Preface

The idea that circuits might be able to explain brain function never really occurred to me until once when, as a Princeton undergraduate, I heard a lecture about a Cornell psychologist, Frank Rosenblatt, and the machine he had designed, the *Percepron*, which could "learn" to recognize letters of the alphabet. The lecture planted a seed, but it took very long before anything grew from it.

I went to Cornell after graduating from Princeton, to work on a PhD in theoretical physics, and made occasional visits to Rosenblatt's lab. I kept telling him how I didn't think the brain worked the way he described, and he countered by pointing to the large number of "theories" I could choose from. Of course I had no idea how the brain worked, but in the process I started reading books on it.

It was hard to know where to start. Physiology, anatomy, psychology, neurology, computer science – these were all subjects directly relevant to what I needed to know, not to mention that data and facts about the human mind are all around us; and observations too common to require laboratory research are not automatically unimportant.

I finished my PhD work (on solid-state plasma oscillations known as *helicons*), and proceeded to a job with United Aircraft (now United Technologies) in Hartford, as a physicist under a contract which permitted me to spend one third of my time on work I chose, which in my case would be brain theory.

From Hartford I often drove up to MIT, where McCulloch had his office, cluttered with papers, and Lettvin his lab, with the name tag "J. Y. Lettvin, experimental epistemology." They were both still active, and gave generously of their time. At United Aircraft I wrote my first brain theory paper, on ignitable neuron groups (Legéndy, 1967). Influenced by the approach of McCulloch, Rosenblatt, and von Neumann, the 1967 paper attempted to explore the duality between subjective perception and the facts of connectivity, and describe a set of "networks with circles" which are more reliable than their individual elements.

The proper career path for me at this point would probably have been to enroll in a graduate school again and study under some established neuroscientist, but by this time I have become aware that there was nobody able to teach me what I needed to know. I was navigating uncharted waters.

My good luck steered me to a small Westinghouse research group, located right on "Tech Square," a block of research buildings across the street from MIT, where

I was offered a job doing pure research in brain theory, justified as potentially leading to computing machines better suited to pattern recognition than digital computers. Tech Square was the world's greatest place to learn about brains. I often spoke to McCulloch, got to know Minsky and Papert, and, on the other side of the river, at Harvard Medical School, got to visit the labs of Hubel, Wiesel, and Palay.

While at Westinghouse, I wrote down many of the principles which underlie the present book; among them were the need to formulate communication inside the brain in terms of "surprising" events of firing, the conceptual linkage between them and "local knowledge," and the idea that neuron groups representing objects can transmit "syntactic" relations between the objects through prearranged relative timing of their outputs (Legéndy, 1970). I also wrote down the idea of "reaching," and some methods the brain uses to achieve it, such as the "trick of retinotopic mapping" and the "trick of small connective fields" (Remarks on the Brain, unpublished internal report, 1970).

After my Westinghouse money ran out, I took a postdoctoral position in Italy, at the lab of Caianiello (at the CNR Laboratory of Cybernetics at Arco Felice), then another one in Germany (at the University of Tübingen), a few minutes away from the Max Planck Institute of Biological Cybernetics, where Valentino Braitenberg had his group. I spent many good hours at Braitenberg's office looking at Golgi slides through his microscope.

I still had much to learn. For one thing, I had felt since my conversations with Hubel that my education would remain incomplete until learned to do experiments of my own. My years in electrical engineering (which I had studied at Princeton) steered me toward electrophysiology, and an opportunity to learn the techniques soon arose at Otto Creutzfeldt's lab in Göttingen, where I had a chance to join a project with Tadaharu Tsumoto on cortico-geniculate correlations. At Göttingen I went on to spend a very busy one and a half years, and wrote, among other things, a program to make "Poisson surprise scans" of some spike trains I recorded. By the time I returned to the New York City, I knew enough to take postdoctoral jobs as a neurophysiologist, first with Alden Spencer at Columbia University, then with Herb Vaughan at Albert Einstein College of Medicine.

But my career in experimental brain research was short-lived, because after a while I lost my NIH funding and had to take jobs in the aerospace and computer industry. It seemed that my brain research days had come to an end. Except that around this time my good luck once again intervened.

