain?

S. Shoham
Faculty of Biomedical Engineering,
Technion-Israel Institute of Technology,
Technion, Haifa 32000, Israel

D. H. O'Connor · R. Segev (✉) · S. Shoham
Department of Molecular Biology,
Princeton University, Princeton, NJ 08540, USA
E-mail: rsegev@princeton.edu
Tel.: +1-609-2580495
Fax: +1-609-2581028

Here, we review evidence that many neurons in the brain are far more silent than commonly thought. These silent neurons could be referred to as “dark neurons,” in analogy to the astrophysical observation that much of the matter in the universe is undetectable, or dark (for a different use of the term “dark matter” in neuroscience, see Binzegger et al. 2004). The presence of neurons that do not fire action potentials is counterintuitive, since electrical activity is central to neuronal signaling and to what distinguishes neurons from other cell types. Moreover, tissue is energetically expensive to construct (Raichle and Gusnard 2002), suggesting that “silent” neurons must play some necessary role in brain function even though they do not fire during typical experiments. But what could this role be?

The purpose of this review is to present a coherent perspective on the issue of silent neurons. To do so, we pull together related information from a number of studies employing a diverse set of recording and analysis methods, and dating back nearly four decades. The topic of silent neurons is particularly timely in light of recent technological advances in neural imaging and recording. In addition to reviewing the evidence for silent neurons, we highlight new techniques that offer the means to record from entire populations of neurons, and that are already providing results that may appear at odds with those of traditional methods.

Accumulating evidence for neural silence

Ever since the pioneering studies in the 1950s by Amassian (1953), Hubel (1957), and Mountcastle et al. (1957), the tool of choice used by in vivo experimental neurophysiologists for cellular recording has been the metal microelectrode. This workhorse of neurophysiology measures action potentials by detecting small electrical signals that occur outside a neuron resulting from induced current flow when the neuron fires an action potential. Consequently, the method is able to detect a neuron when it fires, and does not require physical

contact. However, with this method neurons that do not fire are not detected.

In 1968 David Robinson, in his review of single-unit electrode technology, pointed out a large discrepancy between the number of neurons present near a recording microelectrode and the number of neurons actually detected (Robinson 1968). Using theoretical calculations of the electrical fields surrounding a spiking neuron (Rall 1962), he estimated that along a 2-mm track in macaque monkey brain an electrode should encounter around 100 neurons. He observed, however, that in practice only a few neurons along the track are observed. He noted that “why this is so is a very disturbing question to users of microelectrodes.” His early suspicion has been supported by work in the decades since; now, evidence for such discrepancies comes from many brain areas (Table 1).

For instance, in anesthetized rats, Henze et al. (2000) recorded from hippocampal neurons by simultaneous intracellular and extracellular recording, thus allowing them to know by intracellular recording the exact position of a firing neuron. They could reliably detect neurons from at least 60 μm away, and on average they found six active neurons at a time. However, simple geometrical calculations based on anatomical data show

that within 60 μm of their extracellular recording electrodes, hippocampal area CA1 of rat would contain at least 120 neurons with a detectable electrical signature, a number far larger than six. The existence of a large number of “silent cells” has also been confirmed under these conditions by antidromic activation (Thompson and Best 1989), in which neurons were made to fire by evoking spikes in axons which then propagate back to the cell body. Thompson and Best (1989) also showed that apparent silence is unlikely to arise from the use of barbiturate anesthesia, which in area CA1 actually increases neural activity.

