Chapter 1
Introduction

Two friends, one living in the city and the other on the family farm, describe to one another the experiences of everyday life. The farmer conjures up pastoral images, acres of wheat swaying in a gentle breeze, the sweet smells of spring, and the songs of the birds. The city dweller recounts scenes of thousands of people emerging from the train station, the inescapable odors of traffic, and the throbbing beat of a street musician's drums. It would seem that these sensory experiences are as different as one could imagine, yet they share with all our sensory experiences one crucial feature: In each case, our perception of the world is constructed out of the raw data sent to our brains by our sensory nerves, and in each case these data come in the same standard form—as sequences of identical voltage pulses called action potentials or "spikes."

When we see, we are not interpreting the pattern of light intensity that falls on our retina; we are interpreting the pattern of spikes that the million cells of our optic nerve send to the brain. When we hear, we are not interpreting the patterns of amplitude and frequency modulation that characterize the acoustic waveform; we are interpreting the patterns of spikes from roughly thirty thousand auditory nerve fibers. All the myriad tasks our brains perform in the processing of incoming sensory signals begin with these sequences of spikes. When it comes time to act on the results of these computations, the brain sends out sequences of spikes to the motor neurons. Spike sequences are the language for which the brain is listening, the language the brain uses for its internal musings, and the language it speaks as it talks to the outside world.

If spikes are the language of the brain, we would like to provide a dictionary. We would like to understand the structure of this dictionary, perhaps even providing the analog of a thesaurus. We would like to know if, as in language, there are notions of context that can influence the meaning of the individual words. And of course we would like to know whether our use of the linguistic analogy makes sense. We must travel a long road even to give these questions a precise formulation. We begin at the beginning, more than two centuries ago.

1.1 THE CLASSICAL RESULTS

Our understanding of how the sensory world is represented in the electrical activity of the sensory nerves is limited, first and foremost, by our ability to record this activity. Indeed, the history of experiments on the electrical activity of nerves is intertwined with the history of electrical measurements more generally. The science of electricity as we understand it today began with Galvani and Volta in the 1700s (Pera 1986). Galvani observed that the muscles of a frog could be made to twitch when touched with a piece of metal, and he believed that the metal evoked "animal electricity" in the muscle. Volta suspected that the electricity was generated at the contact point itself, and that similar effects should be observable from a contact between different inorganic materials. Volta was right, and the pursuit of his ideas led him to what we now call a Voltaic pile, the first real battery. The fact that electricity was not the special provenance of animals was one of the first nails in the coffin of vitalism.

Galvani and Volta made macroscopic measurements. Their biological preparations consisted of large hunks of muscle—often the entire muscle—not what we now know to be the single muscle fibers or motor neurons that make up these tissues. The notion that the body is constructed from cells emerged only through the efforts of the nineteenth-century microscopists, which culminated in the beautiful observations of Ramón y Cajal on the cellular nature of the brain itself (Cajal 1909–11). As a more microscopic picture of the nervous system began to take shape, it seemed natural to ask how the activity of individual cells might relate to our perceptions. Müller developed the doctrine of specific nerve energies, according to which the identity of a sensory stimulus is represented by the fact that certain nerve fibers, and not others, are activated by that stimulus (Boring 1942). Helmholtz provided evidence for this view in his analysis of the inner ear, arguing that cells at different locations along the cochlear spiral are sensitive to different frequencies of sound (Helmholtz 1885). These discussions from the late nineteenth century form the foundation for much of our current thinking about the nervous system. When we read about a computational map in the cortex (Knudsen, du Lac, and Esterly 1987), where an array of neurons decomposes incoming signals according to the values of different component features, we are reminded of Helmholtz, who realized that the array of auditory nerve fibers would decompose sound into its component frequencies.

Testing the ideas of Helmholtz and Müller requires the direct observation of electrical activity in *individual* sensory neurons, not just the summed activity

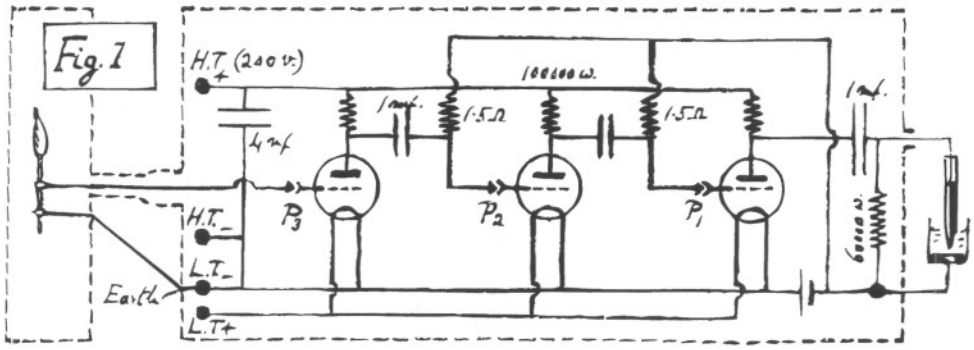

Figure 1.1
Schematic of Adrian's apparatus for recording the electrical activity in a nerve fiber. The fiber itself is at the far left. Adrian placed the fiber across two electrodes and measured the difference in the voltage at these two points along the nerve. The signal was amplified and used to control a mercury column, at the far right. Records were obtained by scanning a piece of film behind the mercury column, an example of which is shown in Fig. 1.3. Redrawn from Adrian (1926).

