Available online at www.sciencedirect.com
SCIENCE DIRECT®
Physics of Life Reviews 1 (2004) 71–102
PHYSICS of LIFE reviews
www.elsevier.com/locate/plrev

Neural coding: computational and biophysical perspectives

Gabriel Kreiman

Center for Biological and Computational Learning, McGovern Institute for Brain Research, Massachusetts Institute of Technology, 45 Carleton Street, MIT E25-201B, Cambridge, MA 02142, USA

Received 7 June 2004; accepted 7 June 2004

Communicated by L. Perlovsky

Abstract

While recognizing a face or kicking a ball may seem to be easy tasks for us, they still constitute challenging problems for even the most sophisticated computer algorithms available nowadays. The brain has evolved complex mechanisms to encode behaviorally relevant information. Here we review the types of codes used by the brain, what their constraints are and how they map the sensory environment or the motor output. We start by defining neural codes and briefly describing some of the current tools available to record activity from the brain. We give several examples of coding strategies used by different systems and multiple organisms and discuss how spiking patterns can be read out. Going beyond correlations between physiology and stimuli, we show what is currently known about the direct causal link between neuronal responses and behavioral output or sensory input. Finally, we identify what we consider to be some of the pressing questions in the field.

Contents

E-mail address: kreiman@mit.edu (G. Kreiman).

1571-0645/$ – see front matter
doi:10.1016/j.plrev.2004.06.001

1. Introduction and scope of the review

It seems effortless for us to see, hear, smell and thus interpret what is outside in the environment. Yet these processes involve complex interactions of large numbers of neurons. A three-year old performs better at face recognition than the best computational algorithms available today. The intricate stimuli that we are constantly subject to are represented by the neuronal activity in the approximately 10^{11} neurons in our brains. Our internal thoughts, emotions and memories are also represented by these neurons and so are the commands sent to the muscles for motor output. One of the main goals of Systems Neuroscience is to try to comprehend how neuronal activity represents sensory input, cognitive processes and motor output. However attractive the idea of a little homunculus in our brains trying to interpret the world and issuing commands may be, our task is to explain these processes using only sensory receptors and neuronal signals that are sent to other neurons, read by neurons and submitted to muscles for output.

Studies about coding in the nervous system have a long history. Perhaps one of the most influential ideas came about with the realization that different brain areas are involved in processing different types of signals [1]. A lucid review of the field of neural coding was published by Perkel and Bullock over thirty years ago [2]. Interestingly, they wondered back then whether the code was about to be broken. There have been major advances in our understanding of coding in the brain as illustrated by many theme-specific reviews related to neural coding [3–14]. However, many questions about coding are more complicated that they thought back then and many mysteries remain.

We will start by asking what we mean by "coding" in the context of the nervous system and what would constitute evidence that we can decipher particular types of codes. We will not devote much effort to enumerating theoretically plausible codes (there are too many). We focus on patterns that are actually observed in the nervous systems of different species and for which we can find neurophysiological evidence. Embedded in the idea of a code is the notion that some aspects of neuronal activity can be discriminated as a signal while other aspects may be due to noise. The distinction between what constitutes a signal and what constitutes noise will be a fundamental one. We are particularly interested in trying to extract general principles that govern the encoding of information in the nervous system, the mechanisms that are responsible for such neuronal patterns of activation and the way that those patterns are read out. However, this does not imply that there is a single answer, a single universal code as in the case of the mapping from nucleotides in DNA to amino acids in proteins [15].

We will draw examples from multiple species and systems throughout the review but we would like to make some of our biases explicit at this point. We are particularly interested in the visual system of primates (particularly the higher stages of the so-called ventral pathway in macaque monkeys and humans). Therefore, we hope the reader will excuse us if there are more examples pertaining to the visual system than from other sensory systems, cognitive processes or the motor side. The examples cited in later sections only constitute a small fraction selected for clarity. The examples do not attempt to constitute an exhaustive list of all possible coding schemes nor of all possible examples for particular coding schemes.

Throughout the review, we will focus on spiking patterns from one or more neurons. As we will discuss below, spikes are not the only way of neuronal communication nor are they the only variables one could study to identify neuronal codes. For example, photoreceptor neurons in the retina do not fire spikes and encode the intensity of light in the form of graded potentials [16,17]. A special section of the review will be devoted to understanding the state of the art of what can be measured and the different trade-offs of different techniques. Coding can be studied at different levels of specificity and we still do not know what the most appropriate levels are. The annual meeting of the Society for Neuroscience is typically attended by approximately 30 000 scientists. Some are concerned with the detailed structure of ionic channels at the atomic level. Others present models and data averaging the activity of millions of neurons. In the same way that it is not necessary to write complex quantum mechanics equations to describe the movement of a ball in free fall, it is likely that there is a most appropriate scale to understand coding in the nervous system.

2. Codes and features

We can generally think of a code as a way of mapping two spaces. In the particular case that we are discussing here, the stimuli in the outside world are represented by the activity of neurons in our brains. Most of the basic questions that one can ask about maps in general are still unanswered in the case of the brain. Is it a one-to-one map (a particular neuronal signature for every stimulus) or many-to-one or one-to-many? To what extent is it a deterministic or stochastic map? What is the domain (in stimulus space and neuronal space) of the map? We start by constraining the definition of codes to the field of Neuroscience.

2.1. What is a code and what needs to be explained?

A code can be a substitution scheme (e.g., a coded language where each letter is substituted by the following one as in "Ofvsptdjfodf"). The code(s) in the nervous system are somewhat different: (i) they are not necessarily static ("n" always maps to "o" in the above example), (ii) it is unclear whether the number of symbols is fixed (and if they are fixed the number of combinations may be very large), (iii) neural codes show considerable levels of noise (e.g., "n" sometimes may map to "o", other times to "n" and other times to "m" with different frequencies). Some of these discrepancies are addressed by Shannon's information theoretic approach [18]. Indeed information theory has been used quite successfully in many examples in Neuroscience [19–21]. One should be cautious, however, about the interpretation of these results. In general, the ability to decode a signal from neuronal spike trains does not necessarily imply that this decoding process is used by the nervous system. While the number of bits conveyed by neurons about sensory stimuli can be very high and approximate the theoretical limits [7], it is still hard to establish a direct link between information transmission rate and behavior.

Encoding the information through multiple stages of processing may be useful for an organism to compress information, get rid of behaviorally irrelevant variables and extract the particular stimulus features that are important to understand the environmental signals. For example, a neuron in inferior temporal cortex might be interested in achieving a representation of complex stimuli (such as a face) in a way that is invariant to changes in position, size, illumination, and even rotation in some cases.

What would it mean to understand the neural code? We would like to be able to predict the neuronal activity from a given stimulus and guess which stimulus was presented based on the neuronal activity. An analogy may be useful. It is common to store digital images in a computer. Given the pixels, there are several different ways of storing the information into the pattern of zeros and ones that computers understand. If we take one particular format, say JPEG, we know how to convert pixels to 0s and 1s and we know how to read the 0s and 1s to re-create the digital image.

The question is more complicated in the case of neural coding. We generally do not know the equivalent of the pixels, that is, exactly what is being encoded in different areas of the brain. And there are potentially several million neurons that may be involved in the representation (it is estimated that the human brain contains on the order of 10^{11} neurons making approximately 10^{14} synapses [22]). The task seems daunting. However, we do not need to write quantum mechanics level equations for a large number of atoms to predict the trajectory of a ball in space. Maybe it is possible to derive similar "classical" approximations for how neurons encode information.

As a first step, we can suspect that a particular area of the brain is relevant for certain tasks. This type of information has initially come from cases of patients that had neurological lesions (see also the comments below about functional imaging). In this way, it was initially established that areas of the occipital cortex were important for visual processing [23]. This type of initial notions then gave way to more detailed studies about the responses of individual neurons. A lesion in area V5 (also called MT) in monkeys impairs the animals' perception of motion [24]. Several investigators reported that the neuronal activity in this area was correlated with the direction of motion of the stimuli. Within this frame it is possible to ask how to predict certain aspects of the stimulus (e.g., its direction of motion in this case) from the neuronal activity.

A distinction must be made between explicit and implicit encoding of information. If we observe a face, the information about that face is encoded in our retinae. This information is then submitted to several stages of visual processing. The retinae contain information about the face in the same way that

the pixels in a computer monitor contain information about a face photograph. However, the retinae know little about faces. If the face is slightly rotated, or if luminosity changes, or if the face moves away from the observer changing its size, most of the neurons in the retinae will completely alter their patterns of firing. The retinae implicitly contain information about the face but faces are not the features encoded in the retinae.

2.2. Which features are encoded?

An area of the cortex called V1 receives visual input from the lateral geniculate nucleus in the thalamus which, in turn, receives visual input from the retina [17]. Investigators have studied the receptive fields and response characteristics of neurons in this area to visual stimuli. Searching for a response to tactile stimuli in this area would not yield much success. This obvious point shows that in order to study how neurons encode information one also needs to understand what kind of information is being encoded. However, it seems that in order to understand what kind of information is being encoded, one needs to understand what to look for in the neuronal response. This chicken and egg dilemma emphasizes the strong link between neural coding and the study of feature selectivity. In fact, the two issues could very well be considered to be part of the same problem. The above example appears trivial but consider the case of recordings in inferior temporal (IT) cortex in the macaque brain. We know that neurons in this area change their firing rate in response to complex stimuli such as faces and hands [25–28]. Does this mean that neurons in IT are actually interested in the whole face or just the eyes, or a particular combination of eyes and mouth, or some other complex features? A model of object recognition posits that selectivity and invariance of neurons in IT can be derived from a hierarchical arrangement of different visual areas with distinct feature selectivities [29]. Still, understanding what particular features are being encoded in IT has proved to be a rather difficult task in spite of over three decades of active research in the field.

Detecting what features are encoded is difficult because the stimulus universe is multidimensional and potentially quite complex and recording time is highly limited. A typical recording session for a given neuron may allow the investigator to present a set of approximately a few hundred stimuli and this is typically not enough to fully characterize or extract what features are encoded. Investigators may parametrically vary the stimuli according to some particular features they have in mind but these parameters may be as relevant to the neurons as the tactile stimuli may be to V1.

2.3. Comparison to coding in modern computers

The brain has often been compared to computers. In computer science, the word “code” is sometimes used to refer to a program or set of instructions that a computer will follow but this is unrelated to our discussion. Closer to neural codes are the ways in which information is encoded in the bits of a computer. Gross estimates of the number of neurons in the human brain give numbers of 10^{11} neurons and on the order of 10^{14} synaptic connections [22]. The Intel chip in the CPU of a modern computer contains on the order of 10^8 transistors. The clock speed may be on the order of 1 ns for the Pentium chip, a number that is several orders of magnitude faster than the type of communications present in the brain. One of the major differences lies in the connectivity; the brain has an average of approximately 10^3 connections per neuron. Another remarkable difference is the reliability. At the single transistor level, the voltage may be very stable whereas the timing of firing of a neuron may vary by several hundred ms from one trial to the next under apparently similar conditions. Even in the most remarkable cases of timing accuracy (see

Section 4.3), transistors would appear to be more precise. Furthermore, transmission of action potentials and neurotransmitter release from action potentials is also very unreliable (see Section 5). While it is unclear how to accurately estimate the storage capacity of the human brain, common sense seems to indicate that in terms of sheer memory storage, computers highly surpass humans. Still, the mundane and apparently simple tasks of recognizing objects, understanding an arbitrary conversation and walking around the house still constitute hard tasks for modern computers.

Suppose an investigator tries to understand how information is stored in a computer. The task of reverse-engineering a computer to find out the codes may by highly nontrivial. We could further assume that his tools are far from ideal. For example, he can start only by monitoring the voltage at a particular transistor within the circuit. Trying to correlate this measurement to the processes going on in the computer or to its output (e.g., what is shown in the screen) may not be easy. This trivial analogy, though flawed in many aspects, illustrates how difficult the task of understanding the codes used by a computer or a brain can be. One way to start is by using large-scale models and measurements (e.g., what happens if this chip is removed?), and then fine-tune the techniques to study the smaller components.

3. What kind of experimental data and resolutions do we need to study neural coding?

As in many other areas of science, what can and what cannot be measured influences the kind of hypotheses that can be directly tested. There are many theoretical ideas and models about coding by large groups of neurons and multiple brain areas. Unfortunately, at this point many of these hypotheses are necessarily highly speculative given the kind of evidence that experimentalists can gather. We will give a brief overview here of the most common type of measurements that neuroscientists use nowadays to interrogate the activity within the brain. We will describe the temporal and spatial resolution of each technique (Fig. 1). The description is brief and the reader is referred to other more comprehensive experimental treatises for further information (see, for example, Refs. [30–32]). The list of techniques indicated here corresponds to some of the most prominent experimental tools used by neuroscientists today; they are not intended to represent theoretical limits of what can be measured. It should be kept in mind, of course, that this is a highly dynamic discipline and that techniques can be dramatically improved over the course of a decade.

3.1. A succinct description of experimental techniques

One of the oldest techniques is electroencephalography (EEG). EEG measurements have a very high temporal resolution (<1 ms), but the spatial resolution is highly limited [16]. Several investigators have tried to improve models used to attempt to derive the precise location of the sources of EEG activity but it is still at best on the order of several mm (several cm in most cases). This is due to the necessity of solving an ill-posed problem in which multiple solutions exist for electric fields inside the head that can give rise to a given pattern of EEG measurements. For a density of approximately 10^5 neurons per mm^3 in cortex, a resolution of 1 cm^3 implies listening to the activity of on the order of millions of neurons. An important advantage of EEGs is that this is a non-invasive technique. This implies that it is possible to work in humans as well as in other species. A related technique is magnetic encephalography to which similar comments apply [33,34].

