

Chapter 14

Simplified Models of Individual Neurons

In the previous thirteen chapters, we meet and described, sometimes in excruciating detail, the constitutive elements making up the neuronal hardware: dendrites, synapses, voltage-dependent conductances, axons, spines, calcium and other second messengers. We saw how, different from electronic circuits in which only very few levels of organization exist, the nervous systems has many, tightly interlocking levels of organization that co-depend on each other in crucial ways. It is now time to put some of these elements together into a functioning whole, a single nerve cell. With such a single nerve cell model in hand, we can ask functional questions, such as at what time scale does it operates, what sort of operations can it carry out and how good is it at encoding information.

We begin this Herculean task by (i) completely neglecting the dendritic tree and (ii) by replacing the conductance-based description of the spiking process (e.g. the Hodgkin-Huxley equations) by one of two canonical descriptions. These two steps dramatically reduce the complexity of the problem of characterizing the electrical behavior of neurons; instead of having to solve coupled, nonlinear partial differential equations, we are left with a single ordinary differential equation. Such simplifications allow us to formally treat networks of large numbers of interconnected neurons, as exemplified in the neural network literature, and to simulate their dynamics. Understanding any complex system always entails choosing a level of description that retains key properties of the system while removing nonessential ones for the purpose at hand. The study of brains is no exception to this.

Numerous simplified single cell models have been proposed over the years, yet most of them can be reduced to just one of two forms. These can be distinguished by the form of their output: *spike* or *pulse* models generate discrete, all-or-none impulses. Their output over time can be treated as a series of delta functions $\sum_i \delta(t - t_i)$. Implicitly this assumes that no information is contained in the spike height or width. The original model of a neuron in McCulloch and Pitts (1943) as well as the venerable *integrate-and-fire* unit are instances of pulse models. In *firing rate* neurons, the output is a continuous firing rate, assumed to be a positive, bounded and stationary function of the input. Examples of these are the units

at the heart of Hopfield's (1984) associative memory network.

Yet before we can delve into more detail we need to introduce the deep issue of the proper output representation of a spiking cell that directly relates to the question of the neuronal code used to transmit information among neurons.

14.1 Rate Codes, Temporal Coding and All of That

How the neuronal output is represented, as a series of discrete pulses or as a continuous firing rate, relates to the question of the *code* used by the nervous system to transmit information between cells. So let us briefly digress and talk about neuronal codes.

In a typical physiological experiment, the same stimulus is presented multiple times to a neuron and its response is recorded (Fig. 14.1). One immediately notices that the detailed response of the cell changes from trial to trial; characterizing and analyzing stochastic component of the neuronal response is important enough that we dedicate the entire following chapter to it.

Given the pulse-like nature of spike trains, the standard procedure to quantify the neuronal response is to count how many spikes arrived within some sampling window Δt and to divide this number by the number of presentations. This yields the conditional probability that a spike occurred between t and $t + \Delta t$ given some particular stimulus.

In the limit of very small sampling windows, such that the probability for more than one spike occurring within Δt is vanishingly small (usually 1 msec), the probability of spiking is given by $f(t)\delta t$, where $f(t)$ is the *instantaneous firing rate* of the neuron (in units of spikes/time). $f(t)$ as a function of time after the onset of a stimulus is known as the *post-stimulus time histogram* (PSTH; Fig. 14.1). Because of the stochastic nature of the neuronal response $f(t)$ will, in general, not become arbitrary sharp as Δt goes to zero but retains a considerable degree of "fuzziness" (but see Fig. 15.10).

It is important to realize the artificial nature of this construct. The nervous system has to make a decision based on a single spike train and not on the average of tens or more of spike trains. A visual neuron in the fly does not have the opportunity to see the hand that is about to swat it approach ten times before it makes the decision to initiate an escape response! A neuron can only observe the $\sum_i \delta(t - t_i)$ from its presynaptic partners, and not $f(t)$.

Under certain conditions this might be different. If a cell, say in cortex, has access to the spiking output of many cells with the same receptive field properties, the temporal average of the single presynaptic neuron can be replaced by an ensemble average over a population of neurons, thereby approximating $f(t)$. In many cases though *population rate coding* cannot occur for lack of a sufficient large population to average over. In the insect, a very small number of clearly identifiable neurons code for particular features of the sensory input and no ensemble averaging occurs.

Given the stochastic nature of spike trains, a common assumption is that the averaged firing rate of a neuron constitutes the primary variable relating neuronal response to sensory experience (Adrian and Zotterman, 1926; Adrian, 1932; Lettvin, Maturana, McCulloch and

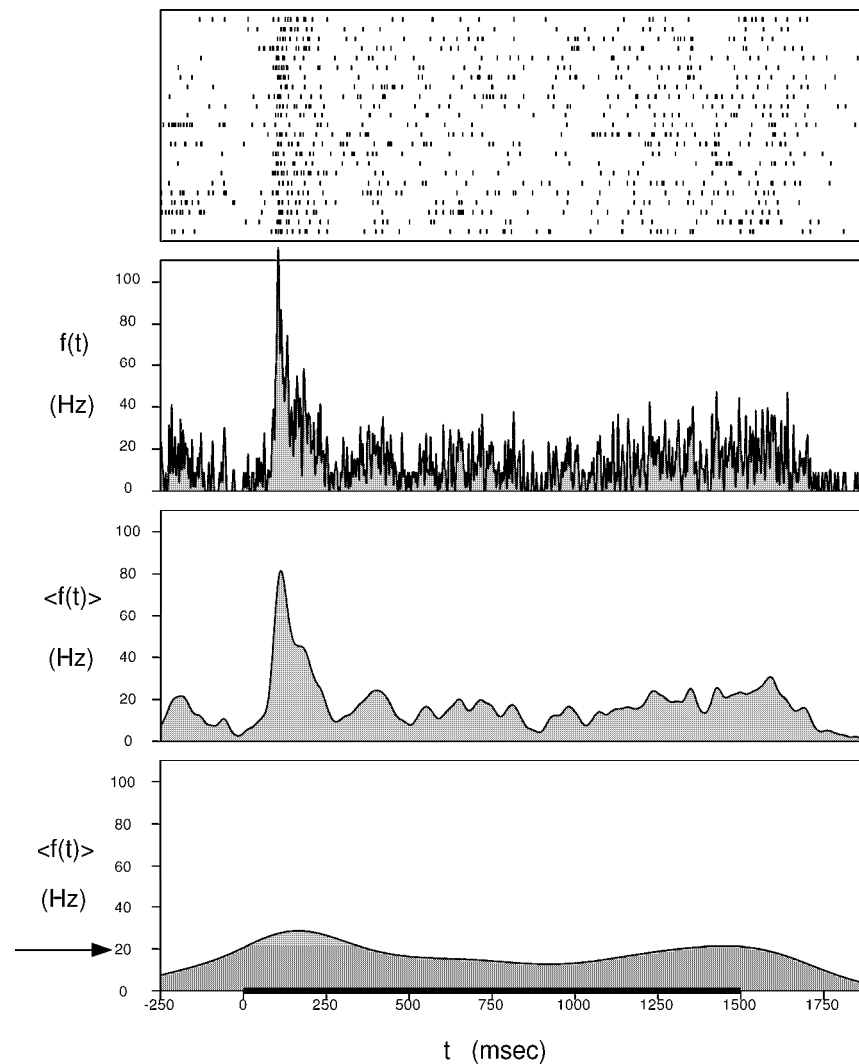


Figure 14.1: WHAT IS THE FIRING RATE

The definition of the *firing rate*. Starting point are numerous trials in which the same stimulus is repeatedly presented to the animal and the spikes generated by some cell recorded. These are shown in the *raster diagram* at the top taken from a cell in cortical area V4 in the awake monkey. The stimulus—a grating—is flashed on at 0 and lasted until 1500 msec. 23 of these trials are averaged, smoothed with a Gaussian of 2 msec standard deviation σ and normalized. This averaging window is so small that it effectively defines the instantaneous firing rate $f(t)$. Due to the stochastic nature of neuronal firing these *post-stimulus time histogram* (PSTH) do not become arbitrarily sharp but remain fuzzy. The two lower plots illustrate an *average firing rate* $\langle f(t) \rangle$ obtained from the raster diagrams using Gaussian smoothing with $\sigma = 20$ and 200 msec. In many experiments only the average number of spikes triggered during each trial, corresponding to $\sigma \rightarrow \infty$ (see arrow at 19.5 Hz) is used to relate the cellular response to the behavior of the animal. It is important to realize that a single neuron only sees spike trains, and not a smoothly varying firing rate. From D. Leopold and N. Logothetis, personal communication.