of a nerve bundle. But the electrical signals from individual cells are very small, at least when seen by an observer outside the cell. To pick up these small signals required a new method of low noise amplification, and this was provided in the first decade of this century by the vacuum tube. Using these new devices at Cambridge University, Lucas (1917) built instruments which allowed the recording of microvolt signals in bandwidths of several kiloHertz. We should remember that these experiments predate the oscilloscope, so even the display of submillisecond signals posed a significant problem. The solution to this problem, together with a general schematic of the instruments, is shown in Fig. 1.1. Lucas, sadly, died young, and the task of using these instruments fell to E. D. Adrian. In the space of roughly a decade, Adrian learned much of what we know to this day about the problem of neural coding. Independently, H. K. Hartline made many of the same discoveries. We follow the line of reasoning laid out by Adrian, and return shortly to some special features of Hartline's observations.

The classic early work of Adrian is contained, primarily, in a series of papers published in 1926 (Adrian 1926; Adrian and Zotterman 1926a, 1926b). Adrian summarized these results and their implications in a (still) very readable monograph, *The Basis of Sensation* (1928). One can trace the evolution of Adrian's thinking in two subsequent books (Adrian 1932, 1947).

Adrian's experiments established three fundamental facts about the neural code. First, he saw that individual sensory neurons produce stereotyped action potentials, or spikes. This is the all-or-none law, which had already been established for muscles and motor neurons: Incoming stimuli either produce action potentials, which propagate long distances along the cell's axon, or they do not; there are no intermediate signaling mechanisms. This means that a single neuron can provide information to the brain only through the arrival times of the spikes.

To make Adrian's observations a bit clearer, we look at a modern version of the same experiment. In Fig. 1.2 we show raw data from a fine tungsten wire electrode which has been placed close to a single neuron in the brain of a fly; the voltage at this electrode is measured relative to that at a reference electrode placed in the body fluids. Although the trace is noisy, there are clear, stereotyped events that can be isolated by appropriate filtering. These are the action potentials or spikes produced by this neuron and seen from outside the cell. The observation of all-or-none responses raises several questions:[1] Why does the nervous system choose this mode of communication? How is the stereotyped action potential waveform selected and stabilized? Is this mechanism universal?

Action potential propagation is an active process—the cell expends energy to produce and transmit a spike, and the energy expenditure increases the farther the spike must travel. In the absence of active processes, the electrical properties of cell membranes are such that a pulse starting at one end of a cell would spread and decay rather than propagating at constant velocity, and the characteristic decay lengths are on the order of one millimeter (Hodgkin and Rushton 1946). Therefore, passive mechanisms are inadequate for sending signals over long distances, such as the roughly one meter from your fingertips to your spinal cord, or even from one area of the cortex to a neighboring area; action potentials provide the means for such long distance communication. On the other hand, cells that send signals only over short distances, such as within the retina or even across the body of a small animal, need not generate action potentials and can, instead, operate entirely with "graded" voltage responses to sensory stimuli (Roberts and Bush 1981); we will see examples of this more continuous mode of neural signalling in section 3.1.4.

1. The experiments and theoretical developments which provided the answers to these questions are by now classic chapters in the history of neuroscience (Aidley 1989). We provide only a brief summary, but we encourage the reader to look at the original papers, as well as the lovely text by Katz (1966). Some of the history is recounted in the essays collected for the one-hundredth anniversary of the Physiological Society (Hodgkin et al. 1977).