Wayne Wicklegren, formerly of MIT and the University of Oregon, moved to New York and joined his wife Norma Graham at Columbia University. I knew both Wayne and Norma from the literature, as they both knew me; in particular Wayne, encouragingly, had a lively interest in my 1967 paper. We talked about the brain on many afternoons for several years; and by the time we stopped, I was doing brain research once again, this time not experiments but theory. Along with my wife, Annemarie, who patiently stuck it out with me throughout, Wayne deserves my deepest thanks for the push he gave me when I most needed it.

During working hours I continued my computer jobs, but before and after work I started filling notebooks with my thoughts on vision; then after retiring I started typing them into a computer, until this book came together.

The book you see here does not deal with the whole brain, only with vision, and within that subject only with one class of large image-determined circuits: the *contour strings*. My original intent was to include another class of large circuits, the "color pools," and I had also hoped to extend the discussion beyond the V1 area of visual cortex, to V2, where the first version of a "stable" cortical image (which does not move around with the retinal image) promises to arise.

However, preliminary work convinced me that both of the other subjects, the *color pools* and the *stable image*, were hugely more complex than the *contour strings* and I decided to leave them out, at least for now.

New York, NY Charles R. Legéndy

Chapter 2
Issues Concerning the Nature of Neuronal Response

2.1 Impressions Gained from Histograms and Raster Displays

A person whose intuition of neuronal reliability is shaped by post-stimulus time (PST) histograms (Gerstein and Kiang, 1960) and dot raster displays (see for example Schmidt et al, 1975) (myself included, in the early papers Legéndy, 1970, 1975) may well be tempted to assume that the single neuron is only reliable in broad statistical terms. To elicit reproducible behavior from a neuron of the visual cortex, for instance, one must sweep the receptive field 10–20 times, combine the sweeps into a PST histogram or a raster display, and examine the way the spikes are distributed. Two or three sweeps do not appear to be enough to get a clear idea of the neuron's behavior, because, as seen in Fig. 2.1, the responses do not repeat spike-for-spike; some sweeps show more spikes, some fewer.

The displays suggest to us that a neuron responds, when it does, by emitting *a number of* spikes. Some typical responses are "brisk," where a number of spikes,

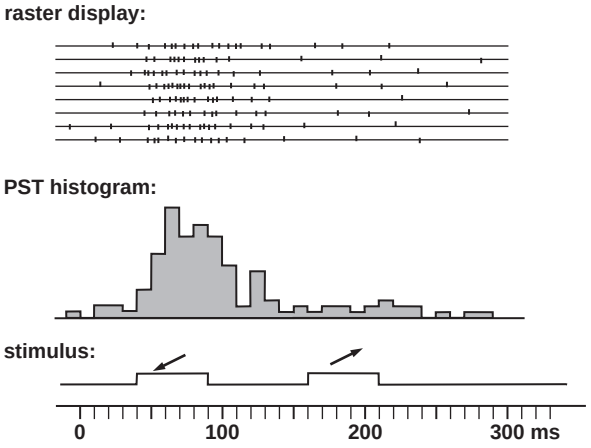

Fig. 2.1 Raster display and post-stimulus time histogram. *Top*: responses of a neuron to eight triggered sweeps of a stimulus (hypothetical neuron and stimulus); *bottom*: spike counts gathered from the same spike trains (10 ms bins)

C.R. Legéndy, *Circuits in the Brain*, DOI 10.1007/978-0-387-88849-1_2,

say 5–10, are crowded into a short time interval, say 20–30 ms; others are more "sluggish" where fewer spikes appear and they are spread over a longer time interval. They do not encourage us to trust the neuron to emit spikes individually timed in ways that fit into the overall processing.

However, recent studies tell us that neuronal outputs are quite reliable, as long as the inputs are. The internal signal conduction within the neuron has been shown to be both reliable and fast (Mainen and Sejnowski, 1995; Ariav et al., 2003), and in particular much faster than had been believed on the basis of the classical resistive-capacitive picture of dendrites (Rall, 1962).

The reliability and speed of neurons seemingly contradicts their irregular mode of firing in behaving animals, and leads to a useful inference concerning the neuronal input stream, as will be seen next.