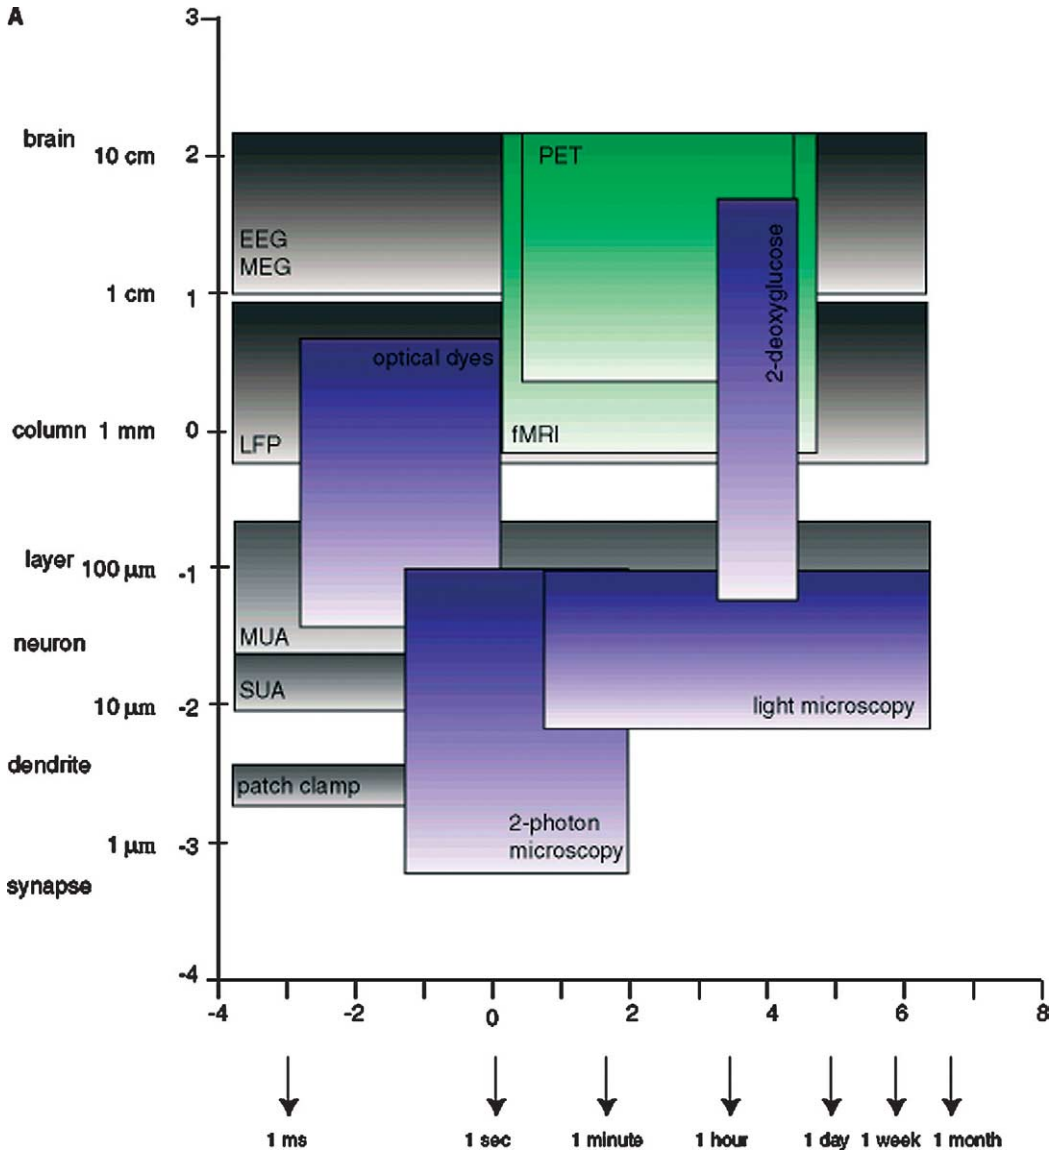

Fig. 1. Experimental techniques used to study neural coding. Schematic illustration of the spatial and temporal resolution of different experimental techniques used in Neuroscience. The resolution limits indicated here are only approximate and may depend on experimental conditions. In some cases, combining different techniques can improve the resolution. The limits shown here only indicate what has been experimentally reported and do not necessarily imply fundamental physical limits in the techniques. We exclude from these diagrams measurements from psychophysics, psychology, and computational models. (A) Techniques that measure neuronal activity directly or indirectly. Electrophysiology-based techniques are shown in gray, optical techniques are shown in blue, PET and fMRI are shown in green. (B) Techniques that interfere with neuronal activity. Lesion techniques are shown in red, stimulation techniques are shown in yellow. Further comments about each technique and references are given in the text (see Section 3.1). EEG = electroencephalography, MEG = magnetoencephalography, LFP = local field potentials, MUA = multiunit activity, SUA = single unit activity, PET = positron emission tomography, fMRI = functional magnetic resonance imaging, TMS = transcranial magnetic stimulation. This is an update of a figure prepared by Churchland and Sejnowski ([145], with permission).

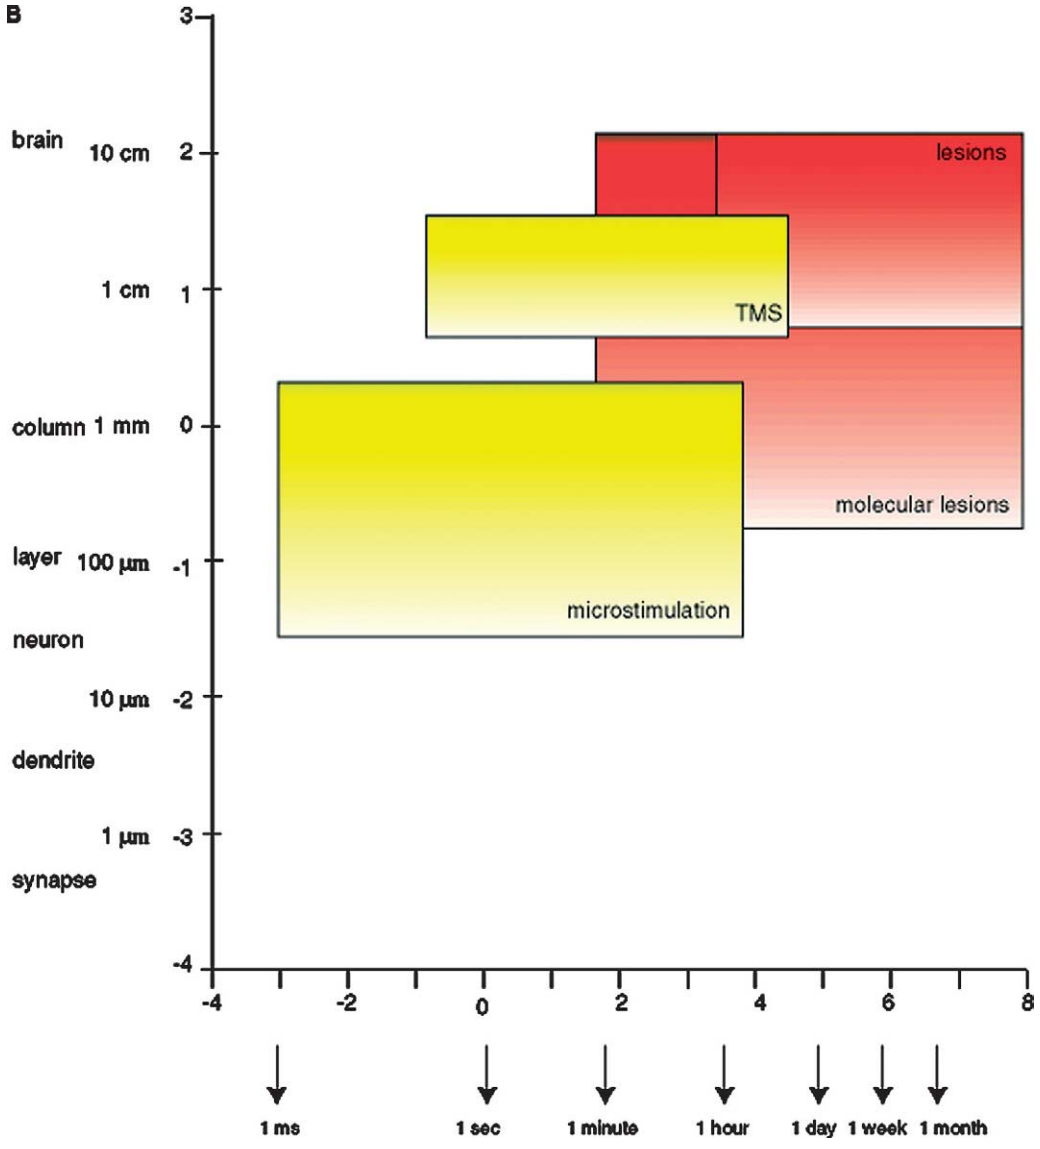

Fig. 1. (Continued).

Another important technique of wide use in Neuroscience is functional magnetic resonance imaging (fMRI). The basic principle is based on the observation made by Linus Pauling several decades ago that the magnetic state of oxygen (O_2), changes when it is bound to hemoglobin [35]. Given that increases in neuronal activity lead to a concomitant increase in blood flow to a particular brain area, it is possible to indirectly infer neuronal activation based on blood flow [32,36]. The same principle is applied in positron emission tomography (PET). The best current reports about spatial resolution come from investigators applying this technique in monkeys. Using magnetic fields of 4.7 Tesla, Logothetis and colleagues report a resolution of $125 \times 125 \times 720$ µm [37]. This technique is also non-invasive and several thousand

papers have been published already using this technique in humans. Another important advantage is that it permits to observe the activity in the whole brain at once. Unfortunately, the temporal resolution is rather poor and is limited by the speed of blood inflow to a particular area, typically on the order of 1 to 2 s [30]. At the present moment, it seems unlikely that the temporal resolution could be improved by several orders of magnitude to reach the ms or sub-ms level of electrophysiology.

It is worth mentioning other imaging techniques including optical imaging and two-photon microscopy. Optical imaging uses the infrared frequency band (typically 600–750 nm) to measure the reflected light with a CCD camera. It is based on the change of absorption with neuronal activity [38]. The signal has a delay of several hundred ms and therefore also does not provide the temporal resolution of electrophysiological recordings. However, it allows the investigators to observe larger areas of cortex [39]. Conventional microscopy has also been an important tool in Neuroscience, providing important anatomical insights but it lacks the temporal resolution required to study neuronal firing. Two-photon microscopy provides sub-neuronal resolution (see, for example, Refs. [40–42]). Optical imaging with voltage sensitive dyes promises a high spatial and temporal resolution while still keeping many of the advantages of imaging [43]. This invasive technique achieves a spatial resolution <1 mm with ms precision in the time domain.

Lesions have provided unique insights about the functioning of the brain (see, for example, [43–47]) and the historical account in [1]). This is, by nature, an invasive technique. In humans, we are limited of course to natural lesions. In animals lesions can also strongly point to the areas of the brain involved in particular behaviors or sensations. Indeed, lesions have provided the foundation for many of the electrophysiological experiments that are described in this review. More restricted types of lesions are also possible in animals. As we will discuss in Section 7, the advent of tools from molecular biology promises to radically change the lesion tools available today. However, the application of molecular tools as high-resolution lesion tools only works for animal models.

Our hunch is that a detailed understanding of the mechanisms of coding and decoding will require very high temporal (ms) and spatial (neurons) resolution. If we ultimately want to understand questions such as how many neurons represent a given stimulus, how the pattern of neuronal firing relates to the stimulus, and what type of neuronal responses a given stimulus will elicit, it seems that we cannot, by definition, rely exclusively on low spatial or low temporal resolution data. The same seems to apply, although this is far more tentative, to being able to predict the stimulus based on the neuronal activity. This is more tentative because it is possible that one could build rather accurate classifiers based on low-resolution non-invasive data at least for some aspects of the stimulus world. For example, it may be possible to easily discriminate based on EEG data or fMRI data whether the subject saw a face or a house [48–51] but it is unclear how to arbitrarily predict exactly what stimulus the subject saw from EEG or fMRI data. It should be emphasized that currently we cannot do this with high spatial resolution data either. Thus, functional imaging (including optical imaging and fMRI), EEG/MEG recordings and lesion studies seem provide a fundamental foundation to guide the search for neuronal coding mechanisms.

The focus of our review will be single neuron electrophysiology. Neurons emit all-or-none electrical signals, called action potentials, or spikes. These spikes can be monitored by inserting an electrode close enough to the neuron. These impulses are on the order of a few μV when monitored extracellularly and last about 1–2 ms. The path to neuroelectrophysiology was opened by the design of amplifiers able to detect such small signals [1,52]. Extracellular recordings are based on measuring electrical changes outside of the neuron. If the electrode is sufficiently close to the soma of one neuron, the signal derives mostly from a single neuron. In some cases, the investigators can also insert the electrode inside the

soma of the neuron to monitor the intracellular potential. However, it is generally difficult to maintain stable intracellular recordings for prolonged periods of time. If the extracellular electrode is farther from the soma of the neuron, it may pick up the activity of multiple nearby neurons. Single unit activity can be obtained from multi-unit recordings using algorithms of spike sorting [53,54]. Spiking activity is obtained by high-pass filtering the raw signal (typically with a corner frequency of 300–600 Hz). If, instead of taking the high-pass band of the extracellular recordings, the investigators use the lower frequency band (from 0.5–300 Hz), the resulting recording is called local field potentials (LFP). LFPs also show sub-ms temporal resolution and monitor the activity of large ensembles of neurons, probably on the order of several mm to cm [37].

