Integrate-and-Fire Neurons and Networks*

Wulfram Gerstner
EPFL, Laboratory of Computational Neuroscience, CH-1015 Lausanne EPFL

Introduction

Most biological neurons communicate by short electrical pulses, called action potentials or spikes. In contrast to the standard neuron model used in artificial neural networks, integrate-and-fire neurons do not rely on a temporal average over the pulses. In integrate-and-fire and similar spiking neuron models, the pulsed nature of the neuronal signal is taken into account and considered as potentially relevant for coding and information processing. In contrast to more detailed neuron models, integrate-and-fire models do not describe explicitly the form of an action potential. Pulses are treated as formal events. This is no real drawback, since, in a biological spike train, all action potentials of a neuron have roughly the same form. The time course of an action potential does therefore not carry any information.

Integrate-and-fire and similar spiking neuron models are phenomenological descriptions on an intermediate level of detail. Compared to other SINGLE-CELL MODELS, they offer several advantages. In particular, coding principles can be discussed in a transparent manner. Moreover, dynamics in networks of integrate-and-fire neurons can be analyzed mathematically. Finally, large systems with thousands of neurons can be simulated rather efficiently. Reviews of integrate-and-fire networks can be found in Maass and Bishop (1998) or in Gerstner and Kistler (2002).

*To appear in The Handbook of Brain Theory and Neural Networks, Second edition, (M.A. Arbib, Ed.), Cambridge, MA: The MIT Press, 2002. http://mitpress.mit.edu

Spiking neuron models

Integrate-and-fire model

In its simplest form an integrate-and-fire neuron i consists of a resistor R in parallel to a capacitor C driven by an external current I_i. The voltage u_i across the capacitor is interpreted as the membrane potential. The voltage scale is chosen so that $u_i = 0$ is the resting potential. The temporal evolution of u_i is

$$\tau_m \frac{du_i}{dt} = -u_i + R\, I_i(t) \tag{1}$$

where $\tau_m = RC$ is the membrane time constant of the neuron. Spikes are formal events. We say that neuron i has fired a spike if u_i reaches at a time $t = t_i^f$ a threshold ϑ. The form of the action potential is not described explicitly. Immediately after spike firing, the potential u_i is simply reset to a value $u_{\rm reset} < \vartheta$. Integration of (1) is then resumed with $u_{\rm reset}$ as initial condition (Stein, 1967). Since the spatial structure of the neuron is neglected, such a model is also called a point model. (SINGLE CELL MODELS).

In a network of neurons, the input I_i to neuron i is due to the spikes of presynaptic neurons j. Detailed models of synaptic input can be found in SYNAPTIC CURRENTS. In the simplest model of a synapse, each presynaptic spike arrival evokes a postsynaptic current with a standard time course α. The total input to neuron i is then

$$I_i = \sum_{j,f} w_{ij}\, \alpha(t - t_j^f) \tag{2}$$

where the sum runs over all firing times t_j^f of all presynaptic neurons. The factor w_{ij} is the synaptic efficacy of a connection from a presynaptic neuron j to a postsynaptic neuron i. Choices for the postsynaptic current include a delayed δ-pulse, $\alpha(s) = \delta(s - \Delta^{\rm ax})$, or a double exponential, $\alpha(s) = [e^{-(s-\Delta^{\rm ax})/\tau_1} - e^{-(s-\Delta^{\rm ax})/\tau_2}]/(\tau_1 - \tau_2)$, where $\Delta^{\rm ax}$ is the axonal transmission delay and τ_1, τ_2 are synaptic time constants.

