

- Michael Recce** *Chapter 4*
 Department of Computer and Information Science
 New Jersey Institute of Technology
 Newark, NJ 07102, USA
 recce@homer.njit.edu
- Barry J. P. Rising** *Chapter 13*
 Department of Computer Science
 Royal Holloway, University of London
 Egham, England, UK
 barry@dcs.rhbc.ac.uk
- Ulrich Roth** *Chapter 9*
 Institut of Microelectronics
 TU Berlin
 Berlin, Germany
 roth@mikro.ee.tu-berlin.de
- Tim Schönauer** *Chapter 9*
 Institut of Microelectronics
 TU Berlin
 Berlin, Germany
 tim@mikro.ee.tu-berlin.de
- Terrence J. Sejnowski** *Foreword*
 The Salk Institute
 La Jolla, CA 92037, USA
 terry@salk.edu
- John S. Shawe-Taylor** *Chapter 13*
 Department of Computer Science
 Royal Holloway, University of London
 Egham, UK
 john@dcs.rhbc.ac.uk
- Philippe Venier** *Chapter 7*
 Advanced Microelectronics Division
 Centre Suisse d'Electronique et de Microtechnique
 Neuchatel, Switzerland
 venier@csemne.ch
- Hermann Wagner** *Chapter 14*
 Institut für Biologie
 Lehrstuhl für Zoologie/Tierphysiologie
 RWTH Aachen
 D-52074 Aachen, Germany
 wagner@tyto.bio2.rwth-aachen.de
- Adrian M. Whatley** *Chapter 6*
 Institut für Neuroinformatik
 Universität Zürich & ETH Zürich
 Zürich, Switzerland
 amw@ini.phys.ethz.ch
- Anthony M. Zador** *Chapter 12*
 The Salk Institute
 La Jolla, CA 92037, USA
 zador@salk.edu

Part I

Basic Concepts and Models

1 Spiking Neurons

Wulfram Gerstner

An organism which interacts with its environment must be capable of receiving sensory input from the environment. It has to process the sensory information, recognize food sources or predators, and take appropriate actions. The difficulty of these tasks is appreciated, if one tries to program a small robot to do the same thing: It turns out to be a challenging endeavor. Yet animals perform these tasks with apparent ease.

Their astonishingly good performance is due to a neural system or 'brain' which has been optimized over the time courses of evolution. Even though a lot of detailed information about neurons and their connections is available by now, one of the fundamental questions of neuroscience is unsolved: What is the code used by the neurons? Do neurons communicate by a 'rate code' or a 'pulse code'?

In the first part of this chapter, different potential coding schemes are discussed. Various interpretations of rate coding are contrasted with some pulse coding schemes. Pulse coded neural networks require appropriate neuron models. In the second part of the chapter, several neuron models that are used throughout the book are introduced. Special emphasis has been put on spiking neurons models of the 'integrate-and-fire' type, but the Hodgkin-Huxley model, compartmental models, and rate models are reviewed as well.

1.1 The Problem of Neural Coding

1.1.1 Motivation

Over the past hundred years, biological research has accumulated an enormous amount of detailed knowledge about the structure and the function of the brain see, e.g., [Kandel and Schwartz, 1991]. The elementary processing units in the brain are neurons which are connected to each other in an intricate pattern. A portion of such a network of neurons in the mammalian cortex is sketched in Figure 1.1. It is a reproduction of a famous drawing by Ramón y Cajal, one of the pioneers of neuroscience around the turn of the century. We can distinguish several neurons with triangular or circular cell bodies and long wire-like extensions. This drawing gives a glimpse of the network of neurons in the cortex. Only a few of the neurons present in the sample have been made visible by the staining procedure. In reality the neurons and their connections form a dense network with more than 10^4 cell bodies and several kilometers of 'wires' per cubic millimeter.

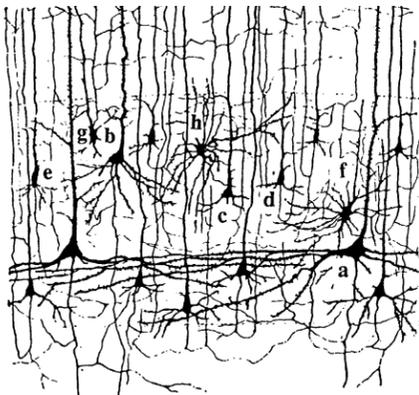


Figure 1.1. This reproduction of a drawing of Ramón y Cajal shows a few neurons in the cortex. Only a small portion of the neurons are shown; the density of neurons is in reality much higher. Cell b is a nice example of a pyramidal cell with a triangularly shaped cell body. Dendrites, which leave the cell laterally and upwards, can be recognized by their rough surface. The axon extends downwards with a few branches to the left and right. From Ramón y Cajal.

In other areas of the brain the wiring pattern looks different. In all areas, however, neurons of different sizes and shapes form the basic elements.

A typical neuron has three parts, called dendritic tree, soma, and axon; see Figure 1.2. Roughly speaking, signals from other neurons arrive onto the dendritic tree and are transmitted to the soma and the axon. The transition zone between the soma and the axon is of special interest. In this area the essential non-linear processing step occurs. If the total excitation caused by the input is sufficient, an output signal is emitted which is propagated along the axon and its branches to other neurons. The junction between an axonal branch and the dendrite (or the soma) of a receiving neuron is called a synapse. It is common to refer to a sending neuron as the presynaptic neuron and to the receiving neuron as a postsynaptic neuron. A neuron in the cortex often makes connections to more than 10^4 postsynaptic neurons. Many of its axonal branches end in the direct neighborhood of the neuron, but the axon can also stretch over several millimeters and connect to neurons in other areas of the brain.

So far, we have stated that neurons transmit signals along the axon to thousands of other neurons – but what do these signals look like? The neuronal signals can be observed by placing a fine electrode close to the soma or axon of a neuron; see Figure 1.2. The voltage trace in a typical recording shows a sequence of short pulses, called action potentials or spikes. A chain of pulses emitted by a single neuron is usually called a spike train – a sequence of stereotyped events which occur at regular or irregular intervals. The duration of an action potential is typically in the range of 1-2 ms. Since all spikes of a given neuron look alike, the form of the action potential does not carry any information. Rather, it is the number and the timing of spikes which matter.

Throughout this book, we will refer to the moment when a given neuron emits an action potential as the firing time of that neuron. The firing time

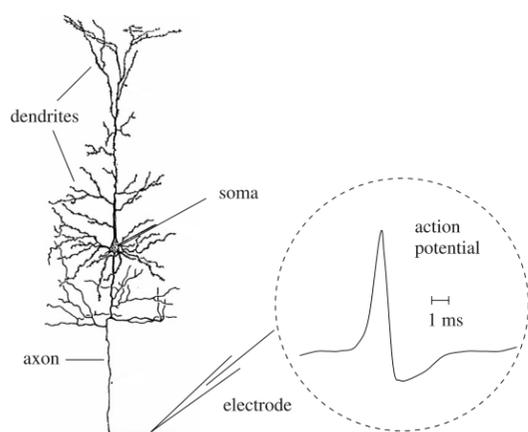


Figure 1.2. A single neuron. Dendrite, soma, and axon can be clearly distinguished. The inset shows an example of a neuronal action potential (schematic). Neuron drawing after Ramón y Cajal. The action potential is a short voltage pulse of 1-2 ms duration.

of neuron i will be denoted by $t_i^{(f)}$. The spike train of a neuron i is fully characterized by the set of firing times

$$\mathcal{F}_i = \{t_i^{(1)}, \dots, t_i^{(n)}\} \quad (1.1)$$

where $t_i^{(n)}$ is the most recent spike of neuron i .

In an experimental setting, firing times are measured with some resolution Δt . A spike train may be described as a sequence of ones and zeros for 'spike' and 'no spike' at times $\Delta t, 2\Delta t \dots$, respectively. The choice of ones and zeros is, of course arbitrary. We may just as well take the number $1/\Delta t$ instead of unity to denote the occurrence of a spike. With this definition, the spike train of a neuron i corresponds to a sequence of numbers $S_i(\Delta t), S_i(2\Delta t), \dots$ with

$$S_i(n\Delta t) = \begin{cases} 1/\Delta t & \text{if } n\Delta t \leq t_i^{(f)} < (n+1)\Delta t \\ 0 & \text{otherwise.} \end{cases} \quad (1.2)$$

Formally we may take the limit $\Delta t \rightarrow 0$ and write the spike train as a sequence of δ -functions

$$S_i(t) = \sum_{t_i^{(f)} \in \mathcal{F}_i} \delta(t - t_i^{(f)}) \quad (1.3)$$

where $\delta(\cdot)$ denotes the Dirac δ function with $\delta(s) = 0$ for $s \neq 0$ and $\int_{-\infty}^{\infty} \delta(s) ds = 1$.

So far we have focused on the spike train of a single neuron. Since there are so many neurons in the brain, thousands of spike trains are emitted constantly by different neurons; see Figure 1.3. What is the information contained in such a spatio-temporal pattern of pulses? What is the code used by the neurons to transmit that information? How might other neurons decode the signal? As external observers, can we read the code, and understand the message of the neuronal activity pattern?

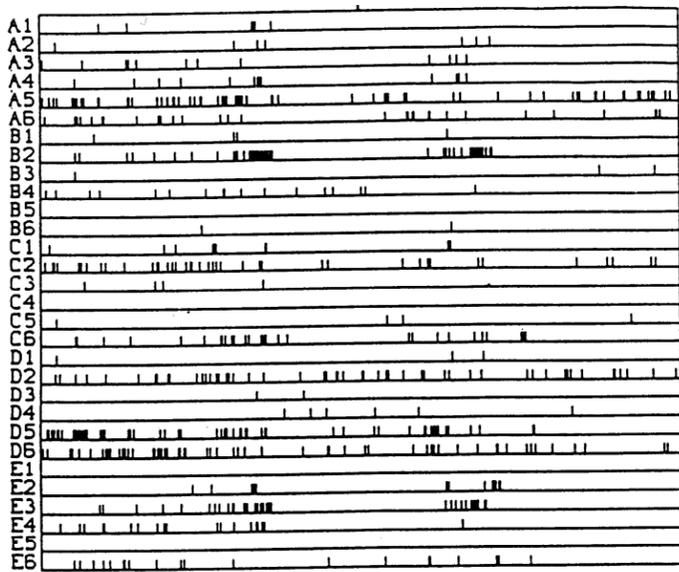


Figure 1.3. Spatio-temporal pulse pattern. The spikes of 30 neurons (A1-E6, plotted along the vertical axes) are shown as a function of time (horizontal axis, total time is 4 000 ms). The firing times are marked by short vertical bars. From [Krüger and Aiple, 1988].

At present, a definite answer to these questions is not known. Traditionally it has been thought that most, if not all, of the relevant information was contained in the mean firing rate of the neuron. The firing rate is usually defined by a temporal average; see Figure 1.4. The experimentalist sets a time window of, let us say, $T = 100$ ms or $T = 500$ ms and counts the number of spikes $n_{\text{sp}}(T)$ that occur in this time window. Division by the length of the time window gives the mean firing rate

$$\nu = \frac{n_{\text{sp}}(T)}{T} \quad (1.4)$$

usually reported in units of s^{-1} or Hz.

The concept of mean firing rates has been successfully applied during the last 80 years. It dates back to the pioneering work of Adrian [Adrian, 1926, 1928] who showed that the firing rate of stretch receptor neurons in the muscles is related to the force applied to the muscle. In the following decades, measurement of firing rates became a standard tool for describing the properties of all types of sensory or cortical neurons [Mountcastle, 1957; Hubel and Wiesel, 1959], partly due to the relative ease of measuring rates experimentally. It is clear, however, that an approach based on a temporal average neglects all the information possibly contained in the exact timing of the spikes. It is therefore no surprise that the firing rate concept has been repeatedly criticized and is subject of an ongoing debate [Abeles, 1994; Bialek et al., 1991; Hopfield, 1995; Shadlen and Newsome, 1994; Softky, 1995; Rieke et al., 1996].

During recent years, more and more experimental evidence has accumulated which suggests that a straightforward firing rate concept based on temporal averaging may be too simple for describing brain activity. One of the main arguments is that reaction times in behavioral experiments are often too short to allow slow temporal averaging [Thorpe et al., 1996]. Moreover, in experiments on a visual neuron in the fly, it was possible to ‘read the neural code’ and reconstruct the time-dependent stimulus based on the neurons firing times [Bialek et al., 1991]. There is evidence of precise temporal correlations between pulses of different neurons [Abeles, 1994; Lestienne, 1996] and stimulus dependent synchronization of the activity in populations of neurons [Eckhorn et al., 1988; Gray and Singer, 1989; Gray et al., 1989; Engel et al., 1991; Singer, 1994]. Most of these data are inconsistent with a naïve concept of coding by mean firing rates where the exact timing of spikes should play no role. In this book we will explore some of the possibilities of coding by pulses. Before we can do so, we have to lay the foundations which will be the topic of this and the next three chapters.

We start in the next subsection with a review of some potential coding schemes. What exactly is a pulse code – and what is a rate code? We then turn to models of spiking neurons (Section 2). How can we describe the process of spike generation? What is the effect of a spike on a postsynaptic neuron? Can we mathematically analyze models of spiking neurons?

The following Chapters 2 and 3 in the ‘Foundation’ part of the book will focus on the computational power of spiking neurons and their hardware implementations. Can we build a Turing machine with spiking neurons? How many elements do we need? How fast is the processing? How can pulses be generated in hardware? Many of these questions outlined in the Foundation chapters will be revisited in the detailed studies contained in the parts II and III of the book. Chapter 4, the last chapter in the Foundation part, will discuss some of the biological evidence for temporal codes in more detail.

1.1.2 Rate Codes

A quick glance at the experimental literature reveals that there is no unique and well-defined concept of ‘mean firing rate’. In fact, there are at least three different notions of rate which are often confused and used simultaneously. The three definitions refer to three different averaging procedures: either an average over time, or an average over several repetitions of the experiment, or an average over a population of neurons. The following three subsections will reconsider the three concepts. An excellent discussion of rate codes can also be found in [Rieke et al., 1996].

1.1.2.1 Rate as a Spike Count (Average over Time)

The first and most commonly used definition of a firing rate refers to a temporal average. As discussed in the preceding section, this is essentially the spike count in an interval T divided by T ; see Figure 1.4. The length of the

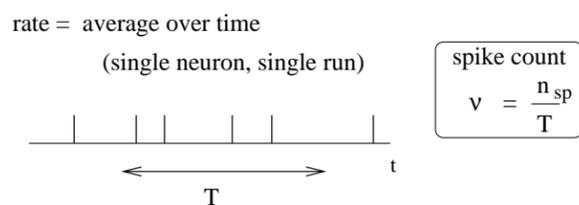


Figure 1.4. Definition of the mean firing rate via a temporal average.

time window is set by the experimenter and depends on the type of neuron recorded from and the stimulus. In practice, to get sensible averages, several spikes should occur within the time window. Values of $T = 100$ ms or $T = 500$ ms are typical, but the duration may also be longer or shorter.

This definition of rate has been successfully used in many preparations, particularly in experiments on sensory or motor systems. A classical example is the stretch receptor in a muscle spindle [Adrian, 1926]. The number of spikes emitted by the receptor neuron increases with the force applied to the muscle. Another textbook example is the touch receptor in the leech [Kandel and Schwartz, 1991]. The stronger the touch stimulus, the more spikes occur during a stimulation period of 500 ms.

These classical results show that the experimenter as an external observer can evaluate and classify neuronal firing by a spike count measure – but is this really the code used by neurons in the brain? In other words, is a neuron which receives signals from a sensory neuron only looking at and reacting to the numbers of spikes it receives in a time window of, say, 500 ms? We will approach this question from a modeling point of view later on in the book. Here we discuss some critical experimental evidence.

From behavioral experiments it is known that reaction times are often rather short. A fly can react to new stimulus and change the direction of flight within 30-40 ms; see the discussion in [Rieke et al., 1996]. This is not long enough for counting spikes and averaging over some long time window. It follows that the fly has to react to single spikes. Humans can recognize visual scenes in just a few hundred milliseconds [Thorpe et al., 1996], even though recognition is believed to involve several processing steps. Again, this leaves not enough time to perform temporal averages on each level.

Temporal averaging can work well where the stimulus is constant or slowly moving and does not require a fast reaction of the organism – and this is the situation usually encountered in experimental protocols. Real-world input, however, is hardly stationary, but often changing on a fast time scale. For example, even when viewing a static image, we perform saccades, rapid changes of the direction of gaze. The retinal photo receptors receive therefore every few hundred milliseconds a new input.