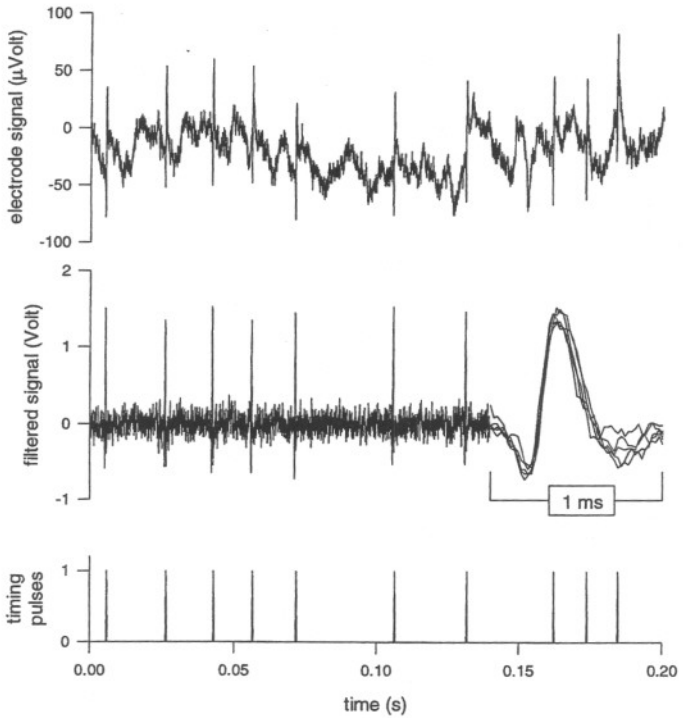

Figure 1.2
All-or-none coding by action potentials. Each action potential generated by the cell has a similar shape. Thus action potentials are the elementary units of the neural code. The top panel shows the difference between the voltage recorded with a fine tungsten wire placed near a cell in the fly's brain and that recorded with a reference electrode placed in the body fluid. The middle panel shows the same voltage after band-pass filtering to separate the relatively high frequency components in the action potential from low frequency noise; after filtering, the shapes of individual action potentials are quite similar. At the right, five action potentials are shown overlaid on an expanded time scale. This gives an impression of the shape and of the reproducibility of the time course. The bottom panel shows timing pulses generated electronically by a threshold discriminator circuit.

The local circuit properties of a cell membrane include active elements, conductances that are modulated by voltage changes and are electrically in series with power supplies (or, effectively, batteries) that are maintained by ion pumps; these pumps in turn are powered by chemical energy from the cell's metabolism. Hodgkin and Huxley (1952a, 1952b, 1952c) analyzed the electrical dynamics of the cell membrane in the giant axon of squid, and showed that these dynamics could be described with relatively simple phenomenological models of conductances that depend on voltage and are selective for different ions. When these local, active elements are assembled into a long cable, such as the axon, the nonlinear dynamics of the conductances select a stereotyped pulse which can propagate at constant velocity, while all other voltage changes eventually decay; the great triumph of this work was to show that this pulse has a shape and speed essentially identical to the observed action potentials (Hodgkin and Huxley 1952d). The mathematics of pulse selection has its roots in the nineteenth century, but a complete theory came much later (Aronson and Weinberger 1978), and the Hodgkin–Huxley equations continue to provide the inspiration for interesting mathematics and physics problems.

Although their analysis was purely phenomenological, the form of the Hodgkin–Huxley equations suggested a microscopic picture in which the conductances selective for different ions correspond to different molecular elements, or channels, in the membrane, and the modulations of the conductance correspond to transitions among discrete states of these channel molecules. Continuing advances in low noise amplification made it possible to resolve the electrical noise generated by spontaneous transitions among the different channel states, and finally to detect the currents flowing through single channel molecules (Sakmann and Neher 1983). Measurements on the properties of individual channel molecules, together with the techniques of modern molecular biology, have made it possible to identify a great diversity of channel types (Hille 1992), but these studies also demonstrate that many features of channel structure and function are strongly conserved throughout the animal kingdom (Jan and Jan 1994). This universality of mechanism at the molecular level harks back to Adrian's observations on the universality of spike encoding. Over the years, Adrian and his colleagues recorded the activity of sensory neurons from an enormous variety of different sensory systems in different animals. Although the quantitative details vary from neuron to neuron, it seems that the principles are universal.

The second of Adrian's fundamental observations was that, in response to a static stimulus such as a continuous load on a stretch receptor, the rate of spiking increases as the stimulus becomes larger. The raw data from Adrian's

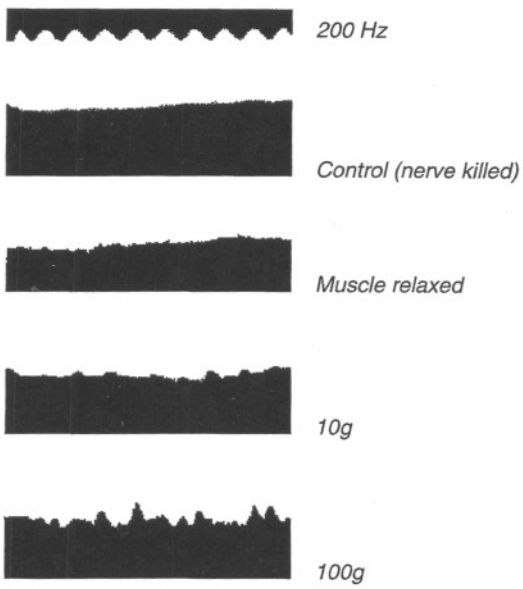

Figure 1.3
Firing rate as a function of stimulus strength, adapted from Adrian and Zotterman (1926a). The spikes in these panels are visible as the fluctuations riding on the black-white interface. A time marker is shown on top. Adrian and Zotterman measured the relation between the force applied to a muscle and the firing rate in a stretch receptor embedded in the muscle. Different forces were generated by hanging weights with different masses from the muscle. This type of experiment established that the frequency of firing in sensory neurons increased with increasing stimulus strength.