3.2. About spikes and codes

The advent of electrophysiological recordings from single fibers made possible by Edgar Adrian and colleagues revolutionized the field of Neuroscience [1]. This technique gave rise to eight decades now where researchers monitor the activity of single neurons in different species and different brain areas. The initial studies of Edgar Adrian were concerned with motor output. He showed that there was a correlation between the number of impulses emitted by a single fiber and the strength of the output [16,52]. As another major example of the application of this technique, Kuffler observed that neurons in the retina enhanced their firing response when the stimulus was within a delimited area of the visual field, called the receptive field [17]. Hubel and Wiesel inserted electrodes in the first cortical stage that receives visual input, the so-called primary visual area or V1, and observed that neurons may prefer (meaning fire more action potentials) bars of particular orientations [55].

What is the business of spikes? Spikes provide a fast way of communicating signals between neurons. Furthermore, their regenerative properties imply that they can travel long distances with little attenuation. Textbooks typically indicate that all spikes from a given neuron are equal in height and duration (except for bursts of spikes). Are they really equivalent from the point of view of coding? An analogy may better illustrate the point. Suppose an investigator is interested in understanding human communication. He could pay attention to what somebody says during a lecture to many students or he could pay attention to what the person sings while he is showering. Even during the lecture, the speaker may emphasize certain points. Similarly, the same neuron may carry different messages depending on which neurons will listen, different messages at different time points, etc.

Spikes are accessible experimentally. What other variables may be relevant? It may be interesting to consider the concentration of neurotransmitter released at the synaptic cleft as well as the Ca^{2+} concentration in the pre-synaptic terminal and dendrites. It is generally harder to get detailed information (see two-photon microscopy in Section 3.1) about these variables in a dynamic fashion with an acceptable temporal resolution.

At least theoretically, there is enormous room to perform very fancy computations with spikes. Perhaps a stimulus may be encoded by a neuron A firing a spike precisely 5.2 ms after another neuron B fired a spike and 14.2 ms before a spike from neuron C. Some investigators have spent considerable efforts trying to study some of these possibilities [11,56]. Experimental evidence for a direct role of any of these complex spiking patterns in coding information is still hard to come by.

There may be a tacit assumption that individual neurons can only code simple features whereas the encoding of complex features requires the coordination of thousands or hundreds of thousands of neurons. For example, a single neuron in the retina can be interested in the intensity in a particular small patch

of the visual field and several of these patches can be combined to form a line [17,55]. The activity of a single neuron representing a face may appear to be more mysterious. However, one can think of complex stimuli as being composed of a certain number of more elementary features [29]. An elegant example of coding complex computations in individual neurons is given by the so-called lobula giant movement detector neuron in locusts which can encode the product of two separate variables that are important to detect a looming object [57].

4. A gallery of examples

By reading the previous sections, it may be possible to assume that we are searching for one single neuronal code that can explain all available data. While indeed we want to find universal principles about encoding and decoding, it is not necessarily true that there should be one single type of representation. The nervous systems of different organisms are the product of very long periods of evolution. This implies that multiple independent ways of representing information may have appeared through time. These coding mechanisms may have remained available if they were in some sense efficient and useful to the survival of the organism. For example, a certain coding scheme can have a selective advantage over others through its efficient energy usage [58], through a higher capacity or through the possibility of separating stimuli that are indistinguishable for other codes.

In this section we describe several coding strategies (Fig. 2), illustrating each one with specific examples. The themes we summarize here are important coding strategies for which experimental support is currently available. However, this does not constitute an exhaustive list of coding schemes. We discuss the continuous distinction between "rate" coding and "time" coding. We show examples of sparse codes and the encoding of time-varying signals. Finally, we discuss the particular properties of how ensembles of neurons may encode information.

4.1. Rate codes

The most common notion of a neural code follows the paradigm laid out by Edgar Adrian (Section 3.2). In a "rate code" the only variable of interest is the total number of action potentials emitted by a neuron in a relatively long time period of several hundred ms or even seconds (Fig. 2A). "Spike timing" codes and "rate" codes are part of a continuum that depends on the size of the time window used to count spikes [7,9,14]. In a "spike timing" code, the precise time at which the spike occurs, at the ms level, is relevant for encoding. In a rate code, the time of occurrence of spikes is considered to be noise; two spike trains with the same number of action potentials are considered to be equivalent regardless of the timing pattern (Fig. 2A, top). For many electrophysiologists, a rate code constitutes the simplest and clearest notion of how neurons encode information.

Several investigators have recorded the neuronal activity in visual area V5 in the monkey brain (also called area MT) while the animal observed motion stimuli or performed a motion discrimination task [59]. By counting the number of spikes in windows of several hundred ms (or even more than 1 s [59]), it has been shown that: (i) neurons in this area are selective to the direction of motion [60], (ii) the spike count correlates well with the motion discrimination performance of the monkey [59,61], (iii) the timing of action potentials with respect to stimulus onset upon repeated presentations of the same stimulus is quite variable [4,62].

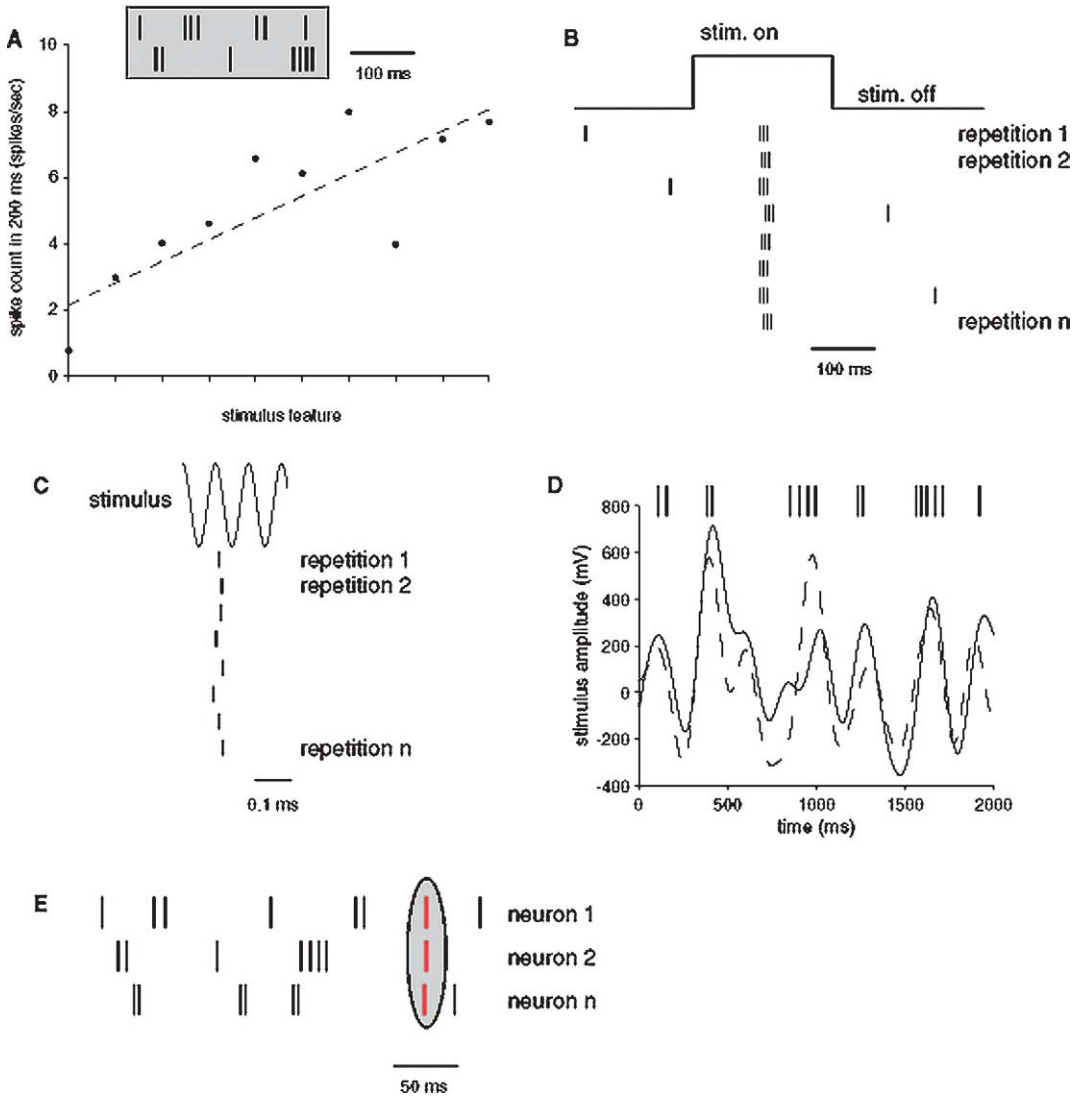

Fig. 2. Schematic illustration of different coding schemes and their corresponding temporal resolutions. (A) In a rate coding scheme, the number of spikes in windows of several hundred ms correlates with some stimulus feature or motor output. The two spike trains shown on the top part are considered to be equivalent for a rate code since they carry the same number of action potentials in spite of the different temporal patterns. Although a linear trend is illustrated here, the relationship between stimulus and the spike count may be non-linear. Examples of rate codes are given in Refs. [59,63]. (B) In a sparse representation, the neuron shows a very low spontaneous activity. The neuron reliably fires a single burst of spikes at a particular time from stimulus onset during multiple repetitions of the same stimulus. Examples of sparse representations are given in Fig. 3 and Refs. [69,71,72]. (C) A neuron shows very precise spike timing, with a trial-to-trial variation which can be less than 1 ms. Examples of remarkable temporal precision in neuronal firing can be seen in Refs. [9,14,79,146]. (D) A time varying signal (solid trace) is represented by a neuron that can follow the rapid changes in the stimulus (top, action potentials). The stimulus can be reconstructed (dashed trace) from the instantaneous firing rate of the neuron (see Refs. [20,21,87] for examples of this type of representation). (E) In this example, the synchronized activity of multiple neurons (symbolized by the spikes marked in red) constitutes the code to represent information. Examples of this type of representation are given in Fig. 4 and Refs. [10,97, 101,103].

Similarly, in the macaque inferior temporal cortex, spike counts correlate well with the identity of the stimulus the monkey is viewing [63]. Charles Gross and colleagues showed that a neuron may enhance the number of spikes with respect to the baseline firing rate upon presenting to the animal a certain stimulus, such as a given face [25,63]. Upon repeatedly presenting the same stimulus, investigators typically observe that the timing of the spikes, with respect to the stimulus onset, is highly variable. The mean of the interspike interval distribution is typically on the order of the standard deviation, which corresponds to the variability observed in a Poisson process [22]. The interpretation of these observations by many investigators is that the variability in spike timing constitutes mostly noise that needs to be "averaged out" and that neurons only care about the mean spike counts in windows of several hundred ms [4].

With an appropriate dynamic rate, a neuron could encode multiple different features in different spike count bands. For example, a V1 neuron could, in principle, signal the presence of a bar of a given orientation with 10–20 spikes/s, a bar of a different orientation with 20–30 spikes/s, etc. In the brain, the mechanism of encoding seems to be different. A V1 neuron may have an orientation preference and the number of spikes per second correlates with how close the actual orientation is to the preferred one.

How are rate codes read out? The "averaging out" of the variability could be performed by a postsynaptic neuron that integrates input from large numbers of neurons. In this scheme, the input neurons are assumed to function independently and communicate a spike rate with some noise [4,64]. This view greatly simplifies the tasks of recording, analyzing and decoding neuronal activity. Furthermore, this suggests that at least to a certain degree it is possible to ignore the complexities of where neurons receive input, dendritic processing, spike timing and correlations between neurons. Rate codes are robust to timing jitter by definition. The degree of robustness to spike failures and spontaneous spikes depends on the variability in spike counts compared to these sources of noise.

4.2. Sparse codes

In stark contrast with neurons that constantly fire many spikes per second and then briefly change their firing rate by several tens of spikes per second, there are some remarkable examples where individual neurons seem to respond selectively to specific stimuli using only a few spikes (Fig. 2B). Such *sparse* representations have caught the attention of many scientists [65,66].

As an example of a sparse representation, Kreiman and colleagues have studied the responses of individual neurons in the human medial temporal lobe (MTL). Subjects are patients who show pharmacologically intractable forms of epilepsy. Electrodes are implanted, typically in the hippocampus, amygdala, entorhinal cortex and parahippocampal gyrus, in order to localize the seizure onset focus [67–69]. Single neurons in the human MTL show selective responses to visual presentation of complex stimuli including faces, objects and spatial layouts [67]. Some neurons showed a sparse response with very low background rate (less than 1 spike/s) and an enhancement of only a few additional spikes in the presence of their preferred stimuli [69,70]. An example of this is shown in Fig. 3A. This neuron, located in the right amygdala, fired a few extra spikes upon presenting to the subject a drawing of Curly, one of the characters in a famous American TV series. This should not be interpreted as indication that this is the only visual stimulus that the neuron would respond to. Recording time is highly limited and the investigators could only present less than 60 different pictures [69]. This observation does not imply either that this is the only neuron in the human brain that would be selective to Curly. Apart from being highly unlikely, such a representation would not be robust.