Despite its shortcomings, the concept of a firing rate code is widely used not only in experiments, but also in models of neural networks. It has led to the idea that a neuron transforms information about a single input variable (the stimulus strength) into a single continuous output variable (the firing rate). In this view, spikes are just a convenient way to transmit the analog output over long distances. In fact, the best coding scheme to transmit

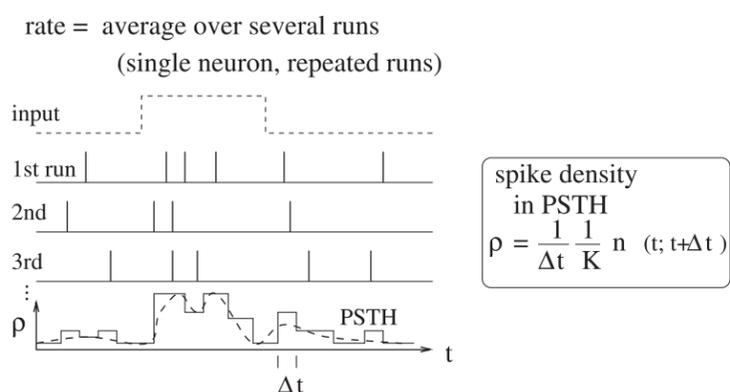


Figure 1.5. Definition of the spike density in the Peri-Stimulus-Time Histogram (PSTH).

the value of the rate ν would be by a regular spike train with intervals $1/\nu$. In this case, the rate could be reliably measured after only two spikes. From the point of view of rate coding, the irregularities encountered in real spike trains of neurons in the cortex must therefore be considered as noise. In order to get rid of the noise and arrive at a reliable estimate of the rate, the experimenter (or the postsynaptic neuron) has to average over a larger number of spikes. A critical discussion of the temporal averaging concept can be found in [Shadlen and Newsome, 1994; Softky, 1995; Rieke et al., 1996].

1.1.2.2 Rate as a Spike Density (Average over Several Runs)

There is a second definition of rate which works for stationary as well as for time-dependent stimuli. The experimenter records from a neuron while stimulating with some input sequence. The same stimulation sequence is repeated many times and the results are reported in a Peri-Stimulus-Time Histogram (PSTH); see Figure 1.5. For each short interval of time $[t, t + \Delta t]$, before, during, and after the stimulation sequence, the experimenter counts the number of times that a spike has occurred and sums them over all repetitions of the experiment. The time t is measured with respect to the start of the stimulation sequence and Δt is typically in the range of one or a few milliseconds. The number of occurrences of spikes $n(t; t + \Delta t)$ divided by the number K of repetitions is a measure of the typical activity of the neuron between time t and $t + \Delta t$. A further division by the interval length Δt yields the spike density of the PSTH

$$\rho(t) = \frac{1}{\Delta t} \frac{n(t; t + \Delta t)}{K}. \quad (1.5)$$

Sometimes the result is smoothed to get a continuous 'rate' variable. The spike density of the PSTH is usually reported in units of Hz and often called the (time-dependent) firing rate of the neuron.

As an experimental procedure, the spike density measure is a useful method to evaluate neuronal activity, in particular in the case of time-dependent

rate = average over pool of equivalent neurons
(several neurons, single run)

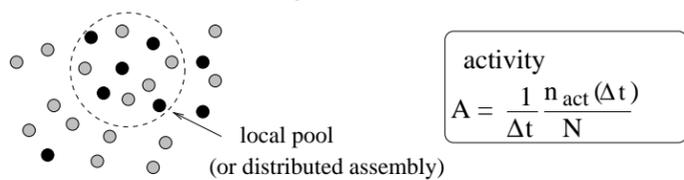


Figure 1.6. Definition of the population activity.

stimuli. The obvious problem with this approach is that it can not be the decoding scheme used by neurons in the brain. Consider for example a frog which wants to catch a fly. It can not wait for the insect to fly repeatedly along exactly the same trajectory. The frog has to base its decision on a single ‘run’ – each fly and each trajectory is different.

Nevertheless, the experimental spike density measure can make sense, if there are large populations of neurons which are independent of each other and sensitive to the same stimulus. Instead of recording from a population of N neurons in a single run, it is experimentally easier to record from a single neuron and average over N repeated runs. Thus, the spike density coding relies on the implicit assumption that there are always populations of neurons and therefore leads to the third notion of a firing rate, viz., a rate defined as a population average.

1.1.2.3 Rate as Population Activity (Average over Several Neurons)

The number of neurons in the brain is huge. Often many neurons have similar properties and respond to the same stimuli. For example, neurons in the primary visual cortex of cats and monkeys are arranged in columns of cells with similar properties [Hubel and Wiesel, 1962, 1977; Hubel, 1988]. Let us idealize the situation and consider a population of neurons with identical properties. In particular, all neurons in the population should have the same pattern of input and output connections. The spikes of the neurons in a population j are sent off to another population k . In our idealized picture, each neuron in population k receives input from all neurons in population j . The relevant quantity, from the point of view of the receiving neuron, is the proportion of active neurons in the presynaptic population j ; see Figure 1.6. Formally, we define the population activity

$$A(t) = \frac{1}{\Delta t} \frac{n_{\text{act}}(t; t + \Delta t)}{N} \quad (1.6)$$

where N is the size of the population, Δt a small time interval, and $n_{\text{act}}(t; t + \Delta t)$ the number of spikes (summed over all neurons in the population) that occur between t and $t + \Delta t$. If the population is large, we can consider the limit $N \rightarrow \infty$ and take then $\Delta t \rightarrow 0$. This yields again a continuous quantity with units s^{-1} – in other words, a rate.

The population activity may vary rapidly and can reflect changes in the stimulus conditions nearly instantaneously [Tsodyks and Sejnowsky, 1995].

Thus the population activity does not suffer the disadvantages of a firing rate defined by temporal averaging at the single-unit level. The problem with the definition (1.6) is that we have formally required a homogeneous population of neurons with identical connections which is hardly realistic. Real populations will always have a certain degree of heterogeneity both in their internal parameters and in their connectivity pattern. Nevertheless, rate as a population activity (of suitably defined pools of neurons) may be a useful coding principle in many areas of the brain. For inhomogeneous populations, the definition (1.6) may be replaced by a weighted average over the population. A related scheme has been used successfully for an interpretation of neuronal activity in primate motor cortex [Georgopoulos et al., 1986].

1.1.3 Candidate Pulse Codes

In this subsection, we will briefly introduce some potential coding strategies based on spike timing. All codes will be discussed in more detail later on and will be referred to throughout the book.

1.1.3.1 Time-to-First-Spike

Let us study a neuron which abruptly receives a new input at time t_0 . For example, a neuron might be driven by an external stimulus which is suddenly switched on at time t_0 . This seems to be somewhat academic, but even in a realistic situation abrupt changes in the input are quite common. When we look at a picture, our gaze jumps from one point to the next. After each saccade, there is a new visual input at the photo receptors in the retina. Information about the time t_0 of a saccade would easily be available in the brain. We can then imagine a code where for each neuron the timing of the first spike to follow t_0 contains all information about the new stimulus. A neuron which fires shortly after t_0 could signal a strong stimulation, firing somewhat later would signal a weaker stimulation; see Figure 1.7.

In a pure version of this coding scheme, only the first spike of each neuron counts. All following spikes would be irrelevant. Alternatively, we can also assume that each neuron emits exactly one spike per saccade and is shut off by inhibitory input afterwards. It is clear that in such a scenario, only the timing conveys information and not the number of spikes.

A coding scheme based on the time-to-first-spike is certainly an idealization. In Chapter 2 it will be formally analyzed by Wolfgang Maass. In a slightly different context coding by first spikes has also been discussed by S. Thorpe [Thorpe et al., 1996]. Thorpe argues that the brain does not have time to evaluate more than one spike from each neuron per processing step. Therefore the first spike should contain most of the relevant information. Using information-theoretic measures on their experimental data, several groups have shown that most of the information about a new stimulus is indeed conveyed during the first 20 or 50 milliseconds after the onset of the neuronal response [Optican and Richmond, 1987; Kjaer et al., 1994; Tovee et al., 1993; Tovee and Rolls, 1995]. Rapid computation during the

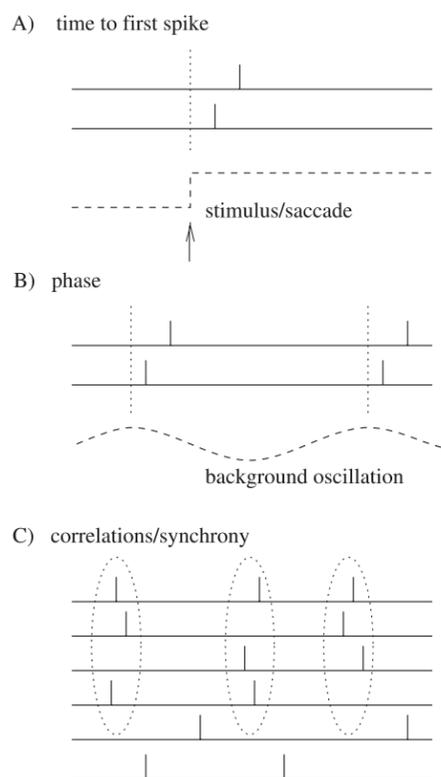


Figure 1.7. Three examples of pulse codes. A) Time-to-first spike. The second neuron responds faster to a change in the stimulus than the first one. Stimulus onset marked by arrow. B) Phase. The two neurons fire at different phases with respect to the background oscillation (dashed). C) Synchrony. The upper four neurons are nearly synchronous, two other neurons at the bottom are not synchronized with the others.

transients after a new stimulus has also been discussed in model studies [Hopfield and Herz, 1995; Tsodyks and Sejnowsky, 1995; van Vreeswijk and Sompolinsky, 1997].

1.1.3.2 Phase

We can apply a coding by 'time-to-first-spike' also in the situation where the reference signal is not a single event, but a periodic signal. In the hippocampus, in the olfactory system, and also in other areas of the brain, oscillations of some global variable (for example the population activity) are quite common. These oscillations could serve as an internal reference signal. Neuronal spike trains could then encode information in the phase of a pulse with respect to the background oscillation. If the input does not change between one cycle and the next, then the same pattern of phases repeats periodically; see Figure 1.7 B.

The concept of coding by phases has been studied by several different groups, not only in model studies [Hopfield, 1995; Jensen and Lisman,

1996; Maass, 1996], but also experimentally [O'Keefe and Recce, 1993]. There is for example evidence that the phase of a spike during an oscillation in the hippocampus of the rat conveys information on the spatial location of the animal which is not accounted for by the firing rate of the neuron alone [O'Keefe and Recce, 1993].

1.1.3.3 Correlations and Synchrony

We can also use spikes from other neurons as the reference signal for a pulse code. For example, synchrony between a pair or a group of neurons could signify special events and convey information which is not contained in the firing rate of the neurons; see Figure 1.7 C. One famous idea is that synchrony could mean 'belonging together' [Milner, 1974; Malsburg, 1981]. Consider for example a complex scene consisting of several objects. It is represented in the brain by the activity of a large number of neurons. Neurons which represent the same object could be 'labeled' by the fact that they fire synchronously [Malsburg, 1981; Malsburg and Buhmann, 1992; Eckhorn et al., 1988; Gray et al., 1989]. Coding by synchrony has been studied extensively both experimentally [Eckhorn et al., 1988; Gray et al., 1989; Gray and Singer, 1989; Singer, 1994; Engel et al., 1991ab; Kreiter and Singer, 1992] and in models [Wang et al., 1990; Malsburg and Buhmann, 1992; Eckhorn, 1990; Aertsen and Arndt, 1993; Koenig and Schillen, 1991; Schillen and Koenig, 1991; Gerstner et al., 1993; Ritz et al. 1993; Terman and Wang, 1995; Wang, 1995]. For a review of potential mechanism, see [Ritz and Sejnowski, 1997]. Coding by synchrony is discussed in Chapter 11.

More generally, not only synchrony but any precise spatio-temporal pulse pattern could be a meaningful event. For example, a spike pattern of three neurons, where neuron 1 fires at some arbitrary time t_1 followed by neuron 2 at time $t_1 + \delta_{12}$ and by neuron 3 at $t_1 + \delta_{13}$, might represent a certain stimulus condition. The same three neurons firing with different relative delays might signify a different stimulus. The relevance of precise spatio-temporal spike patterns has been studied intensively by Abeles [Abeles, 1991; Abeles et al., 1993; Abeles, 1994]. Similarly, but on a somewhat coarse time scale, correlations of auditory neurons are stimulus dependent and might convey information beyond the firing rate [deCharms and Merzenich, 1996].

1.1.3.4 Stimulus Reconstruction and Reverse Correlation

Let us consider a neuron which is driven by a time dependent stimulus $s(t)$. Every time a spike occurs, we note the time course of the stimulus in a time window of about 100 ms immediately before the spike. Averaging the results for several spikes yields the typical time course of the stimulus just before a spike. Such a procedure is called a 'reverse correlation' approach; see Figure 1.8. In contrast to the PSTH experiment sketched in Section 2.2 where the experimenter averages the neuron's response over several trials with the same stimulus, reverse correlation means that the experimenter averages the input under the condition of an identical response, viz., a spike. In other words, it is a spike-triggered average; see,

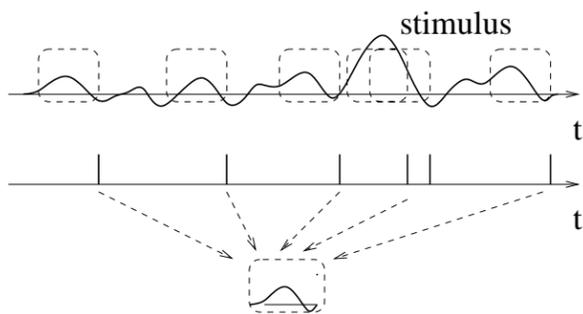


Figure 1.8. Reverse correlation technique (schematic). The stimulus in the top trace has caused the spike train shown immediately below. The time course of the stimulus just before the spikes (dashed boxes) has been averaged to yield the typical time course (bottom).

e.g., [de Ruyter van Steveninck and Bialek, 1988; Rieke et al., 1996]. The results of the reverse correlation, i.e., the typical time course of the stimulus which has triggered the spike, can be interpreted as the ‘meaning’ of a single spike. Reverse correlation techniques have made it possible for example to measure the spatio-temporal characteristics of neurons in the visual cortex [Eckhorn et al., 1993; DeAngelis et al., 1995].

With a somewhat more elaborate version of this approach, W. Bialek and his co-workers have been able to ‘read’ the neural code of the H1 neuron in the fly and to reconstruct a time-dependent stimulus [Bialek et al., 1991; Rieke et al., 1996]. Here we give a simplified version of the argument.

Results from reverse correlation analysis suggest, that each spike signifies the time course of the stimulus preceding the spike. If this is correct, a reconstruction of the complete time course of the stimulus $s(t)$ from the set of firing times $\mathcal{F} = \{t^{(1)}, \dots, t^{(n)}\}$ should be possible; see Figure 1.9.

As a simple test of this hypothesis, Bialek and coworkers have studied a linear reconstruction. A spike at time $t^{(f)}$ gives a contribution $\kappa(t - t^{(f)})$ to the estimation $s^{\text{est}}(t)$ of the time course of the stimulus. Here, $t^{(f)} \in \mathcal{F}$ is one of the firing times and $\kappa(t - t^{(f)})$ is a kernel which is nonzero during some time before and around $t^{(f)}$; see inset of Figure 1.9. A linear estimate of the stimulus is

$$s^{\text{est}}(t) = \sum_{f=1}^n \kappa(t - t^{(f)}). \quad (1.7)$$

The form of the kernel κ was determined through optimization so that the average reconstruction error $\int dt [s(t) - s^{\text{est}}(t)]^2$ was minimal. The quality of the reconstruction was then tested on additional data which was not used for the optimization. Surprisingly enough, the simple linear reconstruction (1.7) gave a fair estimate of the time course of the stimulus [Bialek et al., 1991; Bialek and Rieke, 1992; Rieke et al., 1996]. These results show nicely that information about a time dependent input can indeed be conveyed by spike timing.