original demonstration of this principle is shown in Fig. 1.3, and a quantitative analysis is shown in Fig. 1.4a. Thus the rate, or frequency, of spikes indicates the intensity of the stimulus. To be a bit more precise, the number of spikes in a fixed time window following the onset of a static stimulus represents the intensity of that stimulus. This is the idea of *rate coding.*

The third of Adrian's discoveries was that if a static stimulus is continued for a very long time, the spike rate begins to decline, as illustrated in Fig. 1.4b. This is called *adaptation,* although this term is also used more generally to describe a dependence of the neural response on the history of stimulation. Adrian suggested that this physiological phenomenon corresponds to perceptual phenomena wherein we become gradually unaware of constant stimuli.

As we have tried to find a precise modern formulation for the problem of neural coding, we have been struck by the extent to which the ideas of Adrian

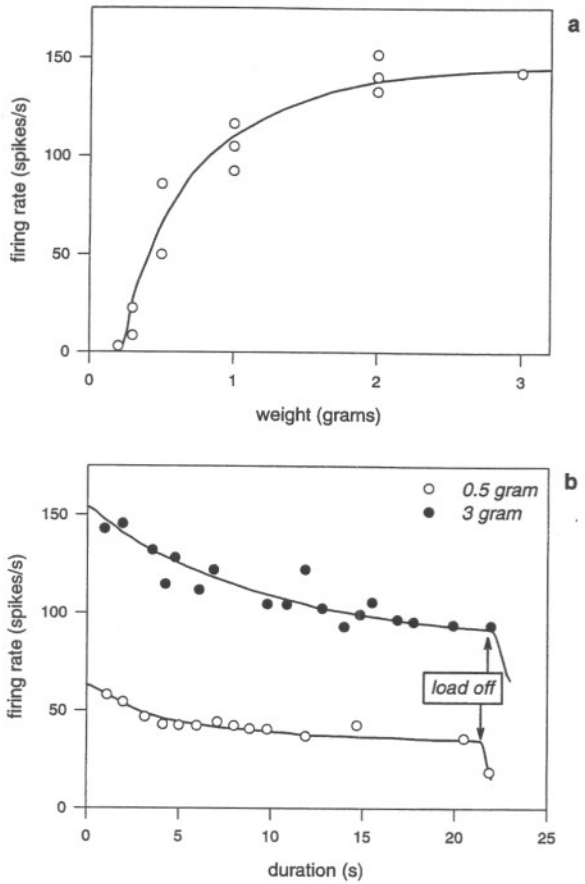

Figure 1.4
Rate coding and adaptation. (a) Average firing rate of a stretch receptor as a function of the weight applied to the muscle, in an experiment similar to that of Fig. 1.3. (b) Decrease in firing rate with time following the onset of a static stimulus at $t = 0$, adapted from Adrian (1926). This desensitization, or *adaptation*, is a general property of neural coding.

and Hartline have formed the paradigm for subsequent exploration of the nervous system. On the one hand this must mean that their early experiments captured essential and universal features of the neural code. On the other hand one must worry that, in following this single line of ideas, some crucial points may have been missed.

In the first experiments on single sensory neurons the stimulus was often defined by a single parameter. This parameter, such as the load on a stretch receptor, was held fixed while the stimulus was on. But naturally occurring stimuli are defined by a much larger number of parameters. In vision, for example, a small region of the visual field may be described by its overall luminance, but also by its contrast relative to the background, the size and shape of any features in the region, the positions and orientations of such features, their color, depth, and so on. By analogy with the Adrian–Hartline observations on spike rate as a function of stimulus intensity, one can plot the responses of a visual neuron as a function of these multiple parameters. This leads to the notion of *feature selectivity,* in which the cell's response depends most strongly on a small number of parameters and is maximal at some optimum value of these parameters.

Precursors to the notion of feature selectivity can be found in the work of Hartline and collaborators, who studied the responses of single neurons from the compound eyes of the horseshoe crab *Limulus polyphemus.* In addition to reproducing Adrian's results concerning rate coding, Hartline found that the stimulus whose strength was coded by one neuron reflected the difference between the light intensity at the location of that cell and the intensity at neighboring cells. Thus the crab retina has an enhanced response to spatial contrast or edges. Hartline, Ratliff, and coworkers suggested that this enhancement is connected to the perceptual phenomenon of Mach bands, shown schematically in Fig. 1.5. The unraveling of the retinal circuitry responsible for contrast enhancement led to a long sequence of now classic papers (Ratliff 1974).