2.2 Cortical Firing Should Be Nearly Periodic – So Why Isn't It?

When a steady current is injected into the interior of a neuron (Fig. 2.2), the neuron emits more or less periodic firing, whose spike rate increases with the current; it does not for instance fire irregularly at a current-dependent average rate. More recently it has been demonstrated (in vitro) that current injection into apical dendrites, and even thin (1.9 micron) basal dendrites, or comparably thin branches of the dendritic tuft in layer 2–3 pyramidal cells, causes similar and more or less regular firing (Larkum et al., 2001; Nevian et al., 2007; Larkum et al., 2007).

The counterintuitive aspect of the observation becomes clear when one notes that, at least in a crude approximation, every spike coming to an excitatory synapse injects a small amount of electric charge into the cell interior. Since typically thousands of spike trains arrive to a cortical neuron through its synapses, their steady rainfall is expected to simulate a DC current injected into the cell interior, and cause more or less steady firing like what is seen in Fig. 2.2, or at least steady firing at a slowly varying rate, as shown in Fig. 2.3(a). However, in fact, typical spike trains in the visual cortex look something like Fig. 2.3(b).

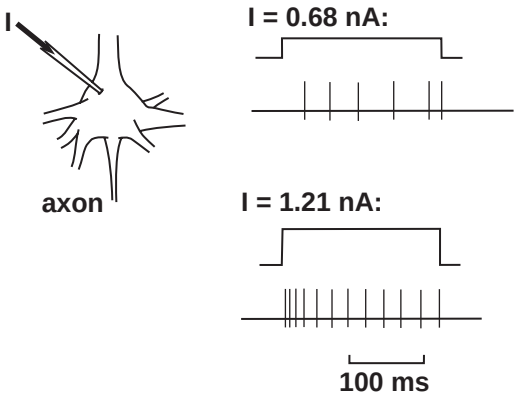

Fig. 2.2 Response of a cell to intracellular injection of DC current. Response (in vivo) of a Betz cell from the motor cortex, when a pulse of steady electrical current, lasting about 200 ms, is injected into its cell body (Anesthetized cat. Re-drawn from Creutzfeldt et al., 1964)

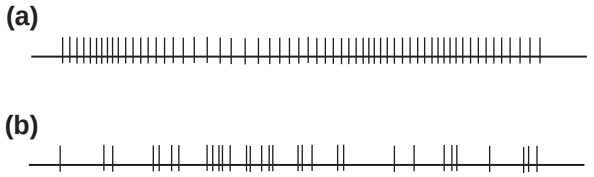

Fig. 2.3 Effect of input correlation on the appearance of neuronal output. (**a**) Appearance of the firing of a cortical neuron on the assumption that its input spike trains are only crudely correlated (to an accuracy of about hundred ms). (**b**) Appearance of firing as seen in a typical recording from the visual cortex of an unanaesthetized cat

The contradiction was first raised in the form of a problem requiring a solution by Perkel and Bullock (1969) who used the mathematical results of Cox and Smith (1954); it was later brought to wider attention by Softky and Koch (1993), and has since then given rise to a fair amount of discussion in the literature. It may be remarked that some of this discussion incorrectly implies that current injection only causes regular spiking in brain slices (in vitro), whereas in fact it also does so in the intact brain (Creutzfeldt et al., 1964; Oshima, 1969; Ahmed et al., 1993).

The irregularity of firing in response to natural inputs has been attributed by some writers to brief pauses in firing caused by volleys of inhibition (Shadlen and Newsome, 1998), and by others to noise in the membrane and synapses (Destexhe et al., 2001). Both groups of authors see the irregular firing as resulting from noise-like random events of one sort or another.

The explanation in terms of random events has never been entirely satisfying. Ahmed et al. (1993), in an in vivo study, point out (and it is also observable in the data of Creutzfeldt et al., 1964) that at low levels of injected current the firing is often irregular, but at stronger currents it becomes more regular. The random noise explanation would predict that during the brief epochs of elevated spike rate (bursts) reported in waking animals (Legéndy and Salcman, 1985), when electric charge is expected to enter the cell interior at an increased rate, the irregularity-causing effects of noise would tend to be overwhelmed by the input current, and the spike discharge would become more uniform (Holt et al., 1996).