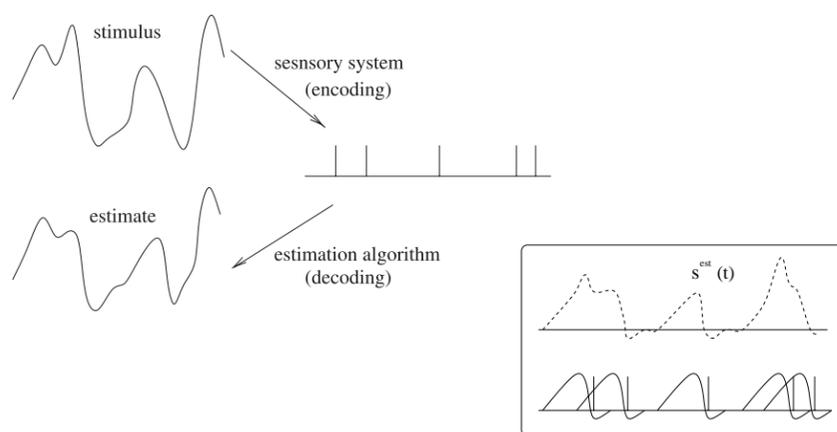


Figure 1.9. Reconstruction of a stimulus (schematic). A stimulus evokes a spike train of a neuron. The time course of the stimulus may be estimated from the spike train. The inset shows the principle of linear stimulus reconstruction. The estimation $s^{\text{est}}(t)$ (dashed) is the sum of the contributions (solid lines) of all spikes. Main figure redrawn after [Rieke et al., 1996].

1.1.4 Discussion: Spikes or Rates?

The dividing line between pulse codes and firing rates is not always as clearly drawn as it may seem at first sight. Some codes which were first proposed as pure examples of pulse codes have later been interpreted as variations of rate codes.

For example the stimulus reconstruction (1.7) with kernels seems to be a clear example of a pulse code. Nevertheless, it is also not so far from a rate code based on spike counts [Theunissen and Miller, 1995]. To see this, consider a spike count measure with a running time window $K(\cdot)$. We can estimate the rate ν at time t by

$$\nu(t) = \frac{\int K(\tau) S(t - \tau) d\tau}{\int K(\tau) d\tau} \quad (1.8)$$

where $S(t) = \sum_{f=1}^n \delta(t - t^{(f)})$ is the spike train under consideration. The integrals run from minus to plus infinity. For a rectangular time window $K(\tau) = 1$ for $-T/2 < \tau < T/2$ and zero otherwise, (1.8) reduces exactly to our definition (1.4) of a rate as a spike count measure.

The time window in (1.8) can be made rather short so that at most a few spikes fall into the interval T . Furthermore, there is no need that the window $K(\cdot)$ be symmetric and rectangular. We may just as well take an asymmetric time window with smooth borders. Moreover, we can perform the integration over the δ function which yields

$$\nu(t) = c \sum_{f=1}^n K(t - t^{(f)}) \quad (1.9)$$

where $c = [\int K(s) ds]^{-1}$ is a constant. Except for the normalization, the generalized rate formula (1.9) is now identical to the reconstruction for-

mula (1.7). In other words, the linear reconstruction is just the firing rate measured with a cleverly optimized time window.

Similarly, a code based on the 'time-to-first-spike' is also consistent with a rate code. If, for example, the mean firing rate of neuron is high for a given stimulus, then the first spike is expected to occur early. If the rate is low, the first spike is expected to occur later. Thus the timing of the first spike contains a lot of information about the underlying rate.

Finally, a code based on population activities introduced in Section 1.1.2 as an example of a rate code may be used for very fast temporal coding schemes [Tsodyks and Sejnowski, 1995]. As discussed later in Chapter 10 the population activity reacts quickly to any change in the stimulus. Thus rate coding in the sense of a population average is consistent with fast temporal information processing, whereas rate coding in the sense of a naïve spike count measure is not.

We do not want to go into the details of the discussion whether or not to call a given code a rate code [Theunissen and Miller, 1995]. What is important, in our opinion, is to have a coding scheme which allows neurons to quickly respond to stimulus changes. A naïve spike count code with a long time window is unable to do this, but many of the other codes are. The name of such a code, whether it is deemed a rate code or not is of minor importance.

In this book, we will explore some of the possibilities of coding and computation by spikes. As modelers – mathematicians, physicists, and engineers – our aim is not to give a definite answer to the problem of neural coding in the brain. The final answers have to come from experiments. One possible task of modeling may be to discuss candidate coding schemes, study their computational potential, exemplify their utility, point out their limitations – and this is what we will attempt to do in the course of the following chapters.

1.2 Neuron Models

Neural activity may be described at several levels of abstraction. On a microscopic level, there are a large number of ion channels, pores in the cell membrane which open and close depending on the voltage and the presence (or absence) of various chemical messenger molecules. Compartmental models, where each small segment of a neuron is described by a set of ionic equations, aim at a description of these processes. A short introduction to this model class can be found in section 1.2.4.

On a higher level of abstraction, we do not worry about the spatial structure of a neuron nor about the exact ionic mechanisms. We consider the neuron as a homogeneous unit which generates spikes if the total excitation is sufficiently large. This is the level of the so-called integrate-and-fire models. In Section 1.2.3, we will discuss this model class in the framework of the 'spike response model'.

The spiking neuron models should be contrasted with the rate models reviewed in Section 1.2.5. Rate models neglect the pulse structure of the neuronal output, and are therefore higher up in the level of abstraction. On a

yet coarser level would be models which describe the activity in and interaction between whole brain areas.

Most chapters in the book will make use of a generic neuron model on the intermediate description level. We therefore devote most of the space to the discussion in Section 1.2.3. For those readers who are not interested in the details, we present the basic concepts of our generic neuron model in a compressed version in the following section 1.2.1.

1.2.1 Simple Spiking Neuron Model

Spike Response Model - definitions

The state of neuron i is described by a state variable u_i . The neuron is said to fire, if u_i reaches a threshold ϑ . The moment of threshold crossing defines the firing time $t_i^{(f)}$; see Figure 1.10. The set of all firing times of neuron i is denoted by

$$\mathcal{F}_i = \{t_i^{(f)}; 1 \leq f \leq n\} = \{t \mid u_i(t) = \vartheta\}. \quad (1.10)$$

For the most recent spike $t_i^{(f)} < t$ of neuron i we write either $t_i^{(n)}$ or, shorter, \hat{t} .

Two different processes contribute to the value of the state variable u_i .

First, immediately after firing an output spike at $t_i^{(f)}$, the variable u_i is lowered or 'reset'. Mathematically, this is done by adding a negative contribution $\eta_i(t - t_i^{(f)})$ to the state variable u_i . An example of a refractory function η_i is shown in Figure 1.10. The kernel $\eta_i(s)$ vanishes for $s \leq 0$ and decays to zero for $s \rightarrow \infty$.

Second, the model neuron may receive input from presynaptic neurons $j \in \Gamma_i$ where

$$\Gamma_i = \{j \mid j \text{ presynaptic to } i\}. \quad (1.11)$$

A presynaptic spike at time $t_j^{(f)}$ increases (or decreases) the state u_i of neuron i for $t > t_j^{(f)}$ by an amount $w_{ij} \epsilon_{ij}(t - t_j^{(f)})$. The weight w_{ij} is a factor which accounts for the strength of the connection. An example of an ϵ_{ij} function is shown in Figure 1.10b. The effect of a presynaptic spike may be positive (excitatory) or negative (inhibitory). Because of causality, the kernel $\epsilon_{ij}(s)$ must vanish for $s \leq 0$. A transmission delay may be included in the definition of ϵ_{ij} ; see Figure 1.10.

The state $u_i(t)$ of model neuron i at time t is given by the linear superposition of all contributions,

$$u_i(t) = \sum_{t_i^{(f)} \in \mathcal{F}_i} \eta_i(t - t_i^{(f)}) + \sum_{j \in \Gamma_i} \sum_{t_j^{(f)} \in \mathcal{F}_j} w_{ij} \epsilon_{ij}(t - t_j^{(f)}). \quad (1.12)$$

An interpretation of the terms on the right-hand side of (1.12) is straightforward. The η_i contributions describe the response of neuron i to its own spikes. The ϵ_{ij} kernels model the neurons response to presynaptic spikes.

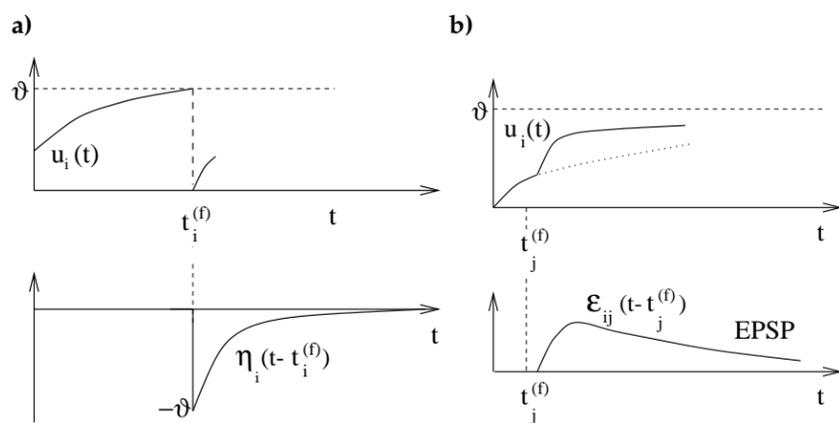


Figure 1.10. a) The state variable $u_i(t)$ reaches the threshold ϑ at time $t_i^{(f)}$. Immediately afterwards $u_i(t)$ is reset to zero. The reset is performed by adding a kernel $\eta_i(t - t_i^{(f)})$. The function $\eta_i(s)$ takes care of refractoriness after a spike emitted at $s = 0$. b) The state variable $u_i(t)$ changes after a presynaptic spike has occurred at $t_j^{(f)}$. The kernels ϵ_{ij} describes the response of u_i to a presynaptic spike at $s = 0$. The postsynaptic potential can either be excitatory (EPSP) or inhibitory (IPSP).

We will refer to (1.10) - (1.12) as the Spike Response Model (SRM). In a biological context, the state variable u_i may be interpreted as the electrical membrane potential. The kernels ϵ_{ij} are the postsynaptic potentials and η_i accounts for neuronal refractoriness.

To be more specific, let us consider some examples of suitable functions η_i and ϵ_{ij} . The kernel $\eta_i(s)$ is usually nonpositive for $s > 0$. A typical form of η_i is shown in Figure 1.10a). A specific mathematical formulation is

$$\eta_i(s) = -\vartheta \exp\left(-\frac{s}{\tau}\right) \mathcal{H}(s) \quad (1.13)$$

where τ is a time constant and $\mathcal{H}(s)$ is the Heaviside step function which vanishes for $s \leq 0$ and takes a value of 1 for $s > 0$. Note that at the moment of firing $u_i(t) = \vartheta$. The effect of (1.13) is that after each firing the state variable u_i is reset to zero. If the factor ϑ on the right-hand side of (1.13) is replaced by a parameter $\eta_0 \neq \vartheta$, then the state variable would be reset to a value $\vartheta - \eta_0 \neq 0$. For a further discussion of the η -kernel, the reader is referred to section 1.2.3.1.

The kernels ϵ_{ij} describe the response to presynaptic spikes; see Figure 1.10b). For excitatory synapses ϵ_{ij} is non-negative and is called the excitatory postsynaptic potential (EPSP). For inhibitory synapses, the kernel takes non-positive values and is called the inhibitory postsynaptic potential (IPSP). One of several potential mathematical formulations is

$$\epsilon_{ij}(s) = \left[\exp\left(-\frac{s - \Delta^{\text{ax}}}{\tau_m}\right) - \exp\left(-\frac{s - \Delta^{\text{ax}}}{\tau_s}\right) \right] \mathcal{H}(s - \Delta^{\text{ax}}). \quad (1.14)$$

where τ_s, τ_m are time constants and Δ^{ax} is the axonal transmission delay. The amplitude of the response is scaled via the factor w_{ij} in (1.12). For

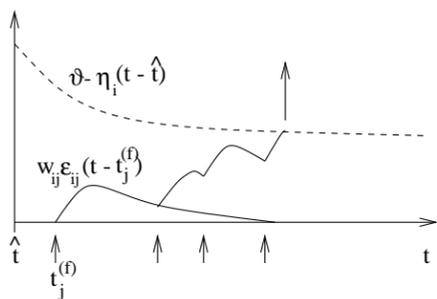


Figure 1.11. Dynamic threshold interpretation. The last firing of neuron i has occurred at $t = \hat{t}$. Immediately after firing the dynamic threshold $\vartheta - \eta_i(t - \hat{t})$ (dashed) is high. The next output spike occurs, when the sum of the EPSPs $\sum_{t_j^{(f)} \in \mathcal{F}_j} w_{ij} \epsilon_{ij}(t - t_j^{(f)})$ caused by presynaptic spikes at times $t_j^{(f)}$ (arrows in the lower part of the figure) reaches the dynamic threshold again.

inhibitory synapses, the kernel ϵ_{ij} would have a negative sign in front of the expression on the right-hand side. Alternatively, we can put the sign in the synaptic efficacy and use $w_{ij} > 0$ for excitatory synapses and $w_{ij} < 0$ for inhibitory synapses.

Equations (1.10) and (1.12) give a fairly general framework for the discussion of neuron models. We will show in Section 1.2.3., that the Spike Response Model (1.10) - (1.12) with kernels (1.13) and (1.14) is equivalent to the integrate-and-fire model. Furthermore, with a different choice of kernels, the Spike Response Model also approximates the Hodgkin-Huxley equations with time-dependent input; see Section 1.2.4. and [Kistler et al., 1997].

Dynamic threshold model

We note that (1.10) - (1.12) may also be formulated in terms of a dynamic threshold model. To see this, consider the threshold condition $u_i(t) = \vartheta$; see (1.10). With (1.12) we get

$$\sum_{j \in \Gamma_i} \sum_{t_j^{(f)} \in \mathcal{F}_j} w_{ij} \epsilon_{ij}(t - t_j^{(f)}) = \vartheta - \sum_{t_i^{(f)} \in \mathcal{F}_i} \eta_i(t - t_i^{(f)}) \quad (1.15)$$

where we have moved the sum over the η_i 's to the right-hand side of (1.15). We may consider the expression $\vartheta - \sum_{t_i^{(f)} \in \mathcal{F}_i} \eta_i(t - t_i^{(f)})$ as a dynamic threshold which increases after each firing and decays slowly back to its asymptotic value ϑ in case of no further firing of neuron i .

Short term memory

There is a variant of the Spike Response Model which is often useful to simplify the analytical treatment. We assume that only the last firing contributes to refractoriness. Hence, we simplify (1.12) slightly and only keep the influence of the *most recent* spike in the sum over the η contributions. Formally, we make the replacement

$$\sum_{t_i^{(f)} \in \mathcal{F}_i} \eta(t - t_i^{(f)}) \longrightarrow \eta(t - \hat{t}_i) \quad (1.16)$$

where $\hat{t}_i < t$ denotes the most recent firing of neuron i . We refer to this simplification as a neuron with short term memory. Instead of (1.12), the membrane potential of neuron i is now

$$u_i(t) = \eta_i(t - \hat{t}_i) + \sum_{j \in \Gamma_i} \sum_{t_j^{(f)} \in \mathcal{F}_j} w_{ij} \epsilon_{ij}(t - t_j^{(f)}). \quad (1.17)$$

The next spike occurs when

$$\sum_{j \in \Gamma_i} \sum_{t_j^{(f)} \in \mathcal{F}_j} w_{ij} \epsilon_{ij}(t - t_j^{(f)}) = \vartheta - \eta(t - \hat{t}_i). \quad (1.18)$$

A graphical interpretation of (1.18) is given in Fig. 1.11.

External input

A final modification concerns the possibility of external input. In addition to (or instead of) spike input from other neurons, a neuron may receive an analog input current $\mathcal{I}^{\text{ext}}(t)$, for example from a non-spiking sensory neuron. In this case, we add on the right-hand side of (1.12) a term

$$h^{\text{ext}}(t) = \int_0^\infty \tilde{\epsilon}(s) \mathcal{I}^{\text{ext}}(t - s) ds. \quad (1.19)$$

Here $\tilde{\epsilon}$ is another kernel, which describes the response of the membrane potential to an external input pulse. As a notational convenience, we introduce a new variable h which summarizes all contributions from other neurons and from external sources

$$h(t) = \sum_{j \in \Gamma_i} w_{ij} \sum_{t_j^{(f)} \in \mathcal{F}_j} \epsilon_{ij}(t - t_j^{(f)}) + h^{\text{ext}}(t). \quad (1.20)$$

The membrane potential of a neuron with short term memory is then simply

$$u_i(t) = \eta(t - \hat{t}_i) + h(t). \quad (1.21)$$

We will make use of (1.21) repeatedly, since it allows us to analyze the neuronal dynamics in a transparent manner.