Spike Response Model

The integrate-and-fire equation (1) with the synaptic current (2) can be integrated, either numerically or analytically. Since it is a linear equation, the

analytical integration can be done for each term in the sum of eq. (2) separately. The total membrane potential is then the sum of all the postsynaptic potentials (PSPs) caused by presynaptic firing plus the refractory effect of a negative reset potential. Given the last firing time $\hat{t}_i$ of neuron i, the result of the integration is therefore of the form ($t > \hat{t}_i$)

$$u_i(t) = \eta(t - \hat{t}_i) + \sum_{j,f} w_{ij}\, \epsilon(t - \hat{t}_i, t - t_j^f)\,. \tag{3}$$

The next firing of i occurs, if the membrane potential u_i approaches the threshold ϑ from below. Eq. (3) defines the dynamics of the Spike Response Model (SRM). It has been introduced above as an integrated version of the integrate-and-fire model, but the SRM is in fact more general (Fig. 1). The function η describes the action potential at $\hat{t}_i$ and the spike afterpotential that follows. The function ϵ describes the voltage response of neuron i to a presynaptic spike at t_j^f. Let us suppose that the last spike of the postsynaptic neuron i was far back in the past ($t - \hat{t}_i \to \infty$). Then $\epsilon(\infty, s)$ as a function of s describes the time course of the PSP caused by a presynaptic spike. If the postsynaptic neuron i has been active in the recent past, then a presynaptic spike is less effective in exciting a postsynaptic response. The first argument of $\epsilon(t - \hat{t}_i, t - t_j^f)$ describes the dependence upon the recent firing history of the postsynaptic neuron. With an appropriate choice of the functions ϵ and η, about ninety percent of the firing times of the Hodgkin-Huxley model with time dependent input can be correctly predicted by the SRM, with a precision of ± 2 ms (Kistler et al., 1997). Moreover the spatial structure of neurons with linear dendritic tree can be incorporated by an appropriate choice of ϵ. For synapses which are farther out on the dendritic tree, the PSP, and hence the function ϵ rises more slowly.

Noise

Biological neurons which are driven by a time-dependent intracellular current exhibit a reliable, (nearly) deterministic behavior, just as the models (1) or (3). On the other hand, neurons which are part of a cortical network emit spikes at irregular intervals. Since the exact spike times cannot be controlled by the experiment, the irregularity is interpreted as noise.

Formally, noise can be introduced into the integrate-and-fire model by adding a fluctuating input $\sigma\,\xi_i(t)$ on the right-hand side of (1) where σ

is a parameter controlling the amplitude of the noise and ξ is a normally distributed random variable with zero mean (DIFFUSION MODELS OF NEURON ACTIVITY). In the presence of noise, we may ask the following question. Given the last firing time $\hat{t}_i$ of neuron i and the input current $I_i(t')$ for $t' > \hat{t}_i$, what is the probability that the next spike occurs around time t? The answer is given by the conditional interval distribution $P(t|\hat{t}_i, I(.))$. The calculation of $P(t|\hat{t}_i, I(.))$ for the diffusion model is equivalent to the solution of a first passage time problem. The general solution to this problem is not known.

Noise can also be introduced in a different manner into spiking neuron models. The voltage $u_i(t)$ is calculated according to (1) or (3). Even before u_i reaches the threshold ϑ, neuron i may fire with an 'escape rate' $\rho(t)$ which depends on the momentary distance from threshold and possibly also on the current input I, viz., $\rho(t) = h(u(t) - \vartheta; I(t))$. In this case, an explicit expression for the conditional interval distribution is known, viz.,

$$P(t|\hat{t}_i, I(.)) = \rho(t) \, \exp\left[- \int_{\hat{t}_i}^{t} \rho(t') \, dt' \right] . \tag{4}$$

With an appropriate choice of the escape function h, the diffusion model can be approximated by the escape model to a high degree of accuracy For a review of noise models, see Gerstner and Kistler (2002), Ch. 5.