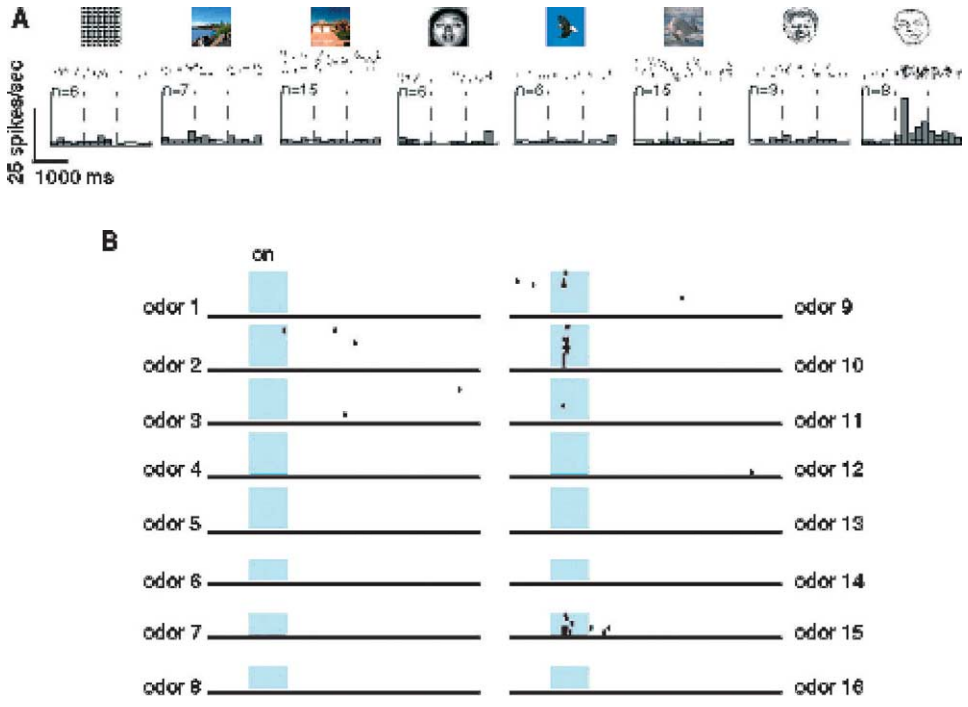

Fig. 3. Examples of sparse neuronal responses. (A) Responses of a neuron in the amygdala of a human epileptic patient implanted with depth electrodes in order to localize the seizure focus [70]. The subject was presented with pictures of complex stimuli (top, only a subset of the images is shown here). Below each image, a raster plot indicates all the spikes aligned to stimulus onset. The post-stimulus time histograms (PSTHs) show the average response of the neuron to multiple repetitions of each image (the number of repetitions of each image is indicated above the PSTH). The dashed vertical lines indicate the stimulus onset and offset respectively. This neuron transiently increased its activity when the subject saw an image of Curly (from [69], with permission). (B) Responses of an olfactory neuron located in the so-called mushroom body of a locust to 16 different odors. The stimulus presentation time (1 s) is denoted by the gray rectangles. Each tick indicates an action potential and multiple repetitions of each odor are shown. The neuron remained silent most of the time, with a baseline firing rate of less than 1 spike/s. The neuron reliably fired a few spikes at a specific time in response to 2 odors only (from [71], with permission). (C) Responses of 10 neurons in an area called the hyperstriatum ventralis pars caudalis (HVC) nucleus of the songbird while the animal was singing. The top part shows the vocalizations as a sonogram (frequency versus time, the intensity is color coded). The song is divided into motifs which are, in turn, composed of different syllables. The bottom part shows the neuronal raster plot, with multiple repetitions for each of the 10 neurons; each neuron is shown in a different color. The spikes are aligned to the onset of the nearest song syllable. The neurons showed a very low rate and fired a brief burst of spikes during a specific syllable. In general, different neurons selectively responded during distinct syllables or parts of a syllable, some neurons did not respond during this song motif (from [72], with permission).

Another example of sparse coding comes from the study of encoding of olfactory information by individual neurons in the mushroom body of locusts [12]. Remarkably, these neurons show, on average, interspike intervals longer than 20 s. Some of these neurons responded specifically to one or a few out of 16 possible odorants by firing only one or two spikes at specific times [71]. One such neuron is shown in Fig. 3B. It is striking to note how the neuron is basically silent most of the time. The baseline firing rate of this type of neurons was 0.025 spikes/s.

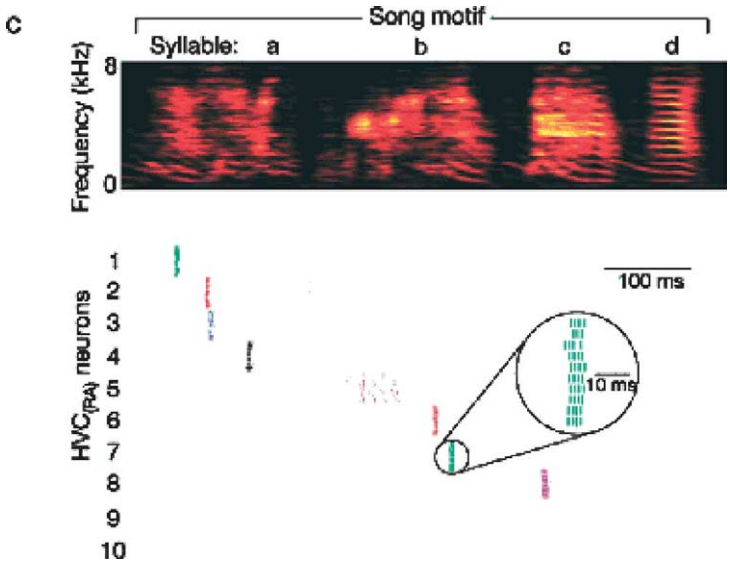

Fig. 3. (Continued).

Another interesting example is the study of the songbird's auditory system. Single neurons in one of the nuclei important in producing songs (the so-called HVC nucleus) show a short burst of spikes at specific time periods with respect to certain syllables of the song (Fig. 3C, [72]). The investigators suggested that these neurons show similar responses to those in the controversial proposal of 'grandmother' cells in the case of object recognition [72]. A "grandmother cell" is a neuron that would be activated exclusively when the subject saw his grandmother [65,73]. This example shows that the representation of motor output, not only that of sensory stimuli, can also be sparse.

How are sparse representations decoded? While many modeling studies have addressed the read-out of rate codes, fewer studies have focused on how to interpret sparse representations. For a representation using very few spikes to be efficiently decoded, the transmission of these spikes should be very reliable, the spontaneous activity should be low and the neurons sending sparse information should make strong synapses. One may also speculate that the post-synaptic neurons receive fewer inputs than the average neuron in cortex.

4.3. Spectacular timing

There are some examples where neurons display a striking temporal precision upon repeatedly presenting the same stimulus (Fig. 2C). We review some of those examples here but we leave those cases related to synchronous interaction of two or more neurons for Section 4.5.2. Multiple other ideas that take into account the timing of action potentials have been suggested including the existence of complex spiking patterns in multiple neurons [11], the precise timing with respect to a given phase of the local field potential [12], encoding based on first spikes or spike latency [74]. For an overview of different spike timing codes and their resolutions, see references [9,14]. In this section, we focus on cases that illustrate that neurons can be precise at the ms or even sub-ms level.

Owls can localize sounds using two cues, the difference in the arrival time of a sound to its right and left ears and the difference in the intensity of the sound at the two ears. The circuit and algorithms involved in this process have been the subject of intensive study [75,76]. The time difference between the two ears is in the sub-ms range. An elegant mechanism involving delay lines from neurons sending phase-locked information from each ear converges on a part of the brain called the nucleus laminaris in owls. Neurons in the nucleus laminaris are sensitive to coincidences in the spike arrival times and can detect interaural time differences on the order of tens to hundreds of μs [76,77].

A similar algorithm is used by electric fish in a behavior used to avoid jamming of frequencies with nearby fish [75,78]. The fish are able to detect phase differences of 400 ns in the signals arriving at different parts of its body. Interestingly, at the neuronal level, the primary afferents in the phase pathway show a response jitter on the order of 30 μs. In contrast to what might be expected, accuracy increases in higher processing stages and phase neurons in the midbrain show a jitter of 11 μs [79]. This shows a striking contrast with the variabilities of tens to hundreds of ms of some of the cortical neurons discussed in Section 4.1.

A somewhat different example comes from adjacent retinal ganglion cells that communicate through gap junctions in the retina. Cross-correlation analysis of the spike trains of ganglion cells in the salamander retina during spontaneous activity shows that a neuron can fire on average within 600 μs of the other neuron [80]. Retinal ganglion cells that share a common input show correlated firing with timescales on the order of 10–50 ms.

What does all this imply? We take these cases as evidence that neurons can show high temporal precision. Furthermore, even the sub-ms precision of spike timing can be relevant to encode information at least in these examples. This seems to be particularly true of situations in which time is an essential component of the signal itself (e.g., auditory time differences or phase of electrical signals). This should not be taken to imply that the timing of all spikes in the nervous system needs to be studied at the μs level but, at the very least, the machinery for timing accuracy is there [81,82]. Therefore, this certainly casts a doubt on arguments suggesting that neurons cannot keep precise timing. It is possible that in order to encode a stimulus that remains present on the order of several hundred ms or more, neurons do not need to show such striking temporal precision. For example, in the case of macaque monkeys, visual stimuli may be stable for 100 ms or more and therefore a different coding strategy may have evolved. This stresses the importance of studying coding of natural rather than artificial stimuli; see for example [83–86]. It is rather interesting and important to question how neurons can show this type of temporal hyperacuity when the spikes last on the order of 1 ms and interspike intervals may last 10 ms or more. How decoding works for signals that have extraordinary precision is still unclear but the detection of coincident firing may play an important role (see Section 4.5.2).

4.4. Time varying signals

Another situation where time is essential is the case of dynamic signals (Fig. 2D). For example, for a fly, estimating motion with relatively high precision may be crucial for its survival. The encoding of time-varying signals imposes some constraints on the types of codes that can be used and precludes the system from integrating spikes over hundreds of ms. All signals in the natural world are dynamic either because things move or because the animal moves or moves his eyes but the speed of change for a monkey deciding whether a face is friend or foe, is much slower than for a fly deciding which way to turn.

An example of the encoding of time-varying information is the work on an identified neuron called H1 in the fly. Bialek and colleagues were able to show that they could quite accurately reconstruct the motion stimuli experienced by the animal from the spikes recorded from a single H1 neuron [7,20]. An information theoretic analysis can put a bound on the number of bits transmitted by such a spiking neuron, reaching values of 3 bits per spike. Interestingly, when considering the flight speed of the animal and the firing rates of the H1 neuron, Bialek and colleagues also concluded that in many instances decisions about direction of motion are based on a very small number of spikes [7,20].

Another example of encoding of time-varying signals comes from a very different system. The phase pathway in the weakly electric fish was already mentioned in this article as one of the most remarkable examples of timing precision. In parallel to this pathway, electrosensory neurons in the electric fish are sensitive to the amplitude of electric field modulations in the environment. Amplitude modulations are used to locate objects and for communication [78]. The amplitude and phase pathways converge in higher brain centers of the electric fish. So-called P-type primary receptor afferents show high firing rates and their activity is modulated by amplitude changes. Up to 80% of the stimulus can be reconstructed by applying a linear filter to the spiking activity of these neurons [87]. Interestingly, the code shows a considerable degree of robustness to spike failures, spontaneous activity and timing jitter. The tolerance for timing jitter depends on the cut-off frequency of the stimulus and can be on the order of 3 ms for rapidly changing signals [21]. Signals from approximately ten of these afferent neurons converge on the next stage of signal processing, the pyramidal neurons in the electrolateral line lobe (ELL). The pyramidal neurons show much lower firing rates and typically fire bursts of 10–20 ms duration. Stimulus reconstruction from pyramidal cells is very poor compared to P-receptors [88]. The possibility of studying both stages allows for investigating how time-varying signals could be decoded. In this case, the pyramidal cells in the ELL do not seem to represent the detailed time course of the amplitude modulations. Instead, they extract behaviorally relevant features (such as upstrokes or downstrokes in the electric field) by firing bursts of spikes [88].

4.5. Coding by multiple neurons

The brain solves a different problem than the decoding of single electrode activity performed by many electrophysiologists. Cognitive processing and decisions about motor output depend on the activity of large numbers of neurons. Therefore, it is important to ask how groups of neurons can encode information in the nervous system. A large fraction (but not all) of the available data has been recorded with single electrodes. Investigators recording from individual neurons still wondered how ensembles of neurons could encode information. Because of the lack of simultaneous recordings, these studies typically assumed independent firing. Here we review some of the ideas about coding by ensembles of independent neurons. We also discuss findings obtained from multiple electrode recordings. In some cases, the independence assumptions seem to hold but there is also evidence that dependencies may be important in several other cases.

4.5.1. Independent neurons

In some cases, investigators recorded from single electrodes on multiple sessions using the same stimuli. From these types of recordings, it is possible to extrapolate to ensembles of neurons only after assuming independence. One such example is the study of Georgopoulos and colleagues where they studied how a population of neurons could encode the direction of movement of the monkey's arm [89].

Each single neuron was only broadly tuned to a specific direction of movement but the population vector obtained by adding the vectors denoting the "votes" of each neuron was a much better indicator of the direction of movement. Using a total of 224 neurons, this population vector was, on average, 15.8° away from the actual direction of movement (the direction of movement in the plane could be any value between 0° and 360°) [89]. A similar example is given by the attempts to assess the capacity of inferior temporal cortex neurons for complex objects of Rolls and colleagues [90]. Other models and analyses have been proposed beyond the simple voting schemes [91,92] but the underlying assumption is that of independence.

The question of independence has been hotly debated for the past two decades. An experimental assessment requires recording from multiple electrodes. One study where investigators attempted to directly evaluate the independence hypothesis is the work of Nirenberg and colleagues in retinal ganglion cells in the isolated mouse retina [93]. By comparing the amount of information conveyed by multiple neurons assuming independent firing versus the information without this assumption, the investigators observed that more than 90% of the information about the natural stimuli could be retrieved without studying concerted firing [93]. In another example, by recording from pairs of pyramidal cells in the electrolateral line lobe of the *Eigenmania* electric fish, Krahe et al. observed that they could consider the cells to fire independently [94]. In other words, the degree of correlation between pairs of neurons was the same after randomly shuffling the trials. This suggests that, for those experimental conditions, simultaneous recordings did not add to the information encoded by pairs of neurons. Fig. 4A shows the probability of misclassifying upstrokes and downstrokes in the electric field amplitude from single neuron recordings and also from pairs of pyramidal cells. Shuffling the trials did not change the classification performance for pairs of neurons, therefore suggesting that the responses of the two neurons could be considered to be independent. Similar results were reported by other investigators (see, for example, [95,96]).