The equations (1.10) - (1.21) will be used in several chapters of this book. The following subsections 1.2.3-1.2.5 will put this generic neuron model in the context of other models of neural activity. Readers who are not interested in the details may proceed directly to Chapter 2 - or continue with the next subsection for a first introduction to coding by spikes in the framework of the above spiking neuron model.

1.2.2 First Steps towards Coding by Spikes

Before we proceed further with our discussion of neuron models, let us take a first glance at the type of computation we can do with such a model. To this end, we will reconsider some of the pulse codes introduced in Section 1.1.3. A full discussion of computation with spiking neurons follows

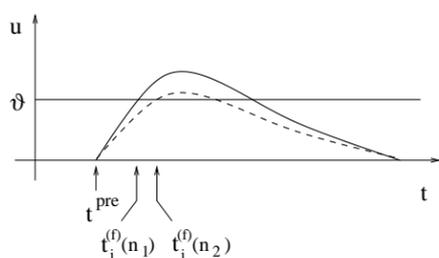


Figure 1.12. Time to first spike. The firing time $t^{(f)}$ encodes the number n_1 or n_2 of presynaptic spikes which have been fired synchronously at t^{pre} . If there are less presynaptic spikes, the potential u rises more slowly (dashed) and the firing occurs later. For the sake of simplicity, the axonal delay has been set to zero.

in Chapter 2 of this book. Here we use simple arguments from a graphical analysis to get a first understanding of how the model works.

Time-to-first-spike

Let us start with a coding scheme based on the ‘time-to-first-spike’. In order to simplify the argument, let us consider a single neuron i which receives spikes from N presynaptic neurons j over synaptic connections which all have the same weight $w_{ij} = w_0$. There is no external input. We assume that the last spike of neuron i occurred long ago so that the spike afterpotential $\eta(\cdot)$ in (1.12) may be neglected.

At $t = t^{\text{pre}}$, a total number of $n_1 < N$ presynaptic spikes are simultaneously generated and transmitted to the postsynaptic neuron i . For $t > t^{\text{pre}}$, the potential of i is

$$u_i(t) = n_1 w_0 \epsilon(t - t^{\text{pre}}). \quad (1.22)$$

An output spike of neuron i occurs whenever u_i reaches the threshold ϑ . We consider the firing time $t_i^{(f)}$ of the first output spike

$$t_i^{(f)} = \min\{t > t^{\text{pre}} \mid u_i(t) = \vartheta\} \quad (1.23)$$

A graphical solution of (1.23) is shown in Figure 1.12. If there are less presynaptic spikes $n_2 < n_1$, then the postsynaptic potential is reduced and the firing occurs later as shown by the dashed line in Figure 1.12. It follows that the time difference $t_i^{(f)} - t^{\text{pre}}$ is a measure of the number of presynaptic pulses. To put it differently, the timing of the first spike encodes the input strength.

Phase coding

Phase coding is possible if there is some periodic background signal which can serve as a reference. We include the background into the external input and write

$$h^{\text{ext}}(t) = h_0 + h_1 \cos(2\pi \frac{t}{T}) \quad (1.24)$$

where h_0 is a constant and h_1 is the amplitude of the T -periodic signal.

Let us consider a single neuron driven by (1.24). There is no input from other neurons. We start from the simplified spike response mode (1.21)

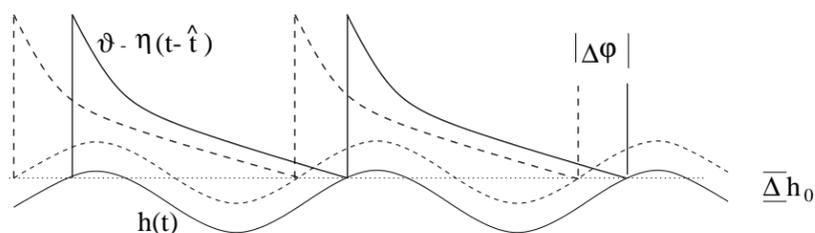


Figure 1.13. Phase coding. Firing occurs whenever the total input potential $h(t) = h_0 + h_1 \cos(2\pi t/T)$ hits the dynamic threshold $\vartheta - \eta(t - \hat{t})$ where \hat{t} is the most recent firing time; cf. Fig. 1.11. In the presence of a periodic modulation $h_1 \neq 0$, a change Δh_0 in the level of (constant) stimulation results in a change $\Delta\varphi$ in the phase of firing.

which yields the membrane potential

$$u(t) = \eta(t - \hat{t}) + h^{\text{ext}}(t). \quad (1.25)$$

As usual \hat{t} denotes the time of the most recent spike. To find the next firing time, (1.25) has to be combined with the threshold condition $u(t) = \vartheta$. We are interested in a solution where the neuron fires regularly and with the same period as the background signal. In this case the threshold condition reads

$$\vartheta - \eta(T) = h_0 + h_1 \cos\left(2\pi \frac{\hat{t}}{T}\right). \quad (1.26)$$

For a given period T , the left-hand side has a fixed value and we can solve for $\varphi = 2\pi \frac{\hat{t}}{T}$. For most combinations of parameters, there are two solutions but only one of them is stable. Thus the neuron has to fire at a certain phase φ with respect to the external signal. The value of φ depends on the level of the constant stimulation h_0 . In other words, the strength h_0 of the stimulation is encoded in the phase of the spike. In (1.26) we have moved η to the left-hand side in order to suggest a dynamic threshold interpretation. A graphical illustration of equation (1.26) is given in Figure 1.13.

Correlation coding

Let us consider two identical uncoupled neurons. Both receive the same constant external stimulus $h^{\text{ext}}(t) = h_0$. As a result, they fire regularly with period T given by $\eta(T) = h_0$ as can be seen directly from (1.26) with $h_1 = 0$. Since the neurons are not coupled, they need not fire simultaneously. Let us assume that the firings of neuron 2 are shifted by an amount δ with respect to neuron 1.

Suppose that, at a given moment t^{pre} , both neurons receive input from a common presynaptic neuron j . This causes an additional contribution $\epsilon(t - t^{\text{pre}})$ to the membrane potential. If the synapse is excitatory, the two neurons will fire slightly sooner. More importantly, the spikes will also be closer together. In the situation sketched in Figure 1.14 the new firing time difference $\tilde{\delta}$ is reduced, $\tilde{\delta} < \delta$. In later chapters, we will analyze this phenomenon in more detail. Here we just note that this effect allows us to encode information using the time interval between the firings of two or more neurons. The reader who is interested in the computational aspects of coding by firing time differences may move directly to Chapter 2. The

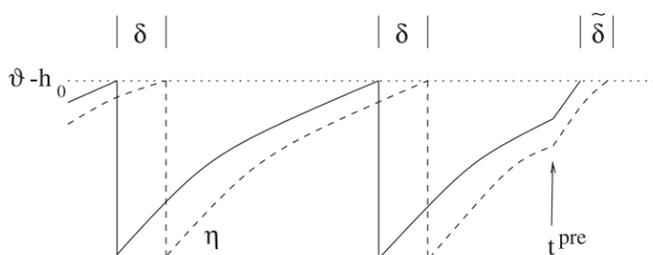


Figure 1.14. The firing time difference δ between two independent neurons is decreased to $\tilde{\delta} < \delta$, after both neurons receive a common excitatory input at time t^{pre} .

above argument also plays a major role in chapters 10 and 11 in the context of neuronal locking.

The remainder of Chapter 1 continues with a discussion of neuron models. Before turning to conductance-based neuron models, we want to put our simple neuron model into a larger context.

1.2.3 Threshold-Fire Models

The simple spiking neuron model introduced in Section 1.2.1 is an instance of a ‘threshold-fire model’. The firing occurs at the moment when the state variable u crosses the threshold. A famous example in this model class is the ‘integrate-and-fire’ model.

In this section we review the arguments that motivate our simple model of a spiking neuron the Spike Response Model introduced in Section 1.2.1. We show the relation of the model to the integrate-and-fire model and discuss several variants. Finally we discuss several noisy versions of the model.

1.2.3.1 Spike Response Model - Further Details

In this paragraph we want to motivate the simple model of a spiking neuron introduced in Section 1.2.1., give further details, and discuss it in a more general context. Let us start and review the arguments for the simple neuron model.

We aim for a simple model which captures some generic properties of neural activity without going into too much detail. The neuronal output should consist of pulses. In real spike trains, all action potentials of a given neuron look alike. The pulses of our model can therefore be treated as stereotyped events that occur at certain firing times $t_i^{(f)}$. The lower index i denotes the neuron, the upper index is the spike number. A spike train is fully characterized by the set of firing times

$$\mathcal{F}_i = \{t_i^{(1)}, \dots, t_i^{(n)}\} \quad (1.27)$$

already introduced in Equation (1.1).

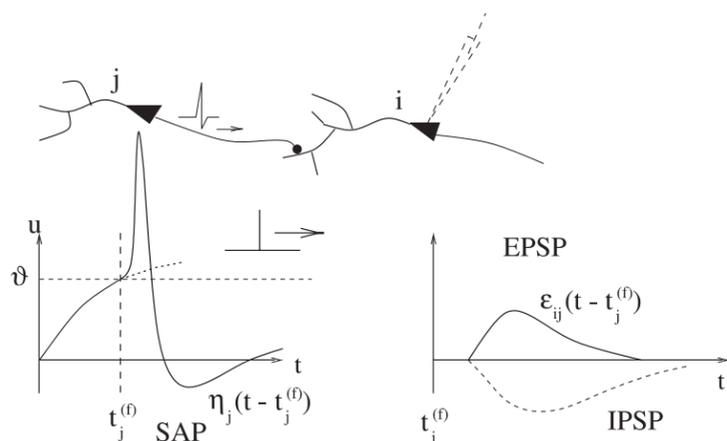


Figure 1.15. The Spike Response Model as a generic framework to describe the spike process. Spikes are generated by a threshold process whenever the membrane potential u crosses the threshold ϑ . The threshold crossing triggers the spike followed by a spike afterpotential (SAP), summarized in the function $\eta_j(t - t_j^{(f)})$. The spike evokes a response of the postsynaptic neuron described by the kernel $\epsilon_{ij}(t - t_j^{(f)})$. The voltage response to an input at an excitatory synapse is called the excitatory postsynaptic potential (EPSP) and can be measured with an electrode (schematically). Spike arrival at an inhibitory would cause an inhibitory postsynaptic potential (IPSP) which as a negative effect (dashed line).

The internal state of a model neuron is described by a single variable u . We will refer to u_i as the membrane potential of neuron i . Spikes are generated when the membrane potential crosses a threshold ϑ from below. The moment of threshold crossing can be used to formally define the firing times. If $u_i(t) = \vartheta$ and $u_i'(t) > 0$, then $t = t_i^{(f)}$. As always, the prime denotes a derivative. The set of firing times therefore is

$$\mathcal{F}_i = \{t \mid u_i(t) = \vartheta \wedge u_i'(t) > 0\}. \quad (1.28)$$

In contrast to (1.10) we have included here explicitly that the threshold must be reached from below ($u_i' > 0$). In the simple model of Section 1.2.1, this condition was automatically fulfilled, since the state variable could never pass threshold: as soon as u_i reached ϑ , the state variable was reset to a value below threshold. In the following we want to be slightly more general.

After a spike is triggered by the threshold process, a whole sequence of events is initiated. Ion channels open and close, some ions flow through the cell membrane into the neuron, others flow out. The result of these ionic processes is the action potential, a sharp peak of the voltage followed by a long lasting negative afterpotential. As mentioned before, the forms of the spike and its afterpotential are always roughly the same. Their generic time course will be described in our model by a function $\eta_i(s)$, where $s = t - t_i^{(f)} > 0$ is the time since the threshold crossing at $t_i^{(f)}$. A typical form of $\eta_i(\cdot)$ is sketched in Figures 1.15 and 1.16a. Also, information about the spike event is transmitted to other postsynaptic neurons; see

Figure 1.15. The response of these neurons is described by another function $\epsilon_{ij}(s)$ which will be discussed further below. Let us concentrate on the function η_i first.

Since the form of the pulse itself does not carry any information, the exact time course during the positive part of the spike is irrelevant. Notice, however, that while the action potential is quickly rising or steeply falling, emission of a further spike is impossible. This effect is called absolute refractoriness. Important for refractoriness is also the spike afterpotential (SAP). A negative spike afterpotential means that the emission of a second spike immediately after the first pulse is more difficult. The time of reduced sensitivity after a spike is called the relative refractory period.

In an attempt to simplify the neuron model, we may therefore replace the initial segment of η_i by an absolute refractory period and concentrate on the negative spike after potential only. This is shown in Figure 1.16b. Here the positive part of the spike is reduced to a pulse of negligible width. Its sole purpose is to mark the firing time $s = 0$. Absolute and negative refractoriness may be modeled by

$$\eta_i(s) = -\eta_0 \exp\left(-\frac{s - \delta^{\text{abs}}}{\tau}\right) \mathcal{H}(s - \delta^{\text{abs}}) - K \mathcal{H}(s) \mathcal{H}(\delta^{\text{abs}} - s) \quad (1.29)$$

with a constant $K \rightarrow \infty$ in order to ensure absolute refractoriness during the time δ^{abs} after the firing. The Heaviside step function $\mathcal{H}(s)$ is unity for $s > 0$ and vanishes for $s \leq 0$. The constant η_0 is a parameter which scales the amplitude of relative refractoriness. If we are interested in a question of spike timing (but not in the form of the action potential), then a model description with such a simplified kernel may be fully sufficient.

Let us now consider two neurons connected via a synapse. If the presynaptic neuron j fires at time $t_j^{(f)}$, a pulse travels along the axon to the synapse where it evokes some response of the postsynaptic neuron i . In our model, we disregard all details of the transmission process and concentrate on the effect that the pulse has on the membrane potential at the soma of neuron i . This response is a measurable function called the postsynaptic potential (PSP) and can be positive (excitatory) or negative (inhibitory). The typical time course of an excitatory postsynaptic potential (EPSP) is sketched in Figure 1.15. In our model, the time course is described by a function $\epsilon_{ij}(s)$ where $s = t - t_j^{(f)}$ is the time which has passed since the emission of the presynaptic pulse. The kernel $\epsilon_{ij}(s)$ vanishes for $s \leq \Delta^{\text{ax}}$. We refer to Δ^{ax} as the axonal transmission delay. We often approximate the time course for $t - t_j^{(f)} > \Delta^{\text{ax}}$ by a so-called α -function $\propto x e^{-x}$ where $x = (t - t_j^{(f)} - \Delta^{\text{ax}})/\tau_s$ and τ_s some time constant. Another possibility is to describe the form of the response by the double exponential introduced in (1.14).

So far we have restricted our discussion to a pair of neurons. In reality, each postsynaptic neuron i will receive input from many different presynaptic neurons $j \in \Gamma_i$. All inputs cause some postsynaptic response and contribute to the membrane potential of i . In our model, we assume that the total membrane potential u_i of neuron i is the linear superposition of

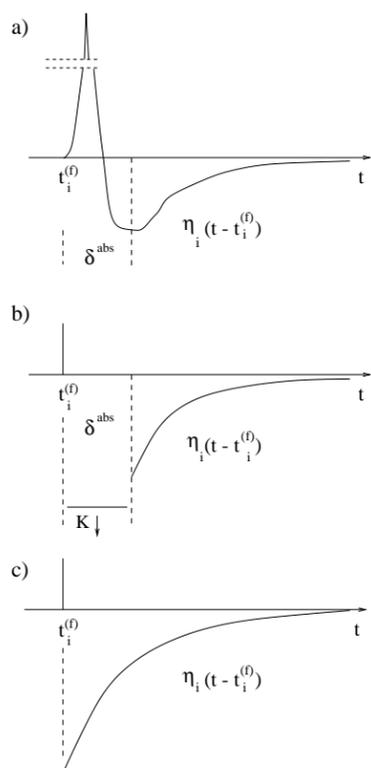


Figure 1.16. a) Action potential and spike after potential (schematic). The peak of the pulse is out of scale; cf. Fig.1.25. During the time δ^{abs} , emission of a further action potential is practically impossible. b) A simplified kernel η_i which includes an absolute refractory period of δ^{abs} followed by an exponential decay. The form of the action potential is not described explicitly. The firing time $t_i^{(f)}$ is marked by a vertical bar. c) As a further simplification of η_i , absolute refractoriness may be neglected.

all contributions

$$u_i(t) = \sum_{t_i^{(f)} \in \mathcal{F}_i} \eta_i(t - t_i^{(f)}) + \sum_{j \in \Gamma_i} \sum_{t_j^{(f)} \in \mathcal{F}_j} w_{ij} \epsilon_{ij}(t - t_j^{(f)}). \quad (1.30)$$

As above, the kernel η_i describes the neuron's response to its own firing. (In the following we often omit the subscript i .) The kernel ϵ_{ij} describes the generic response of neuron i to spikes from each presynaptic neurons $j \in \Gamma_i$. The weight w_{ij} gives the amplitude of the response. It corresponds to the synaptic efficacy of the connection from j to i . For the sake of simplicity, we often assume that the response has the same form for any pair ij of neurons, except for an amplitude factor w_{ij} . This means that we may suppress the index ij of ϵ_{ij} in (1.30) and write ϵ instead of ϵ_{ij} . Since the membrane potential in (1.30) is expressed in terms of response kernels, we will refer to the above description of neuronal activity as the Spike Response Model [Gerstner, 1991; Gerstner and van Hemmen, 1992; Gerstner et al., 1996]. Equation (1.30) is exactly the simplified neuron model introduced already in Section 1.2.1.