A comparison of extracellular recordings with patch-clamp recordings from the same brain regions suggests that extracellular recordings reflect a selection bias that inflates estimates of activity. In layers IV and II/III of rat barrel cortex, extracellularly recorded neurons appear to fire spontaneously or in response to principal whisker stimulation at a rate of 0.8–1.5 spikes per second (Brecht et al. 2003; Brecht and Sakmann 2002; Diamond et al. 1993). In contrast, whole-cell or cell-attached patch-clamp recording methods reveal firing rates nearly 10-fold lower, 0.05–0.15 spikes per second (Brecht et al. 2003; Brecht and Sakmann 2002; Margrie et al. 2002). The same research group has found a similar

Table 1 Evidence for silence or highly sparse activity among neurons

Brain area	Method	Evidence for sparse activity
Zebrafinch nucleus HVC	Extracellular recording with antidromic stimulation	Neurons projecting to song nucleus RA fire at most one burst per song motif (Hahnloser et al. 2002)
Rat hippocampus, area CA1	Comparison of intracellular and extracellular recordings	10- to 100-fold fewer neurons observed than expected from neuron density (Henze et al. 2002)
Rat hippocampus, area CA1	Single unit	Sixty-five percent of neurons active during anesthesia become silent during waking behavior (Thompson and Best 1989)
Rat barrel neocortex, layers 4 and 2/3	Blind selection of neurons by whole-cell patch-clamp recording	~10-fold lower spike rate measured with blind patch versus single-unit methods (Brecht et al. 2003; Brecht and Sakmann 2002)
Rat primary auditory neocortex	Blind cell-attached patch-clamp recording	An average of 1 spike/cell evoked in response to a tone (DeWeese et al. 2003)
Rat olfactory bulb	Whole-cell patch-clamp recording	10- to 50-fold lower spike rate measured with blind patch versus extracellular single-unit methods (Margrie et al. 2002)
Rat cerebellar granule cells	Blind patch	Average spontaneous firing rate of 0.5 Hz (Chadderton et al. 2004)
Rabbit primary motor neocortex	Extracellular with antidromic stimulation	Seventy-five percent of neurons silent during walking (Beloozerova et al. 2003)
Rabbit primary somatosensory neocortex	Extracellular with antidromic stimulation	Most neurons projecting within neocortex could not be made to spike with sensory stimulation (Swadlow and Hicks 1996)
Cat primary visual neocortex	Microelectrode array	~10-fold fewer neurons observed (Blanche et al. 2005)
Monkey neocortex	Comparison of recordings with known anatomy	Far fewer neurons are seen in extracellular recordings than expected from the spread of extracellular potential fields (Robinson 1968)
Human neocortex	Theoretical calculation of metabolic costs	Energetics limits average firing rate to 0.16 spikes/s, constraining the number of highly active cells to about 1 in 100 (Lennie 2003)

discrepancy in rat olfactory bulb mitral cells, in which patch-clamp recordings yield lower spike rates than do extracellular unit recordings (Margrie et al. 2002). These patch-clamp measurements may have their own selection biases, in particular toward the larger neurons that are more likely to be encountered during “blind” penetrations, and due to possible differences in the ability of the recording pipette to seal with the cellular membrane. However, patch-clamp methods do not select neurons on the basis of activity, and therefore may offer a more accurate estimate of true firing rate.

Another approach to measuring neuronal activity levels has been provided recently by a method that records from many neurons at once. Blanche et al. (2005) recorded activity from cat primary visual neocortex using a silicon microelectrode array. Arrays had recording sites spaced $\sim 50\text{--}75\ \mu\text{m}$ apart and were placed between neocortical layers II and VI. These arrays allowed the authors to record from 20 to 100 neurons simultaneously. How many neurons should have been observable by the array? The density of neurons in layers II–V of cat primary visual neocortex is $60,000\text{--}80,000\ \text{per mm}^3$ (Sholl 1956). Approximately $1\ \text{mm} \times 0.15\ \text{mm}$ array should cover, and therefore report activity from, about 700 neurons within the array boundaries (assuming a recording distance of $50\ \mu\text{m}$), and even more if significant numbers of neurons are detectable at a larger distance from the electrode array. In practice, however, only about 60 neurons were routinely observed. This suggests that the remaining neurons, over 90% of the total, were damaged or silent. Careful anatomical studies, and the fact that it was possible to record from the same neurons upon repeated advances and retractions of the electrode, suggest the latter.