The degree of independence may depend on several experimental parameters including the stimulus itself and the state of the animal (e.g., anesthetized versus awake). Furthermore, the correlation in the firing of two neurons may be a function of the distance between them. Inserting two electrodes very close to each other may not be easy and in some cases the connectivity may be such that it may be very difficult to find two neurons that are connected or have common input.

4.5.2. Synchrony

In spite of these caveats, other investigators have shown cases where the independence assumption breaks and it is important to consider synchronous firing. The group of Singer and colleagues has shown several examples of synchronous firing in the cat visual cortex [10,97]. In one such example, the investigators recorded from multiple neurons in the cat visual cortex and observed that two neurons synchronized their firing only when their activating stimulus belonged to a single object [98]. In many of these papers the idea of synchronous firing has been linked to the solution to the so-called binding problem. In brief, if you see a red apple falling from a tree, some color neurons may be detecting the red, some movement neurons may be detecting the motion and some object neurons may be interested in the shape of the apple. How does the subject distinguish the red apple falling from the tree from the red car moving in the street, the red apple hanging in the tree, and the green leaf falling next to the apple? This is the so-called binding problem. One proposal (oversimplified here) has been that the key lies in the synchronous firing between "red" neurons, "apple" neurons and "downward" movement neurons [10]. This notion is still hotly debated (see, for example, [99,100]).

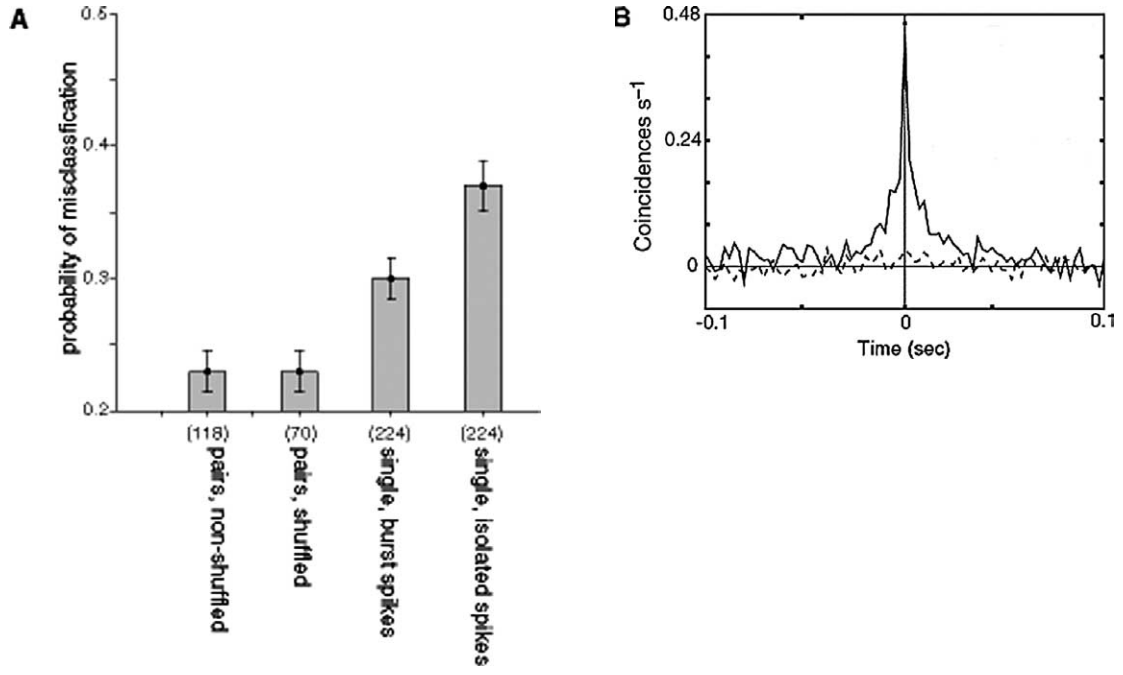

Fig. 4. Examples of correlations in spike timing between two neurons. (A) Feature extraction performance by pyramidal cells in the anterior lateral line lobe of the *Eigenmania* electric fish. Pyramidal neurons signal changes in the amplitude of the electric field around the fish. The classification of upstrokes and downstrokes in the electric field is characterized here by the probability of misclassifying the stimulus, p_e, ranging from 0.5 (chance performance) to 0 (perfect classification). The diagram shows that bursts of spikes from pyramidal cells are better indicators of environmental signals than isolated spikes. Synchronized spikes from two pyramidal neurons are even better at classifying the changes in electric field. The lack of a distinction between shuffled and non-shuffled trials suggests that neurons can be considered to fire independently under these experimental conditions. Results shown here correspond to the I-type neurons. Error bars represent standard errors of the mean. The numbers below the bars give the overall number of stimulus conditions for all cells or cell pairs analyzed (from [94], with permission). (B) Change in synchrony between an attended (solid line) and unattended (dashed line) condition between two neurons recorded from macaque SII motor area. A monkey was trained to switch attention between two different tasks, a visual task and a tactile task. Steinmetz et al. found that the attentional state was correlated with the level of synchronous firing between neurons. Synchrony was assessed by the degree of correlated firing after subtracting coincidences expected by changes in firing rate. The x-axis indicates the time delay between the firing of the two neurons and the y-axis shows the number of coincidences normalized to coincidences/s; bin size = 50 ms Shuffling the trials significantly affected the results; this indicates that, in contrast to the example shown in part A, the neurons could not be considered to respond independently (reproduced from [101], with permission).

In a different example of synchronous interactions between pairs of neurons, Steinmetz and colleagues showed that the degree of attention to different stimuli is correlated with the degree of synchrony between neurons in the somatosensory cortex of awake monkeys (Fig. 4B). Their careful statistical analysis showed that the synchrony observed between neurons is independent of changes in the firing rate of the neurons involved [101]. Other examples of synchronous interactions have been observed between pairs of neurons in the lateral geniculate nucleus in monkeys and also in primary visual cortex [8,102].

The most dramatic demonstration to date that synchrony can have a direct role in the behavior of the organism has been the study of Stopfer and colleagues in the olfactory system of honeybees [103]. The

authors showed that inhibiting synchronous firing by injecting picrotoxin, a $GABA_A$ channel antagonist caused poor performance in an olfactory discrimination task. This constitutes the first direct evidence that synchrony can introduce an important dimension to encode different stimuli.

5. The biophysical point of view

We have given several examples of different types of representations. These codes need to be decoded to ultimately exert any effects on behavior. At the processing stages near the motor output, the decoders can be the muscle cells themselves. Muscle cells have to interpret the commands by the pre-synaptic neurons and transform those into appropriate behavior. For other processing stages more remote from the output, a set of post-synaptic neurons is in charge of decoding and recoding. This scheme is an oversimplification. For example, it is very common in the nervous system to observe that higher stages in turn project back to earlier processing modules.

The question of how neurons can "read" the representation made by other neurons leads us to ask: How are action potentials generated, propagated and transmitted to other neurons? We do not attempt to give a detailed description of the biophysics of action potentials. Instead, here we will argue why some particular patterns of firing may be more easily transmitted to post-synaptic neurons than other patterns. For a detailed description of experimental and computational studies of the biophysics of action potentials, see [22,31].

5.1. Action potential propagation and neurotransmitter release

One of the first questions to ask is whether the action potential can reach the pre-synaptic terminal. Branch points where there is an impedance change can pose a challenge to the propagation of the spikes [22]. Given that an axon may reach several thousand targets, it is important to understand what fraction of these sites will be reached by the spike. An important study by Cox and colleagues recently showed that, at least in the large initial branches, action potentials propagate quite reliably in neocortical pyramidal cells [40]. Data from the smaller branches is very hard to acquire and therefore the question of the reliability of spike propagation still remains open.

When an action potential arrives at the pre-synaptic terminal, there is an influx of calcium to small microdomains that causes neurotransmitter vesicles to release their contents in a quantal fashion to the synaptic cleft. The number of vesicles released depends on the number of available vesicles and the release probability. According to the standard Katz model, the number of vesicles released can be approximated by a binomial distribution [16]. It has been suggested that the probabilistic nature of neurotransmitter release constitutes the main cause of failures in synaptic transmission [104].

Are there particular patterns of spikes that are more likely to be transmitted than others? This is a topic that has undeservingly received rather little attention in the field. The best studied such specific pattern is a burst of action potentials. In some cases, a neuron may fire several action potentials within a short period of 10–30 ms. These bursts of spikes are typically manifested by a bimodal interspike interval distribution with a narrow peak at short ISIs for spikes belonging to a burst and a broader peak at longer intervals for isolated spikes [88,105]. Bursts can be triggered by depolarization caused by calcium influx to the cell. Bursts show a higher probability of being transmitted to the post-synaptic neuron (see Fig. 5C

and [105–108]). Furthermore, it has been shown that bursts can outperform isolated spikes in conveying information about stimuli (see Fig. 4A and [72,88,94,108–110]).

In addition to failures in spike propagation or neurotransmitter release, another potential source of noise is the spontaneous occurrence of action potentials. A look at most areas of the brain at any given time shows a bewildering amount of firing activity (see, for example, [69,80,111]). Any decoding mechanism at the post-synaptic level must have incorporated somehow the probability that some of the spikes (or some patterns of spikes) may have occurred spontaneously and bear little information about the stimulus. There is still little information about the reliability of propagation and neurotransmitter release of spontaneous spikes in comparison with spikes emitted during a particular task.

5.2. Generation of action potentials

Hodgkin and Huxley pioneered the research into the ionic conductances responsible for the generation and propagation of action potentials [16,22,112]. This has given us a considerable degree of understanding of the processes that may occur when the voltage, in most cases near the soma of the neuron, exceeds a certain threshold. In order to understand how signals are decoded at the neuronal level, we also need to study the map that relates input spikes (or input neurotransmitter concentrations) to the generation of spikes. Unfortunately, it is not easy to construct this map with the available data. How neuronal input relates to neuronal output constitutes the core of extensive debate that directly parallels the questions about coding by groups of spikes [22,113–115].

According to some models, a neuron may act as a noisy integrator of large numbers of excitatory post-synaptic potentials (EPSPs [64], see however [113]). Other studies suggest that the probability of generating an action potential might depend on the pattern of EPSPs arriving at different dendrites or even different locations within the same dendrite [116]. An important factor that may play a role in achieving the right amount of depolarization to trigger an action potential is the relative timing of multiple inputs [117,118]. One indication of how sensitive a neuron can be to timing differences is given by the study of changes in synaptic strength. Bi and Poo have shown that differences of less than 10 ms between an EPSP and a post-synaptic spike can dramatically change the direction of change in strength of a synapse (see Fig. 5A, [119]). Another study also suggests that the relative timing of inputs may play an important role. As discussed in Section 4.1, in many cases the standard deviation of the interspike interval distribution is very close to the mean value; this is typical of a Poisson process. What kind of inputs can give rise to such variable responses? Zador and Stevens showed that synchronous input, but not purely excitatory or mixed excitatory and inhibitory input, can give rise to responses with a level of variability that is similar to that observed in vivo (Fig. 5B).

5.3. Where do action potentials go?

Another important aspect of the decoding process is who receives the information. By and large, where neurons project to has been largely ignored in many electrophysiological experiments. Wiring, however, can be part of the code. As an example, it has been shown in several areas of the brain that neurons form topographical maps of the environment (see, for example, [16,75,94,120]). Part of the experimental difficulty is that an electrophysiologist registering the activity of a neuron with an extracellular electrode rarely knows with precision where the neuron may project. One experimental approach towards finding