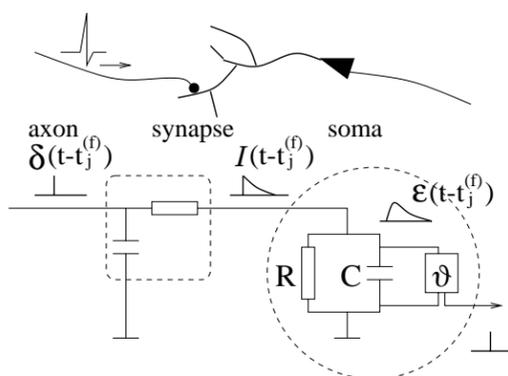


Figure 1.17. Integrate-and-fire neuron. The basic module is the RC circuit shown inside the circle on the right-hand side of the diagram. The circuit is charged by an input current I . If the voltage across the capacitor reaches a threshold ϑ the circuit is shunted and a δ -pulse is transmitted to other neurons (lower right). A δ -pulse sent out by a presynaptic neuron and travelling on the presynaptic axon (left), is low-pass filtered first (middle) before it is fed as a current pulse $I(t - t_j^{(f)})$ into the integrate-and-fire circuit. The voltage response of the RC circuit to the presynaptic pulse is the postsynaptic potential $\epsilon(t - t_j^{(f)})$

Equation (1.30) is a linear equation for the membrane potential. All contributions to u_i are caused by the firing events $t_i^{(f)}$ and $t_j^{(f)}$. The essential nonlinearity of the neuronal dynamics is given by the threshold condition (1.28) which defines the firing times. The Spike Response Model is defined by the combination of (1.28) and (1.30) and is the starting point for the analysis of spike based computation in Chapter 2. It is also used in some other chapters, e.g., Chapters 9 and 10.

1.2.3.2 Integrate-and-Fire Model

An important example in the class of ‘threshold-fire models’ is the integrate-and-fire neuron. The basic circuit of an integrate-and-fire model consists of a capacitor C in parallel with a resistor R driven by a current $I(t)$; see Figure 1.17. The driving current splits into two components, one charging the capacitor, the other going through the resistor. Conservation of charge yields the equation

$$I(t) = \frac{u(t)}{R} + C \frac{du}{dt} \quad (1.31)$$

where u is the voltage across the capacitor C . We introduce the time constant $\tau_m = RC$ of the ‘leaky integrator’ and write (1.31) in the standard form

$$\tau_m \frac{du}{dt} = -u(t) + RI(t). \quad (1.32)$$

We refer to u as the membrane potential and to τ_m as the membrane time constant of the neuron.

Equation (1.32) is a first-order linear differential equation and cannot describe full neuronal spiking behavior. To incorporate the essence of pulse

emission, (1.32) is supplemented by a threshold condition. A threshold crossing $u(t^{(f)}) = \vartheta$ is used to define the firing time $t^{(f)}$. The form of the spike is not described explicitly. Immediately after $t^{(f)}$, the potential is reset to a new value u_r ,

$$\lim_{\delta \rightarrow 0} u(t^{(f)} + \delta) = u_r. \quad (1.33)$$

For $t > t^{(f)}$ the dynamics is again given by (1.32), until the next threshold crossing occurs. The combination of leaky integration (1.32) and reset (1.33) defines the basic integrate-and-fire model.

To see how the model works, let us consider a model neuron with constant input current I_0 and reset potential $u_r = 0$. We assume that a first spike has occurred at $t = t^{(0)}$. The trajectory of the membrane potential can be found by integrating (1.32) with the initial condition $u(t^{(0)}) = u_r = 0$. The solution is

$$u(t) = R I_0 \left[1 - \exp\left(-\frac{t - t^{(0)}}{\tau_m}\right) \right]. \quad (1.34)$$

For $R I_0 < \vartheta$ no further spike can occur. For $R I_0 > \vartheta$, the membrane potential reaches the threshold ϑ at time $t^{(1)}$, which can be found from the threshold condition

$$\vartheta = R I_0 \left[1 - \exp\left(-\frac{t^{(1)} - t^{(0)}}{\tau_m}\right) \right]. \quad (1.35)$$

Solving (1.35) for the time interval $T = t^{(1)} - t^{(0)}$ yields

$$T = \tau_m \ln \frac{R I_0}{R I_0 - \vartheta}. \quad (1.36)$$

After the spike at $t^{(1)}$ the membrane potential is again reset to $u_r = 0$ and the integration process starts again. We conclude that for a constant input current I_0 , the integrate-and-fire neuron fires regularly with period T given by (1.36).

Refractoriness

It is straightforward to include an absolute refractory period. After a spike at $t^{(f)}$, we force the membrane potential to a value $u = -K < 0$ and keep it there during a time δ^{abs} . At $t^{(f)} + \delta^{\text{abs}}$ we restart the integration (1.32) with the initial value $u = u_r$.

As before, we can solve the dynamics for a constant input current I_0 . If $R I_0 > \vartheta$, the neuron will fire regularly. Due to the absolute refractory period the interval between firings is now longer by an amount δ^{abs} compared to the value in (1.36). Instead of giving the interval T between two spikes, the result is often stated in terms of the mean firing rate $\nu = 1/T$, viz.,

$$\nu = \left[\delta^{\text{abs}} + \tau_m \ln \frac{R I_0}{R I_0 - \vartheta} \right]^{-1}. \quad (1.37)$$

The firing rate of the integrate-and-fire neuron as a function of the constant input current is plotted in Figure 1.18.

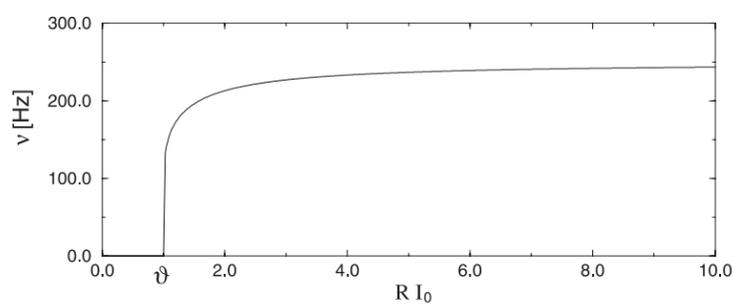


Figure 1.18. Gain function of an integrate-and-fire neuron with absolute refractoriness.

Synaptic currents

If the integrate-and-fire model is part of a network of spiking neurons, then the input current $I(t)$ must be generated somehow from the pulses of other neurons. A simple possibility is to describe the spikes of a presynaptic neuron j as Dirac δ -pulses and feed them directly into the postsynaptic neuron i . The input *current* to unit i is then

$$I_i(t) = \sum_{j \in \Gamma_i} c_{ij} \sum_{t_j^{(f)} \in \mathcal{F}_j} \delta(t - t_j^{(f)}). \quad (1.38)$$

The factor c_{ij} is a measure of the strength of the connection from j to i and corresponds directly to the charge deposited on the capacitor C by a single presynaptic pulse of neuron j . The parameter c_{ij} is, of course, proportional to the synaptic efficacy w_{ij} as we will see later on.

More generally, we can say that each presynaptic spike generates a *current* pulse of finite width and with time course $\alpha(t - t_j^{(f)})$ for $t > t_j^{(f)}$. In this case, the input current to neuron i should be written as

$$I_i(t) = \sum_{j \in \Gamma_i} c_{ij} \sum_{t_j^{(f)} \in \mathcal{F}_j} \alpha(t - t_j^{(f)}). \quad (1.39)$$

This is not too far from reality, since an input spike arriving at a synapse from j to i indeed evokes a current through the membrane of the postsynaptic neuron i . From measurements it is known that the form of the postsynaptic current (PSC) can often be approximated by

$$\alpha(s) = \frac{s - \Delta^{\text{ax}}}{\tau_s^2} \exp\left(-\frac{s - \Delta^{\text{ax}}}{\tau_s}\right) \mathcal{H}(s - \Delta^{\text{ax}}) \quad (1.40)$$

where τ_s is a synaptic time constant in the millisecond range and Δ^{ax} is the axonal transmission delay. As usual, $\mathcal{H}(x)$ denotes the Heaviside step function which vanishes for $x \leq 0$ and has a value of one for $x > 0$. For a yet more realistic description of the synaptic input current the reader should consult the Section 1.2.4 on conductance-based neuron models in this chapter. In passing we remark that in the literature, a function of the form $x \exp(-x)$ is often called an α -function. While this has motivated our

choice of the symbol α for the synaptic input current, $\alpha(\cdot)$ in (1.39) could in fact stand for any form of an input current pulse.

In Figure 1.17 we have sketched the situation where $\alpha(s)$ consists of a simple exponentially decaying pulse

$$\alpha(s) = \frac{1}{\tau_s} \exp\left(-\frac{s}{\tau_s}\right) \mathcal{H}(s). \quad (1.41)$$

(1.41) is a first approximation to the low-pass characteristics of a synapse. Since the analytic expressions are simpler, (1.41) is often used instead of the more complicated expression (1.40). In simulations, it is convenient to generate the exponential pulse (1.41) by a differential equation for the *current*. The total postsynaptic current I_i into neuron i can be described by

$$\tau_s \frac{d}{dt} I_i(t) = -I_i + \sum_{j \in \Gamma_i} c_{ij} \sum_{t_j^{(f)} \in \mathcal{F}_j} \delta(t - t_j^{(f)}). \quad (1.42)$$

The differential equation (1.42) replaces then the input current (1.39). In order to see the relation between (1.42) and (1.39) more clearly, we integrate (1.42), which yields

$$I_i(t) = \sum_{j \in \Gamma_i} c_{ij} \sum_{t_j^{(f)} \in \mathcal{F}_j} \frac{1}{\tau_s} \exp\left(-\frac{t - t_j^{(f)}}{\tau_s}\right) \mathcal{H}(t - t_j^{(f)}). \quad (1.43)$$

Comparison of (1.43) with (1.39) and a current pulse according to (1.41) shows that the two formulations for the input current, the differential formulation (1.42) or the ‘integrated’ formulation (1.39) are indeed equivalent. In both cases, the resulting current is then put into (1.32) to get the voltage of the integrate-and-fire neuron.

Relation to the Spike Response Model

In this paragraph we show that the integrate-and-fire model discussed so far is in fact a special case of the Spike Response Model. To see this, we have to note two facts. First, (1.32) is a linear differential equation and can therefore easily be integrated. Second, the reset of the membrane potential after firing at time $t_i^{(f)}$ is equivalent to an outgoing current pulse of negligible width

$$I_i^{\text{out}}(t) = -C(\vartheta - u_r) \sum_{t_i^{(f)} \in \mathcal{F}_i} \delta(t - t_i^{(f)}) \quad (1.44)$$

where $\delta(\cdot)$ denotes the Dirac δ -function. We add the current (1.44) on the right-hand side of (1.32)

$$\tau_m \frac{du_i}{dt} = -u_i(t) + R I_i(t) + R I_i^{\text{out}}(t). \quad (1.45)$$

Let us check the effect of the last term. Integration of $\tau_m du/dt = R I_i^{\text{out}}$ yields at time $t_i^{(f)}$ indeed a reset of the potential from ϑ to u_r , as it should be.

Equation (1.45) may be integrated and yields

$$u_i(t) = \sum_{t_i^{(f)} \in \mathcal{F}_i} \eta(t - t_i^{(f)}) + \sum_{j \in \Gamma_i} w_{ij} \sum_{t_j^{(f)} \in \mathcal{F}_j} \epsilon(t - t_j^{(f)}) \quad (1.46)$$

with weights $w_{ij} = R c_{ij} / \tau_m$ and

$$\eta(s) = -(\vartheta - u_r) \exp\left(-\frac{s}{\tau_m}\right) \mathcal{H}(s) \quad (1.47)$$

$$\epsilon(s) = \int_0^\infty \exp\left(-\frac{s'}{\tau_m}\right) \alpha(s - s') ds'. \quad (1.48)$$

If $\alpha(s)$ is given by (1.41), then the integral on the right-hand side of (1.48) can be done and yields

$$\epsilon(s) = \frac{1}{1 - (\tau_s / \tau_m)} \left[\exp\left(-\frac{s}{\tau_m}\right) - \exp\left(-\frac{s}{\tau_s}\right) \right] \mathcal{H}(s). \quad (1.49)$$

Note that (1.46) is exactly the equation (1.12) or (1.30) of the Spike Response Model [except for the trivial replacement $\epsilon_{ij}(s) \rightarrow \epsilon(s)$]. We remark that the Spike Response Model is slightly more general than the integrate-and-fire model, because kernels in the Spike Response Model can be chosen quite arbitrarily whereas for the integrate-and-fire model they are fixed by (1.47) and (1.48).

In section 1.2.1 where we introduced a simple version of the Spike Response Model, we suggested a specific choice of response kernels, viz. (1.13) and (1.14). Except for a different normalization, these are exactly the kernels (1.47) and (1.49) that we have found now for the the integrate-and-fire model. Thus, we can view the basic model of section 1.2.1 as an alternative formulation of the integrate-and-fire model. Instead of defining the model by a differential equation (1.31), it is defined by its response kernels (1.47) and (1.49).

1.2.3.3 Models of Noise

Noise is omnipresent in biological systems due to nonzero temperature and finite numbers of molecules and ion channels. The effects of noise include failures in the synaptic transmission and different responses of a neuron to the same input current.

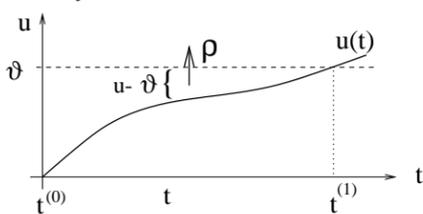
There are various ways to introduce noise into threshold-fire models. Here we briefly mention three possibilities: a noisy threshold, a noisy reset, or a noisy integration. In all cases, we are interested in the effect of the noise on the firing period. More precisely, let us assume that the last firing has occurred at a time $t^{(0)}$. In the simple version (1.21) of the Spike Response Model with short-term memory, the membrane potential for $t > t^{(0)}$ would be given by $u(t) = \eta(t - t^{(0)}) + h(t)$ and the next spike would occur at $t^{(1)}$ given by the threshold condition $u(t^{(1)}) = \vartheta$. In the presence of noise, we can no longer predict the exact time of firing. Instead we ask the following question. What is the probability that the next spike occurs between t and

$t + \Delta t$, given that the last spike occurred at $t^{(0)}$ and that the total input potential in the noiseless case is $h(t)$? For $\Delta t \rightarrow 0$, this defines the *probability density* for firing

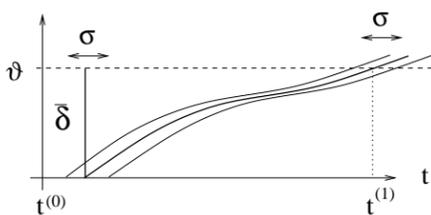
$$P_h(t | t^{(0)}) \quad (1.50)$$

which we would like to calculate for each of the three noise models. We can interpret $P_h(t | t^{(0)})$ as the distribution of interspike intervals in the presence of an input potential h . The lower index h is intended to remind the reader that the distribution depends on the time course of $h(t')$ for $t^{(0)} < t' < t$. We now discuss each of the three models of noise in turn.