In contrast to this large discrepancy in neocortex between the number of neurons that should be recordable and the number actually seen, recent recordings in retina (Segev et al. 2004) indicate that comparable silicon arrays with a planar multielectrode arrangement can in fact record 80–100% of neurons present near the array. In these recordings, the spacing of the multielectrode array was closely matched to the properties of the target ganglion cell layer tissue. In the tiger salamander retina where these recordings were performed, this layer is practically two dimensional. The planar nature of the tissue and the array electrode spacing resulted in each ganglion cell generating a unique pattern of activation on the array. The use of a special spike sorting technique then allowed the efficient detection and sorting of spikes into their unique neuronal origins. The ganglion cells were also counted directly after retrograde labeling with a fluorescent dye placed at the stump of the optic nerve. Comparison of these values with the number of cells present in the extracellular signal indicates that in the retina, nearly all ganglion cells are active and detectable by extracellular recording. On the one hand, this study in retina suggests that electrode arrays may be capable of recording all the neurons within the expected recordable volume and that the silence seen in neocortical

array studies therefore evidences truly silent neurons. On the other hand, there may be important differences in tissue–electrode interactions between retina and neocortex that make extrapolation between the two recording locations difficult.

Why so many silent neurons?

What explains the large discrepancy between the number of neurons expected and the number seen in so many extracellular recordings? In one scenario, neurons might be missed simply because of technical limitations. For instance, neurons could be silenced by recording electrode-induced damage or shielded from the electrode by glial walls that impede electrical current flow (Robinson 1968). Even the method of analysis could, to some degree, cause the spurious appearance of silence. For instance, during extracellular recording the raw signal needs to be sorted into single unit activity (Towe and Harding 1970). In many cases, this is done by dimension-reduced cluster analysis, for example, after plotting spike width versus spike amplitude. With these methods neurons that fire at very low rates usually do not form substantial clusters and are likely to be missed. In such cases, there is a sampling bias caused by the method of analysis.

However, these technical limitations are unlikely to offer a complete explanation of apparent neural silence. In their hippocampal recordings, Henze et al. (2000) pointed out that in area CA1 the ratio of recorded pyramidal cells to recorded interneurons, 6 to 1, is much lower than the relative abundance of these neurons, 33 to 1, determined from anatomical studies. Thus, pyramidal cells appear to be disproportionately likely to be missed. Assuming that most interneurons are recorded (a possibility, since hippocampal interneurons fire tonically at consistently high rates; Freund and Buzsaki 1996), this discrepancy implies that 6/33, or less than 20% of pyramidal neurons, are detected by extracellular recording. This estimate is consistent with the abundance of silent neurons observed using intracellular recording.

A second class of explanations assumes that there is indeed a large population of very quiet neurons. These silent neurons may be silent over long or short periods of time. Silent neurons may simply have sparse stimulus selectivities (Olshausen and Field 2004); that is, respond only to very specific stimuli. Such stimuli may be sufficiently specific that they have not been routinely used in experiments. Silent neurons could also be available for recruitment to higher activity levels through plasticity mechanisms.

Neuronal silence may well be the norm in the neocortex. Theoretical calculations have been used to compare total neocortical energy consumption to the cost of single spikes (Lennie 2003; Olshausen and Field 2004). Such analysis suggests that human neocortical neurons fire at very low baseline rates of up to 0.16

spikes per second on average (Lennie 2003). This low average rate places a strong bound on the number of highly active neurons such as those commonly studied in perceptual and cognitive tasks, which have been reported to fire tens of spikes per second. Although such an energetic calculation can be sensitive to assumptions, the discrepancy suggests that highly active neurons could constitute as few as 1 in 100 neurons.

Why we need to know how silent the brain is

Studies suggesting the existence of silent neurons pose a fundamental challenge to our understanding of brain physiology. Have neuroscientists routinely been recording from only the most active neurons, which constitute a small minority? And, if so, how representative are these neurons aside from their higher firing rates? Our understanding of how information is represented in the brain depends critically on the answers to these questions.