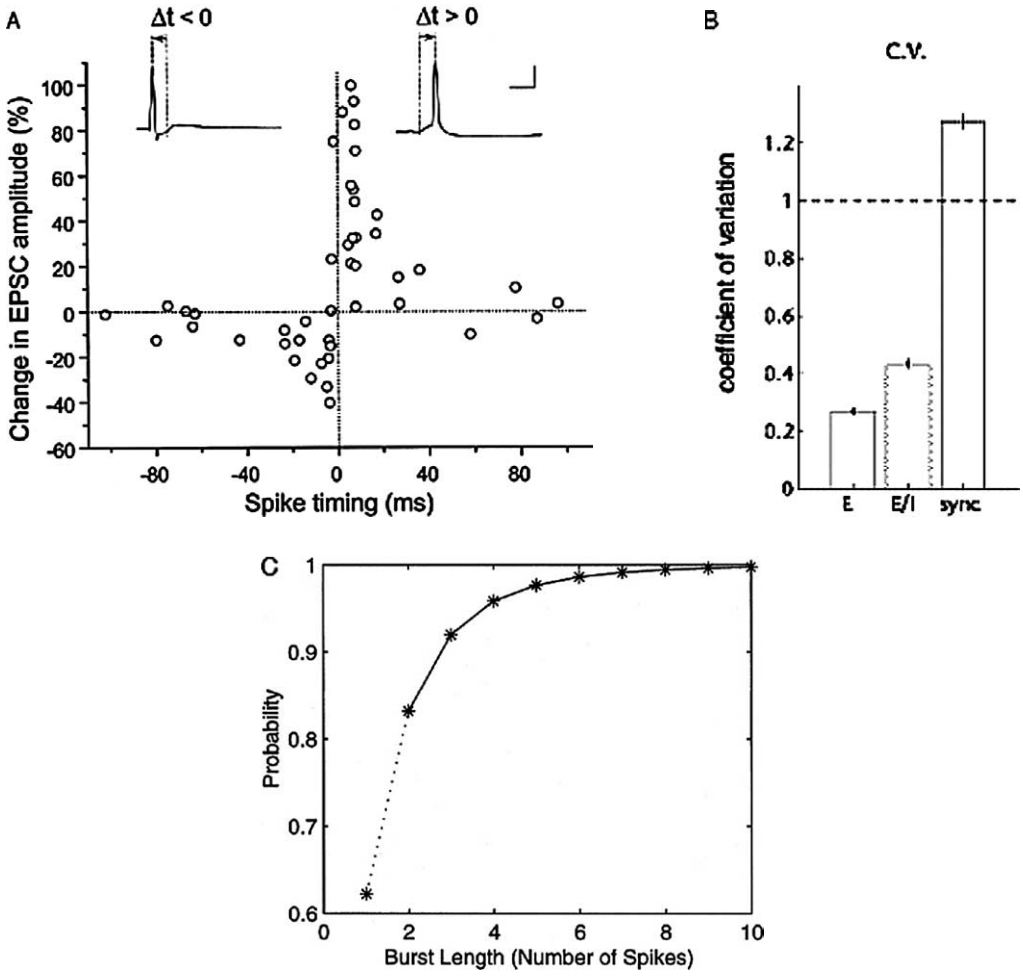

Fig. 5. Spiking patterns that may lead to enhanced transmission and neurotransmitter release. (A) Small differences in timing between a pre-synaptic spike and a post-synaptic EPSP can exert a major influence in synaptic strength. The figure shows the change in excitatory post-synaptic current (EPSC) measured 30 minutes after stimulation of hippocampal neurons in culture as a function of the relative timing between excitatory post-synaptic potentials (EPSP) and a post-synaptic action potential. The scheme on the top shows the relative timing between EPSP and action potential. A positive change in EPSC indicates synaptic potentiation whereas a negative change shows synaptic depression. Note that there is a narrow band around $\Delta t = 0$ ms that can change the direction of change from potentiation to depression (reproduced from [119], with permission). (B) Synchronous firing leads to irregular spike trains. This figure shows the effect of different kinds of input (E = purely excitatory input, E/I = mixed excitatory and inhibitory input, sync. = synchronous input) on the degree of irregularity of the post-synaptic spike train. Irregularity of the spike train is assessed by the coefficient of variation (CV) of the interspike interval distribution. A Poisson process shows CV = 1 (dotted line); CV values very close to 1 are observed in electrophysiological recordings in vivo in cortex. Error bars indicate standard errors of the mean (reproduced from [114], with permission). (C) Effectiveness of spikes in bursts to elicit a post-synaptic action potential. Cumulative probability distribution of generating a post-synaptic action potential as a function of the number of spikes per burst in the pre-synaptic neuron. This study was based on electrophysiological recordings in the cat visual cortex using sine wave gratings as stimuli. Putative connectivity was defined by a shifted and short-latency peak in the cross-correlogram of responses between the pre- and post-synaptic neurons. The data show that bursts of several spikes are more reliable than isolated spikes in eliciting spiking activity in the next information processing stage (reproduced from [105], with permission).

out the projections of a neuron is to inject a tracer [94]. This may be possible, though laborious, for some species, but it is not always feasible in the context of electrophysiology in some animals like macaques.

At least in principle, the message conveyed by a neuron may depend on who receives the signal. A given spiking pattern by a neuron in inferior temporal cortex may be decoded differently by the target neurons in the amygdala than by the target neurons in the prefrontal cortex. A deeper understanding of the decoding process could arise from knowledge of the detailed anatomy of neuronal connections.

6. Correlations and cause

The examples of coding schemes discussed in Section 4 show correlations between neurophysiological variables (such as particular patterns of action potentials) and actions or percepts. However, physiological measurements per se do not establish a causal link between neuronal activity and perception or action. The distinction between cause and correlation is an important one because it may help us better understand the functional role of a particular coding scheme. Unfortunately, establishing a causal link between physiology and perception or behavior constitutes, in general, an extremely difficult task given the methods and technologies available today.

One important clue towards understanding the relationship between activity in a given brain area and perception comes from brain lesions. For example, ablation of area IT in the macaque monkey brain causes impairment in the monkey's ability to visually discriminate between objects [46]. As discussed in Section 3.1, current lesion techniques still involve very large numbers of neurons (Fig. 1B). This is particularly true in humans where neurological lesions are poorly defined and rarely involve exclusively one area. Therefore, lesions can provide fundamental data and can direct attention to the area of the brain to study, but they do not provide mechanistic details about coding at neuronal resolution.

Another line of evidence to establish a causal link comes from stimulation studies. The most important non-invasive method to stimulate the human brain is transcranial magnetic stimulation (TMS, see, for example, Refs. [121,122]). The spatial resolution of this technique is on the order of 1 cm, which means that TMS probably interferes with the activity of at least hundreds of thousands of neurons (Fig. 1B). The invasive nature of electrical microstimulation makes it very difficult to directly stimulate the human brain. However, under particular circumstances, it has been possible to invasively stimulate the human brain while the subjects are conscious. Some classical and intriguing stimulation studies in the human brain were performed in epileptic patients by Penfield and colleagues [123]. In some striking examples, patients would recall faces, events or places during stimulation of the temporal lobe. Recent examples of the usage of this technique on epileptic patients can be seen in the work of Libet, Ojeman, Fried and colleagues [124–126]. The difficult nature of these experiments precludes from drawing strong statistical conclusions. However, it is reassuring and suggestive that it is possible to elicit complex perceptual states by stimulation of groups of neurons in the temporal lobe.

It is possible to perform more detailed and elaborate microstimulation studies in monkeys. A series of landmark studies by Newsome and colleagues has revealed that stimulation within the MT cortical area in macaques can bias the performance of the monkey in a motion discrimination task [127]. In a typical situation, the animal was presented with a series of randomly moving dots. If a certain percentage of dots move coherently in the same direction instead of randomly, the subject can discriminate this dominant direction of motion and motion-selective neurons in MT are strongly activated. The task becomes trivial for 100% coherence and performance is at chance levels for 0% coherence. Newsome's group showed

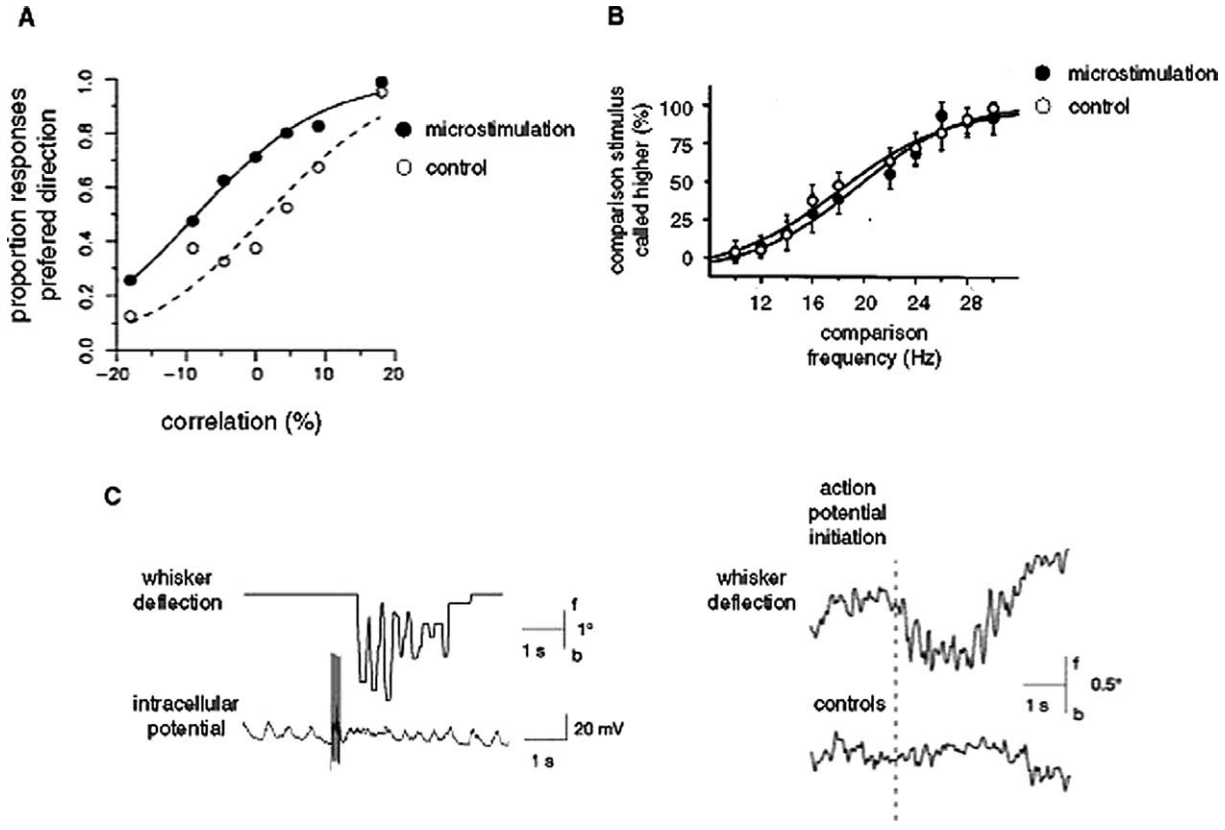

Fig. 6. Microstimulation of cortical neurons to assess causality. (A) Microstimulation in visual area MT can bias the decision of the monkey in a motion discrimination task. A monkey was shown dots moving randomly in the screen. When a given percentage of dots moved coherently in one direction, the monkey had to discriminate the direction of motion; performance increased with higher percentages of coherent dots. For a correlation of 0%, the decision is arbitrary. The investigators recorded extracellular neuronal activity; performance here denotes the proportion of times that the animal reported the direction of motion aligned to the preferred direction of the recorded neurons. Open symbols indicate trials without electrical microstimulation whereas solid symbols indicate trials where electrical current was injected through the recording electrode. Microstimulation shifts the psychometric curve up suggesting that it biased the monkey's perception of motion direction (reproduced from [127], with permission). (B) Microstimulation in area 3b within motor area S1 in the monkey brain. In this experiment, the monkey was trained to discriminate the frequency of tactile vibration. A first stimulus was presented at 20 Hz and then a second comparison stimulus was presented at a different frequency. The task was to assess which frequency was higher. The solid circles indicate trials where the second stimulus was an actual tactile stimulus whereas the open circles represent cases where the second stimulus was electrical microstimulation in the absence of any tactile stimulation. Remarkably, performance during microstimulation trials was indistinguishable from that during real stimulation trials (reproduced from [132], with permission). (C) Electrical stimulation of a single layer 6 neuron in the rat motor cortex can elicit deflection of the rat's whisker. The left plot shows the displacement of the whisker and the intracellular potential of an individual neuron in a single trial. Large deviations of the intracellular potential correspond to action potentials. The right plot shows the average of 30 single neuron stimulation trials and 30 control trials. The dashed line indicates the onset of stimulation (reproduced from [134] with permission).

that extracellularly stimulating in an area near neurons that prefer a certain direction of motion biases the performance of the monkey in that direction (Fig. 6A). These remarkable observations provide an

important link between neuronal activity and the perception of motion (the assumption being here that perception is correlated with the behavior of the animal [128,129]). This work was followed by several other microstimulation studies (see, for example, [130,131]).

Would it be possible to elicit a perceptual state in the absence of any stimulus, just by electrical microstimulation? Romo and colleagues showed that microstimulation in somatosensory area S1 in the absence of any sensory input can elicit a perception that seems to be indistinguishable from that occurring in the presence of a real stimulus (Fig. 6B, [132]). Monkeys were trained to discriminate between two trains of tactile stimulation by indicating which one had higher frequency. When the second tactile stimulus was replaced by electrical microstimulation in area S1, the animals could still perform the task at a level that was statistically not different from the one with a real stimulus.

The exact extent and spread of microstimulation is unclear [133]. It seems likely that the studies reported above involve the activation of large ensembles of neurons (perhaps on the order of several cortical columns). While these studies provide important insights about the putative causal relationship between neuronal activity and perception or behavior, they still do not directly answer what the coding mechanisms are. It is extremely hard to disambiguate different coding schemes such as the ones described in Section 4 by using microstimulation. Furthermore, it is possible to argue that the actual effect of microstimulation is due to the activation of another area which is elicited by microstimulation only as a secondary effect. The involvement of large numbers of neurons precludes the direct study of how spikes encode information. Recently, Brecht and colleagues have provided the most direct evidence to date that stimulation of an individual neuron can be related to, in this case, eliciting a particular motor output [134]. By recording intracellularly from a neuron in the motor cortex of rats, they were able to induce whisker deflections by injecting enough current to drive the neuron to fire multiple action potentials (Fig. 6C). Furthermore, the whisker motor output depended on the pattern of electrical stimulation and on the stimulation layer. This kind of experiment is almost impossible to conduct in humans. Establishing such a direct link between perception and the activity of an individual neuron may be difficult in most cases given that the behavioral repertoire in which animals are tested is limited.

7. Directions

We finish this review by highlighting some questions that we consider will be important in advancing the field. While some of these are areas of active research others seem to be largely neglected, partly because of the experimental difficulties involved.

(i) As discussed in Section 5, action potentials need to propagate through axons, reach the pre-synaptic terminals and lead to neurotransmitter release. It will be important to understand the process that leads from the generation of action potentials to neurotransmitter release in further detail. For example, do certain spike patterns and certain interspike intervals (in addition to bursts) show higher probability of eliciting neurotransmitter release than others? Under what circumstances will isolated spikes lead to neurotransmitter release? Is there a particular processing or signaling mode that can distinguish spontaneous spikes from stimulus related spikes? Do different coding schemes require different types of neurons and synapses? To what extent are the coding mechanisms determined by the type of input?