Pitts, 1959; Barlow, 1972). This belief is supported by the existence of a quantitative relationship between the averaged firing rate of single cortical neurons and psychophysical judgments made by a monkey. That is, the animal's behavior in a visual discrimination task can be statistically predicted by counting spikes over a long interval (typically one second or more) in a single neuron in visual cortex (Werner and Mountcastle, 1963; Barlow, Kaushal, Hawken and Parker 1987; Newsome, Britten and Movshon, 1989; Vogels and Orban, 1990; Zohary, Hillman and Hochstein, 1990; Britten, Shadlen, Newsome and Movshon, 1992).

In these experiments, the rate is estimated by averaging over a window that is large compared to the time in which the sensory stimulus itself changes:

$$\langle f(t) \rangle = \frac{1}{n} \sum_{j=1}^n \frac{1}{\Delta T} \int_t^{t+\Delta T} \sum_{i=1}^{n_j} \delta(t' - t_{ij}) dt', \quad (14.1)$$

where the first sum j is executed over the n identical trials and the second sum over all n_j spikes at time t_{ij} during the j -th trial. Notice that the δ terms have the dimension of spikes per δt , so that the average $\langle f(t) \rangle^1$ has the correct units associated with a rate. Instead of a rectangular window, a frequent alternative is to smooth using a Gaussian convolution kernel with standard deviation σ :

$$\langle f(t) \rangle = \frac{1}{n} \sum_{j=1}^n \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{+\infty} e^{-\frac{(t-t')^2}{2\sigma^2}} \sum_{i=1}^{n_j} \delta(t' - t_{ij}) dt', \quad (14.2)$$

(Fig. 14.1). Because of this success in linking $\langle f(t) \rangle$ with behavior, it has been assumed by some that only the mean rate, averaged over a significant fraction of a second or longer, is the relevant code in the nervous system and that the detailed time course of spikes is not.

The past decade has witnessed a revival in the question to what extent an average rate coding neglects information (for an expose of these ideas, see the superb textbook by Rieke, Warland, van Steveninck and Bialek, 1996). On the basis of signal-processing and information-theoretical approaches, we now know that individual motion-selective cells in the fly (Bialek, Rieke, van Steveninck and Warland, 1991), single afferent axons in the auditory system in the bullfrog (Rieke, Warland and Bialek, 1993) and in the electrosensory system in weakly electric fish (Wessel, Koch and Gabbiani, 1996) can carry between 1 and 3 bits of sensory information per output spike, amounting to rates of up to 300 bits per second. This information is encoded using the instantaneous rate with a resolution of 5 msec or less. And the elegant experiments of Markram and his colleagues (1997), demonstrating the effect a short delay between a presynaptic and a postsynaptic spike arriving at a synapse can have on its "weight" (section 13.6.2), provides a biophysical rationale for why timing at the 10 msec level is crucial for synaptic plasticity.

We summarize this vast body of work (see Rieke *et al.*, 1996) by concluding that in many instances $\langle f(t) \rangle$ —averaged over a 5 to 10 msec time frame—appears to be the relevant code used to transmit information:

¹Sometimes also written as $\langle f(t) \rangle_T$ to express its dependency on the size of the averaging window.

More complex neuronal codes do exist and are frequently referred to under the catch-all term of *temporal coding* (for an exhaustive listing of possible codes, see Perkel and Bullock, 1968). However, because of its implication that rate codes do not preserve detail timing information, we prefer a term coined by Larry Abbott (personal communication) *correlation coding*.

In an instantaneous firing rate code, the generation of each spike does not depend on other spikes in the trains (neglecting the refractory period and bursting). In a correlation code this assumption is abandoned in favor of coupling among pairs, triplets or higher-order groupings of spikes.

To give one example of a correlation code, let $f(t)$ in response to some stimulus be a maintained response. We assume that this cell is very noisy with distinct spike patterns on individual trials. Averaging over them all leads to a flat response of amplitude f_c . In a rate code, f_c is the only information available to a postsynaptic neuron. Closer inspection of the microstructure of spiking reveals that the intervals between consecutive spikes are not independent of each other but that two short spike intervals are always followed by a long one. Any code that fails to exploits these higher-order correlations amongst four consecutive spikes would miss something. For experimental evidence of such codes see (Segundo, Moore, Stensaas and Bullock, 1963; Chung, Raymond and Lettvin, 1970; Eskandar, Richmond and Optican, 1992; Richmond and Optican, 1992; Optican and Richmond, 1987).

A generic problem with the assumption of correlation codes is the question of decoding. It is unclear what sort of biophysical mechanisms are required to exploit the information hidden in such correlations among spikes. They might be prohibitively complicated to implement at the membrane level.

So far we have said little about *population coding*. Once again, one can distinguish two broad types of codes, correlated one and non-correlated ones (Abbott, 1994). The latter are straightforward: here the information from numerous neurons is combined into a population code but without taking account of any correlations among neurons (Knight, 1972a,b). There is plenty of good evidence for such codes in a great variety of different sensory and motor systems, ranging from the four cricket cercal interneurons encoding the direction the wind is blowing from (Theunissen and Miller, 1991) to the larger ensembles encoding the direction of sound in the barn owl (Knudsen, du Lac and Esterly, 1987; Konishi, 1992) and eye movements in the mammalian superior colliculus (Lee, Rohrer and Sparks, 1988) to the posterior parietal cortex in the monkey encoding our representation of space (Pouget and Sejnowski, 1997).

Correlation population codes exploit the exact temporal relationships among streams of action potentials. One way to discover such codes is to record from two or more neurons simultaneously and to measure their cross-correlation function. For instance, it may be that two presynaptic neurons always generate spikes within one or two millisecond of each other. Much technological advance has occurred in this area in recent years, so that recordings from multi-unit recordings have now become routine.

In a variety of sensory systems in both invertebrates and vertebrates, physiological evidence indicates that cross-correlations among groups of cells appears to encode various

stimulus features (Freeman, 1975; Abeles, 1982, 1990; Strehler and Lestienne, 1986; Eckhorn, Bauer, Jordan, Brosch, Kruse, Munk and Reitboeck, 1988; Gray, König, Engel and Singer, 1989; Bialek, Rieke, van Steveninck and Warland, 1991; Eskandar, Richmond and Optican, 1992; Konishi, 1992; Singer and Gray, 1995; Decharms and Merzenich, 1996; Wehr and Laurent, 1996). The best evidence to date linking neuronal synchronization directly to behavior comes from the bee's olfactory system (Stopfer, Bhagavan, Smith and Laurent, 1997). When pharmacological agents were used to block cellular synchronization—without interrupting neuronal activity *per se*—fine sensory discrimination was disrupted.

Much theoretical work using pulse-coded neural networks has focused on the idea that spike coincidence across neurons encodes information (Sejnowski, 1977a; Abeles, 1982, 1990; Amit and Tsodyks, 1991; Koch and Schuster, 1992; Griniasty, Tsodyks and Amit, 1993; Abbott and van Vreeswijk, 1993; Zipser, Kehoe, Littlewort and Fuster, 1993; Softky, 1995; Hopfield, 1995; Maass, 1996). Indeed, it has been proposed that temporal correlation among groups of neurons is a crucial signal for a number of perceptual processes, including figure-ground segregation and the binding of different attributes of an object into a coherent, unitary percept (the so-called *binding problem*; Milner, 1974; von der Malsburg, 1981; von der Malsburg and Schneider, 1986), selective visual attention (Niebur and Koch, 1994), and even the neuronal basis of awareness (Crick and Koch, 1990; for a review see Koch, 1993).

From our point of view as biophysicists, a correlation population code based on coincidence detection has the significant advantages that it is straightforward to implement at the membrane level (witness Fig. 21.2).