Many current models of cortical neurons are based on experimental observations of broad ‘tuning curves’ where an individual neuron responds similarly to a broad range of stimuli and experimental conditions, gradually modulating its response as a function of the experimentally controlled variables (Georgopoulos et al. 1982; Maunsell and Van Essen 1983; O’Keefe and Dostrovsky 1971). These observations have led to the popularity of population coding models (Georgopoulos et al. 1986; Pouget et al. 2003), where a reliable representation by a population replaces the coarse noisy coding of an individual neuron. However, the apparent sparseness of neuronal activity—that is, the apparent presence of neurons that are silent most of the time or fire at extremely low rates—may render greater plausibility to other well-known models where information is represented in the selection of which small subset of neurons is active (out of a large set of silent neurons). Examples of sparsely active models are common in the artificial pattern recognition literature (Rosenblatt 1962) and include, for example, Marr’s model of efficient encoding by cerebellar Purkinje neuron activation (Marr 1969). In the neocortex, encoding of complex stimulus features by single neocortical neurons (Gross 2002) led to a suggestion by Barlow (1972) that “at progressively higher levels in sensory pathways information about the physical stimulus is carried by progressively fewer active neurons.” A compelling example of this behavior was demonstrated recently by Quiroga et al. (2005) who described neurons in the human brain that were highly selective to particular categories of stimuli, such as to pictures of the actress Halle Berry. Interestingly, Barlow did not mention the experimental difficulty that was described earlier by Robinson (1968), and a link between sparse neural codes (i.e. neural codes that rely on neurons being extremely selective to one special feature) and experimental recording bias has only recently been drawn by Olshausen and Field (2005).

The possibility that in some brain regions only a sparse subset of neurons are active at once also opens up new possibilities for how information might be stored. In neocortex, synaptic connectivity is highly nonrandom (Song et al. 2005), and particular sequences of neural activity can be highly favored (Yuste et al. 2005). This opens the possibility that which specific connections are made may be a means of storing information (Stepanyants et al. 2002). Selection of possible connections from a large space of possibilities may thus be a means of learning patterns (Marr 1969) and sequences (Fiete et al. 2004). Speculatively, silent neurons may be neurons that are currently not in use or rarely used, but awaiting potential future use.

A better understanding of baseline activity patterns is also a crucial component of our understanding of brain energetics (Lennie 2003) and of the basis of collective signals like the blood oxygenation level dependent (BOLD) signal in functional MRI (Heeger et al. 2000). Finally, if indeed most neurons are silent or mostly silent, what are the implications for our understanding of the anomalous activity patterns during epileptic seizures? Does recruitment of silent neurons play as much of a role in the pathology as hyper-excitability of the firing neurons? Does the proportion of silent neurons play a role in the predisposition of different brain regions to epilepsy?

Towards a definitive answer

The existing experimental evidence indicates that the true overall average level of action potential activity can sometimes be far lower than estimates based on extracellular measurements. This discrepancy could be explained if extracellular neuronal recording detects mainly a small but active minority of the neurons present. How can such a bias be safely avoided in the future? A variety of techniques are becoming available.

Evoking activity in silent neurons

One way to eliminate the bias toward highly active cells is to force neurons to fire. This has been done using antidromic stimulation (Hahnloser et al. 2002; Swadlow 1998). This method allows extracellular recordings to be established from neurons regardless of whether they are active or not, as long as they are activated by an action potential evoked in their axons.

Interestingly, the use of antidromic stimulation has revealed populations of silent neurons (and therefore highly sparse activity patterns) in motor and somatosensory areas. In a series of recordings from primary motor neocortex of walking rabbits, Swadlow and colleagues (Beloozerova et al. 2003) have observed anatomically segregated populations of neurons in which up to 66% of the cells were essentially silent at rest (less than one spike per minute), and up to 75%

were silent during walking. Similar results have been seen in the primary somatosensory neocortex of awake rabbits (Swadlow and Hicks 1996). Highly sparse activity patterns have also been found in nucleus HVC (high vocal center) of singing zebra finches following antidromic identification of neurons projecting to RA (the robust nucleus of the arcopallium) (Hahnloser et al. 2002).