(ii) Few experimenters insert multiple electrodes in spite of the fact that the basic technology is a couple of decades old [2,135]. This implies that the studies are restricted to studying the responses of individual neurons or assuming independence in the responses of separate neurons. However, as discussed

in Section 4.5.2, several investigators have shown that correlated firing can increase the amount of information conveyed by groups of neurons. The degree to which correlated firing represents information that is independent of changes in firing rate still remains controversial. Whether synchronous firing constitutes a mechanism to enhance post-synaptic response or to bind specific features into a whole also remains a matter of heated debate. What proportion of neurons show concerted firing? How does concerted firing depend on the distance between the neurons and on the neuronal types? Are there other mechanisms of population coding? It is likely that novel and important insights will emerge from recording from larger ensembles of neurons.

(iii) An important coding dimension is the spatial arrangement of selective neurons. In many cases, investigators have observed a topographical map where adjacent neurons code for similar features in stimulus space [16,55,76]. While we think that the ultimate details about coding will depend on the study of spikes and single neurons, it may be important to study in further detail the activity from local field potentials. The simple observation that LFPs show selectivity [36,91] suggests that the average activity of large numbers of neurons within a small region is not just noise. The selectivity from LFPs is not entirely redundant with the information carried by spikes, suggesting that there may be emergent channels of information that may become apparent only by studying multiple neurons. Furthermore, in some cases, it has been reported that the relationship between spikes and the LFP is important (e.g., the phase between the two signals [12,136]).

(iv) What constitutes noise to some investigators may be an important part of the signal for others. It has been observed that repeated presentations of the same stimulus in higher cortical areas lead to widely different times of action potentials (Section 4.1 and Ref. [4]). This is partly what led to the idea of a rate code. However, some other investigators have suggested that this timing noise would turn out to be a very important signal if looked at appropriately (e.g., in conjunction with the responses of other neurons).

(v) Although several investigators have studied the structure of neural circuits, efforts in neuroanatomy have substantially diminished. In order to fully understand coding, we will need to know who sends information to the neurons under study and whom these neurons talk to (see Section 5.3). Details about connectivity are not easy to acquire and require laborious efforts. For example, consider the case of a neuron in inferior temporal cortex. In which layer is this neuron located? Does it receive feed-forward input from pyramidal neurons in V4? Does it receive inhibitory connections from interneurons? What kind of interneurons? Any feedback from frontal cortex or other areas? Any direct input from other ventral or dorsal visual areas? Does it project to pyramidal neurons? What type of frontal cortex neurons does it project to? Does it also send the information to the entorhinal cortex and/or to the amygdala? To complicate matters further, some investigators have suggested that the particular location of synaptic spines (specializations within dendrites that receive a large fraction of excitatory synapses) may be relevant [137].

(vi) Not all neurons are equal. The responses of a pyramidal cell in layer 4 may be completely different from those of an interneuron in layer 2 in terms of the spiking patterns, the input and output projections, etc. It is still unclear how many different types of neurons there are [138]. It will be interesting to be able to characterize neurons in terms of their firing patterns, their functional selectivities, and also their gene expression, morphology and connectivity.

(vii) The distinction between correlation and causation is important to separate coding from epiphenomena or indirect correlations (Section 6). We need novel tools to study the causal relationship (if any) between particular spiking patterns of certain neurons and perception or action. Microstimulation will

need to be refined to better understand how many neurons are being activated (and what types of neurons and how). It will also be interesting to stimulate from multiple nearby electrodes.

(viii) Current lesion techniques have provided fundamental insights but they lack the specificity required to understand mechanistic aspects of coding at the neuronal level. It seems that specificity in ablating specific neuronal types or neuronal networks will come from molecular biology tools. For example, we may ask: what would change in the selectivity of a pyramidal neuron in layer 4 in inferior temporal cortex if we could temporarily silence all the GABA-ergic interneurons that project to it? This would require a selective molecular marker of those interneurons and then neuronal silencing techniques [139,140]. One difficulty with this line of ideas is that molecular biology tools are well established in flies and mouse (e.g., making transgenic or knock-out mice). However, some cognitive questions are hard to study in mice (e.g., the visual system of mice is quite poor compared to primates). On the other hand, developing the molecular biology tools for macaque monkeys seems to be quite expensive and complicated. One promising technique to circumvent these difficulties is the usage of virus that can be topically applied [141].

This constitutes only a small sample of important questions that come up in the study of neuronal codes. Progress in some of these areas is taking place rapidly. For example, methods are becoming more common to record from large ensembles of neurons [142] and neuroscientists are becoming ever more interested in applying molecular tools, e.g., [143,144].

8. Summary

We have described several different strategies that neurons use in different systems to encode information about the environment, internal processing or motor output. These strategies include rate coding, sparse coding, encoding by spike timing and encoding by concerted firing of multiple neurons. For any type of neuronal representation, there should also be a corresponding mechanism that can interpret and process the corresponding spike patterns. This establishes a direct relationship between encoding and the biophysics of generation and propagation of action potentials. Evidence of correlations between spike patterns and sensory stimuli or motor output still does not provide evidence for the functional significance of a coding scheme. A full understanding of neural coding would require proving that the spiking activity can be causally linked to the stimulus or behavior. The combination of tools from molecular biology, biophysics, electrophysiological recordings and microstimulation promises to provide many novel insights in the near future about how neurons encode information.

Acknowledgements

We would like to thank Tommy Poggio, Mariela Zirlinger, Thomas Serre, Srinivas Turaga, Ryan Rifkin and David Cox for discussions and comments on the manuscript, all the authors who have given permission to reproduce their figures and the CBCL at MIT for support. G.K. is supported by a Whiteman fellowship.

References

[1] Finger S. Minds behind the brain: A history of the pioneers and their discoveries. New York: Oxford Univ. Press; 2000.
[2] Perkel D, Bullock T. Neural coding. Neurosci. Res. Prog. Sum. 1968;3:405–527.
[3] Meister M. Multineuronal codes in retinal signaling. Proc. Natl. Acad. Sci. USA 1996;93:609–14.
[4] Shadlen MN, Newsome WT. The variable discharge of cortical neurons: Implications for connectivity, computation and information coding. J. Neurosci. 1998;15:3870–96.
[5] Abbott LF. Decoding neuronal firing and modelling neural networks. Q. Rev. Biophys. 1994;27:291–331.
[6] Koch C, Laurent G. Complexity and the nervous system. Science 1999;284:96–8.
[7] Rieke F, Warland D, van Steveninck R, Bialek W. Spikes. Cambridge, MA: MIT Press; 1997.
[8] Usrey WM, Reid RC. Synchronous activity in the visual system. Annu. Rev. Physiol. 1999;61:435–56.
[9] Bair W. Spike timing in the mammalian visual system. Curr. Oppin. Neurobiol. 1999;9:447–53.
[10] Engel A, Singer W. Temporal binding and the neural correlates of sensory awareness. Trends Cogn. Sci. 2001;5:16–25.
[11] Abeles M. Corticonics. Cambridge: Cambridge Univ. Press; 1991.
[12] Laurent G. Olfactory network dynamics and the coding of multidimensional signals. Nat. Rev. Neurosci. 2002;3:884–95.
[13] deCharms R, Zador A. Neural representation and the cortical code. Annu. Rev. Neurosci. 2000;23:613–47.
[14] Lestienne R. Spike timing, synchronization and information processing on the sensory side of the central nervous system. Progr. Neurobiol. 2001;65:545–91.
[15] Alberts B, Bray D, Lewis J, Raff M, Roberts K, Watson JD. Molecular Biology of the Cell. Third ed. New York: Garland; 1994.
[16] Kandel E, Schwartz J, Jessell T. Principles of Neural Science. Fourth ed. New York: McGraw–Hill; 2000.
[17] Wandell BA. Foundations of Vision. Sunderland, MA: Sinauer; 1995.
[18] Shannon CE. A Mathematical theory of communication. Bell Syst. Tech. J. 1948;27:379–423.
[19] Optican LM, Richmond BJ. Temporal encoding of two-dimensional patterns by single units in primate inferior temporal cortex: III. Information theoretic analysis. J. Neurophysiol. 1987;57:162–78.
[20] Bialek W, Steveninck R, Warland D. Reading a neural code. Science 1991;252:1854–7.
[21] Kreiman G, Krahe R, Metzner W, Koch C, Gabbiani F. Robustness and variability of neuronal coding by amplitude sensitive afferents in the weakly electric fish *Eigenmannia*. J. Neurophys. 2000;84:189–204.
[22] Koch C. Biophysics of Computation. New York: Oxford Univ. Press; 1999.
[23] Glickstein M. The discovery of the visual cortex. Sci. Am. 1988;September:84–91.
[24] Newsome W, Pare E. A selective impairment of motion perception following lesions of the middle temporal visual area (MT). J. Neurosci. 1988;8:2201–11.
[25] Gross CG. How inferior temporal cortex became a visual area. Cereb. Cortex 1994;5:455–69.
[26] Tanaka K. Neuronal mechanism of object recognition. Science 1993;262:685–8.
[27] Logothetis NK, Pauls J, Poggio T. Shape representation in the inferior temporal cortex of monkeys. Curr. Biol. 1995;5:552–63.
[28] Logothetis NK, Sheinberg DL. Visual object recognition. Annu. Rev. Neurosci. 1996;19:577–621.
[29] Riesenhuber M, Poggio T. Hierarchical models of object recognition in cortex. Nat. Neurosci. 1999;2:1019–25.
[30] Logothetis NK. The underpinnings of the BOLD functional magnetic resonance imaging signal. J. Neurosci. 2003;23:3963–71.
[31] Johnston D, Wu SM. Foundations of Cellular Neurophysiology. Cambridge, MA: MIT Press; 1995.
[32] Raichle M. Visualizing the mind. Sci. Am. 1994:58–64.
[33] Halgren E, Raij T, Marinkovic K, Jousmaki V. Hari R. Cognitive response profile of the human fusiform face area as determined by MEG. Cereb. Cortex 2000;10:69–81.
[34] Srinivassan R, Russell D, Edelman G, Tononi G. Increased synchronization of neuromagnetic responses during conscious perception. J. Neurosci. 1999;19:5435–48.
[35] Pauling L, Coryell C. The magnetic properties and structure of hemoglobin, oxyhemoglobin and carbonmonoxyhemoglobin. Proc. Natl. Acad. Sci. USA 1936;22:210–6.
[36] Logothetis N, Pauls J, Augath M, Trinath T, Oeltermann A. Neurophysiological investigation of the basis of the fMRI signal. Nature 2001;412:150–7.
[37] Logothetis N. The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. Philos. Trans. R. Soc. London B Biol. Sci. 2002;357:1003–37.

[38] Mrsic-Flogel T, Hübener T, Bonhoeffer T. Brain mapping: New wave optical imaging. Curr. Biol. 2003;13:778–80.
[39] Chapman B, Stryker M, Bonhoeffer T. Development of orientation preference maps in ferret primary visual cortex. J. Neurosci. 1996;16:6443–53.
[40] Cox C, Denk W, Tank D, Svoboda K. Action potentials reliably invade axonal arbors of rat neocortical neurons. Proc. Natl. Acad. Sci. USA 2000;97:9724–8.
[41] Goldberg J, Tamas G, Aronov D, Yuste R. Calcium microdomains in aspiny dendrites. Neuron 2001;40:807–21.
[42] Potter SM. Two photons are better than one. Curr. Biol. 1996;6:1595–8.
[43] Shoham D, Glaser D, Arieli A, Grinvald A. Imaging cortical dynamics at high spatial and temporal resolution with novel blue voltage-sensitive dyes. Neuron 1999;24:791–802.
[44] Scoville WB, Milner B. Loss of recent memory after bilateral hippocampal lesions. J. Neurol. Neurosurg. Psych. 1957;20:11–21.
[45] Weiskrantz L. Blindsight revisited. Curr. Oppin. Neurobiol. 1996;6:215–20.
[46] Holmes E, Gross C. Stimulus equivalence after inferior temporal lesions in monkeys. Behav. Neurosci. 1984;98:898–901.
[47] Sperry R. Some effects of disconnecting the cerebral hemispheres. Science 1982;217:1223–6.
[48] Kanwisher N, McDermott J, Chun MM. The fusiform face area: A module in human extrastriate cortex specialized for face perception. J. Neurosci. 1997;17:4302–11.
[49] Epstein R, Harris A, Stanley D, Kanwisher N. The parahippocampal place area: Recognition, navigation, or encoding?. Neuron 1999;23:115–25.
[50] Allison T, Ginter H, McCarthy G, Nobre AC, Puce A, Luby M, Spencer DD. Face recognition in human extrastriate cortex. J. Neurophys. 1994;71:821–5.
[51] McCarthy G, Puce A, Belger A, Allison T. Electrophysiological studies of human face perception: II. Response properties of face-specific potentials generated in occipitotemporal cortex. Cereb. Cortex 1999;9:431–44.
[52] Adrian E. The impulses produced by sensory nerve endings: Part 2. The response of a single end-organ. J. Physiol. 1926;61:151–71.
[53] Quiroga R, Nadasdy N, Ben-Shaul Y. Unsupervised spike sorting with wavelets and superparamagnetic clustering. Neural Comput. 2004. In Press.
[54] Lewicki MS. A review of methods of spike sorting: The detection and classification of neural action potentials. Network: Comput. Neural Syst. 1998;9:R53–78.
[55] Hubel DH, Wiesel TN. Early exploration of the visual cortex. Neuron 1998;20:401–12.
[56] Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, Yuste R. Synfire chains and cortical songs: temporal modules of cortical activity. Science 2004;304:559–64.
[57] Gabbiani F, Krapp H, Koch C, Laurent G. Multiplicative computation in a visual neuron sensitive to looming. Nature 2002;420:320–4.
[58] Laughlin SB, van Steveninck RRR, Anderson JC. The metabolic cost of neural information. Nat. Neurosci. 1998;1:36–41.
[59] Britten KH, Shadlen MN, Newsome WT, Movshon JA. The analysis of visual motion: A comparison of neuronal and psychophysical performance. J. Neurosci. 1992;12:4745–65.
[60] Baker JF, Petersen SE, Newsome WT, Allman JM. Visual response properties of neurons in four extrastriate visual areas of the owl monkey (*Aotus trivirgatus*): A quantitative comparison of medial, dorsomedial, dorsolateral, and middle temporal areas. J. Neurophysiol. 1981;45:397–416.
[61] Celebrini S, Newsome W. Neuronal and psychophysical sensitivity to motion signals in extrastriate area MST of the macaque monkey. J. Neurosci. 1994;14:4109–24.
[62] Shadlen M, Britten K, Newsome W, Movshon J. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. J. Neurosci. 1996;16:1486–510.
[63] Desimone R, Albright T, Gross C, Bruce C. Stimulus-selective properties of inferior temporal neurons in the macaque. J. Neurosci. 1984;4:2051–62.
[64] Shadlen MN, Newsome WT. Noise, neural codes and cortical organization. Curr. Oppin. Neurobiol. 1994;4:569–79.
[65] Barlow H. Single units and sensation: A neuron doctrine for perception. Perception 1972;1:371–94.
[66] Poggio T. A theory of how the brain might work. Cold Spring Harb. Sym. Quantit. Biol. 1990;55:899–910.
[67] Kreiman G, Koch C, Fried I. Category-specific visual responses of single neurons in the human medial temporal lobe. Nat. Neurosci. 2000;3:946–53.