Network dynamics and population equations

In many areas of the brain, neurons are organized in groups of cells with similar properties, e.g., pools of motoneurons or columns in the visual cortex. Instead of looking at the firings of individual neurons, we may simply be interested in the fraction of neurons which are active in the population. In each small time window Δt, let us count the number of spikes $n_{\rm sp}(t; t + \Delta t)$ which are emitted across the population and divide by the number N of neurons and Δt. This procedure defines the population activity or population rate

$$A(t) = \lim_{\Delta t \to 0} \frac{n_{\rm sp}(t; t + \Delta t)}{N \, \Delta t} = \frac{1}{N} \sum_{j,f} \delta(t - t_j^f) \, , \tag{5}$$

where δ is the Dirac δ function and the sum runs over all spikes of all neurons in the population. The population activity has units of one over time and can

be seen as the rate at which the total spike count increases. Note that the definition of the population rate (5) does not involve a temporal, but only a spatial average. What is the temporal evolution of $A(t)$ in a (homogeneous) network of spiking neurons?

The state of each neuron depends on its input *and* on the time $\hat{t}$ of its last spike, see (3). We define a *homogeneous* population by the conditions that (i) lateral coupling has a fixed value $w_{ij} = w_0/N$, and (ii), external inputs $I^{\rm stim}(t)$ are the same for all neurons. The total input to any neuron in the network is therefore

$$I(t) = w_0 \int_0^\infty \alpha(s)\, A(t-s)\, ds + I^{\rm stim}(t) \tag{6}$$

Even though they all receive the same input, different neurons will, in general, have different firing times $\hat{t}$. A neuron which has fired its last spike at $\hat{t}$ and has received an input $I(t')$ for $t' > \hat{t}$ will contribute with weight $P(t|\hat{t}, I(.))$ to the population activity at time t. Hence the expected value of the population activity at time t is

$$A(t) = \int_{-\infty}^{t} P(t|\hat{t}, I(.))\, A(\hat{t})\, d\hat{t}\,. \tag{7}$$

For spiking neurons with escape noise $\rho(t)$, $P(t|\hat{t}, I(.))$ is given by eq. (4) and therefore highly non-linear. Equation (7) is implicitly contained in (Wilson and Cowan, 1972; Knight, 1972) and formally derived in (Gerstner, 2000) for a homogeneous, fully connected network of spiking neurons in the limit of $N \to \infty$.

In their 1972 paper, Wilson and Cowan proposed to transform the integral equation (7) into a differential equation of the form

$$\tau \frac{d}{dt} A(t) = -A(t) + g\left[w_0 \int_0^\infty \alpha(s)\, A(t-s)\, ds + I^{\rm stim}(t)\right] \tag{8}$$

where τ is a time constant, w_0 is the neuronal coupling strength, $I^{\rm ext}(t)$ is a stimulus, and g a nonlinear transfer function. One of the problems of eq. (8) is that the time constant τ is the result of a process of 'time coarse-graining' which is necessary for the transition from eq. (7) to (8). Sine the time window of course graining has to be defined somewhat arbitrarily, the time constant τ is basically ad hoc. Because of the problems inherent in eq. (8), it is preferable to work directly with eq. (7).

For the diffusion noise model, eq. (7) is valid but not very useful, because the conditional interval distribution $P(t|\hat{t}, I(.))$ is not known. As an alternative to eq. (7), the state of the population can be described by the distribution of membrane potentials $P(u,t)$ (Abbott and van Vreeswijk, 1993; Brunel 2000; Nykamp and Tranchina 2000). At each moment of time $P(u,t)\,\Delta u\,N$ gives the number of neurons in the population with a membrane potential between u and $u+\Delta u$. The equation of the integrate-and-fire model (1) with additive diffusion noise $\sigma\,\xi(t)$ can be transformed into a Fokker-Planck equation for the distribution of membrane potentials:

$$\tau\,\frac{\partial P(u,t)}{\partial t} = \frac{\sigma^2}{2\tau}\,\frac{\partial^2 P(u,t)}{\partial u^2} + \frac{\partial}{\partial u}\,\{[u - R\,I(t)]P(u,t)\}\;. \tag{9}$$

The threshold is treated as an absorbing boundary so that the probability density vanishes for $u \geq \vartheta$. The probability current across threshold equals the population activity

$$A(t) = \frac{\sigma^2}{2\tau^2}\,\frac{\partial P(u,t)}{\partial u}|_{u=\vartheta}\,. \tag{10}$$

Since the membrane potential of active neurons is immediately reset to $u_{\rm reset}$, the population activity $A(t)$ acts a source of probability current at $u = u_{\rm reset}$. For a review, see Gerstner and Kistler (2002), Ch. 6.