The question of rate versus correlation coding remains with us. Yet this stark either-or dichotomy is not very useful. Clearly, the timing of spikes, at least at the 10 msec level, is important. And in some systems, detailed information across groups of cells will also prove to be of relevance, although for which properties and at what time scale is an open question.

It therefore behooves us to study how accurately and reliably neurons can generate individual action potentials and how robust these are to noise. We here lay the groundwork by introducing pulse neurons as well as firing rate models and describing their basic properties. We believe that such single cell models represent the most reasonable tradeoff between simplicity and faithfulness to key neuronal attributes. The following chapter will deepen our discussion of stochastic aspects of neuronal firing. We will also discuss firing rate models.


14.2 Integrate-and-Fire Models

We turn to a very simple, but quite powerful, model of a spiking cell with a long and distinguished history, first investigated by Lapique (1907, 1926) before anything specific was known about the mechanisms underlying impulse generation. In its vanilla flavored version, it is known as the *integrate-and-fire* model (Stein, 1967a,b; Knight, 1972a; Jack *et al.*, 1975; Tuckwell, 1988b; frequently referred to as *voltage threshold* or *Lapique's model* in the older literature). In its simplicity it rivals physics' linear oscillator model, yet it does encapture the two key aspects of neuronal excitability: a passive, integrating subthreshold domain and the generation of stereotypical impulses once a threshold has been exceeded.

Perfect Integrate-and-Fire Unit



Leaky Integrate-and-Fire Unit



Adapting Integrate-and-Fire Unit




Figure 14.2: INTEGRATE-AND-FIRE MODELS

Three basic variants of integrate-and-fire units. Common to all are passive integration within a single compartment for the subthreshold domain and a nonlinear element, spike generation accomplished with a voltage threshold V_{th} . Whenever the membrane potential $V(t)$ reaches V_{th} , a pulse is generated and the unit is short-circuited. For a duration t_{ref} following spike generation, any input $I(t)$ is shunted to ground (corresponding to an absolute refractory period). The *perfect* or *non-leaky* integrate-and-fire model contains but a capacitance. The *leaky* or *forgetful integrate-and-fire* unit accounts for the decay of the membrane potential by an additional component, a leak resistance R . The *adapting integrate-and-fire unit* with six free parameters (eqs. 14.12 and 13) shows firing rate adaptation via the introduction of g_{adapt} , corresponding to an calcium-dependent potassium conductance (in addition to the absolute refractory period). Each spike increments the amplitude of the conductance; its value decays exponentially between spikes.

The *non-leaky* or *perfect integrate-and-fire* unit consists but of a single capacitance for integrating the charge delivered by synaptic input in addition to a fixed and stationary voltage threshold V_{th} for spike initiation (Fig. 14.2). The *leaky* or *forgetful integrate-and-fire*

model includes a resistance, accounting for leakage currents through the membrane. While integrate-and-fire models don't incorporate the detailed time course of the action potential, the effect of adaptation can be incorporated; indeed, the current-frequency relationship of such an integrate-and-fire cell with a handful of parameters can be very close to that of a much more complex, conductance-based cell model.

14.2.1 The Perfect or Non-Leaky Integrate-and-Fire Unit

We will be considering a number of variants of integrate-and-fire "units". All are characterized by a subthreshold domain of operation and a voltage threshold V_{th} for spike generation. The *perfect integrate-and-fire* unit deals with subthreshold integration via a single capacitance C . While unphysiological, it is mathematical tractable, which is why it is frequently invoked for pedagogical purposes. For the sake of mathematical convenience we assume the input to be a current $I(t)$, arising either from synaptic input or from an intracellular electrode. The generalization to a conductance based input is straightforward.

The voltage trajectory of the perfect integrator is governed by the first-order differential equation:

$$C \frac{dV(t)}{dt} = I(t). \quad (14.3)$$

Together with an initial condition eq. 14.3 specifies the subthreshold time course of the membrane potential.

Once the potential reaches V_{th} , a pulse is triggered and the charge that has accumulated on the capacitance is shunted to zero (through the closed switch in Fig. 14.2). This would normally be accomplished by the various conductances underlying spiking. Sweeping the charge to "ground" has the effect of instantaneously resetting $V(t)$ to zero. Because the model has no pretense of mimicking the currents involved in shaping the action potential, spike generation itself is not part of the model. Formally, we model the action potential by assuming that at the instant t' at which $V(t') = V_{th}$ (or the first time $V(t)$ exceeds V_{th} for models with instantaneously rising EPSPs) an output pulse, described by a delta function $\delta(t-t')$, is generated. The successive times, t_i , of spike occurrence are determined recursively from the equation

$$\int_{t_i}^{t_{i+1}} I(t) dt = CV_{th}. \quad (14.4)$$

As discussed in section 6.4, a canonical way in which experimentalists characterize a cell's behavior is by determining its *discharge, frequency-current* or $f-I$ curve, the relationship between the amplitude of an injected current step and the average firing frequency (defined over an interval longer than the interspike interval; as in eq. 14.4, $\langle f \rangle$ is computed as the inverse of the interspike interval).

In response to a maintained current, the membrane potential will charge up the capacitance until V_{th} is reached and V is reset to zero. The larger the current, the smaller the

intervals between spikes and the higher the firing rate according to

$$\langle f \rangle = \frac{I}{CV_{th}}. \quad (14.5)$$

Three features are worthwhile here. (i) The firing rate is linearly related to the input current (Fig. 14.3B). (ii) Arbitrarily small input currents will eventually lead to a spike, since no input is forgotten. (iii) The output spike train is perfectly regular. Of course, real neurons rarely, if ever, respond to a sustained current injection with a regular discharge of spikes but instead show substantial variability in the exact timing of the spikes. This is particularly true of neurons recorded *in vivo* (Holt, Softky, Koch and Douglas, 1996). The following chapter will deal with this situation.

The dynamic firing range of nerve cells is limited by the fact that the sodium current responsible for spiking must recover from inactivation. Potassium currents furthermore limit the peak firing range. The effect of a refractory period is mimicked by postulating that following spike generation, the membrane potential is set to zero for a fixed duration t_{ref} ; any current arriving within this window is shunted away. This introduces a non-linear saturation into the f-I curve of the perfect integrator,

$$\langle f \rangle = \frac{I}{CV_{th} + t_{ref}I}. \quad (14.6)$$

The output of such an integrator to an arbitrary input current consists of a series of impulses, $\sum_i \delta(t - t_i)$, all of which are spaced at least t_{ref} apart.

14.2.2 The Forgetful or Leaky Integrate-and-Fire Unit

The model considered so far will sum linearly multiple subthreshold inputs irrespectively of their temporal relationship because no account is made of a leak. A more realistic behavior is obtained by incorporating a leak resistance into the subthreshold domain (Fig. 14.2),

$$C \frac{dV(t)}{dt} + \frac{V(t)}{R} = I(t). \quad (14.7)$$

Having first meet this equation in Chapter 1 (eq. 1.6), we know that the evolution of the subthreshold voltage is completely characterized by convolving $I(t)$ with the associated Green's function, $e^{-t/\tau}$ (with $\tau = RC$). The time-course of the membrane potential of the leaky integrate-and-fire unit to a step of constant current I , switched on at $t = 0$ and remaining on, can be obtained by solving eq. 14.7,

$$V(t) = IR(1 - e^{-t/\tau}) + V(t=0)e^{-t/\tau}. \quad (14.8)$$

The membrane charges exponentially up to its stationary value $V = IR$.

The integrator model will only follow this equation as long as the voltage remains below V_{th} , since upon reaching the threshold a spike is initiated and the voltage is reset to zero

(Fig. 14.3A). The minimal sustained current necessary to trigger an action potential, that is the threshold current, is

$$I_{th} = \frac{V_{th}}{R}. \quad (14.9)$$

For any current I larger than I_{th} , an output impulse will be generated at time T_{th} , such that $IR(1 - \exp(-T_{th}/\tau)) = V_{th}$ holds. Inverting this relationship yields the time to spike as

$$T_{th} = -\tau \log \left(1 - \frac{V_{th}}{IR} \right). \quad (14.10)$$

Solving this equation for the minimal duration needed for a sustained current of a fixed amplitude to generate a spike yields what is known as *strength-duration curve* (Jack *et al.*, 1975; Noble and Stein, 1966). Since the voltage is reset following an impulse and if we assume that the input current persists, the membrane will again charge up to the threshold, triggering the next spike $T_{th} + t_{ref}$ later (Fig. 14.3A).