The concept of feature selectivity was clearly enunciated by Barlow (1953a, 1953b), who was Adrian's student. Recording from retinal ganglion cells in the frog, he showed that the response of these cells to a spot of light at first grows with the area of the spot, but then declines if the spot exceeds a critical size, as summarized in Fig. 1.6a. The portion of the visual world that can influence the activity of a neuron is called the receptive field of that cell, and Barlow's results can be described as a "center–surround" organization of the receptive field: spots within a small region (the center) excite the cell, but spots just outside this region (in the surround) inhibit the cell (Fig. 1.6b).

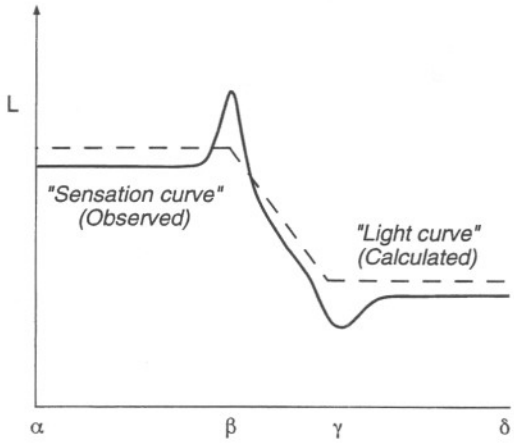

Figure 1.5
Mach bands at the edge of a shadow. The "Light curve" is based on physical calculations of the luminance at the edge of a shadow. The point α is in the fully illuminated space, the point β is at the outer edge of the shadow, the point γ is at the inner edge of the shadow, and the point δ is in the full shadow. The thicker line represents the apparent luminance, or "Sensation curve" actually observed. The maximum and minimum of these curves correspond to the light and dark Mach bands that arise from differencing mechanisms in the visual system that enhance contrast. Redrawn from Ratliff (1974).

To a good approximation these receptive fields are circularly symmetric. Essentially identical receptive fields were found in cat retinal ganglion cells by Kuffler (1953). In many cases the excitation and inhibition are balanced so that spatially uniform stimulation produces no response. Another interpretation is that these cells are tuned to objects of a given apparent size, perhaps that of the bugs the frog is hunting. The picture of frog retinal ganglion cells as specialized "bug detectors" was emphasized by Lettvin and coworkers (1959). In the limiting case this view presents sensory neurons as yes/no devices, signaling the presence or absence of certain elementary features.

The importance of feature selectivity was strongly supported by the observations of Kuffler's colleagues Hubel and Wiesel (1962). They found that many cells in cat visual cortex are selective not only for the size of objects (e.g., the width of a bar) but also for their orientation. As in the Barlow–Kuffler experiments, Hubel and Wiesel observed this selectivity by counting the number of spikes the cell produced in response to the presentation of a static stimulus or in response to the motion of the stimulus through the cell's receptive field. Hubel and Wiesel presented a scenario for how this orientation selectiv-

