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.

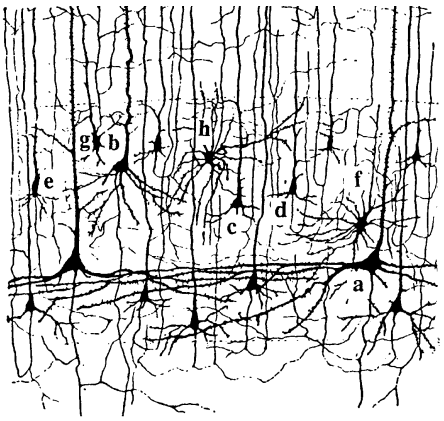

Figure 1.1. This reproduction of a drawing of Ramón y Cajal shows a few neurons in the cortex. Only a small portion of the neurons are shown; the density of neurons is in reality much higher. Cell b is a nice example of a pyramidal cell with a triangularly shaped cell body. Dendrites, which leave the 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 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

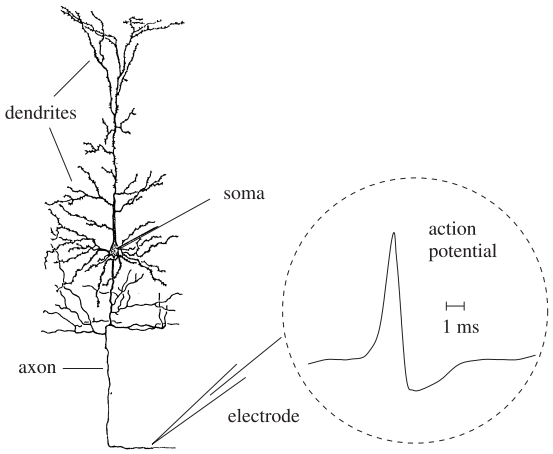

Figure 1.2. A single neuron. Dendrite, soma, and axon can be clearly distinguished. The inset shows an example of a neuronal action potential (schematic). Neuron drawing after Ramón y Cajal. The action potential is a short voltage pulse of 1-2 ms duration.

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

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

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 $\Delta t, 2\Delta t \ldots$, 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 a sequence of numbers $S_i(\Delta t), S_i(2\Delta t), \ldots$ with

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

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

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

where $\delta(.)$ 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?

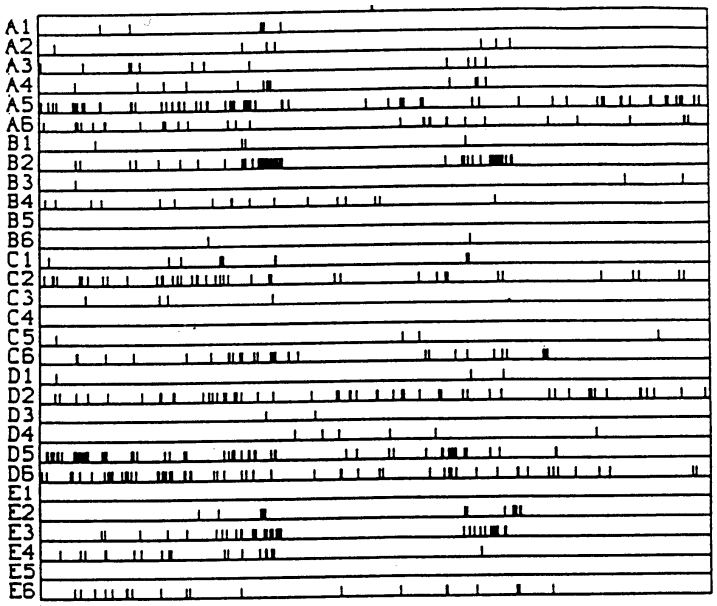

Figure 1.3. Spatio-temporal pulse pattern. The spikes of 30 neurons (A1-E6, plotted along the vertical axes) are shown as a function of time (horizontal axis, total time is 4 000 ms). The firing times are marked by short vertical bars. From [Krüger and Aiple, 1988].