If we take proper account of the refractory period by assuming that for t_{ref} following each spike, all input current is simply lost (due to the shunting effect of the conductances underlying the afterhyperpolarization), the continuous firing rate as a function of the injected current will be (Fig. 14.3B)

$$\langle f \rangle = \frac{1}{T_{th} + t_{ref}} = \frac{1}{t_{ref} - \tau \log \left(1 - \frac{V_{th}}{IR} \right)}. \quad (14.11)$$

For currents below I_{th} no spike is triggered and at $I = I_{th}$, the slope of the f-I curve is infinite. For large currents, the firing rate saturates at the inverse of the refractory period (Fig. 14.3B). In the absence of a refractory period, the slope of the f-I curve levels off to a constant value of $1/(V_{th}C)$, identical to the slope of the non-leaky unit². Its steepness can be increased by reducing the threshold voltage or by decreasing the membrane capacitance. Due to the refractory period, the f-I curve gently bends over to level off at $f_{max} = 1/t_{ref}$ for (unphysiological) high current levels.

14.2.3 Other Variants

Besides the generic version of the integrate-and-fire model discussed above, a number of variants are in usage.

- In order to better account for the 50 to 100 msec time-course of adaptation, Wehmeier and colleagues (1989) introduced a purely time-dependent shunting conductance g_{adapt} (with a reversal potential equal to the resting potential, here assumed zero). Each spike increases this conductance by a fixed amount G_{inc} ; between spikes, g_{adapt} decreases exponentially with a time constant τ_{adapt} to zero. Such an effective calcium-dependent

²This can be seen upon developing the log term in eq. 14.11 into a Taylor series, with $\log(1+x) \approx x - x^2/2$ for $|x| \ll 1$.

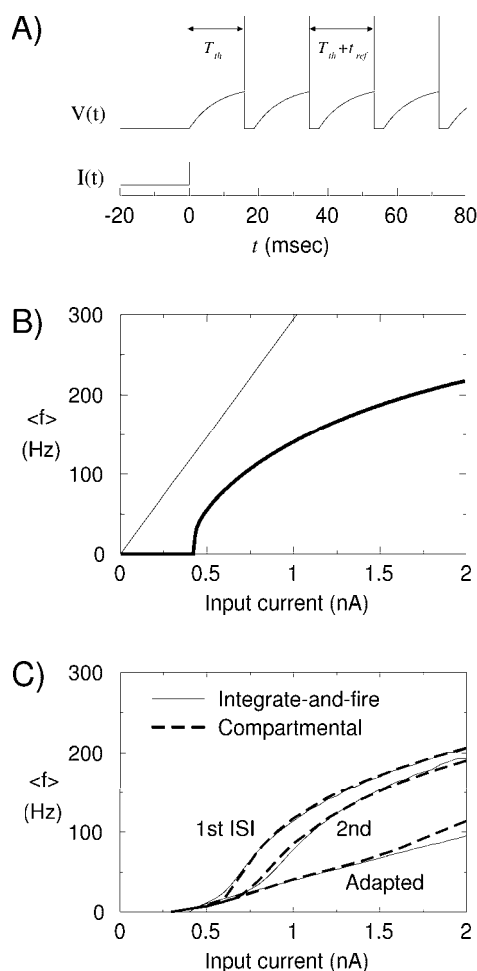


Figure 14.3: SPIKING IN A LEAKY INTEGRATE-AND-FIRE MODEL

The average firing frequency, determined as the inverse of the interspike interval, as a function of the amplitude of a maintained current input, for the leaky integrate-and-fire units of Fig. 14.2. **(A)** Exemplar trace of such a unit receiving a current step with $I = 0.5$ nA. Before the membrane potential has time to reach equilibrium, the unit spikes. $V_{th} = 16.4$ mV, $C = 0.207$ nF, $R = 38.3$ M Ω and $t_{ref} = 2.68$ msec. **(B)** The f-I or *discharge curve* for the same leaky unit with refractory period (eq. 14.11). At threshold, $I_{th} = V_{th}/R$, the slope is infinite. The firing rate saturates at $1/t_{ref}$. For comparison, the f-I curve of the non-leaky unit without refractory period with constant slope $1/(V_{th}C)$, is superimposed. **(C)** An adapting conductance (with $G_{inc} = 20.4$ nS and $\tau_{adapt} = 52.3$ msec) is added to the leaky integrate-and-fire unit from the upper panel (see Fig. 14.2) and the resulting f-I curve is compared against the discharge curve of the biophysical detailed compartmental model of the layer 5 pyramidal cell (Fig. 17.10). The degree of matching between the simple and the very complex models are quite remarkable and support our contention that suitable modified integrate-and-fire models do mimic numerous aspects of the behavior of neurons. Adaptation is already evident when considering the first interspike interval (between the first and the second spike); adaptation linearizes the very steep f-I curve around I_{th} (compare with B).

potassium conductance imitates both the absolute as well as the relative refractory period following spike initiation. We will refer to such a unit as an *adapting integrate-and-fire* model (Fig. 14.2C). Note that a refractory period t_{ref} is still necessary in order to mimic the very short-term aspect of adaptation. In the subthreshold domain, this unit is described by

$$C \frac{dV}{dt} = -\frac{V(1 + Rg_{adapt})}{R} + I \quad (14.12)$$

$$\tau_{adapt} \frac{dg_{adapt}}{dt} = -g_{adapt} \cdot \quad (14.13)$$

If V reaches V_{th} at time t' , a spike is generated at this point in time and $g_{adapt}(t')$ is incremented by G_{inc} . This model is completely characterized by six parameters: V_{th} , C , R , t_{ref} , G_{inc} and τ_{adapt} .

- An alternative to this output-dependent membrane conductance is to increase the voltage threshold following each spike in a deterministic manner, for instance using the rule

$$V_{th}(t) = V_{th,0}(1 + \alpha e^{-(t-t')/\tau_{adapt}}), \quad (14.14)$$

where $t - t'$ is the time from the last impulse, $V_{th,0}$ the threshold in the absence of any adaptation and α the maximal normalized voltage threshold (Calvin and Stevens, 1968; Holden, 1976).

- In order to account for those neurons that do not show any profound afterhyperpolarization following spiking, the membrane potential can be reset to a value closer to V_{th} (e.g. 20% of V_{th}) instead of zero. This is equivalent to resetting the potential to zero but adding a constant current and can have a considerable effect on the jitter in pulse timing (Troyer and Miller, 1996).
- In a strategy to imitate the seemingly random nature of spike times, some authors resort to drawing the voltage threshold from some probability density distribution (Holden, 1976; Gestri, Masterbroek and Zaagman, 1980). Yet in real neurons, the spiking mechanism itself appears to be quite reliable (Calvin and Stevens, 1968; Mainen and Sejnowski, 1995). In the case of the perfect integrator, a random threshold and constant input can be shown to be equivalent to a random input and a constant threshold. Frequently, the former situation is both mathematically and computationally easier to deal with than the latter (for more details, see Gabbiani and Koch, 1997 and Appendix B).

What all of these models share is a mechanism for passive integration of synaptic inputs, a voltage threshold for spike initiation and a lack of specific spiking currents.

In chapter 17 we will learn that an action potential in a full-blown model of a pyramidal cell (with eight voltage-dependent conductances) is, indeed, generated whenever the somatic membrane potential exceeds -49 mV. This is because the synaptic current flowing into the soma—caused by rapid EPSPs on the time-scale of milliseconds—primarily charges up the

capacitance. Relative to this rapid charging current, the ionic currents in the subthreshold regime change on a much slower time-scale.