A) Noisy threshold



B) Noisy reset



C) Noisy integration

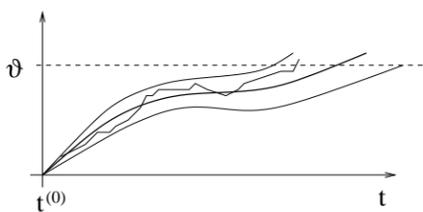


Figure 1.19. Noise in threshold neurons. a) Noisy threshold; a neuron may fire [with probability density $\rho(u - \vartheta)$] even though the membrane potential u has not yet reached the threshold ϑ . b) Noisy reset; a stochastic change in the absolute refractory period shifts the trajectory horizontally. c) Noisy integration; a stochastic contribution in the input *current* of an integrate-and-fire neuron causes the membrane potential to drift away from the reference trajectory.

Noisy threshold

In this first noise model, we assume the neuron can fire even though the formal threshold ϑ has not been reached yet. To do this consistently, we

introduce an ‘escape rate’ ρ which depends on the distance between the momentary value of the membrane potential and the threshold,

$$\rho = f(u - \vartheta). \quad (1.51)$$

In the mathematical literature, the quantity ρ would be called a ‘stochastic intensity’. The choice of the function f is arbitrary. A plausible assumption is an exponential dependence

$$\rho = \frac{1}{\tau_0} \exp[\beta(u - \vartheta)] \quad (1.52)$$

which can be motivated by the Arrhenius formula for chemical reaction rates. β and τ_0 are parameters. Note that the escape rate ρ is implicitly time-dependent, since the membrane potential $u(t) = \eta(t - t^{(0)}) + h(t)$ varies over time. In addition, we may also include an explicit time dependence, e.g., to account for a reduced spiking probability immediately after the spike at $t^{(0)}$.

Let us now calculate $P_h(t^{(1)} | t^{(0)})$, the probability density of having a spike at $t^{(1)}$ given that the last spike occurred at $t^{(0)}$, and in the presence of an input potential $h(t)$ for $t > t^{(0)}$. At each moment of time, the value $u(t)$ of the membrane potential determines the escape rate $\rho(t) = f[u(t) - \vartheta]$. In order to emit the next spike at $t^{(1)}$, the neuron has to ‘survive’ the interval $(t^{(0)}, t^{(1)})$ without firing and then fire at $t^{(1)}$. Given the escape rate $\rho(t)$, the probability of survival from $t^{(0)}$ to $t^{(1)}$ without a firing is

$$S_\rho(t^{(1)} | t^{(0)}) = \exp\left(-\int_{t^{(0)}}^{t^{(1)}} \rho(t) dt\right). \quad (1.53)$$

The probability density of firing at time $t^{(1)}$ is $\rho(t^{(1)})$, thus with (1.53) we have

$$P_h(t^{(1)} | t^{(0)}) = \rho(t^{(1)}) \exp\left(-\int_{t^{(0)}}^{t^{(1)}} \rho(t) dt\right) \quad (1.54)$$

which is the desired result. A more detailed derivation of (1.54) can be found in [Gerstner and van Hemmen, 1994].

Noisy Reset

In this noise model, firing is given by the exact threshold condition $u(t^{(f)}) = \vartheta$. Noise is included into the formulation of reset and refractoriness.

Let us consider the integrate-and-fire model with absolute refractoriness. The duration of refractoriness δ^{abs} is not fixed, but chosen stochastically from a distribution $p(\delta^{\text{abs}})$ with mean $\bar{\delta}$. Naturally we have to require that $p(x)$ vanishes for $x < 0$. In the formulation of the Spike Response Model, this procedure is equivalent to replacing the term $\eta(t - \hat{t})$ in (1.21) by $\eta(t - \hat{t} - \delta^{\text{abs}})$.

As with the preceding noise model, we are interested in calculating the firing density $P_h(t^{(1)} | t^{(0)})$ given some input potential $h(t)$. For the sake of simplicity, let us assume that the input is constant $h(t) = h_0$. The first spike

has occurred at $t^{(0)}$. The firing time $t^{(1)}$ of the next spike can now be found from the threshold condition

$$\vartheta = \eta(t^{(1)} - t^{(0)} - \delta^{\text{abs}}) + h_0. \quad (1.55)$$

In the absence of noise and for a refractory period $\bar{\delta}$ the next firing would occur at $t^{(1)} = t^{(0)} + T$ where T is the interspike interval. If, due to the noise, the value of the refractory period is $\delta^{\text{abs}} \neq \bar{\delta}$, then the interval is $t^{(1)} - t^{(0)} = T + \delta^{\text{abs}} - \bar{\delta}$. The firing *density* in the presence of noise is therefore

$$P_h(t^{(1)} | t^{(0)}) = p(t^{(1)} - t^{(0)} - T + \bar{\delta}) \quad (1.56)$$

where $p(\cdot)$ is the distribution of the refractory period introduced above. Graphically, the result (1.56) is easy to understand. A change in the absolute refractory period shifts the trajectory horizontally. A stochastic component in the refractory period generates a stochastic shift of the firing time. Although a stochastic component in the refractoriness probably is not very realistic, it is a convenient way to introduce noise in a system. As shown in (1.56), noise can be treated analytically without causing too much problems.

Noisy integration

The final and most popular way of introducing noise into an integrate-and-fire model is by adding on the right-hand side of (1.32) a stochastic noise current I_{noise}

$$\tau_m \frac{du}{dt} = -u + RI(t) + RI_{\text{noise}} \quad (1.57)$$

with vanishing mean and finite variance. This choice of a noise term could be motivated either by spontaneous openings of ion channels, or else by stochastic arrival of excitatory and inhibitory inputs on dendrites of cortical neurons.

The noise causes the actual membrane trajectory to drift away from the noiseless reference trajectory. To get the distribution $P_h(t^{(1)} | t^{(0)})$ we have to solve the first passage time problem of (1.57) with initial value u_r and absorbing boundary condition at ϑ . Although (1.57) looks simple, it turns out that the first passage time problem for arbitrary input current $I(t)$ is rather difficult to solve analytically. Discussions of the first passage time problem for constant input $I(t) = I_0$ can be found in many textbooks; see, e.g., [Tuckwell, 1988].

1.2.4 Conductance-Based Models

In this section we briefly discuss a class of rather detailed neuron models which are known as conductance based models. At the origin of these models are the equations of Hodgkin and Huxley which we describe first. We then discuss the relation of the Hodgkin-Huxley model to the Spike Response Model introduced in the previous subsection. Finally, we mention the detailed compartmental models which focus on the spatial structure of the neuron as well as on the dynamics of various ion currents.

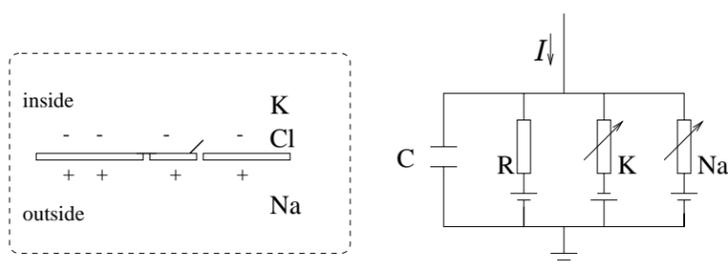


Figure 1.20. Schematic diagram for the Hodgkin-Huxley model.

1.2.4.1 Hodgkin-Huxley Model

The classic description of neuronal spiking dates back to Hodgkin and Huxley [1952] who summarized their extensive experimental studies on the giant axon of the squid with four differential equations. The first describes the conservation of electric charge on a piece of membrane of capacitance C under the influence of some charging currents

$$C \frac{du}{dt} = - \sum_k I_k + I(t) \quad (1.58)$$

where u is the voltage, $I(t)$ is an external driving current and $\sum_k I_k$ is the sum of the ionic currents through the cell membrane. In the case of the Hodgkin-Huxley model, there are three types of ion current indexed respectively by Na, K, and L,

$$\sum_k I_k = g_{\text{Na}} m^3 h (u - V_{\text{Na}}) + g_{\text{K}} n^4 (u - V_{\text{K}}) + g_{\text{L}} (u - V_{\text{L}}). \quad (1.59)$$

The parameters g are conductances and m, n, h are additional variables to be discussed below. Further parameters are the constants $V_{\text{Na}}, V_{\text{K}},$ and V_{L} . They are called reversal potentials since the direction of a current I_k changes when u crosses V_k .

To understand the two equations (1.58) and (1.59) consider the diagram in Figure 1.20. The semipermeable cell membrane separates the interior of the cell from the extracellular liquid. Due to the membrane's selective permeability and also because of active ion transport through the cell membrane, the ion concentrations inside and outside the cell are quite different. As a result, in the absence of external input the interior of the cell has a slightly negative potential with respect to the outside. The cell membrane acts like a capacitor which has been charged by a battery.

If an input current $I(t)$ is injected into the cell, it both charges the capacitor, and leaks through the channels in the cell membrane. This is the essence of (1.58). In the Hodgkin-Huxley model, three types of ion channels are accounted for. There is a sodium channel with subscript Na, a potassium channel with subscript K and an unspecific leakage channel with lower index L. The leakage channel is described by a voltage-independent conductance g_L ; the conductances of the other other ion channels are voltage dependent. If the channels are fully open, they transmit currents with a

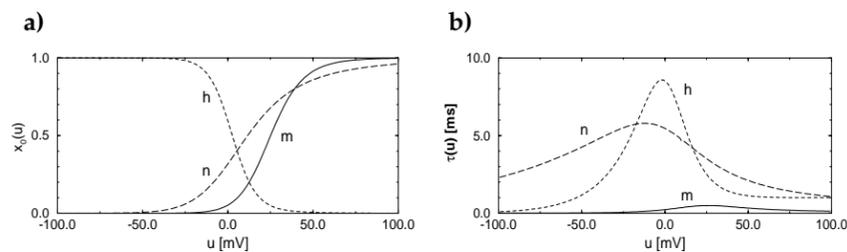


Figure 1.21. Equilibrium function (a) and time constant (b) for the three variables m , n , h in the Hodgkin-Huxley model.

maximum conductance g_{Na} or g_{K} , respectively. Normally, however, the channels are partially blocked. The removal of the block is voltage dependent and is described by the additional variables m , n , and h . The combined action of m and h controls the Na channels. The K gates are controlled by n .

The three variables m , n , and h evolve according to the differential equations

$$\begin{aligned} \dot{m} &= \alpha_m(u)(1-m) - \beta_m(u)m \\ \dot{n} &= \alpha_n(u)(1-n) - \beta_n(u)n \\ \dot{h} &= \alpha_h(u)(1-h) - \beta_h(u)h \end{aligned} \quad (1.60)$$

with $\dot{m} = dm/dt$, and so on. The α and β are empirical functions of u that have been adjusted to fit the data of the giant axon of the squid. Reversal potentials and conductances are also empirical parameters. In the appropriate units (of mS/cm^2), the factors are $g_{\text{Na}} = 120$, $g_{\text{K}} = 36$, and $g_{\text{L}} = 0.3$. The reversal potentials are $V_{\text{Na}} = 115$ mV, $V_{\text{K}} = -12$ mV, and $V_{\text{L}} = 10.6$ mV and the membrane capacity is $C = 1$ $\mu\text{F}/\text{cm}^2$. Eqs. (1.58) - (1.60) define the Hodgkin-Huxley model.

To get an idea of how the model works it is more convenient to write (1.60) in the form

$$\dot{x} = -\frac{1}{\tau(u)}[x - x_0(u)] \quad (1.61)$$

where x stands for m , n , or h . If we force the voltage to a constant value u , the variable x would approach the value $x_0(u)$ with time constant $\tau(u)$. The asymptotic value $x_0(u)$ and the time constant $\tau(u)$ are given by the transformation $x_0(u) = \alpha_x(u)/[\alpha_x(u) + \beta_x(u)]$ and $\tau(u) = [\alpha_x(u) + \beta_x(u)]^{-1}$. Using the parameters given by Hodgkin and Huxley [1952], we have plotted the functions $x_0(u)$ and $\tau(u)$ in Figure 1.21. The function $x_0(u)$ has a sigmoidal shape with maximum slope at some 'threshold' θ_x ; the time constant τ is significant only in a limited range around θ_x . Note that m and n increase with u whereas h decreases. Thus, if some external input causes the membrane voltage to rise, the ion conductance of sodium (Na) increases due to increasing m and sodium flows into the cell. This raises the membrane potential further and further and an action potential is initiated. At high values of u the sodium conductance is shut off due to the factor h . Also potassium (K) outflow sets in which lowers the potential. Due to the longer time constant $\tau_n(u)$, the potassium concentration reaches its

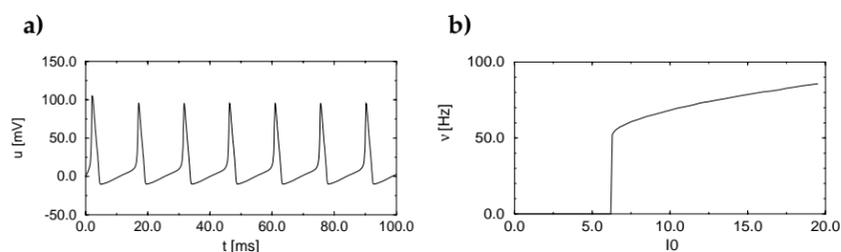


Figure 1.22. Spike train for constant input current (a) and gain function (b) of the Hodgkin-Huxley model.

equilibrium potential only slowly. The overall effect is a short action potential followed by a negative overshoot which decays slowly back to zero as shown in Figure 1.22a. Numerical integration of Eqs. (1.58)-(1.60) shows that a constant input $I(t) = I_0$ larger than a critical value I_θ results in a regular spike train. If the number of spikes during a large interval T is counted and divided by T , a firing rate results. The firing rate as a function of the input I_0 is called the *gain function* of the Hodgkin-Huxley model. It is plotted in Figure 1.22b.

Using the above equations and an appropriate set of parameters, Hodgkin and Huxley were able to describe an enormous amount of data from experiments on the giant axon of the squid. Due to its success in this special system, there have subsequently been several attempts to generalize the model in order to describe other experimental situations as well (for a review see, e.g., [Jack et al., 1975]).

Whereas the model had originally been designed to describe the form and temporal change of an action potential during *axonal* transmission, a set of equations completely analogous to Eqs. (1.58) to (1.60) has been also used to describe spike generation at the *soma* of the neuron [Bernander et al., 1991; Bush and Douglas, 1991; Ekeberg et al., 1991; Rapp et al., 1992; Traub et al., 1991; Wilson et al., 1989; Yamada et al., 1989]. The main difference is that additional ion channels have to be included, in particular those that account for Ca^{2+} and the slow components of the potassium current. For each type of ion channel i , a current $I_i = g_i x_i^{n_i} (u - V_i)$ is added. Here x is yet another variable with dynamics (1.61). The conductance parameters g_i , the exponents n_i , as well as the functions $x_0(u)$ and $\tau(u)$ are adjusted to fit experimental data. This approach leads to detailed compartmental models briefly discussed at the end of this section.

Signal transmission and integration on dendrites have traditionally been described by passive electric processes, most prominently in the dendritic cable theory of Rall; see e.g. [Rall, 1964]. Such an approach can even account for some nonlinear dendritic phenomena [Abbott, 1991; Bernander et al., 1991; Rapp et al., 1992]. A different approach to dendritic integration is based again on Hodgkin-Huxley type equations. The only change with respect to the theory of *axonal* signal transmission is that a different set of parameters for the conductivity of various channels are used on the dendrites. If a few spatial compartments of the neuron are put together, a set of more than 20 coupled nonlinear differential equations results (see

e.g. [Traub et al., 1991]). Numerical solution of the system of equations shows good agreement with experiments. This is an important indication that a Hodgkin-Huxley type analysis is a useful tool in understanding the properties of single neurons. It is however obvious that such an approach is too detailed, if we want to describe large network of neurons.