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 $n_{\rm sp}(T)$ that occur in this time window. Division by the length of the time window gives the mean firing rate

$$\nu = \frac{n_{\rm sp}(T)}{T} \tag{1.4}$$

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 a is 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 second part 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

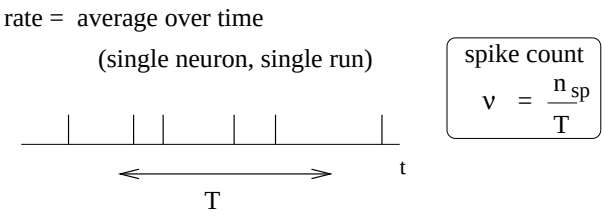

Figure 1.4. Definition of the mean firing rate via a temporal average.

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

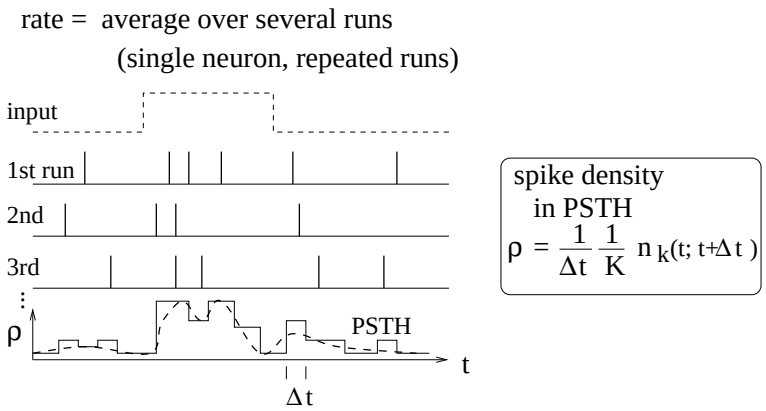

Figure 1.5. Definition of the spike density in the Peri-Stimulus-Time Histogram (PSTH).

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

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

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

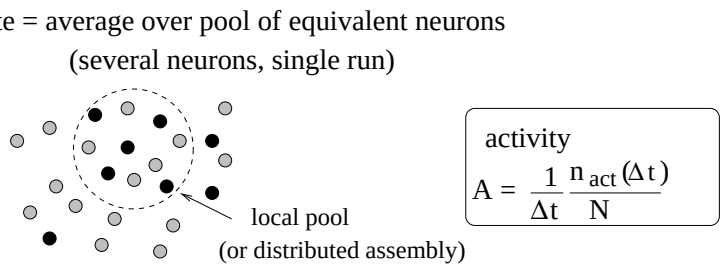

Figure 1.6. Definition of the population activity.

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

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

where N is the size of the population, Δt a small time interval, and $n_{\rm act}(t; t+\Delta t)$ the number of spikes (summed over all neurons in the population) that occur between t and $t+\Delta t$. population is large, we can consider the limit $N \to \infty$ and take then $\Delta t \to 0$. 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