Application to Coding

Integrate-and-fire models can be used to discuss potential principles of coding and dynamics in a transparent manner (Maass and Bishop, 1998, Chs. 1,2,10-14). Before we turn to networks, let us start with two examples of coding on the single-neuron level.

Signal encoding by single neurons

Coherent input is more efficient than incoherent spikes in driving a postsynaptic neuron. To see why, let us consider the SRM (3). For the sake of simplicity, we assume that the postsynaptic neuron i was inactive in the recent past ($t < 0$) and receives, for $t > 0$, input from two presynaptic neurons $j = 1, 2$, both firing at 100 Hz. We set $w_{i1} = w_{i2} = w_0$. According to eq. (3),

each input spike evokes a postsynaptic potential $\epsilon(-\infty, t - t_j^f)$ where t_j^f is one of the firing times of neuron j. If the two spike trains are out of phase, the summed postsynaptic potential is lower than in the synchronous case (Fig. 2). By an appropriate choice of the threshold ϑ, an output spike of the postsynaptic neuron i occurs therefore only in the coherent (or 'coincident') case. Quite generally, coincidence detection is possible if the threshold of the postsynaptic neuron is slightly above the mean value, the membrane potential would take for asynchronous input (Koenig et al., 1996; Kempter et al., 1998). In the auditory system, it is commonly accepted that coincidence detection is used for the localization of sound sources. On the other hand, it is an open question whether *cortical* neurons operate in the regime of coincidence detection (Koenig et al., 1996). (SINGLE CELL MODELS)

Coding by homogeneous populations

Spiking neurons connected to each other by excitatory or inhibitory synapses exhibit non-trivial dynamical properties. The population may respond rapidly to external signals. The network activity may explode or die away. Neurons may spontaneously develop a tendency to fire synchronously or in groups. All of these phenomena which can potentially be the basis of various coding schemes, can be understood from an analysis of equations (6) - (10). Some of the fundamental questions are highlighted in the following.

First, is it possible, in the absence of an external stimulus, to stabilize a population of spiking neurons at a reasonable level of spontaneous activity? For $N \to \infty$, spontaneous activity corresponds to a stationary solution $A(t) \equiv A_0$ of the population dynamics (7) or (10). Spontaneous asynchronous firing seems to be a generic feature of cortical tissue, but its role is still unclear. A stability analysis shows that without noise asynchronous firing is never stable. Thus the apparent noisiness of cortical neurons is a necessary feature of the system.

Even in the presence of noise, neurons often tend to synchronize their firings and develop collective oscillations. This observation leads to the second question: How is the frequency of collective oscillations related to neuronal parameters? It turns out that there are different oscillatory regimes depending on the form of the postsynaptic potential, the axonal delay, and the value of the threshold (Abbott and van Vreeswijk, 1993; Gerstner, 2000; Brunel, 2000). The frequency of the collective oscillation may be low (about that

of individual neurons) or several times faster. Collective oscillations and synchronization (Maas and Bishop, 1998, Chs. 10,11; Gerstner and Kistler, 2002, Ch. 12) have been suggested as potential coding schemes in cortex and hippocampus (SYNCHRONIZATION, BINDING, AND EXPECTANCY).

Third, how rapidly does the population activity $A(t)$ respond to changes in the input? An analysis of (7) shows that the response time is not limited by the membrane time constant of the neurons, but can be much faster (Gerstner, 2000). The fast response is due to the fact that, during spontaneous activity, there are always some neurons with a membrane potential just below threshold. A slight increase in the input will make those neurons fire immediately. The fast response of populations of spiking neurons to a new input could be important for an explanation of reaction time experiments (Thorpe et al., 1996). The same type of arguments also shows that populations of spiking neurons can reliably transmit signals which vary on a time scale that is short compared to the interspike intervals of a neuronal spike train, as is for example the case in the auditory pathway.