Sampling both silent and active neurons

A different sampling method that selects by a criterion other than firing is patch-clamp recording. Instead, patch recording selects for neurons that can seal to the recording electrode (Brecht et al. 2003). This selection criterion opens the possibility of recording from neurons that fire no spikes before the establishment of the recording (though, as noted above, the method may be biased towards cells with big somas). Patch-clamp recordings have shown that sensory input can evoke very small numbers of spikes in the rat primary auditory neocortex (DeWeese et al. 2003), and in the granule cells of the cerebellum (Chadderton et al. 2004).

Recording from all neurons at once

A very direct approach would be simply to record from many neurons in a region at once. If we can




Fig. 1 Combining electrophysiology and imaging may help solve the dark matter problem. The multi-electrode array used by Segev et al. (2004) in conjunction with anatomical imaging to record from 80 to 100% of the ganglion cells in a patch of the retina. The electrodes appear as a grid of *black circles*. Fluorescently labeled ganglion cells and their optic nerve axons are shown in *green*. Directly observing the neurons present in a patch of neural tissue while recording activity may yield a more accurate picture of population activity

record from almost all the neurons in a specific region, then we can directly verify or disprove the existence of a population of silent neurons. In order to know that one is monitoring all the neurons in an area, one must be able to visualize all neurons independent of their activity. This can be done by imaging neurons labeled with fluorescent dyes (Brecht et al. 2004), or by combining dense electrode array recordings and fluorescence imaging (Segev et al. 2004; see also Fig. 1 for more details). With activity-dependent indicators such as calcium-sensitive dyes (Brecht et al. 2004), both spontaneous and evoked neural activity may be sampled across populations of neurons. Thus far, an effective means of labeling neurons *in vivo* has been bulk loading with the acetoxymethyl-(AM)-ester form of synthetic indicators followed by two-photon fluorescence scanning microscopy (Stosiek et al. 2003). This method has been applied to monitor activity in primary visual cortex (Ohki et al. 2005) and cerebellum (Sullivan et al. 2005). In the case of primary visual cortex, this method has shown that under a variety of visual stimulus patterns, only 25–75% of rat and 63% of cat dye-loaded neurons showed measurable calcium signals (Ohki et al. 2005), indicating that this sensory region's neural activity may have an intermediate level of sparseness. Expression of calcium-sensing proteins based on green fluorescent protein (Wang et al. 2003) may improve labeling and allow cell-type specificity (Hasan et al. 2004).

A related approach is to combine electrode array recordings with fluorescence imaging in order to match the extracellularly recorded neurons with their physical location. This could be done by ablating cells one by one. For example, the activity of the whole population could be recorded extracellularly, followed by the targeted ablation of one fluorescent cell. A second phase of recording from the whole population would then be used to find which cell is absent. Repeating this procedure many times would allow matching all the recorded cells with their locations in space. Such a method would also reveal the presence of silent neurons that appose the array but do not produce spikes.

Another important issue that can be solved is whether there are active cells in the multi-electrode array catchment area that are not recorded by the array due to technical limitations. This can be addressed by combining intracellular electrode recordings with array recordings to verify that the cells that appear to be silent are indeed not active.

Recording with complex stimuli in awake animals

As pointed out by Olshausen and Field (2005) with respect to visual cortex, many studies use a combination of reduced stimuli useful for characterizing simple (usually linear) systems, along with a sequential rather than continuous recording method, often in anesthetized animals. Although these features of experimental design

have been chosen for good reasons, the net result may be a biased sample of neuronal activity. We expect that fluorescence imaging techniques and dense multi-electrode arrays, used to record continuously from a large fraction of specific neuronal populations in awake animals, will ultimately answer the question of how active the brain is. For the task of inferring neural activity levels in the behaving animal, it will be of obvious importance to use natural stimuli and complex behavioral paradigms with awake animals rather than a more reduced preparation.