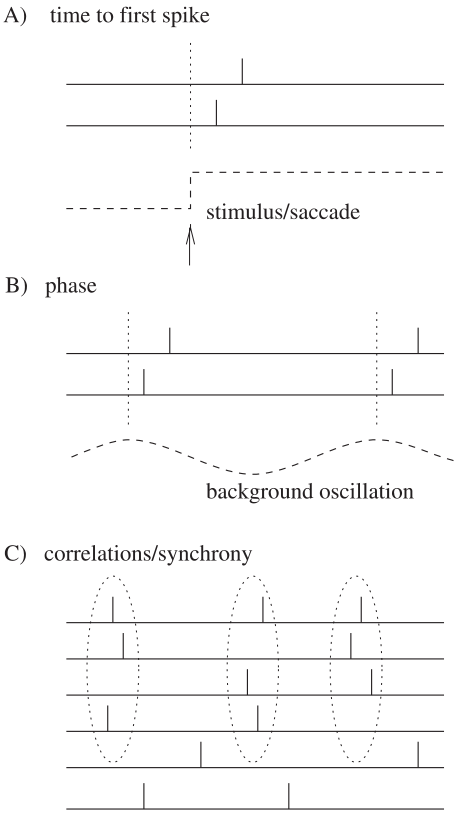

Figure 1.7. Three examples of pulse codes. A) Time-to-first spike. The second neuron responds faster to a change in the stimulus than the first one. Stimulus onset marked by arrow. B) Phase. The two neurons fire at different phases with respect to the background oscillation (dashed). C) Synchrony. The 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,

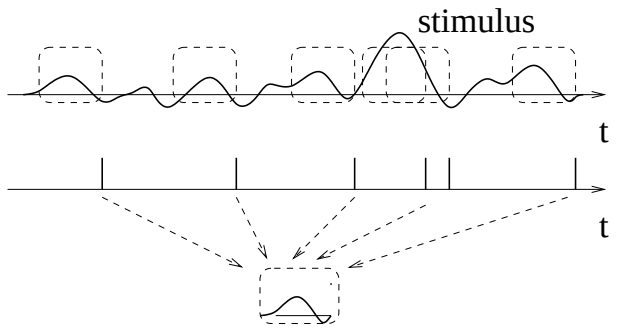

Figure 1.8. Reverse correlation technique (schematic). The stimulus in the top trace has caused the spike train shown immediately below. The time course of the stimulus just before the spikes (dashed boxes) has been averaged to yield the typical time course (bottom).

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 $\mathcal{F} = \{t^{(1)}, \ldots t^{(n)}\}$ 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 $s^{\rm est}(t)$ 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

$$s^{\rm est}(t) = \sum_{f=1}^{n} \kappa(t - t^{(f)}) \,. \tag{1.7}$$

The form of the kernel κ was determined through optimization so that the average reconstruction error $\int dt[s(t) - s^{\rm est}(t)]^2$ 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.

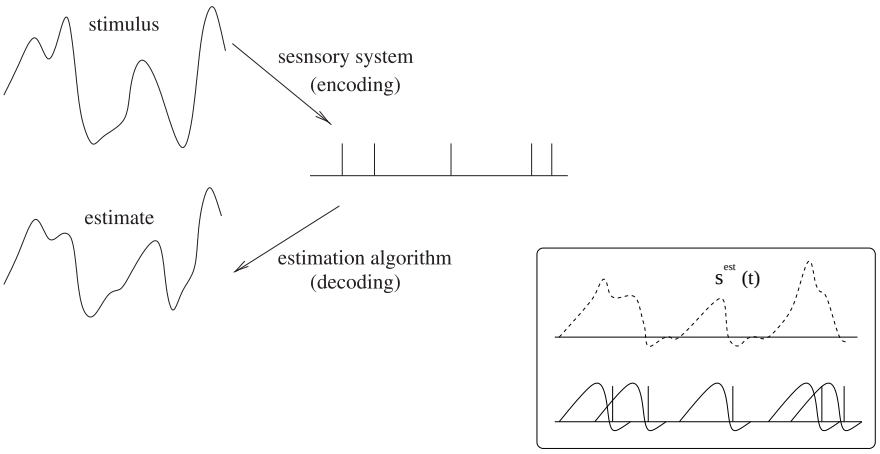

Figure 1.9. Reconstruction of a stimulus (schematic). A stimulus evokes a spike train of a neuron. The time course of the stimulus may be estimated from the spike train. The inset shows the principle of linear stimulus reconstruction. The estimation $s^{\rm est}(t)$ (dashed) is the sum of the contributions (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 $K(.)$. We can estimate the rate ν at time t by

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

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 $K(.)$ 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

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

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