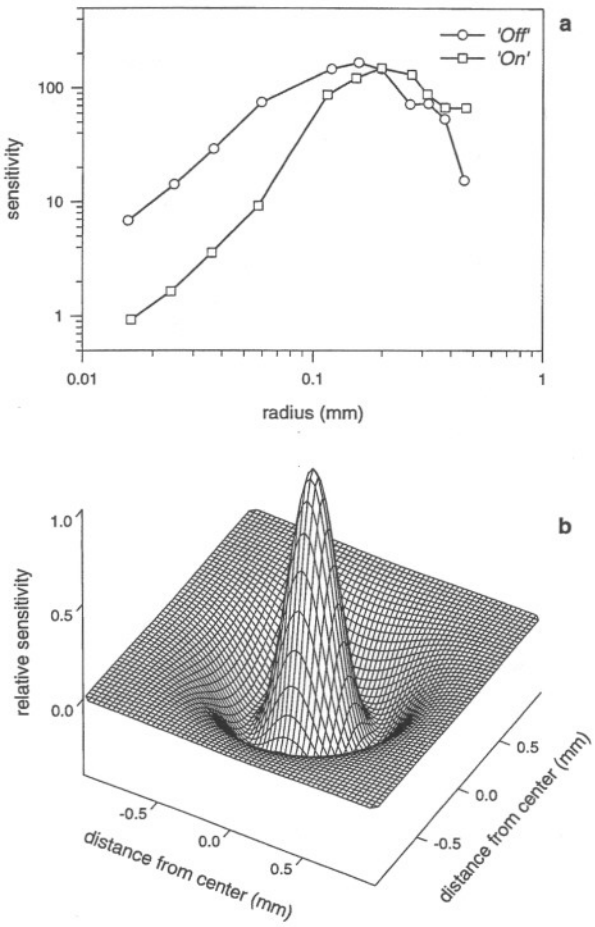

Figure 1.6
Center–surround receptive fields in retinal ganglion cells. (a) Sensitivity of retinal ganglion cells in the frog as a function of the radius of the light stimulus; sensitivity is defined as the light intensity required to elicit a fixed number of spikes. As the stimulus size is increased, the sensitivity initially increases, but then begins to decrease when the stimuli are larger than 0.2 mm in radius. This behavior was seen in both "on" ganglion cells, which respond to an increase in light intensity in the central region of their receptive field, and in "off" ganglion cells, which respond to a decrease in light intensity. (b) Receptive field organization suggested by Barlow to explain measurements such as those in (a). Light falling within the central excitatory region of the cell's receptive field causes an increase in the number of spikes, while light falling in the inhibitory surround causes a decrease in the number of spikes, indicated here as a negative sensitivity. Maximal response to a spot of light is achieved when the stimulus just covers the entire receptive field center.

ity could be built out of center–surround neurons in lower levels of the visual system, making explicit the intuitive notion that higher percepts are built out of elementary features. Finally, they found that neighboring neurons are tuned to neighboring orientations, so that feature selectivity is mapped over the surface of the cortex. This notion of cortical mapping, presaged by Mountcastle's (1957) observations on the responses of cells in the somatosensory cortex, revealed order amid the seemingly impenetrable mass of cortical circuitry. This discovery led, in turn, to the investigation of how this order develops out of the more amorphous circuitry of the embryonic brain. The ideas of feature selectivity, cortical maps, and self-organization of maps during development have dominated the exploration of cortex ever since (Hubel and Wiesel 1977).

If we return to the original Adrian–Hartline experiments on sensory neurons, we see that one could extend the description of the neural code in two very different directions. One direction is to study the coding of multiparameter stimuli, which has been followed extensively in the exploration of the visual system. A second direction is to use stimuli with realistic time dependencies. In a natural environment, sensory inputs are not broken into discrete presentations, and they are not simply turned on and off. More complex dynamic signals have been used in the study of the auditory system, where the main issues concern recognition and classification of temporal waveforms. But even in these experiments there is a tendency to approximate real dynamic signals with more elaborate but still essentially stationary signals. For example, the coding of vowel sounds has often been studied using continuous, periodic stimuli whose power spectra approximate those of real vowels.

A primary concern in this book is to understand how the nervous system represents signals with realistic time dependencies. The problem of coding for nearly static stimuli is very different from the problems faced by the brain under more natural conditions. In particular, the focus on time-dependent signals forces us to think about the significance of much smaller numbers of spikes. But we are getting ahead of ourselves.

Do the ideas of rate coding, feature selectivity, and cortical mapping tell us what we want to know about the neural code? Certainly the fact that neurons in deeper layers of the brain are selective for more complex features tells us something about the kinds of computations that are carried out as sensory signals are passed from one stage of processing to the next, although it is dangerous to take a hierarchical or sequential view of sensory processing too literally. The idea of rate coding leaves open the question of whether other features of the spike train—generally grouped under the catch phrase *timing*—carry meaningful information, and indeed this question has been central to

many discussions of neural coding. The idea of mapping leads us to think about the representation of the sensory world in arrays of neurons; it also leads to the concepts of ensemble or population coding, which are active topics of current research.

The classical results on the neural code suggest many avenues for exploration. We cannot do justice to all the different paths taken by different investigators. In the following section we hope to make precise a more limited set of questions, which, with luck, we can answer in the space of the remaining text.

1.2 DEFINING THE PROBLEM

What would it mean to say that we "understand" the neural code in a particular region of the nervous system? How do we quantify the notion that the spike train of a single cell "conveys information" about the sensory world? In what sense is a particular sequence of spikes the "right answer" to some computational problem faced by the brain? We search for sharper versions of these questions by forcing ourselves to adopt a more precise and more mathematical language. In talking about the nervous system we routinely make colloquial use of terms such as code, information, and reliability. All of these words can be given precise mathematical definitions, and we hope, through the remainder of the text, to convince the reader that these definitions provide a clear guide to the design and analysis of new experiments. In striving for precision we shall see the emergence of some new ideas. We begin, however, by revisiting an old idea, the homunculus.

The homunculus is an often derided concept in discussions of the brain. We recall that this metaphor conjures up a little man—or, in a lovely variant by Michael Land and Simon Laughlin (Fig. 1.7), a little fly—who observes the responses of his own sensory neurons and finally forms the percepts that the organism experiences. The problem with this picture is that it never gets to the essence of what it means to perceive and to experience the world. On the other hand, as explorers of the nervous system we place ourselves, inevitably, in the position of the homunculus—we observe the responses of sensory neurons and try to decide what these responses could mean to the organism. This problem of assigning meaning to the activity of sensory neurons is the central issue in our discussion of the neural code.