Conclusions

Is the existing body of evidence sufficient to provide a reasonable estimate of the fraction of silent cells? Table 1 suggests that such proportions may vary widely among different brain regions and preparations, a notion which is consistent with hierarchical, increasingly sparse neural coding schemes. Conservative estimates may, however, be possible by considering those parameters of the neuron–electrode interface that affect the detection of unit signals: signal magnitudes very close to the neuron membrane, the signal decay constant, and the signal-to-noise level above which one can reliably identify neurons. Extracellular signals around open-field neurons appear to drop roughly exponentially with distance, a prediction (Rall 1962) that is supported by observations in different systems (Gray et al. 1995; Segev et al. 2004). The signal is characterized by $V(r) = V_0 e^{-r/\lambda}$, where $V(r)$ is the signal at distance r from the cell soma, V_0 is the signal very close to the soma, and λ is the decay constant. Interestingly, the decay constant was measured to be 28.42 μm in the (three dimensional) cortex (Gray et al. 1995) and 28 μm in the (nearly two-dimensional) retina (Segev et al. 2004), so it appears reasonably safe to assume that 28 μm is a representative value. Likewise, both theoretical calculations (Holt and Koch 1999) and experimental measurements from a number of different systems (e.g. Gray et al. 1995; Henze et al. 2000; Segev et al. 2001, 2004) suggest that V_0 generally has a value larger than 1 mV. As a conservative estimate we will use $V_0 \approx 0.5$ mV. Finally, a conventional conservative choice for a sufficient signal level (3–4 times above the noise floor) appears to be in the 50–60 μV range (Abeles 1991; Henze et al. 2000; Segev et al. 2004).

Combining these parameters together leads to the following estimate for the ‘recordable radius’ around an electrode: $R = -\lambda \log(60 \mu\text{V}/500 \mu\text{V}) \sim 59 \mu\text{m}$. This radius is in good agreement with the 50–140 μm range suggested by Henze et al. using combined intracellular and extracellular recordings in the rat hippocampus. Using this radius, we can estimate the fraction of recordable neurons by an extracellular electrode: $F = N_e / (4/3\pi R^3 D)$, where D is the neuron density of the specific

brain region under consideration, and N_e is the number of observed neurons on the extracellular electrode. Applying this formula to cat primary visual cortex, where the neuron density is $> 60,000$ cells/ mm^3 , we should expect to find consistently at least 50 neurons in the recorded signal from a single microelectrode, suggesting a silent fraction of at least 90%.

To summarize, the existence of large populations of silent neurons has been suggested recently by experimental evidence from diverse systems. Only some regions and neuron types show this phenomenon: as counterexamples, interneurons and cerebellar Purkinje cells are active most or all of the time. Nonetheless, the diversity of cases in which many neurons appear to be silent includes major neuron types in the mammalian neocortex and hippocampus, the cerebellum, and the zebra finch song system. Silent neurons may be a recurring principle of brain organization.

The likelihood that the brain contains a large fraction of silent neurons suggests that information processing may depend in part on which subset of neurons is active. In this scenario many neurons are silent most of the time, and a given stimulus triggers spikes only in a highly selective, small subset of neurons. While sparse activity patterns could in principle be interpreted using vector-type population decoding schemes (Georgopoulos et al. 1986; Pouget et al. 2003), neural mechanisms that are tuned to synchrony (Abeles 1991), or to co-spiking timing patterns (Abeles and Gerstein 1988; Rieke 1997) are likely to be more powerful for reading sparse activity patterns, and provide a substrate for binding different features of a particular stimulus (Singer 1999). From existing evidence, it seems likely that the coding strategy—and level of sparseness—may vary among brain areas. Thus, determining the level of silence in a given area may provide hints to how information is coded in that area. Although sparse neural representations pose a special challenge to neurophysiologists, recent technical advances may shed light on the question of how silent the brain is, and thereby advance our understanding of how nervous systems encode information.

Note added in proof: A recent study (Kerr et al. PNAS 2005;102(39):14063-8) reported extremely sparse spontaneous activity patterns in rat motor cortical neurons using AM-ester bulk loading and electrical patch recordings. Although nearly 90% of the imaged neurons appear to be spontaneously active, over half of these neurons averaged less than one action potential every 20 seconds, which was the overall average rate for the population. It is highly likely that most of these neurons will go unnoticed in typical electrophysiological recording sessions.