How well does this f-I curve compare against curves obtained from much more detailed and sophisticated models? Presaging section 17.5, we plot the f-I curve of the layer 5 pyramidal cell, including the effect of firing rate adaptation in Fig. 14.3C. Notice the very low slope of the f-I curve around threshold, in contrast to the infinite slope of the f-I curve of the integrate-and-fire unit. If, an adapting conductance is incorporated into the leaky integrator (eqs. 14.12 and 13 and Fig. 14.2C), it is surprising how well this single cell model (with just six degrees of freedom) resembles the much more detailed compartmental model based on membrane conductances.

Due to the presence of the leak term, integrate-and-fire models have been difficult to fully characterize analytically but have also been surprisingly successful in describing neuronal excitability. They have been applied to model the firing behavior of numerous cell types: neurons in the limulus eye (Knight, 1972b), α motoneurons (Calvin and Stevens, 1968), neurons in the visual system of the housefly (Gestri, Masterbroek and Zaagman, 1980), cortical cells (Softky and Koch, 1993; Troyer and Miller, 1996) and others.

While singing the praise of integrator models, it must be pointed out that many cells do not behave like integrate-and-fire units. For instance, cerebellar Purkinje cells (Jaeger, De Schutter and Bower, 1997) or the many types of oscillating neurons that constitute the central pattern generators found throughout the animal kingdom (Marder and Calabrese, 1996), have such strong inherent nonlinearities, generated by powerful intrinsic currents, that any attempts to directly map their behavior onto this class of models would fail miserably. Approximations are possible, though. For instance, bursting (Chapter 16) could be treated by letting the rapid Na^+ spikes be handled by the integrate-and-fire threshold mechanism. The slow dynamics governing when the fast sequences of spikes are generated by incorporating voltage-dependent conductances into the unit.

14.2.4 The Response Time of Integrate-and-Fire Units

When a spiking, non-adapting membrane (such as the squid axon) receives a sustained, suprathreshold current input, its membrane potential never reaches an equilibrium but moves along a limit cycle. That is, it undergoes periodic changes in its state variables (chapter 7). When the integrate-and-fire neuron spikes, and the state variable is reset, it loses memory of the previous input current, and begins to respond to the new current by charging toward the threshold. It follows from these considerations that if there is a step change in current, the integrate-and-fire neuron must converge to its limit cycle by the end of the *first* interspike interval after the change, because everything during this first interval is exactly the same as during subsequent intervals.

Figure 14.4 compares the step response of an integrate-and-fire unit, a compartmental model of a cortical pyramidal neuron, an experimental record derived from a neuron in cat visual cortex and a mean-rate neuron. The first interspike interval already reflects the new firing rate—the convergence occurs on as short a time interval as can be defined (*i.e.* the

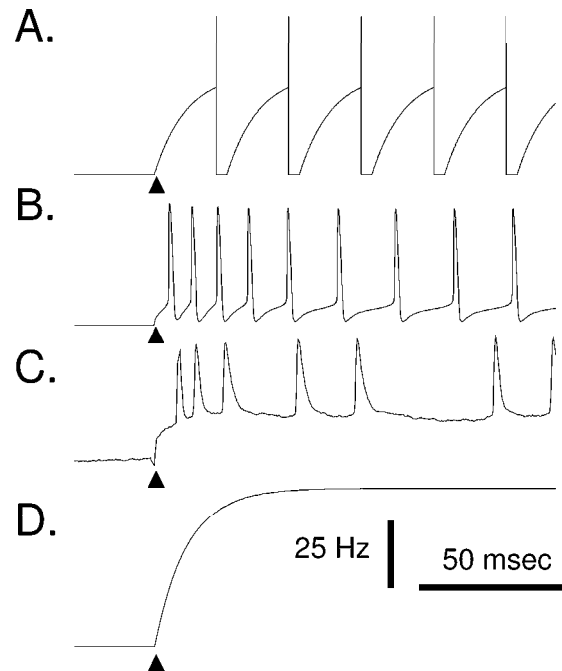


Figure 14.4: SPIKING CELLS

Sample spike rasters in response to a step current injection. **(A)** Leaky integrate-and-fire unit with refractory period spiking in response to a current step of 1.6 nA (for parameters, see legend to Fig. 14.3A). The arrows indicate the time at which the current injection commenced. **(B)** Somatic membrane potential in the layer 5 pyramidal cell model in response to a 1.5 nA current input. **(C)** Response of a cell in primary visual cortex of the anesthetized adult cat to a 0.6 nA current injection (from Ahmed *et al.*, 1993). The firing rate does not increase gradually; the effect of the change in current is fully visible in the first interspike interval. **(D)** The output of a non-adapting firing rate model with $\tau = 20$ msec. In the linear regime of the cell's f-I curve, the firing rate can be considered to be a low-pass filtered version of the step input.

interspike interval). The inclusion of adaptation currents (Fig. 14.3B) does not substantially affect this analysis. Cortical cells (Fig. 14.4C) also reach their maximum firing rate by the first interspike interval. Thereafter, the firing rate decreases slowly due to the temporal dynamics of adaptation.

Just because the subthreshold dynamics is governed by $\tau = RC$ does not imply that the neuron must respond to a suprathreshold input with the same dynamics. Returning to expression 14.10 for the time-to-spike, T_{th} , we notice that the larger the injected current, the sooner the cell spikes (jump ahead to Fig. 14.5A). Furthermore, T_{th} actually decreases as the input resistance, and therefore τ , increases (Fig. 14.5B). This can easily be explained by recalling that T_{th} is the time it takes for the membrane potential to reach the fixed threshold V_{th} . Increasing the input resistance will shorten this time, even if overall it would have taken the membrane longer to reach its ultimate steady-state value RI (which is never reached

since a spike is triggered and the membrane potential reset once V hits V_{th}).

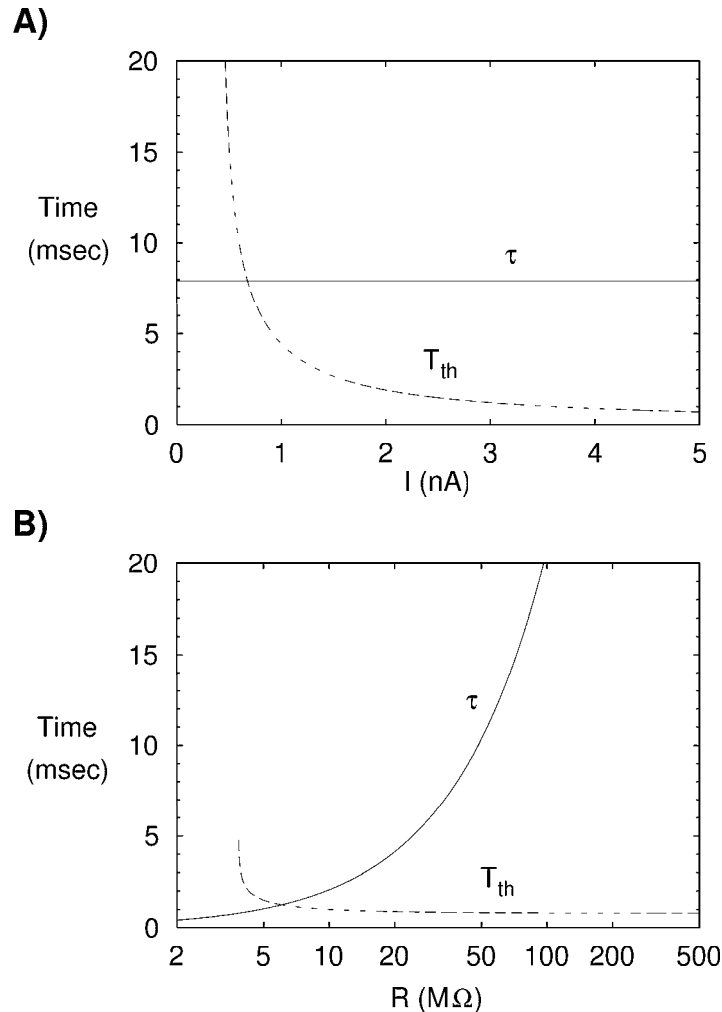


Figure 14.5: INTEGRATE-AND-FIRE UNITS CAN RESPOND MUCH FASTER THAN τ
 The membrane time constant $\tau = RC$ of a leaky integrate-and-fire unit and the time, T_{th} , it takes such a unit, starting at $V = 0$, to reach V_{th} and spike (eq. 14.10). **(A)** As the amplitude of the injected current increases, the unit can spike very rapidly. The true time to spike is $\leq T_{th}$, since the unit's initial state is usually $V \geq 0$. The parameters are as in Fig. 14.3A. **(B)** The input resistance R is varied over two orders of magnitude. As τ diverges, T_{th} converges to CV_{th}/I , the time it takes for the voltage across a capacitance to reach V_{th} . While this may be counterintuitive, it follows from the fact that T_{th} is the time it takes to reach a fixed threshold value, while τ dictates the approach toward a steady-state value beyond V_{th} . The injected current $I = 4.3$ nA. All else as in the upper panel. For small values of R , the current fails to reach I_{th} and T_{th} diverges.