It is easy to imagine that the task of the homunculus is trivial—after all, he just watches a projected image of the world as it flashes through the brain. But this projected image is *encoded* in the patterns of action potentials generated by the sensory neurons. It is not at all clear what the homunculus would have

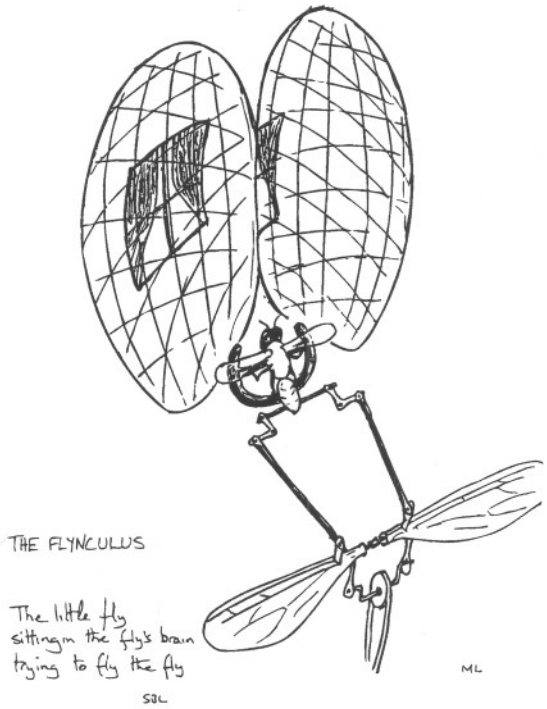

Figure 1.7
The Flynculus. Doodle by M. F. Land, quotation, "The little fly sitting in the fly's brain trying to fly the fly," from S. B. Laughlin, with permission.

to do, even in principle, to make sense out of these encoded data. We propose that "understanding the neural code" means that we would know how to make sense out of the bewildering array of spike trains streaming in from the sense organs: If we understand the code, we can function as the homunculus.

When we ask what a spike train means, or what it can tell us about the world, we need to set some boundaries for the question, or, equivalently, a context for the answer. If we live in a world with only two possible sensory stimuli, we can ask how the homunculus could best use the spike train data to make a decision about which stimulus in fact occurred. This decision rule would constitute a complete understanding of the neural code, assuming that the world offers just two possible signals.

In many psychophysical discrimination experiments (Green and Swets 1966), a world of two alternatives is created artificially, and the subject must

solve the problem of choosing between these alternatives. This binary decision problem provides a convenient context for asking questions about the reliability of our perceptions, and we shall see that it is also useful for investigating the reliability of neurons. But it is not enough to build a homunculus that functions in a world of two alternatives; we want to ask our question about the meaning of spike trains in a context that approaches the complexity of the natural world.

Under natural conditions, the stimulus that will appear in the next brief time window is not known to us in advance. Instead the stimulus is chosen from an infinite set of alternatives. On the other hand, these alternatives are not all equally likely. While there are blue spruce trees, green trees do not suddenly turn blue (or red or yellow either). Natural stimuli develop in time, and these dynamics have some underlying regularity or structure. This structure has a deterministic component, as when a leaf falls downward according to Newton's laws. But since we do not know all the forces that shape the dynamics of sensory stimuli, some aspects of these stimuli are unpredictable, as when the falling leaf is deflected by a gust of wind. The result is that natural signals are presented to us at random, but these signals have correlations that reflect their origins in deterministic physical processes.

Rather than inhabiting a world of two alternatives, we thus inhabit a world of random but correlated time dependent signals. The time dependence is crucial, because it means that we cannot wait forever to decide what we are looking at. Not only does biology press for quick decisions—we must catch our prey and not be caught by predators—the physics of our environment is such that any simple averaging for long periods of time will average away the very signals that interest us. The task of the homunculus, then, is not to create a static image of the sensory world from the input spike trains, but rather to give a sort of running commentary or simultaneous translation. We emphasize that this running commentary need not be, and most likely cannot be, a comprehensive reconstruction of the world around us.

To give meaning to the spike trains nonetheless requires that we recreate at least some aspects of the continuous time dependent world that is encoded in discrete sequences of spikes. From our experience in the laboratory we know that when forced to interpret rapidly changing signals we are very susceptible to noise; usually we try to combat noise by averaging in time or by averaging over repeated presentations of the same signal. But the homunculus is not free to set arbitrary averaging times, and he certainly cannot ask for a second, identical copy of the immediate past. On the contrary, the homunculus (and

the animal as well!) has to reach conclusions about the world from just one example of the spike train in each of his sensory neurons.