Acknowledgements We thank Sam Wang for fruitful discussions during the preparation of this manuscript, and two anonymous reviewers for their helpful comments. We also thank Joe Goodhouse and Michael J. Berry II for allowing us to use the figure.

References

- Abeles M (1991) *Corticonics: neural circuits of the cerebral cortex*. Cambridge University Press, Cambridge
- Abeles M, Gerstein GL (1988) Detecting spatiotemporal firing patterns among simultaneously recorded single neurons. *J Neurophysiol* 60:909–924
- Amassian VE (1953) Evoked single cortical unit activity in the somatic sensory areas. *Electroencephalogr Clin Neurophysiol Suppl* 5:415–438
- Barlow HB (1972) Single units and sensation: a neuron doctrine for perceptual psychology? *Perception* 1:371–394
- Beloozerova IN, Sirota MG, Swadlow HA (2003) Activity of different classes of neurons of the motor cortex during locomotion. *J Neurosci* 23:1087–1097
- Binzegger T, Douglas RJ, Martin KA (2004) A quantitative map of the circuit of cat primary visual cortex. *J Neurosci* 24:8441–8453
- Blanche TJ, Spacek MA, Hetke JF, Swindale NV (2005) Poly-trodes: high-density silicon electrode arrays for large-scale multiunit recording. *J Neurophysiol* 93:2987–3000
- Brecht M, 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, Roth A, 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, Fee MS, Garaschuk O, Helmchen F, Margrie TW, Svoboda K, Osten P (2004) Novel approaches to monitor and manipulate single neurons in vivo. *J Neurosci* 24:9223–9227
- Chadderton P, Margrie TW, Hausser M (2004) Integration of quanta in cerebellar granule cells during sensory processing. *Nature* 428:856–860
- Diamond ME, Armstrong-James M, Ebner FF (1993) Experience-dependent plasticity in adult rat barrel cortex. *Proc Natl Acad Sci USA* 90(5):2082–2086
- DeWeese MR, Wehr M, Zador AM (2003) Binary spiking in auditory cortex. *J Neurosci* 23:7940–7949
- Fiete IR, Hahnloser RH, Fee MS, Seung HS (2004) Temporal sparseness of the premotor drive is important for rapid learning in a neural network model of birdsong. *J Neurophysiol* 92:2274–2282
- Freund TF, Buzsáki G (1996) Interneurons of the hippocampus. *Hippocampus* 6:347–470
- Georgopoulos AP, Kalaska JF, Caminiti R, Massey JT (1982) On the relations between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J Neurosci* 2:1527–1537
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233:1416–1419
- Gray CM, Maldonado PE, Wilson M, McNaughton B (1995) Tetraodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *J Neurosci Methods* 63:43–54
- Gross CG (2002) Genealogy of the “grandmother cell”. *Neuroscientist* 8:512–518
- Hahnloser RH, Kozhevnikov AA, Fee MS (2002) An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature* 419:65–70
- Hasan MT, Friedrich RW, Euler T, Larkum ME, Giese G, Both M, Duebel J, Waters J, Bujard H, Griesbeck O, Tsien RY, Nagai T, Miyawaki A, Denk W (2004) Functional fluorescent Ca²⁺ indicator proteins in transgenic mice under TET control. *PLoS Biol* 2:e163
- Heeger DJ, Huk AC, Geisler WS, Albrecht DG (2000) Spikes versus BOLD: what does neuroimaging tell us about neuronal activity? *Nat Neurosci* 3:631–633
- Henze DA, Borhegyi Z, Csicsvari J, Mamiya A, Harris KD, Buzsáki G (2000) Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *J Neurophysiol* 84:390–400
- Holt GR, Koch C (1999) Electrical interactions via the extracellular potential near cell bodies. *J Comput Neurosci* 6:169–184
- Hubel DH (1957) Tungsten microelectrode for recording from single units. *Science* 125:549–550