All of the above results hold true for homogeneous networks with either excitatory or inhibitory coupling. Formally the theory is valid for full connectivity in the limit of $N \to \infty$. It yields also an excellent approximation for networks with random connectivity if the density of connections is either very high or very low. An extension to mixed excitatory/inhibitory populations as found in the cortex is possible (Brunel, 2000).

Coding in structured networks

Structure in neuronal networks may arise from a spatial arrangement of neurons or from specific patterns stored in a distributed manner in the network.

In networks with local (or distance-dependent) excitatory connections, traveling waves may occur. In two-dimensional sheets of neurons, wave fronts may have planar or spiral shapes, similar to the ones found in reaction-diffusion systems. Collective oscillations and asynchronous firing are other possible network states. These effects can be described by a direct generalization of the theory of homogeneous systems to a spatially continuous population. Replace $A(t)$ in (7) by $A(x,t)$ where x is the spatial location. Instead of (6) we use $I(x,t) = \int dx' w(|x-x'|) \int ds\, \alpha(s)\, A(x', t-s)$ where $w(.)$ is the distance dependent coupling strength. Activity waves have been reported in slice cultures. It has also been suggested that similar activity

waves could account for some of the trial-to-trial variability in cortical spike train recordings.

In the previous example, neurons that are strongly connected are located next to each other. Activity spreads from one group of neurons to its neighbors which is easily recognizable by an external observer as a travelling wave of activity. Let us now keep the connections between the same neurons as before but move all neurons to a new random location on the two-dimensional sheet. Apart from the fact that connection lines are longer, nothing has changed. What used to be a propagating wave in the original spatial arrangement, now looks like asynchronous firing of neurons all over the sheet. Nevertheless, it is a specific nearly deterministic spatio-temporal spike pattern. These 'hidden' waves of activity has been termed a SYNFIRE CHAINS (Abeles 1991). While the existence and stability of synfire chains can be shown by simulation or analysis of model networks, this does not necessarily imply that real brains make use of synfire chains for coding.

Discussion

What is the code used by cortical neurons? What is signal, what is noise in neuronal spike trains? While the final answers to these questions have to come from additional experiments, modeling on the level of integrate-and-fire networks can contribute to answering, because models allow researchers to explore potential coding schemes and to identify relevant operating regimes.

In populations of integrate-and-fire neurons, a rate code can be a very fast code, if rate is defined by a population average ('population activity') rather than by a temporal average (Knight, 1972; Gerstner, 2000). In contrast to a widespread belief, the speed of signal transmission is not limited by the membrane time constant of the neuron. Moreover, with appropriate spike-based learning rules (Maass and Bishop, ch. 14), spiking neurons can work, in principle, at a very high temporal precision (Abeles, 1991). Large-scale simulations of integrate-and-fire networks provide a link between theory and experiments.

One of the points that has been stressed in recent models of integrate-and-fire neurons is the relevance of the subthreshold regime. If neuronal and network parameters are chosen so that the mean membrane potential stays just below threshold, then several interesting properties emerge. First,

neurons act as coincidence detectors. They are sensitive to fluctuations in the input and can therefore 'read out' the coherent aspects of the input signal (Koenig et al., 1996; Kempter et al., 1998). Second, neurons in this regime respond rapidly to changes in the input (Gerstner, 2000). This might be relevant to explain fast reaction times (Thorpe et al., 1996). Third, to stabilize a highly recurrent network of spiking neurons in the subthreshold regime, a certain amount of 'noise' is necessary (Abbott and van Vreeswijk, 1993; Gerstner, 2000). From that point of view, it comes as no surprise that cortical neurons appear to be 'noisy'. Whether this apparent noisiness is due to intrinsic noise sources in the neuronal dynamics, to noise in the synaptic transmission, or the result of deterministic chaos in a network, is not clear. Model studies have shown that 'noise' itself can arise as a network effect if neurons are in the subtreshold regime. While individual neurons behave more or less deterministically, the same neurons show large firing variability, when part of a random network of excitatory and inhibitory neurons with sparse connectivity (Brunel, 2000). Such networks can represent past input in their spatio-temporal firing pattern (TEMPORAL INTEGRATION IN RECURRENT CIRCUITS). Thus the study of integrate-and-fire networks may shed new light on burning questions of brain theory.