14.3 Firing Rate Models

The potential in a continuous firing rate unit, such as those at the heart of most neural networks, has the same dynamics as in the leaky integrate-and-fire unit, *i.e.*

$$C \frac{dV(t)}{dt} = -\frac{V(t)}{R} + I(t). \quad (14.15)$$

A subtle but far-reaching difference is that the instantaneous output of this unit, $f(t)$, is a continuous but nonlinear function of $V(t)$:

$$f = g(V). \quad (14.16)$$

Most commonly g is a *sigmoidal* function, that is a monotonic increasing, positive and saturating function. A popular choice is

$$f = \frac{1}{1 + e^{-2\beta V}}, \quad (14.17)$$

where β is some constant (Fig. 14.6). Other functions have also been used (e.g. $\tanh V$ or V^2). V is sometimes identified with the so-called *generator potential*, the somatic membrane potential when the spikes are disabled (e.g. by blocking the fast sodium conductance). Under certain conditions, the generator potential correlates well with the firing rate when spikes are blocked (Katz, 1950; Granit, Kernell and Shortess, 1963; Stein, 1967a; Dodge, Knight and Toyoda, 1968). If f is thought of as the firing rate of the unit, then the function g can be identified with the f-I curve of the cell under investigation and can be directly fitted against experimental data. Given the continuous and smooth nature of g , no true current threshold exists.

An electronic circuit implementation of a rate neuron has been drawn in Fig. 14.6. Here the ideal operational amplifier is assumed to draw no input current (*i.e.* it has an infinite input impedance), converting the difference between V and ground into the potential $f = g(V)$.

In this class of models the firing rate changes smoothly in response to a rapid change in I . For small steps in input current $g(V)$ is approximately linear. In the linear regime, the firing rate is given by convolving the input current with a first-order low-pass filter with time constant $\tau = RC$, and is therefore a smoothed version of the input.

14.3.1 Comparing the Dynamics of a Spiking Cell with a Firing Rate Cell

As we mentioned above, the voltage in an integrate-and-fire unit in response to a sustained suprathreshold input moves along a limit cycle. This is not true for the potential in a firing rate model. V reaches its equilibrium value in a time dictated by τ ; the firing rate follows V without further delay.




Figure 14.6: A CONTINUOUS FIRING RATE SINGLE CELL MODEL

As in the integrate-and-fire models (Fig. 14.2), the input current in a firing rate neuron charges up a RC element. The instantaneous firing rate is obtained by passing the membrane potential V through a smooth, stationary nonlinearity g . In the circuit idiom used here, this nonlinearity is implemented by an ideal operational amplifier. The associated adapted f-I curve, specified by eq. 4.17 $g(V) = g(RI)$ is a continuous, positive and saturating function with no threshold.

If τ is increased in either model by increasing the neuron input resistance R , the rate of change of the subthreshold voltage also increases. This allows the spiking neuron to reach threshold more quickly (Fig. 14.7). In contrast, it will take longer to reach $1-1/e$ of the equilibrium voltage because the equilibrium voltage is also increased, implying that the dynamics in the firing rate model slows down as τ increases (Fig. 14.7). In the extreme case, when $\tau \rightarrow \infty$ and the leaky integrate-and-fire model converges to a perfect integrator model, the firing rate model does not even asymptotically approach an equilibrium, while the spiking unit settles into its steady-state limit cycle faster than it does for a finite τ . Similarly, decreasing τ will not make a spiking neuron respond faster.

A spiking mechanism can therefore speed up a neuron's response to a step change in the input. In a population of spiking cells with uniformly distributed initial conditions, there will be cells whose T_{th} will be arbitrary close to zero. Knight (1972a) made this point, proving rigorously that an infinite population of either perfect or leaky integrate-and-fire units uniformly distributed in phase will respond instantaneously to any suprathreshold stimulus. This is true for increases or decreases in input. In theory, how fast a neuron that

receives input from such a population could detect a change in its inputs is limited only by the number of statistically independent inputs (Panzeri, Biella, Rolls, Skaggs and Treves, 1996). This is unlike the response of one or even of a population of firing-rate neurons. Because of the low-pass filtering stage, their firing rate cannot change instantaneously.

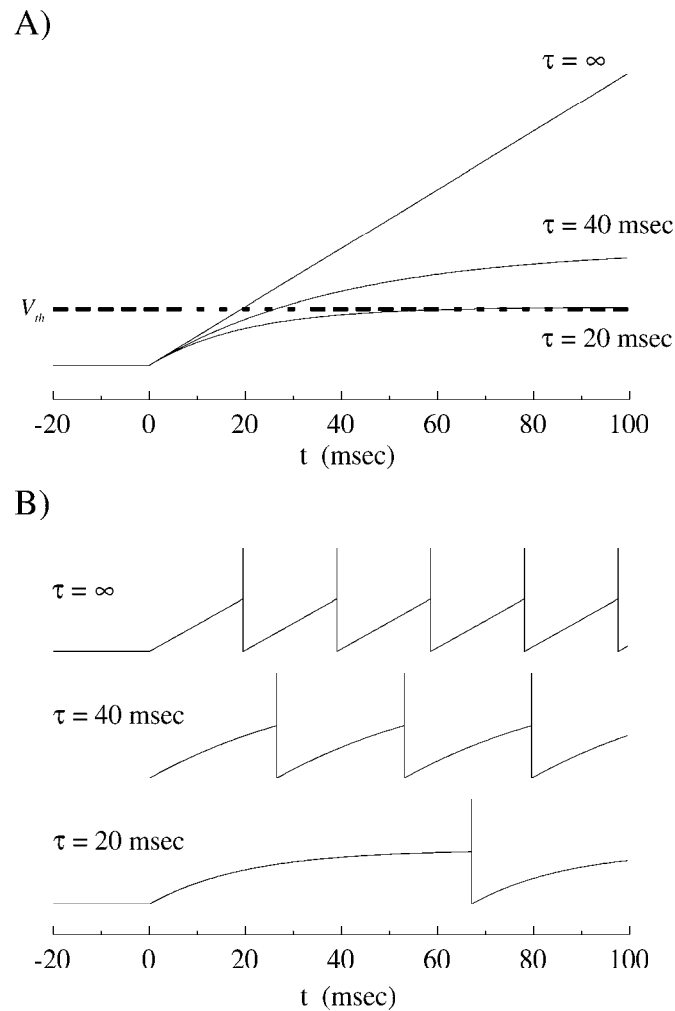


Figure 14.7: RESPONSE TIME IN SPIKING AND FIRING RATE MODELS

Effect of changing the membrane time constant on (A) a non-spiking (or firing rate) neuron and (B) on an integrate-and-fire neuron. The input current to the cell changes from zero to 0.85 nA at time 0. Subthreshold parameters were the same in both panels ($C = 1$ nF; R was varied from 20 M Ω to infinity; no refractory period). The subthreshold voltage of the integrate-and-fire model (B) is exactly the same as the non-spiking model (A) until the threshold V_{th} (dashed line in A) is crossed. Increasing R increases the rate of change of voltage, but also increases the equilibrium voltage. Non-spiking neurons therefore converge more slowly as the time constant increases. As $\tau \rightarrow \infty$ (when the leaky integrate-and-fire unit turns into a perfect one) the integrate-and-fire model responds earlier.