- Lennie P (2003) The cost of cortical computation. *Curr Biol* 13:493–497
- Margrie TW, Brecht M, Sakmann B (2002) In vivo, low-resistance, whole-cell recordings from neurons in the anaesthetized and awake mammalian brain. *Pflügers Arch* 444:491–498
- Marr D (1969) A theory of cerebellar cortex. *J Physiol* 202:437–470
- Maunsell JH, Van Essen DC (1983) Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *J Neurophysiol* 49:1127–1147
- Mountcastle VB, Davies PW, Berman AL (1957) Response properties of neurons of cat’s somatic sensory cortex to peripheral stimuli. *J Neurophysiol* 20:374–407
- Ohki K, Chung S, Ch’ng YH, Kara P, Reid RC (2005) Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex. *Nature* 433:597–603
- O’Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* 34:171–175
- Olshausen BA, Field DJ (2004) Sparse coding of sensory inputs. *Curr Opin Neurobiol* 14:481–487
- Olshausen BA, Field DJ (2005) How close are we to understanding v1? *Neural Comput* 17:1665–1699
- Pouget A, Dayan P, Zemel RS (2003) Inference and computation with population codes. *Annu Rev Neurosci* 26:381–410
- Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435:1102–1107
- Raichle ME, Gusnard DA (2002) Appraising the brain’s energy budget. *Proc Natl Acad Sci USA* 99:10237–10239
- Rall W (1962) Electrophysiology of a dendritic neuron model. *Biophys J* 2(2Pt 2):145–167
- Rieke F (1997) *Spikes: exploring the neural code*. MIT Press, Cambridge
- Robinson DA (1968) The electrical properties of metal microelectrodes. *Proc IEEE* 56:1065–1071
- Rosenblatt F (1962) *Principles of neurodynamics*. Spartan, New York
- Segev R, Shapira Y, Benveniste M, Ben-Jacob E (2001) Observations and modeling of synchronized bursting in two-dimensional neural networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 64:011920
- Segev R, Goodhouse J, Puchalla J, Berry MJ II (2004) Recording spikes from a large fraction of the ganglion cells in a retinal patch. *Nat Neurosci* 7:1154–1161
- Sholl DA (1956) *The organization of the cerebral cortex*. Methuen, London
- Singer W (1999) Neuronal synchrony: a versatile code for the definition of relations? *Neuron* 24:49–65, 111–125
- Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB (2005) Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol* 3:e68
- Stepanyants A, Hof PR, Chklovskii DB (2002) Geometry and structural plasticity of synaptic connectivity. *Neuron* 34:275–288
- Stosiek C, Garaschuk O, Holthoff K, Konnerth A (2003) In vivo two-photon calcium imaging of neuronal networks. *Proc Natl Acad Sci USA* 100:7319–7324
- Sullivan MR, Nimmerjahn A, Sarkisov DV, Helmchen F, Wang SS (2005) In vivo calcium imaging of circuit activity in cerebellar cortex. *J Neurophysiol* 94:1636–1644
- Swadlow HA (1998) Neocortical efferent neurons with very slowly conducting axons: strategies for reliable antidromic identification. *J Neurosci Methods* 79:131–141

- Swadlow HA, Hicks TP (1996) Somatosensory cortical efferent neurons of the awake rabbit: latencies to activation via supra- and subthreshold receptive fields. *J Neurophysiol* 75:1753–1759
- Thompson LT, Best PJ (1989) Place cells and silent cells in the hippocampus of freely-behaving rats. *J Neurosci* 9:2382–2390
- Towe AL, Harding GW (1970) Extracellular microelectrode sampling bias. *Exp Neurol* 29:366–381
- Wang JW, Wong AM, Flores J, Vossell LB, Axel R (2003) Two-photon calcium imaging reveals an odor-evoked map of activity in the fly brain. *Cell* 112:271–282
- Yuste R, Maclean JN, Smith J, Lansner A (2005) Opinion: the cortex as a central pattern generator. *Nat Rev Neurosci* 6:477–483