[68] Kreiman G, Koch C, Fried I. Imagery neurons in the human brain. Nature 2000;408:357–61.
[69] Kreiman G. On the Neuronal Activity in the Human Brain During Visual Recognition, Imagery and Binocular Rivalry, Ph. D. thesis. Pasadena: California Institute of Technology; 2002.
[70] Kreiman G, Fried I, Koch C. Single neuron correlates of subjective vision in the human medial temporal lobe. Proc. Natl. Acad. Sci. USA 2002;99:8378–83.
[71] Perez-Orive J, Mazor O, Turner GC, Cassenaer S, Wilson RI, Laurent G. Oscillations and sparsening of odor representations in the mushroom body. Science 2002;297:359–65.
[72] Hahnloser RH, Kozhevnikov AA, Fee MS. An ultra-sparse code underlies the generation of neural sequences in a songbird. Nature 2002;419:65–70.
[73] Gross C. Genealogy of the "Grandmother Cell". The Neuroscientist 2002;8:512–8.
[74] vanRullen R, Thorpe S. Surfing a spike wave down the ventral stream. Vis. Res. 2002;42:2593–615.
[75] Konishi M. Deciphering the brain's codes. Neural Comput. 1991;3:1–18.
[76] Konishi M. Coding of auditory space. Annu. Rev. Neurosci. 2003;26:31–55.
[77] Pena J, Konishi M. Auditory spatial receptive fields created by multiplication. Science 2001;292:249–52.
[78] Heiligenberg W. Neural Nets in Electric Fish. Cambridge, MA: MIT Press; 1991.
[79] Carr C, Heiligenberg W, Rose G. A time comparison circuit in the electric fish midbrain. I. Behavior and physiology. J. Neurosci. 1986;6:107–10.
[80] Brivanlou I, Warland D, Meister M. Mechanisms of concerted firing among retinal ganglion cells. Neuron 1998;20:527–39.
[81] Marsalek P, Koch C, Maunsell J. On the relationship between synaptic input and spike output jitter in individual neurons. Proc. Natl. Acad. Sci. USA 1997;94:735–40.
[82] Koch C, Rapp M, Segev I. A brief history of time constants. Cereb. Cortex 1996;6:93–101.
[83] Reinagel P. How do visual neurons respond in the real world?. Curr. Oppin. Neurobiol. 2001;11:427–42.
[84] Stanley G, Li F, Dan Y. Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus. J. Neurosci. 1999;19:8036–42.
[85] Warzecha AK, Egelhaaf M. Variability in spike trains during constant and dynamic stimulation. Science 1999;283:1927–30.
[86] van Steveninck R, Lewen GD, Strong SP, Koberle R, Bialek W. Reproducibility and variability in neural spike trains. Science 1997;275:1805–8.
[87] Wessel R, Koch C, Gabbiani F. Coding of time-varying electric field amplitude modulations in a wave-type electric fish. J. Neurophys. 1996;75:2280–93.
[88] Gabbiani F, Metzner W, Wessel R, Koch C. From stimulus encoding to feature extraction in weakly electric fish. Nature 1996;384:564–7.
[89] Georgopoulos A, Schwartz A, Kettner R. Neuronal population coding of movement direction. Science 1986;233:1416–9.
[90] Rolls ET, Treves A, Tovee MJ. The representational capacity of the distributed encoding of information provided by populations of neurons in primate temporal visual cortex. Exp. Brain Res. 1997;114:149–62.
[91] Mehring C, Rickert J, Vaadia E, Cardosa de Oliveira S, Aertsen A, Rotter S. Inference of hand movements from local field potentials in monkey motor cortex. Nat. Neurosci. 2003;6:1253–4.
[92] Pouget A, Dayan P, Zemel R. Inference and computation with population codes. Annu. Rev. Neurosci. 2003;26:381–410.
[93] Nirenberg S, Carcieri S, Jacobs A, Latham P. Retinal ganglion cells act largely as independent encoders. Nature 2001;411:698–701.
[94] Krahe R, Kreiman G, Gabbiani F, Koch C, Metzner W. Stimulus encoding and feature extraction by multiple pyramidal cells in the hindbrain of weakly electric fish. J. Neurosci. 2002;22:2374–82.
[95] Zohary E, Shadlen MN, Newsome WT. Correlated neuronal discharge rate and its implications for psychophysical performance. Nature 1994;370:140–3.
[96] Reich D, Mechler F, Victor J. Independent and redundant information in nearby cortical neurons. Science 2001;294:2566–8.
[97] Fries P, Roelfsema PR, Engel AK, Konig P, Singer W. Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. Proc. Natl. Acad. Sci. USA 1997;94:12699–704.

[98] Engel A, Konig P, Singer W. Direct physiological evidence for scene segmentation by temporal coding. Proc. Natl. Acad. Sci. USA 1991;88:9136–40.
[99] Riesenhuber M, Poggio T. Are cortical models really bound by the "binding problem"?. Neuron 1999;24:87–93.
[100] Shadlen M, Movshon J. Synchrony unbound: A critical evaluation of the temporal binding hypothesis. Neuron 1999;24:67–77.
[101] Steinmetz P, Roy A, Fitzgerald P, Hsiao S, Johnson K, Niebur E. Attention modulates synchronized neuronal firing in primate somatosensory cortex. Nature 2000;404:187–90.
[102] Alonso J, Usrey W, Reid R. Precisely correlated firing in cells of the lateral geniculate nucleus. Nature 1996;383:815–9.
[103] Stopfer M, Bhagavan S, Smith BH, Laurent G. Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. Nature 1997;390:70–4.
[104] Allen C, Stevens CF. An evaluation of causes for unreliability of synaptic transmission. Proc. Natl. Acad. Sci. USA 1994;91:10380–3.
[105] Snider RK, Kabara JF, Roig BR, Bonds AB. Burst firing and modulation of functional connectivity in cat striate cortex. J. Neurophysiol. 1998;80:730–44.
[106] Gabbiani F, Krahe R. Burst firing in sensory systems. Nat. Rev. Neurosci. 2004;5:13–24.
[107] Csicsvari J, Hirase H, Czurko A, Buzsaki G. Reliability and state dependence of pyramidal cell-interneuron synapses in the hippocampus: An ensemble approach in the behaving rat. Neuron 1998;21:179–89.
[108] Lisman J. Bursts as a unit of neural information: Making unreliable synapses reliable. Trends Neurosci. 1997;20:38–43.
[109] Martinez-Conde S, Macknik SL, Hubel DH. Microsaccadic eye movements and firing of single cells in the striate cortex of macaque monkeys. Nat. Neurosci. 2000;3:251–8.
[110] Reinagel P, Godwin D, Sherman S, Koch C. Encoding of visual information by LGN bursts. J. Neurophys. 1999;81:2558–69.
[111] Leinekugel X, Khazipov R, Cannon R, Hirase H, Ben-Ari Y, Buzsaki G. Correlated bursts of activity in the neonatal hippocampus in vivo. Science 2002;296:2049–52.
[112] Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. J. Physiol. 1952;117:500–44.
[113] Softky W, Koch C. The highly irregular firing of cortical-cells is inconsistent with temporal integration of random EPSPs. J. Neurosci. 1993;13:334–50.
[114] Stevens CF, Zador AM. Input synchrony and the irregular firing of cortical neurons. Nat. Neurosci. 1998;1:210–7.
[115] Koch C, Segev I. The role of single neurons in information processing. Nat. Neurosci. 2000;3:1171–7. (Supplement).
[116] Poggio T, Torre V, Koch C. Computational vision and regularization theory. Nature 1985;317:314–9.
[117] Markram H, Lubke J, Frotscher M, Sakmann B. Regulation of synaptic efficacy by coincidence of post-synaptic APs and EPSPs. Science 1997;275:213–5.
[118] Magee J, Johnston D. A synaptically controlled associative signal for Hebbian plasticity in hippocampus. Science 1997;257:209–13.
[119] Bi GQ, Poo MM. Synaptic modifications in cultured hippocampal neurons: Dependence on spike timing, synaptic strength, and post-synaptic cell type. J. Neurosci. 1998;18:10464–72.
[120] Kaas J. Topographic maps are fundamental to sensory processing. Brain Res. Bull. 1997;44:107–12.
[121] Kamitani Y, Shimojo S. Manifestation of scotomas created by transcranial magnetic stimulation of human visual cortex. Nat. Neurosci. 1999;2:767–71.
[122] Pascual-Leone A, Walsh V, Rothwell J. Transcranial magnetic stimulation in cognitive neuroscience—virtual lesion, chronometry, and functional connectivity. Curr. Oppin. Neurobiol. 2000;10:232–7.
[123] Penfield W, Perot P. The brain's record of auditory and visual experience: A final summary and discussion. Brain 1963;86:595–696.
[124] Libet B. Brain stimulation in the study of neuronal functions for conscious sensory experiences. Hum. Neurobiol. 1982;1:235–42.
[125] Fried I, Mateer C, Ojemann G, Wohns R, Fedio P. Organization of visuospatial functions in human cortex. Brain 1982;105:349–71.
[126] Fried I, Wilson C, MacDonald K, Behnke E. Electric current stimulates laughter. Nature 1998;391:650.
[127] Salzman CD, Murasugi CM, Britten KH, Newsome WT. Microsimulation in visual area MT: Effects on direction discrimination performance. J. Neurosci. 1992;12:2331–55.
[128] Rees G, Kreiman G, Koch C. Neural correlates of consciousness in humans. Nature Rev. Neurosci. 2002;3:261–70.

[129] Crick F, Koch C. Consciousness and neuroscience. Cereb. Cortex 1998;8:97–107.
[130] DeAngelis GC, Newsome WT. Perceptual "read-out" of conjoined direction and disparity maps in extrastriate area MT. PLoS Biol. 2004;2:E77.
[131] Ditterich J, Mazurek M, Shadlen M. Microstimulation of visual cortex affects the speed of perceptual decisions. Nat. Neurosci. 2003;6:792–3.
[132] Romo R, Hernandez A, Zainos A, Salinas E. Somatosensory discrimination based on cortical microstimulation. Nature 1998;392:387–90.
[133] Tehovnik EJ. Electrical stimulation of neural tissue to evoke behavioral responses. J. Neurosci. Methods 1996;65:1–17.
[134] Brecht M, Schneider M, Sakmann B, Margrie T. Whisker movements evokes by stimulation of single pyramidal cells in rat motor cortex. Nature 2004;427:704–10.
[135] Perkel D, Gerstein G, Moore G. Neuronal spike trains and stochastic point processes II: Simultaneous spike trains. Biophys. J. 1967;7:419–40.
[136] Fries P, Reynolds J, Rorie A, Desimone R. Modulation of oscillatory neuronal synchronization by selective visual attention. Science 2001;23:1560–3.
[137] Holthoff K, Tsay D, Yuste R. Calcium dynamics of spines depend on their dendritic location. Neuron 2002;33:425–37.
[138] Nelson S. Cortical microcircuits: Diverse or cannonical?. Neuron 2002;36:19–27.
[139] Slimko E, McKinney S, Anderson D, Davidson N, Lester H. Selective electrical silencing of mammalian neurons in vitro by the use of invertebrate ligand-gated chloride channels. J. Neurosci. 2002;22:7373–9.
[140] Lechner HAE, Lein ES, Callaway E. Selective and quickly reversible silencing of mammalian neurons using an insect G protein-coupled receptor. J. Neurosci. 2002;22:5287–90.
[141] Lois C, Hong E, Pease S, Brown E, Baltimore D. Germline transmission and tissue-specific expression of transgenes delivered by lentiviral vectors. Science 2002;295:868–72.
[142] Chapin JK, Moxon KA, Markowitz RS, Nicolelis MAL. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. Nat. Neurosci. 1999;2:664–70.
[143] Heintz N. Bac to the future: The use of bac transgenic mice for neuroscience research. Nat. Rev. Neurosci. 2001;2:861–70.
[144] Zirlinger M, Kreiman G, Anderson D. Amygdala-enriched genes identified by microarray technology are restricted to specific amygdaloid sub-nuclei. Proc. Natl. Acad. Sci. USA 2001;98:5270–5.
[145] Sejnowski T, Koch C, Churchland P. Computational neuroscience. Science 1988;241:1299–306.
[146] Mainen ZF, Sejnowski TJ. Reliability of spike timing in neocortical neurons. Science 1995;268:1503–6.