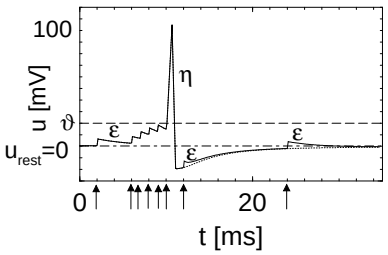

Fig. 1. Each input current pulse (arrows) evokes a postsynaptic potential with time course ϵ. If the sum of the postsynaptic potentials reaches the threshold ϑ, and action potential with time course η is triggered. An input current pulse immediately after the action potential evokes a reduced response because of refractoriness.

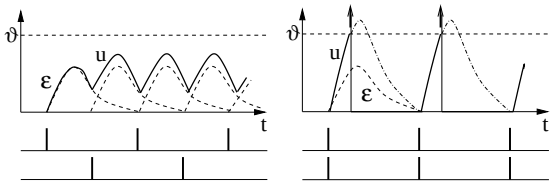

Fig. 2. Left: Spike trains from two different presynaptic neurons are phase shifted with respect to each other. The summed potential u does not reach the threshold. Right: Spikes from the same presynaptic neurons arrive synchronously so that u reaches the threshold ϑ and evokes the generation of output spikes (arrows). Afterwards u is reset (schematic figure).

References

Abbott, L.F. and van Vreeswijk, C. (1993). Asynchronous states in a network of pulse-coupled oscillators. Phys. Rev. E, 48:1483–1490.

(*) Abeles, M. (1991). Corticonics. Cambridge University Press, Cambridge.

Brunel, N. (2000). Dynamics of sparsely connected networls of excitatory and inhibitory neurons, Computational Neuroscience, 8, 183-208.

Gerstner, W. (2000). Population dynamics of spiking neurons: fast transients, asynchronous states and locking. Neural Comput., 12:43-89.

(*) Gerstner, W. and Kistler, W.M. (2002). Spiking neuron models: single neurons, populations, plasticity Cambridge University Press, Cambridge

Kempter,R., Gerstner, W., van Hemmen, J.L. and Wagner, H. (1998). Extracting oscillations: Neuronal coincidence detection with noisy periodic spike input. Neural Comput., 10:1987–2017.

Kistler, W.M., Gerstner, W. and van Hemmen, J.L. (1997). Reduction of Hodgkin-Huxley equations to a single-variable threshold model. Neural Comput., 9:1015–1045.

Knight, B.W. (1972). Dynamics of encoding in a population of neurons.
J. Gen. Physiology, 59:734–766.

(*) König, P., Engel, A.K., and Singer, W. (1996). Integrator or coincidence detector? the role of the cortical neuron revisited. TINS, 19(4):130–137.

(*) Maass, W. and Bishop, C. editors (1998). Pulsed Neural Networks. MIT-Press

Nykamp, D. and Tranchina, D. (2000). A population density approach that facilitates large-scale modeling of neural networks: Analysis and application to orientation tuning. J. Computational Neuroscience, 8:19-50.

Stein, R.B. (1967). Some models of neuronal variability. Biophys. J., 7:37–68.

Thorpe, S., Fize, D., and Marlot, C. (1996). Speed of processing in the human visual system. Nature, 381:520–522.

Wilson, H.R. and Cowan, J.D. (1972). Excitatory and inhibitory interactions in localized populations of model neurons. Biophysical J, 12:1–24.