It has been argued that the temporal dynamics of neurons embedded within a network is primarily determined by the time course of the synaptic currents. This has as a consequence that the subthreshold RC time constant should more properly be replaced by one or more synaptic time constants (e.g. Amit and Tsodyks, 1991; Amit and Brunel, 1993; Burkitt, 1994; Suarez, Koch and Douglas, 1995; Brunel, 1996). A very simple, yet physiological correct way to express the firing rate is by

$$f = h_{ss}(I(t)), \quad (14.18)$$

where h_{ss} is the steady-state f-I discharge curve of the cell and $I(t)$ is the total current flowing into the soma. This equation dispenses with the subthreshold voltage since it only plays a very limited role for suprathreshold currents. The subthreshold domain will come into play for inputs hovering just around threshold (e.g. the lower trace in Fig. 14.7B).

In order to turn this into a complete single cell model, two additional ingredients are needed (Abbott, 1994).

- (i) How to obtain the current $I(t)$? The simplest manner is to use the standard neural network “linear sum over all inputs” formulation (see eq. 14.20). We can be more accurate and incorporate the nonlinear effects occurring in the dendritic tree (synaptic interaction, synaptic saturation and so on). This is discussed at length in section 18.4. Almost any degree of biophysical complexity can be accounted for. What can’t be readily included are back-propagating action potentials and the like (chapter 19) since we assume that the coupling between dendrites and the spike-initiation zone is one-way only.
- (ii) What is the relationship between the instantaneous firing rate $f(t)$ and the steady-state rate? The net effect of synaptic time constants and adaptation currents can all be lumped into a single, low-pass temporal filter that can, from a phenomenological point-of-view, be described by a first-order differential equation

$$\tau_{eff} \frac{df(t)}{dt} = h_{ss}(I(t)) - f(t). \quad (14.19)$$

where τ_{eff} is the above mentioned effective time constant (in the 20 to 30 msec range, in particular if one wants to account for both NMDA as well as adaptation currents). This time-constants is not related to the passive membrane time constant τ_m . This constitutes an alternative procedure to eqs. 14.15 and 16 to define a firing rate model (Abbott, 1994). It gives rise to the same phenomenological equation but with a different physiological interpretation.

We conclude that because the discharge rate is often quite linear over the relevant range of firing rates (Granit, Kernell and Shortess, 1963; Ahmed *et al.*, 1993), a *linear* threshold unit with no threshold may often be a satisfactory approximation for the firing rate of a real neuron.

14.4 Neural Networks

Neural networks consist of a large number of neurons connected using a scalar *synaptic weight* w_{ij} . Almost always, the single cell model used is a mean rate one, from the earliest publications of Steinbuch (1961) to Wilson and Cowan (1972), Hopfield (1984) and more recent work (summarized in Arbib, 1995).

14.4.1 Linear Synaptic Interactions are Common to Almost All Neural Networks

The evolution of the network, often termed *neurodynamics*, is governed by a coupled system of single cell equations. For any one cell i (out of n such cells) it takes the form of

$$C \frac{dV_i(t)}{dt} = -\frac{V_i(t)}{R} + I_i + \sum_{j=1}^n w_{ij} f_j(t), \quad (14.20)$$

where f_j is the firing rate of the j -th neuron and I_i corresponds to the external current injected into the cell. f_i is related to V_i via eq. 14.16, that is via a stationary nonlinearity

$$f_i(t) = g(V_i(t)). \quad (14.21)$$

In our usual circuit idiom (Fig. 14.8) in which both f_i and V_i are voltages, the synaptic coupling between the two neurons, *i.e.* w_{ij} , has the dimension of a conductance. Negative, hyperpolarizing synapses are implemented by inverting the output of the amplifier. The evolution of the circuit in Fig. 14.8 can be expressed by an equation of the form shown in eq. 14.20.

Different from the biophysics of synaptic conductance inputs (section 4.5), a change δf_j in the firing activity of the j -th presynaptic neuron leads to a change $w_{ij} \delta f_j$ in the *current* delivered to the operational amplifier. In effect, the synaptic inputs act as current sources and no nonlinear interaction among synaptic inputs exists.

Firing-rate models incorporating a low-pass filter to capture the passive properties of the underlying membrane have been applied widely in abstract neural network analysis (Wilson and Cowan, 1972; Cohen and Grossberg, 1983; Hopfield, 1984; Arbib, 1995; an excellent textbook covering this area in a relative intuitive manner is the one by Hertz, Krogh and Palmer, 1991) and in the analysis of the dynamics of real neurobiological networks (Knight, Toyoda and Dodge, 1970; Abbott, 1991; Traub and Miles, 1991; Amit and Tsodyks, 1991; Carandini and Heeger, 1994).

The study of the evolution of networks of spiking neurons—usually either of the integrate-and-fire or of the Hodgkin-Huxley variety—has only began of late due to the discontinuous nature of their output and the attendant mathematical difficulties (Hansel and Mato, 1993; van Vreeswijk and Abbott, 1993; Usher, Stemmler, Koch and Olami, 1994; van Vreeswijk and Sompolinsky, 1996; Maass, 1996). Given the revival of interest in the role of spike timing in computation, more progress is likely to be just around the corner. The equations




Figure 14.8: TWO INTERACTING FIRING CELLS IN A NEURAL NETWORK

Circuit model of two interacting mean rate neurons. Such continuous output units constitute the standard working horse of neural networks. Common to all is that the coupling among neurons is characterized by a scalar w_{ij} that can take on any real value, depending on whether the synapse is inhibitory or excitatory. The interaction among synaptic inputs is strictly linear. Local learning rules of the type discussed in section 13.6 are used for determining the amplitude of the w_{ij} 's. A qualitatively very similar model of linear synaptic interactions has been used in the neural network community for studying the computational power of networks of spiking units.

of motion of such a network are remain governed by linear synaptic interactions that usually take the form of

$$C \frac{dV_i(t)}{dt} = -\frac{V_i(t)}{R} + I_i + \sum_{j=1}^n w_{ij} \sum_k \delta(t - t_{jk}), \quad (14.22)$$

with an auxiliary equation for the voltage to deal with the refractory period (or adaptation) following spike generation (Usher *et al.*, 1994). The second sum on the right hand side includes all times t_{jk} at which the j -th neuron generated a spike. Each time this happens all of the postsynaptic targets of j receive a current pulse of amplitude w_{ij} (propagation delays can easily be incorporated into this notation).

14.4.2 Multiplicative Interactions and Neural Networks

The bottom line of the previous section is that the vast majority of neural networks have been built on the assumption of linearity of synaptic interactions. The nonlinearity that is necessary for any true information processing to occur resides solely in the firing mechanism at the output of the cell (V_{th} for spiking models and g for rate neurons).

In view of the many different types of nonlinear interactions that can occur in the dendritic tree, including AND-NOT interactions (section 5.1), NMDA synapses (section 5.2) and the voltage-dependent calcium and sodium membrane conductances found in the dendritic tree as discussed at length in chapter 19, it appears wise to cast around for some canonical single cell model that captures some of these nonlinearities yet is simple enough to be still amenable to analysis.

From a computational point of view, the simplest nonlinearity is *multiplication*, as in

$$\mathcal{M}(x_1, x_2) = \alpha x_1 \cdot x_2, \quad (14.23)$$

(with $\alpha \neq 0$). A substantial body of evidence supports the presence of multiplicative-like operations in the nervous system (Koch and Poggio, 1992). Physiological and behavioral data strongly suggest that the optomotor response of insects to moving stimuli is mediated by a correlation-like operation. Psychophysical work in humans from a number of independent groups strongly supports models of the correlation type, albeit with spatio-temporal filters different from those in insects. Simple cells recorded in the primary visual cortex of the cat (Emerson, Bergen and Adelson, 1992) appear to encode some of the stages in the most popular of these algorithms, the *spatio-temporal energy* model of motion perception (Adelson and Bergen, 1985). From a mathematical point of view, all these motion algorithms can be implemented by multiplication among two inputs (Buchner, 1984; Hildreth and Koch, 1987; Poggio, Yang and Torre 1989; Koch and Poggio, 1992).