However, in behaving animals the spiking is just as irregular during the bursts as it is outside the bursts (Legéndy and Salcman, 1985). In must be added that Stevens and Zador (1998), in a careful statistical study, eliminate both noise and inhibition as reasons for the observed irregularity, and show that their effect is, at least in an in vitro preparation, insufficient for causing the observed behavior. While the role of inhibition in causing the irregular firing is likely to be important, the conclusion remains that the irregular firing of the neurons in behaving animals implies synchronized and correlated synaptic input.

The present model goes along with the latter conclusion and assumes that a significant number of spikes arriving to the neurons are sharply synchronized.

The idea of synchrony sharp enough to dictate the timing of output spikes suggests that the incoming spikes arrange a "rendezvous," in some way, on the receptive surfaces of neurons, which is a seemingly outlandish idea and may account for some of the resistance to it in the literature.

As will be seen, the concept becomes much more plausible when viewed from a different perspective (Sect. 4.2), but before getting to that, let me make a brief remark on the effect of volleys of spikes on neurons, and then survey some of the literature on plasticity.

2.3 Sensitivity of Neurons to Synchronized Volleys of Spikes

The form of "rendezvous" which comes up most often in this book is one which sends a volley of nearly simultaneous spikes to neurons, along the lines shown in Fig. 2.4. Such volleys clearly stand out above the background noise, especially when

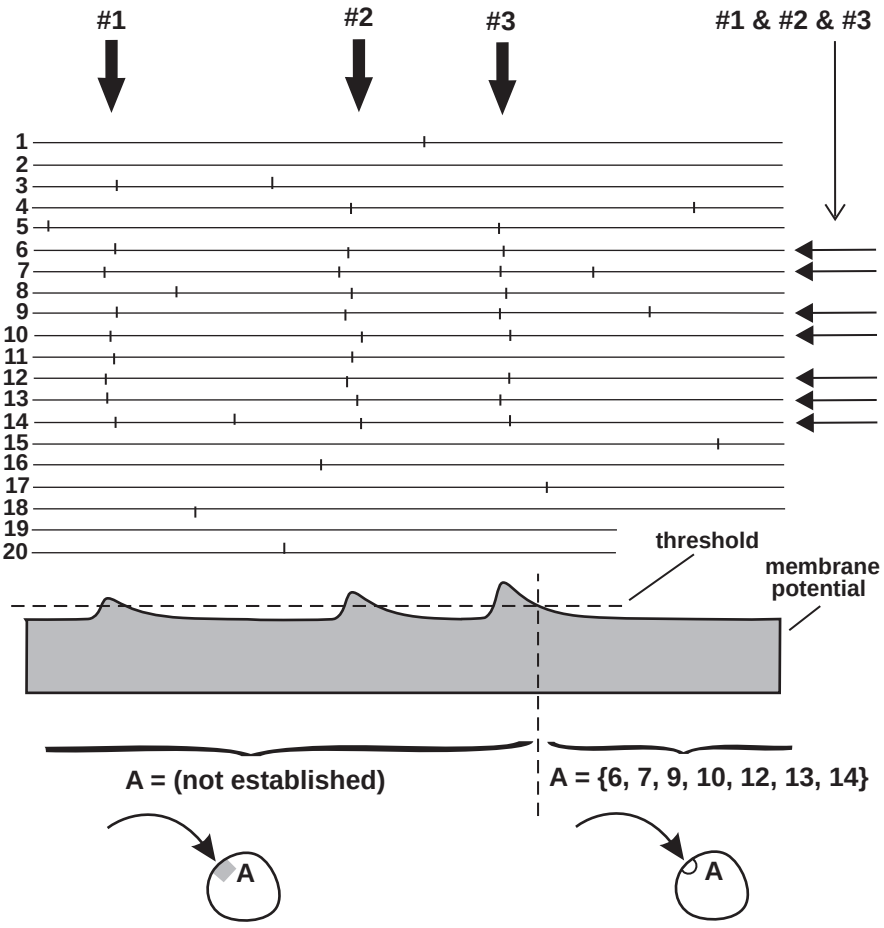

Fig. 2.4 (continued)

repeated, and there is evidence that if they arrive repeatedly they can cause plastic change, marking the synapses on which they arrive.