1.2.4.2 Relation to the Spike Response Model

The system of equations proposed by Hodgkin and Huxley is rather complicated. It consists of four coupled nonlinear differential equations and as such is difficult to analyze mathematically. For this reason, several simplifications of the Hodgkin-Huxley equations have been proposed. The most common reduces the set of four differential equations to a system of two equations [FitzHugh, 1961; Nagumo, 1962; Rinzel and Ermentrout, 1989; Abbott and Kepler, 1990]. Two important approximations are made. First, the m dynamics which has a faster time course than the other variables (see the plot for τ_m in Figure 1.21b) is considered to be instantaneous, so that m can be replaced by its equilibrium value $m_0(u)$. Second, the equations for n and h which have according to Figure 1.21b roughly the same time constants are replaced by a single effective variable. [Rinzel and Ermentrout, 1989] and [Abbott and Kepler, 1990] have shown how to make such a reduction systematically. The resulting two-dimensional model is often called the Morris LeCar model or the FitzHugh-Nagumo Model. More generally, the two-dimensional set of equations is also called a Bonhoeffer/Van der Pol oscillator. The advantage of a two-dimensional set of equations is that it allows a systematic phase plane analysis. For a review of the methods and results see the article of [Rinzel and Ermentrout, 1989] in the review collection [Koch and Segev, 1989]. For a further reduction of the two-dimensional model to an integrate-and-fire model, see [Abbott and Kepler, 1990].

In this section, we will follow a somewhat different approach. We would like to relate the Hodgkin-Huxley equations to our generic threshold model, the Spike Response Model [Kistler et al., 1997]. To do so, we have to determine the following three terms which appear in the equations of the Spike Response Model (1.28) and (1.30): (i) the kernel ϵ which describes the response to incoming pulses, (ii) the kernel η which describes the response to spike emission, and (iii) the value of the threshold ϑ .

We start with the kernel ϵ . In the absence of input the membrane potential u is at some resting value u_{rest} . To find the kernel ϵ we perform a simulation with a short square current pulse as the input

$$I(t) = \frac{q}{\Delta} \quad \text{for } 0 < t < \Delta \quad (1.62)$$

and zero otherwise. Here q is some small unit charge and $\Delta \ll 1$ ms is the duration of the pulse. (Formally, we consider the limit $\Delta \rightarrow 0$.) The voltage response of the Hodgkin-Huxley model to this sub-threshold current pulse defines the kernel ϵ ,

$$q \epsilon(t) = u(t) - u_{\text{rest}}. \quad (1.63)$$

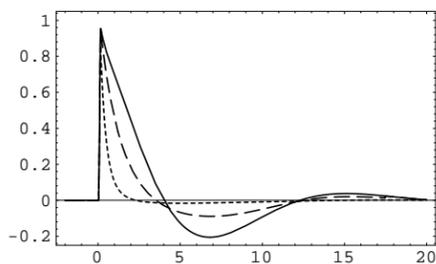


Figure 1.23. The voltage response of the Hodgkin-Huxley model to a short sub-threshold current pulse defines the kernel ϵ (solid line). If the same input current pulse occurs a few milliseconds after an output action potential, the duration of the response is reduced due to refractoriness (dashed line, input spike $\Delta t = 10.5$ ms after the output spike; dotted line $\Delta t = 6.5$ ms) x -axis, time in ms; y -axis voltage in mV; taken from [Kistler et al., 1997].

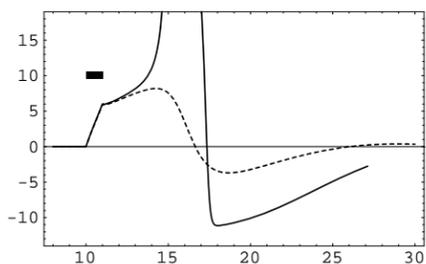


Figure 1.24. Threshold behavior of the Hodgkin-Huxley model. An input current pulse of 1 ms duration (thick bar) can evoke either an action potential (solid line) with amplitude 100 mV (out of scale) or a small response of about 10 mV (dashed) depending on the precise value of the input current. x -axis, time in ms; y -axis potential in mV. Solid line, input current $6.9 \mu\text{Acm}^{-2}$; dashed line, input current $7.0 \mu\text{Acm}^{-2}$; taken from [Kistler et al., 1997].

Here $t > 0$ is the time since the initiation of the pulse. The result is shown in Figure 1.23. Since the input current pulse delivers a unit charge during a very short amount of time $\Delta < 0.1$ ms, the ϵ -kernel jumps almost instantaneously at time $t = 0$ to a value of 1 mV. For synaptic input currents of finite width, the ϵ -kernel would have a finite rise time.

The kernel η is found by a similar procedure. We take a square current pulse as in (1.62) but with a charge q large enough to evoke a spike. The principle is indicated in Fig 1.24. We consider a series of current pulses of increasing q but the same duration of 1 ms. At a critical value of q the responses show an abrupt change from a response amplitude of about 10 mV to an amplitude of nearly 100 mV. If q is increased even further, the form of the pulse remains nearly the same. The kernel η allows us to describe the standard form of the spike and the spike afterpotential in a convenient manner. In order to define the kernel η , we set

$$\eta(s) = u(s) - u_{\text{rest}} \quad (1.64)$$

where $u(s)$ is the voltage trajectory after the supra-threshold current pulse. The kernel $\eta(s)$ is shown in Figure 1.25. The time $s = 0$ denotes the moment of spike initiation.

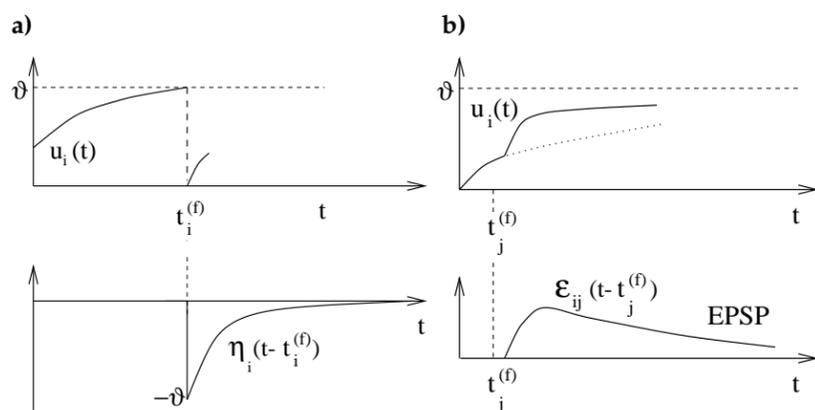


Figure 1.25. *a)* The state variable $u(t)$ reaches the threshold ϑ at time $t_i^{(f)}$. Immediately afterwards $u(t)$ is reset to zero. The reset is performed by adding a kernel $\eta_i(t - t_i^{(f)})$. The function $\eta_i(s)$ takes care of refractoriness after a spike emitted at $s = 0$. *b)* The state variable $u(t)$ changes after a presynaptic spike has occurred at $t_j^{(f)}$. The kernels ϵ_{ij} describes the response of u_i to a presynaptic spike at $s = 0$. The figure shows an excitatory postsynaptic potential (EPSP). In case of an inhibitory postsynaptic potential (IPSP) the response would be negative. Note that the response occurs only after some delay.

The third term to be determined is the threshold ϑ . Even though Figure 1.24 suggests, that the Hodgkin-Huxley equation exhibits some type of threshold behavior, the threshold is not well-defined [Koch et al., 1995] and it is fairly difficult to estimate a voltage threshold directly from a single series of simulations. We therefore take the threshold as a parameter which will be adjusted by a procedure discussed below.

In Spike Response Model (SRM) approximation, we replace the full Hodgkin-Huxley model by the equation

$$u^{SRM}(t) = \eta(t - \hat{t}) + \int_0^\infty \epsilon(s) I(t - s) ds \quad (1.65)$$

where \hat{t} is the most recent firing time and the kernels ϵ and η are defined by (1.63) and (1.64), respectively. The next output spike occurs when the threshold ϑ is approached from below, i.e., $u^{SRM}(t) = \vartheta$ and $\frac{d}{dt}u^{SRM}(t) > 0$.

To test the quality of the approximation (1.65) we compare the performance of the reduced model (1.65) with that of the full Hodgkin-Huxley model (1.58) - (1.60). Since it is our general assumption that timing is important, we study the difficult case of a time-dependent input current $I(t)$. The same current is applied to both models and we compare the firing times of the full Hodgkin-Huxley model with the spike times generated by (1.65). Whenever the firing times differ by more than 2 ms, we record an error in the performance of the SRM. Details may be found in [Kistler et al., 1997].

A time-dependent input current was generated stochastically by the following procedure. Every 2 ms a new target value for the input was drawn from a Gaussian distribution with zero mean and variance $\sigma = 3\mu\text{A}/\text{cm}^2$.

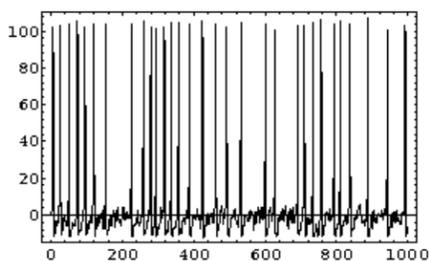


Figure 1.26. Spike train of the Hodgkin-Huxley model driven by a time dependent input current. The action potentials occur irregularly. x -axis, time in ms; y -axis, voltage in mV. Taken from [Kistler et al., 1997].

To get a continuous input current, a linear interpolation was used between the target values. The resulting time-dependent input current was then applied to the Hodgkin-Huxley model (1.58); a sample of the response is shown in Figure 1.26. The spike train looks irregular with a broad distribution of interspike intervals as is commonly found for cortical neurons.

The same input was then applied to the Spike Response Model (1.65). The value of the threshold ϑ was optimized so that the total number of spikes generated by (1.65) was roughly that of the Hodgkin-Huxley model; see [Kistler et al., 1997] for details. The spike times were then compared. About 70 per cent of the spikes of the Spike Response Model occurred within 2 ms of the action potentials of the Hodgkin-Huxley model. Moreover, the voltage trace is generally well approximated by (1.65) except just after a spike; see Figure 1.27. The inset in the lower left of Figure 1.27 shows that in an interval where there is no spike, the voltage of the Spike Response Model (dashed) closely follows that of the Hodgkin-Huxley model (solid line). Immediately after a spike, the approximation is poor (inset lower right).

Is there a simple way to improve the approximation? Obviously the problem is that (1.65) provides an incomplete description of refractoriness. In (1.65) refractoriness appears only in form of the spike afterpotential described by the kernel η . But refractoriness also has some other effects. During and immediately after a spike, the response to incoming current pulses is reduced. This is shown in Figure 1.23 by the dashed lines. To take this effect into account, we allow the kernel ϵ to depend on the time since the last output spike and replace (1.65) by

$$u^{SRM}(t) = \eta(t - \hat{t}) + \int_0^\infty \epsilon(t - \hat{t}, s) I(t - s) ds. \quad (1.66)$$

As before, \hat{t} denotes the firing time of the most recent output spike. In order to find the kernel $\epsilon(t - \hat{t}, s)$ we determine numerically the response to a short input current pulse (1.62) under the condition that the last output spike was at \hat{t} . The original kernel ϵ in (1.65) corresponds to $\epsilon(\infty, s)$, that is, to the case in which the most recent output spike was long ago.

To test the performance of the improved approximation (1.66), we repeat the simulations with the same stochastic input current $I(t)$ as before and

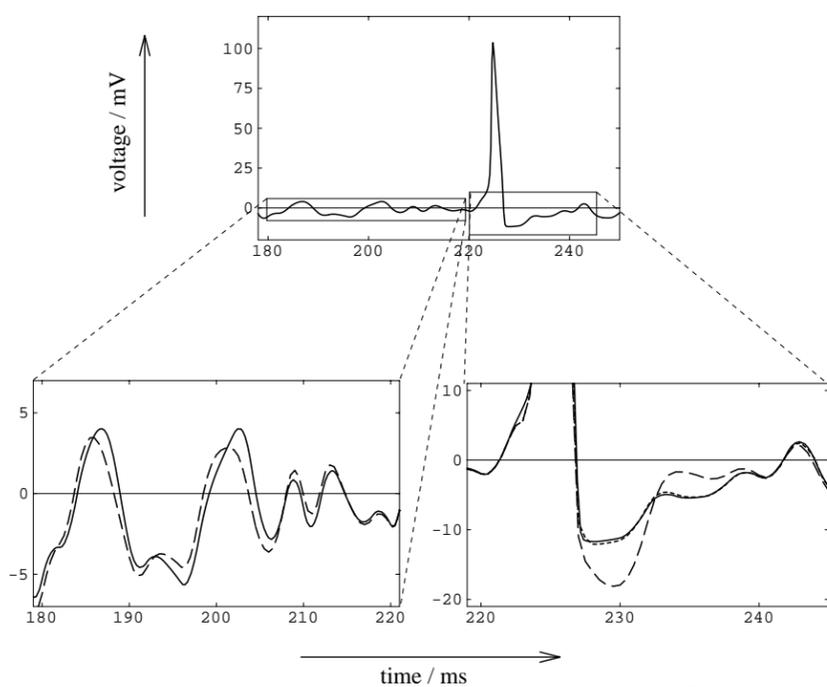


Figure 1.27. A segment of the spike train of the previous figure. The inset in the lower left corner shows the voltage of the Hodgkin-Huxley model (solid) together with the approximation of the Spike Response Model (dashed) during a period where no spike occurs. The inset on the lower right shows the situation during and after a spike. The approximation by the dashed line is not perfect, but the improved approximation (dotted) defined by (1.66) yields a very good fit; taken from [Kistler et al., 1997].

readjust the value of ϑ . The approximation (1.66) increased the percentage of firing times that coincided with the Hodgkin-Huxley action potentials from 70 to 90 percent. This shows that threshold-fire models like the (improved) Spike Response Model can indeed provide a simple and yet useful description of neuronal dynamics.

1.2.4.3 Compartmental Models

The neuron models described so far did not include any spatial structure. A neuron has been considered to be a point-like element and the main focus has been on the process of spike generation. Compartmental models provide a more detailed description of neuronal dynamics by explicitly taking into account the spatial structure of the dendritic tree and by modeling the synaptic transmission at a greater level of detail. Additionally, other ion currents beyond the Na and K currents incorporated in the Hodgkin-Huxley model are included in the description. Of particular interest have been slow currents like Ca^{2+} and Ca-mediated K currents which are related to neuronal adaptation [Yamada et al., 1989]. The basic idea of a compartmental model is indicated in Figure 1.28. Each small segment of the dendritic tree is characterized by ionic and electrical properties, de-

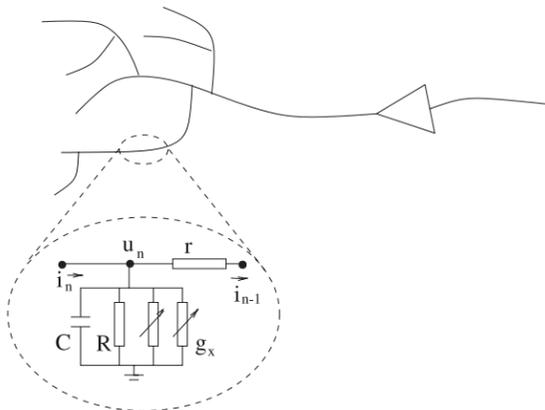


Figure 1.28. In a compartmental neuron model, each small segment of the dendritic tree is modelled by an equivalent electrical circuit. Many more ionic conductances g_x could be present. For the sake of simplicity, the batteries representing the reversal potentials have been suppressed in the drawing.

scribed by a capacitance C and a set of ionic conductances g_x with reversal potential V_x . Neighboring compartments are connected by a longitudinal resistance r . For details, we refer the reader to the vast literature on compartmental models [Bernander et al., 1991; Busch et al., 1991; Ekeberg et al., 1991; Rapp et al., 1992; Traub et al., 1991; Wilson et al., 1989; Yamada et al., 1989], in particular to the reviews in [Koch and Segev, 1989] and [Bower and Beeman, 1995].

Here we discuss a simplified linear model of dendritic structure as introduced by Rall [Rall, 1964, 1989]. The dendritic tree is described by a chain of compartments. In principle, we can allow for branching points, but for the sake of simplicity we focus on a single chain. Each compartment consists of a capacitor C in parallel with a resistor R which describe the linear properties of a patch of membrane. As we have seen above, a more detailed model would also contain non-linear conductances for various ionic currents – these are neglected in our simplified linear model. Each compartment is connected via the axial resistor r to adjacent compartments. The compartment with index $n = 1$ will be referred to as the somatic compartment.