In generating a running commentary on the meaning of spike trains we shall have to deal with whatever level of noise is present in these data. Ideally, our interpretation of the spike trains should be as reliable as possible given this noise, and the statistically sophisticated homunculus would report confidence levels on his estimates of what is happening in the world. If understanding the neural code means building a homunculus, we can compare two different candidate homunculi—two different candidates for the structure of the neural code—by comparing the accuracy of their inferences about events in the sensory world.

We are closing in, then, on a more precise definition of the problems in understanding the neural code. We place ourselves in the position of the homunculus, monitoring the spike trains of sensory neurons as stimuli vary in time along some unknown trajectory. We must generate a running commentary on the identity of these stimuli, using only the spike train data as input. Our inferences about events in the world will have some limited accuracy, and we shall have to quantify this accuracy. Out of many possible homunculi, there is one that tells us as much as possible about the world given the noise in the spike train data itself. The performance of this best homunculus will reflect a compromise between averaging in time to combat noise and responding quickly to keep up with the dynamics of the world, and we shall have to be precise about these time scales.

The construction of a complete homunculus, or even the complete flynculus of Fig. 1.7, is a daunting task. In the fly, visual signals stream in along thousands of parallel paths reflecting the array of lenses in the compound eye, and in ourselves and our primate cousins the corresponding numbers are three orders of magnitude larger. There are a few special cases, such as the moths discussed in section 4.1.1 (Roeder 1963), for which it might be possible to monitor all of the spike trains that encode one sensory modality, but in general this is hopeless. As we have noted, however, there is a long tradition of trying to make sense out of the responses of single neurons, always recognizing that one cell can tell us about only a small piece of the sensory environment. In this tradition, most of this book is about the problem of an impoverished homunculus who looks at the spike train of just one neuron at a time; we take a brief look at the problem of multiple neurons in section 5.1. We thus have a clear question, amenable to experimental investigation: What can the spike train of this one neuron tell us about events in the world?

1.3 CENTRAL CLAIMS OF THIS BOOK

Nearly seventy years ago, Adrian summarized the first generation of experiments on neural coding (Adrian 1928). We have argued that, even today, this classic work contains a large fraction of what we know about the language of the brain. Forty years later, Perkel and Bullock (1968) provided an encyclopedic summary of the state of the field, a handbook of diverse candidate coding strategies in different systems. What can we add after all these years?

We believe that there has been substantial progress in both the formulation and the resolution of three major issues regarding coding by single neurons. These three points form the core of our presentation:

1. Representation of time-dependent signals. In a variety of sensory systems, single neurons produce on the order of one spike per characteristic time of stimulus variations—a sparse temporal representation. This is in direct contradiction to a simple, intuitive implementation of the rate coding idea, since the rate is an average quantity not available from a single spike. Sparse temporal codes can be decoded by simple algorithms, even when the encoding is a complex nonlinear process. Thus the problem of *decoding*—the problem solved by our homunculus—may be simpler than the classical problem of encoding.

2. Information rates and coding efficiency. The focus on signals with realistic time dependencies leads to the demonstration that single neurons can transmit large amounts of information, on the order of several bits per spike. In at least one case, signals with more natural temporal correlations are coded more efficiently, so that the spike train provides more information with roughly the same number of spikes. These high rates come close to saturating the fundamental physical limits to information transmission.

3. Reliability of computation. Understanding the reliability of the nervous system requires that we understand the code which the system uses to represent the answers to its computational problems; the study of neural coding is thus tied to much broader issues of neural computation. In several systems there is agreement between at least two of three fundamental quantities: The reliability of behavior, the reliability of single neurons, and the fundamental physical limits to reliability imposed by noise in the sense data itself. It is clear that the approach to the physical limits is closest for the more natural tasks of processing time-dependent signals.

These three ideas provide, we hope, a clear answer to the questions formulated in section 1.2. Decoding the spike train provides a literal construction of

the "running commentary" that we require from the homunculus, the measurement of information transmission rates quantifies how much our impoverished homunculus can tell us by looking at just one neuron, and the observations on reliability place this information on a meaningful scale relative to the capabilities of the whole organism.

In exploring these three issues, we will refer to experimental results from many different systems, obtained by many different groups over a period of several decades. The common thread running through these diverse studies is the attempt to *quantify* the behavior of neurons, specifically under conditions that approximate the function of the nervous system in the life of the organism. Much of the text is also concerned with methodology, reviewing several theoretical approaches that have been proposed as guides to the design and analysis of quantitative experiments. Many of our readers may reasonably wonder whether the effort of building up this more mathematical framework will be rewarded. One reason for persevering is that the quantitative analysis of neural coding leads to surprising results. As devices for transmitting and processing information, neurons are doing much more than one might have expected, and in a precise sense they are doing almost as much as is physically possible. Even simple quantitative questions—how many spikes carry a meaningful signal?—have surprising answers. Thus we claim that the results of a quantitative approach are sufficiently extreme that they begin to alter our qualitative conception of how the nervous system works.