Another instance of a multiplication-like operation in the nervous system is the modulation of the receptive field location of neurons in posterior parietal cortex by the eye and head position of the monkey (Andersen, Essick and Siegel, 1985; Zipser and Andersen, 1988; Van Opstal and Hepp, 1995; Brotchie, Andersen, Snyder and Goddman, 1995). This operation serves to transform the image from a retinal coordinate system into one that takes eye and body position into account. A final example is the output of an identified neuron—the *descending contralateral movement detector*—in the visual pathway of the locust that signals rapidly approaching objects. Its firing rate can be accurately described as the product of the angular image velocity and a term that depends exponentially on the angular size of the approaching object on the animal's retina (Hatsopoulos, Gabbiani and Laurent, 1995).

The generalization of multiplicative algorithms are *polynomial* ones. The output of a polynomial cell consists of the sum of contributions from a set of products:

$$\mathcal{P}(\mathbf{x}) = a_1 + b_1 x_1 + b_2 x_2 + c_1 x_1^2 + c_2 x_1 x_2 + \cdots d_1 x_1^2 x_2 \cdots \quad (14.24)$$

where $\mathbf{x} = (x_1, x_2, \dots)$ represents the synaptic input and the scalar $\mathcal{P}(\mathbf{x})$ the output of the unit. Single cell models that implement such a function are known as *sigma-pi* or *higher-*

*order*³ units (Feldman and Ballard, 1982; Volper and Hampson, 1987; Mel, 1992). If the highest order product in eq. 14.24 is quadratic, that is only terms such as $x_i x_j$ and x_i^2 are represented and d_1 and all higher coefficients are all identical to zero, the neuron is known as a *second-order* or *multiplicative* unit.

The evolution of the membrane potential is identical to the one used in the linear threshold ones. For a second-order multiplicative unit it is

$$C \frac{dV_i(t)}{dt} = -\frac{V_i(t)}{R} + I_i + \sum_{j=1}^n w_{ij} f_j(t) + \sum_{j=1}^n \sum_{k=1}^n w'_{ijk} f_j(t) f_k(t). \quad (14.25)$$

Associated with this single cell model are second-order “synapses” w'_{ijk} . Such a synapse only contribute to the postsynaptic potential if both inputs j and k are simultaneously active. In principle, up to n^2 synapses exist per neuron. For real neurons with dendritic trees, the connectivity matrices w_{ij} and w'_{ijk} are not independent of each other but are functions of the cable properties and of the specific synaptic architecture used. In general, they are also functions of time.

As an illustrative example, let us re-interpret the interaction occurring between an excitatory and a silent inhibitory synapse in a passive dendritic tree of a direction-selective cell—assuming that the conductance inputs g_e and g_i are small—in terms of a multiplicative unit (section 5.1). We assume that the excitatory postsynaptic conductance change g_e is proportional to the presynaptic firing frequency f_e of the neuron e and that the amplitude of the silent inhibition g_i is proportional to the firing frequency f_i of an inhibitory cell synapsing onto the cell. With $E_i = 0$, we can re-express the steady-state potential of the direction selective cell j in eq. 5.13 as

$$V_j = w_{je} f_e - w_{jee} f_e^2 - w_{jei} f_e f_i. \quad (14.26)$$

The synaptic strengths are specified in terms of the input and transfer resistances and batteries as $w_{je} = E_e \tilde{K}_{es}$, $w_{jee} = E_e \tilde{K}_{es} \tilde{K}_{ee}$ and $w_{jei} = E_e \tilde{K}_{is} \tilde{K}_{ie}$. Note the unconventional nature of the dimensions of f and w . Equation 14.26 directly illustrates the connection between the nonlinear interaction among synapses and multiplicative neural network units.

Mel (1992, 1993, 1994) argues persuasively that NMDA synaptic input in combination with voltage-dependent sodium and calcium conductances distributed throughout the dendritic tree implements something akin to eq. 14.24 (section 5.2). Synaptic input distributed in spatial clusters throughout the dendritic tree approximates a polynomial, in the sense that simultaneous excitation of m (where m can be large) neighbouring synapses causes a larger somatic response than activation of a similar number of synapses distributed in a diffuse fashion throughout the tree (Figs. 5.7–5.9). Different from eq. 14.26, this operation is both more robust and less specific, since the absence of any one particular input will have little effect on the overall output (as indicated by the broad peak in Fig. 5.8B).

It is easy to see the power of polynomial units in the case of binary, Boolean functions and circuits. For instance, a single such unit can implement an exclusive-or function, something

³*Order* refers to the maximal number of variables multiplied in each term (three in the case of eq. 14.24).

that a single-layer neural network of linear threshold units is unable to do (for a rigorous investigation of this, see Bruck, 1990; Bruck and Smolensky, 1992; for background material see Koch and Poggio, 1992). Of course, neurons always have a second type of nonlinear operation available to them in the form of a nonlinear mapping g (eq. 14.16) or h_{ss} (eq. 14.18).

14.5 Recapitulation

We started off this chapter by defining the instantaneous firing frequency $f(t)$. It is a fictive variable that can be obtained by averaging the spiking response of a single neuron to multiple presentations of the same stimulus. Due to the stochastic nature of the neuronal response the exact microstructure of spike trains are rarely reproducible from trial to trial. This is why the temporal average of the firing rate, $\langle f(t) \rangle$, is the most common variable measured during neurophysiological experiments. In a handful of experiments $\langle f \rangle$, evaluated over fractions of a second or longer, has been directly related to the behavior of the animal.

There is no question that firing rate codes that preserve temporal information at the 5 to 10 msec level are used in the nervous system. To what extent more complex correlation codes—exploiting information encoded among n -tuples of spikes from one neuron, or spiking information across multiple neurons—exist remains a subject of considerable debate.

What is the simplest model of a single neuron that captures some of its key operations? Two family of models are in common use today: integrate-and-fire and continuous firing rate models. While the former retains the timing information of individual action potentials the latter assumes that it is only the average or mean firing rate of a neuron that matters to its postsynaptic targets. Both neglect the dendritic tree and both eliminate the complex time course of the sodium and potassium membrane conductances underlying spiking.

The key insight behind the various guises of the integrate-and-fire model is that from a phenomenological point of view the neuron possesses two domains of operation, a sub- and a supra-threshold one. In the subthreshold domain, synaptic inputs are integrated and decay away; their temporal evolution is governed by the time constant τ . Once the voltage threshold is reached, a pulse is generated and the membrane potential is reset. Different versions of integrate-and-fire models, incorporating various mechanisms to account for adaptation, can be well fitted to the discharge curves of cortical and other cells. It will be argued in section 17.3 that firing in response to fast synaptic input in complex, conductance-based single cell models is, indeed, initiated whenever a voltage threshold is exceeded.

In response to a suprathreshold stimulus, these units, in accordance with their biological counterparts, can spike in a time $T_{th} \ll \tau$. A network of integrate-and-fire units can respond almost instantaneously to a stimulus. The take home lesson is that the dynamics of the subthreshold domain do not carry over into the suprathreshold domain.

In firing rate neurons the continuous output variable is an instantaneous function of the voltage. Since the evolution of the voltage is dictated by a time constant, the firing rate will always be low-pass filtered with respect to the input current, distinct from the response of real neurons and different from integrate-and-fire units. If one would like to retain the

continuous nature of the firing rate model, a more physiologically correct way to achieve this would be to make the steady-state firing rate a function of the total current (synaptic, dendritic or otherwise) at the cell body. The output of such a neuron can be interpreted as the firing rate associated with a population of spiking cells.

At the heart of the vast majority of neural networks lies the assumption that synaptic inputs interact in a linear manner. The nonlinearity that is necessary for computation is relegated to the firing mechanism at the output. A biophysical more faithful and more complex model that incorporates multiplicative interactions among synaptic inputs is the polynomial or sigma-pi unit.

Multiplication is a key operation underlying many neuronal operations. Chapter 5 treated the evidence in favor of the view that a dendritic tree endowed with NMDA synapses and voltage-dependent membrane conductances (see chapter 19 for this) can implement a robust version of such a polynomial unit. The nonlinear operations underlying the polynomial interactions do *not* depend on the threshold occurring at the cell body but precede it. The computational power of such neurons is considerably beyond that of their feeble-minded linear threshold counterparts.