In a chain of N neurons, conservation of current at each compartment yields the equations

$$\frac{u_{n+1} - u_n}{r} - \frac{u_n - u_{n-1}}{r} = C \frac{du_n}{dt} + \frac{u_n}{R} \quad (1.67)$$

for the interior compartments $2 \leq n \leq N - 1$ and similar equations for $n = 1$ and $n = N$. The terms on the left-hand side are the longitudinal currents i_n and i_{n-1} leading into and out of the compartment; see Fig 1.29. The difference between the two currents branches off into compartment n and either leaks through the membrane resistor R or charges the capacitor C . We introduce the membrane time constant $\tau_m = RC$ and reorder the terms yielding

$$\tau_m \frac{du_n}{dt} = - \left(1 + 2 \frac{R}{r} \right) u_n + \frac{R}{r} (u_{n+1} + u_{n-1}) \quad (1.68)$$

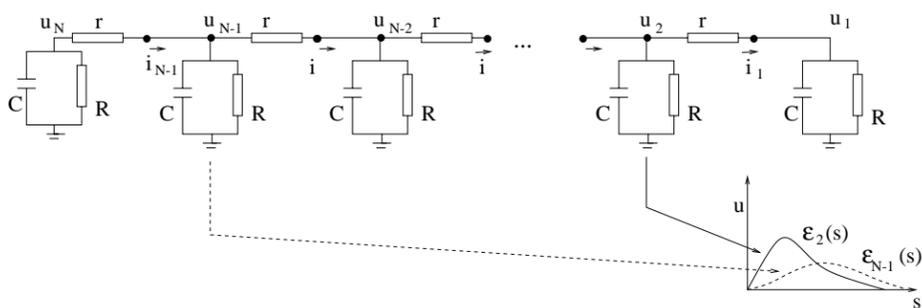


Figure 1.29. A linear compartmental model of the dendritic tree. A current input at compartment 2 at time $s = 0$ gives a strong and quickly rising voltage response ϵ_2 at the somatic compartment 1 (solid line); the same current injected at compartment $N - 1$ gives a weaker and broader response at the soma (ϵ_{N-1} , dashed line). Schematic drawing.

for the interior compartments $2 \leq n \leq N - 1$. At the two ends we have

$$\tau_m \frac{du_1}{dt} = - \left(1 + \frac{R}{r} \right) u_1 + \frac{R}{r} u_2 \quad (1.69)$$

$$\tau_m \frac{du_N}{dt} = - \left(1 + \frac{R}{r} \right) u_N + \frac{R}{r} u_{N-1} \quad (1.70)$$

Since (1.68) - (1.70) are linear equations, they can be solved exactly [Rall, 1989; Tuckwell, 1988; Abbott et al., 1991; Bressloff and Taylor, 1993]. For example we can ask the following important question. What is the voltage response of compartment number $n = 1$ (the soma) to a short input current pulse into compartment number n (somewhere on the dendritic tree)? If the input is close to the soma, e.g., $n = 2$, then the response is a fast rising function with a well pronounced maximum; see Figure 1.29. If the input is far away from the soma, the voltage at the soma rises more slowly. Thus the main effect of a linear dendritic tree is that the shape of the postsynaptic potential (as measured at the soma) depends on the location of the input. We can account for this effect in the generic Spike Response Model by making the form of the response kernel ϵ depend on the label of the compartment where the synapse is located.

The situation becomes more complicated if the synaptic currents are modeled at a greater level of detail. First, synaptic currents are not infinitely short but are pulses of finite width. If the dendrite were linear and described by (1.68) - (1.70) then the only difference would be that the postsynaptic potential $\epsilon(s)$ is somewhat broader. If the form and amplitude of the synaptic current were always identical, then we could simply add the individual postsynaptic potentials to get the total effect of the input as it is seen at the soma.

Reality is somewhat more difficult, however, since the amplitude of the synaptic input current itself depends on the voltage at the compartment where the synapse is located. In detailed models, each presynaptic action potential evokes a change in the synaptic *conductance* with standard time course $g(t - t^{(f)})$ where $t^{(f)}$ is the arrival time of the presynaptic pulse. The

synaptic input current is modeled as

$$I(t - t^{(f)}) = g(t - t^{(f)}) [u_{syn}(t) - u_{rev}] \quad (1.71)$$

where $u_{syn}(t)$ is the voltage at the location of the synapse and u_{rev} is called the reversal potential of the synapse.

The level of the reversal potential depends on the type of synapse. For excitatory synapses, u_{rev} is larger than the resting potential. The synaptic current then shows saturation. The higher the voltage u_{syn} at the synaptic compartment, the smaller the amplitude of the input current. A smaller input current, however, evokes a postsynaptic potential of reduced amplitude. The total postsynaptic potential is therefore not simply the sum of independent contributions. Nevertheless, since the reversal potential of excitatory synapses is usually significantly above the firing threshold, the factor $[u_{syn} - u_{rev}]$ is nearly constant and saturation can be neglected.

For inhibitory synapses, the reversal potential is close to the resting potential and saturation plays an important role. This is sometimes described as the 'shunting' phenomenon of inhibition. An action potential arriving at an inhibitory synapse pulls the membrane potential towards the reversal potential u_{rev} which is close to u_{rest} . Thus, if the neuron is at rest, inhibitory input hardly has any effect. If the membrane potential is instead considerably above the resting potential, then the same input has a strong inhibitory effect.

1.2.5 Rate Models

Before we end the chapter, we want to mention the formal neuron model that is most widely used for the analysis of learning and memory in artificial neural networks - the sigmoidal unit. In standard neural network theory, neural activity is described in terms of rates. The rate ν_i of neuron i is an analog variable which depends nonlinearly upon the excitation u_i of the neuron,

$$\nu_i = g(u_i) \quad (1.72)$$

where $g(\cdot)$ is usually taken as a sigmoidal function with $g(u) \rightarrow 0$ for $u \rightarrow -\infty$ and $g(u) \rightarrow 1$ for $u \rightarrow \infty$; see Figure 1.30. The excitation is given by a linear sum over all input connections

$$u_i = \sum_{j \in \Gamma_i} w_{ij} \nu_j \quad (1.73)$$

where ν_j is the output rate of a presynaptic neuron j . The sum runs over all neurons which send signals to neuron i . The parameter w_{ij} , called synaptic efficacy, is the weight attributed to the connection from j to i .

Equations (1.72) and (1.73) can be summarized in a single equation

$$\nu_i = g \left(\sum_{j \in \Gamma_i} w_{ij} \nu_j \right) \quad (1.74)$$

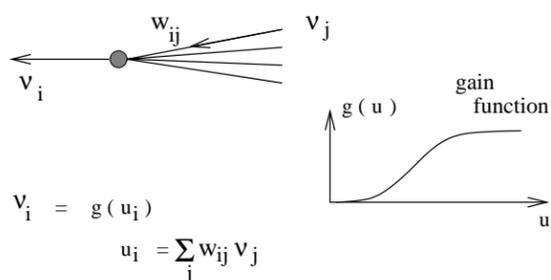


Figure 1.30. The rate model used in standard neural network theory.

which is the starting point of standard neural network theory [Hertz et al., 1991; Rojas, 1996; Haykin, 1994; Bishop, 1995].

Equation (1.74) is a static equation. It applies to situations where a stationary input (a set of firing rates ν_j) is mapped to a stationary output (the rate ν_i). Is it possible to make the equation time-dependent? A straightforward way to introduce dynamics into the rate equation (1.74) is to replace (1.74) by a differential equation [Cowan, 1968]

$$\tau \frac{d\nu_i}{dt} = -\nu_i + g \left(\sum_{j \in \Gamma_i} w_{ij} \nu_j \right) \quad (1.75)$$

with some time constant τ . For stationary input and output, the left-hand side of (1.75) vanishes and (1.75) reduces to (1.74). In other words, the fixed-point solutions of (1.75) are given by (1.74).

Equation (1.75) provides a convenient way to introduce some time dependence in the rate model (1.74), but can it be considered a realistic description of neuronal activity? As we have discussed in Section 1.1.2, an analog variable defined by a spike count measure requires a long temporal averaging window. It can therefore be used only if the input and the output change on a slow time scale. Considering the fact that, for example, the visual input changes due to saccades every 200-500 ms, a slowly changing input can not always be assumed.

It has therefore been argued that the rate equation (1.75) refers to a population average rather than to a temporal average. To make this clear in our notation, we rewrite (1.75) as

$$\tau \frac{dA_k}{dt} = -A_k + g \left(\sum_l J_{kl} A_l \right) \quad (1.76)$$

where A_k is the activity of a population k and the sum in the brackets runs over all other populations l which send signals to k . Again we may ask the question, whether (1.76) can be considered a realistic description of the population dynamics. More specifically, what determines the time constant τ which limits the response time of the system? Is it given by the membrane time constant of a neuron? Is τ really a constant or does it depend on the input or activity of the system?

We will see later in the Chapter 10 on population dynamics that the population activity of a group of spiking model neurons can react instantaneously

to changes in the input. This suggests that the ‘time constant’ τ in (1.76) is, at least in some cases, extremely short. The theory of population dynamics in Chapter 10 does not make use of the differential equation (1.76), but works in a more general setting.

We are now at the end of our review of neuron models. We will return to conductance based models and compartmental models in the context of hardware implementations in Chapters 5 and 6. Most of the other chapters will take one of the simple threshold-fire neurons as the computational unit of the neural network using either the integrate-and-fire neuron or the Spike Response Model. As we have seen, the two formulations are roughly equivalent. In the following chapter, we will study the computational capabilities of simple networks of spiking neurons. For this purpose a formulation with response kernels turns out to be most appropriate.

1.3 Conclusions

How do neurons encode information in a sequence of pulses? Even though it is a fundamental question the problem of neural coding is still not fully resolved. In this chapter, we have reviewed three concepts of rate codes, viz. spike count over some time window, spike density in a histogram, and population activity in an ensemble of neurons. All three concepts have been successfully used in experimental data analysis, but may be criticized on principal grounds. A constructive critic of rate codes may come from a presentation of some candidate pulse codes, if their usefulness in terms of computational power or ease of implementation can be shown. This endeavor is the program of the book.

We have seen that it is often difficult to draw a clear border line between pulse and rate codes. Whatever the name of the code, it should offer a neural system the possibility to react quickly to changes in the input. This seems to be a minimum requirement if fast behavioral reaction times are to be accounted for.

If pulse coding is relevant, neural network theory must be based on spiking neurons. Several pulsed neuron models have been reviewed in this chapter. Conductance based neuron models operate on a detailed level of description. If we want to investigate nonlinear interactions on the dendrite, conductance based neuron models are a suitable level of description.

The classic example of a conductance based neuron model is the Hodgkin-Huxley model. We have seen that the Hodgkin-Huxley model can be approximated by the Spike Response Model, if we use a suitable set of response kernels. This suggests that models in the class of threshold-fire models probably capture some important features of spiking neurons. Most of the theoretical investigations in chapters 10-14 will make use of this model class and use either the integrate-and-fire formulation or the formalism of the Spike Response Model.

References

- [Abbott, 1991] Abbott, L. F. (1991). Realistic synaptic inputs for model neural networks. *Network*, 2:245–258.
- [Abbott et al., 1991] Abbott, L. F., Fahri, E., and Gutmann, S. (1991). The path integral for dendritic trees. *Biol. Cybern.*, 66:49–60.
- [Abbott and Kepler, 1990] Abbott, L. F., and Kepler, T. B. (1990). Model neurons: from Hodgkin-Huxley to Hopfield. *Statistical Mechanics of Neural Networks*. L. Garrido, ed., Springer, Berlin.
- [Abeles, 1991] Abeles, M. (1991). *Corticonics*. Cambridge University Press, Cambridge.
- [Abeles, 1994] Abeles, M. (1994). Firing rates and well-timed events. *Models of Neural Networks 2*, E. Domany, K. Schulten, and J. L. van Hemmen, eds., Springer, New York, chapter 3, 121–140.
- [Abeles et al., 1993] Abeles, M., Bergman H., Margalit E., and Vaadia, E. (1993). Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J. of Neurophysiology*, 70:1629–1638.
- [Adrian, 1926] Adrian, E. D. (1926). The impulses produced by sensory nerve endings. *J. Physiol. (London)*, 61:49–72.
- [Adrian, 1928] Adrian, E. D. (1928). *The Basis of Sensation*. W. W. Norton, New York.
- [Aertsen, 1993] Aertsen, A., and Arndt, M. (1993). Response synchronization in the visual cortex. *Current Opinion in Neurobiology*, 3:586–594.
- [Bernander et al., 1991] Bernander, Ö., Douglas, R. J., Martin, K. A. C., and Koch C. (1991). Synaptic background activity influences spatiotemporal integration in single pyramidal cells. *Proc. Natl. Acad. Sci. USA*, 88:11569–11573.
- [Bialek et al., 1991] Bialek, W., Rieke, F., de Ruyter van Stevenick, R. R., and Warland, D. (1991). Reading a neural code. *Science*, 252:1854–1857.
- [Bialek and Rieke, 1992] Bialek, W., and Rieke, F. (1992). Reliability and information transmission in spiking neurons. *Trends in Neurosciences*, 15(11):428–433.
- [Bishop, 1995] Bishop, C. M. (1995). *Neural Networks for Pattern Recognition*. Clarendon Press, Oxford.
- [Bower, 1995] Bower, J. M., and Beeman, D. (1995). *The Book of Genesis*. Springer, New York.
- [Bressloff and Taylor, 1993] Bressloff, P. C., and Taylor, J. G. (1993). Compartmental-model response function for dendritic trees. *Biol. Cybern.*, 70:199–207.
- [Bush and Douglas, 1991] Bush, R. C., and Douglas, R. J. (1991). Synchronization of bursting action potential discharge in a model network of neocortical neurons. *Neural Computation*, 3:19–30.
- [Cowan, 1968] Cowan, J. D. (1968). *Statistical Mechanics of Nervous Nets*. Proc. 1967 NATO conference on Neural Networks, Springer, Berlin.

- [de Ruyter van Stevenick and Bialek, 1988] de Ruyter van Stevenick, R. R., and Bialek, W. (1988). Real-time performance of a movement-sensitive neuron in the blowfly visual system: coding and information transfer in short spike sequences. *Proc. R. Soc. B*, 234:379–414.
- [deCharms and Merzenich, 1996] deCharms, R. C., and Merzenich, M. M. (1996). Primary cortical representation of sounds by the coordination of action-potential timing. *Nature*, 381:610–613.
- [DeAngelis et al., 1995] DeAngelis, G. C., Ohzaw I., and Freeman, R. D. (1995). Receptive-field dynamics in the central visual pathways. *Trends in Neurosci.*, 18:451–458.
- [Eckhorn et al., 1988] Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., and Reitboeck, H. J. (1988). Coherent oscillations: A mechanism of feature linking in the visual cortex? *Biol. Cybern.*, 60:121–130.
- [Eckhorn et al., 1993] Eckhorn, R., Krause, F., and Nelson, J. L. (1993). The rf-cinematogram: a cross-correlation technique for mapping several visual fields at once. *Biol. Cybern.*, 69:37–55.
- [Eckhorn et al., 1990] Eckhorn, R., Reitboeck, H. J., Arndt, M., and Dicke, P. (1990). Feature linking via synchronization among distributed assemblies: Simulations of results from cat visual cortex. *Neural Computation*, 2:293–307.
- [Ekeberg et al., 1991] Ekeberg, O., Wallen, O., Lansner, A., Traven H., Brodin, L., and Grillner, S. (1991). A computer based model for realistic simulations of neural networks. *Biol. Cybern.*, 65:81–90.
- [Engel et al., 1991a] Engel, A. K., König, P., and Singer, W. (1991). Direct physiological evidence for scene segmentation by temporal coding. *Proc. Natl. Acad. Sci. USA*, 88:9136–9140.
- [Engel et al., 1991b] Engel, A. K., König, P., Kreiter, A. K., and Singer, W. (1991). Interhemispheric synchronization of oscillatory neural responses in cat visual cortex. *Science*, 252:1177–1179.
- [FitzHugh, 1961] FitzHugh, R. (1961). Impulses and physiological states in models of nerve membrane. *Biophys. J.*, 1:445–466.
- [Georgopoulos et al., 1986] Georgopoulos, A. P., Schwartz, A., and Keizer, R. E. (1986). Neuronal population coding of movement direction. *Science*, 233:1416–1419.
- [Gerstner, 1991] Gerstner, W. (1991) Associative memory in a network of 'biological' neurons. *Advances in Neural Information Processing Systems*, vol. 3, Morgan Kaufmann Publishers, San Mateo, CA, 84–90.
- [Gerstner and van Hemmen, 1992] Gerstner, W., and van Hemmen, J. L. (1992). Associative memory in a network of 'spiking' neurons. *Network*, 3:139–164.
- [Gerstner et al., 1993] Gerstner, W., Ritz, R., and van Hemmen, J. L. (1993). A biologically motivated and analytically soluble model of collective oscillations in the cortex: I. theory of weak locking. *Biol. Cybern.*, 68:363–374.