In Fig. 2.4, repeated volleys arrive to a set of synapses (arrows at right), envisioned as being located sufficiently close together on a dendrite to cause the indicated excitatory postsynaptic potentials (EPSPs) at the synapses. The gradually increasing EPSP amplitude (exaggerated) is intended to imply a mechanism of short-term and long-term potentiation (STP and LTP). The *threshold* marked with broken line stands for the minimum level of membrane potential needed, in conjunction with presynaptic activity, for inducing synaptic potentiation.

The insets at the bottom of Fig. 2.4 introduce a graphical notation (see also Sect. 12.3), to be used in the functional diagrams of later chapters for sets of synapses on neurons which undergo LTP, or are "marked," as will be said from now on. A set of synapses, initially unassigned and only known to be a subset of the synapses from a biologically distinct neuron pool (rectangular shading at left), becomes "marked" (semicircle at right) by repeatedly arriving volleys. The amorphously drawn blob shape in these drawings stands for the typical neuron from among a *set* of similar neurons, or alternatively for the neuron set itself, and the arrow for a set of fibers coming to the neuron and bringing the signals which do the marking.

2.4 Notes on Plastic Change at the Synaptic Level

The classic Hebbian scheme of synapse modification (Hebb, 1949) does not permit plastic change to occur as a result of time-concentrated presynaptic events alone; it requires the postsynaptic neuron also to fire, right after the presynaptic firing, and until recently all experimental findings agreed. It appeared that only through firing, and the resulting back-propagated action potential, could a neuron send all its participating synapses an intense enough and fast enough wave of depolarization to induce them to long-term change.

However, recently, Remy and Spruston (2007) showed that the postsynaptic firing is not strictly necessary, as long as the synaptic invasion is intense enough to give rise to *dendritic spikes*. It may be added that, in practice, when volleys are intense enough to cause dendritic spikes, they are also expected, after some repetitions, to cause action potentials, since the level of current injection required for the latter is not much greater (Nevian et al., 2007). (For the purposes of their demonstration,

Fig. 2.4 (continued) Marking a set of synapses participating in repeated volleys. Traces 1–20 stand for spike trains arriving to some of the synapses on a (hypothetical) neuron. At times #1, #2, #3, (*arrows on top*) the spike trains form "volleys"; seven of the spike trains (*arrows at right*) participate in all three of the volleys. The corresponding membrane potential (idealized) is shown below the traces. *Bottom*: synapse set notation symbolizing a whole synapse pool (A) on the neuron receiving input (*left, shaded rectangle*), and a "synapse set" marked out from among members of the same synapse pool as a result of the recurring volleys (*right, semicircle*)

Remy and Spruston, 2007, suppressed the action potentials and back-propagation by means of tetrodotoxin.)

According to the available evidence, induction of LTP requires presynaptic spikes in conjunction with either back-propagated action potentials (Magee and Johnston, 1997) or locally induced dendritic spikes (Remy and Spruston, 2007). Glutamate released in the synaptic transmission binds to the N-methyl D-aspartate (NMDA) receptors and opens up their calcium channels which are, because of the brief change in membrane potential, momentarily freed of the Mg2+ ions usually blocking them, and permit Ca2+ ions into the cell interior (Ascher and Nowak, 1988; Bliss and Collingridge, 1993).

It will be emphasized that Fig. 2.4 implies something not yet addressed by experimental data, namely that a set of synapses can be individually selected by the input arriving to them, even when they are not all adjacent; in other words, the synapses reinforced can be interspersed among synapses not reinforced.

From a functional point of view the capacity for such individual synapse selection is expected to be a requirement, because the arriving volleys come from cooperative neuron groups which, although functionally connected at the moment, are not similarly connected at the earlier time when axonal growth decides the location of their potential synapses. Accordingly they cannot be expected to have their synapses all located right next to each other.

The recently developed technique of two-photon glutamate uncaging at preselected sets of synaptic spines (Losonczy and Magee, 2006; Gasparini and Magee, 2006; Losonczy et al., 2008), now makes it feasible to address the issue of individual synapse selection. The results available at this point indicate that the synapses do not need to be concentrated on one portion of a dendrite but can be scattered throughout its length, but do not demonstrate the possibility that some synapses do and other nearby ones do not undergo potentiation. The issue is still open at the time of this writing.