Spiking Neurons

Pietro Mele (pietromele@yahoo.com)

June 12, 2005

Abstract

Some notes about spiking neurons to be used in a future project (Neuro3).

A few data about the human brain

- 10^4 cell bodies per cubic millimiter.

- Several kilometers of "wires" per cubic millimiter.

- A neuron in the cortex often makes connections to more than 10^4 postsynaptic neurons.

- Many axonal branches end in the direct neighborhooud of the neuron, but the axon can also stretch over several millimiters to connect to neurons in other areas of the brain.

- The duration of an action potential is typically in the range of 1-2 ms.

- Since the spikes of a given neuron look alike, the form of the action potential does not carry any information.

- The information is contained in the number and the timing of the spikes.

- It can happen that a presynaptic and a postsynaptic neuron are connected by more than one synapse (e.g. 5 or more).

- A single synapse generates a binary response to each spike of its presynaptic neuron.

Notation

A set of neurons: $I \quad (i, j \in I)$.

Firing time of neuron i: $t_i^{(f)}$

Spike train of neuron i: $\mathcal{F}_i = \left\{ t_i^{(1)}, ..., t_i^{(n)} \right\}$ where $t_i^{(n)}$ is the most recent spike.

Resolution in time measurement: Δt

Possible representations of a spike train:

- as a sequence of ones (spike) and zeros (no spike) at times Δt, $2\Delta t$, $3\Delta t$, ...
- ...

Presynaptic neurons j of neuron i: $\Gamma_i \subseteq I$

Delay between presynaptic spike arrival and postsynaptic firing: $s = t_j^{(f)} - t_i^{(f)}$

Synaptic weight: w_{ij}

Axonal transmission delay: Δ_{ij}

Response functions that model the postsynaptic potentials: $\epsilon_{ij} \left(t - t_i^{(f)} \right)$

Refractory function: $\eta_i \left(t - t_i^{(f)} \right)$

Heaviside step function: $\mathcal{H}(s) = 0$ for $s \leq 0$ and $\mathcal{H}(s) = 1$ for $s > 0$

Dirac δ function: $\delta(s) = 0$ for $s \neq 0$ and $\int_{-\infty}^{+\infty} \delta(s) ds = 1$

Coding strategies based on spike timing

Time-to-First-Spike

Consider a neuron that abruptly receives a new input at time t_0. With this hypothesys, the time the neuron takes to generate the first spike contains all the information about the stimulus; i.e. if the neuron fires shortly after t_0 could mean a strong stimulation, and viceversa. All following spikes would be irrelevant. This model is an idealization... or maybe not: there might not be enough time to evaluate more than one spike from each neuron per time step.

Phase

This coding strategy is identical to the time-to-first-spike, but, instead of an input signal at time t_0, there is background oscillation that provides an internal reference signal. The information is encoded in the phase of the spikes with respect to this oscillation.

Correlations and Synchrony

In this case, spikes from other neurons are used as the reference signal. Synchrony between a group of neurons could convey information which is not contained in the firing rates of the single neurons.

Stimulus reconstruction

A stimuls evokes a spike train of a neuron. This stimulus can be reconstructed (with a certain approximation) from the spike train itself. A simple, and good, way to do it is with a linear reconstruction with a specific kernel.

Spike Response Model (SRM)

In this section the "time-to-first-spike" model will be examined.

The state of neuron i is described by a state variable u_i. The neuron fires if u_i reaches a threshold ϑ from below. The firing time $t_i^{(f)}$is the moment when the threshold is reached. The set of all firing times of neuron i is:

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

where $t_i^{(n)}$ is the most recent spike (also indicated with $\hat{t}$).

The state variable u_i is influenced by:

- the excitatory or inhibitory input from presynaptic neurons, whose effect is described by the kernel ϵ_{ij};

- the spikes generated by neuron i itself: immediately after a spike, u_i is reset (adding a refractory function η).

So u_i can be described by the following equation:

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

One possible form of the kernels η and ϵ is:

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

$$\epsilon_{ij}(s) = \left[\exp\left(-\frac{s-\Delta_{ij}}{\tau_m}\right) - \exp\left(-\frac{s-\Delta_{ij}}{\tau_s}\right)\right]\mathcal{H}\left(s-\Delta_{ij}\right)$$

where τ, τ_m and τ_s are time constants, $\mathcal{H}(s)$ is the Heaviside step function, and Δ_{ij} is the axonal transmission delay. The synaptic weight w_{ij} scales the effect of the spikes, and is positive for excitatory synapses and negative for inhibitory synapses (consider that an excitatory synapse cannot become inhibitory, and viceversa). Choosing different kernels, it is possible to approximate other models (e.g. Hodgkin-Huxley).

Short term memory

It is possible to simplify the equation that gives u_i assuming that only the last firing contributes to refractoriness:

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

External input

A neuron may receive an analog input current $\mathcal{I}^{ext}(t)$ from a sensory input; considering a specific kernel $\tilde{\epsilon}$, the following term must be added in the expression that gives $u_i(t)$:

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

Learning

A spiking neuron can learn by updating the following parameters:
- its synaptic weights w_{ij}
- its axonal transmission delays Δ_{ij}
- its threshold ϑ.

Feeding and Linking Inputs

A variation of the standard spiking neuron has two different types of input signals:

- feeding input, coming from the input layer or from layers placed before the current one;
- linking input, coming thruogh feedback connections from higher layers of the network (e.g. an associative memory, ...).

These two inputs are controlled by different time constants, and are multiplied to produce the potential $u_i(t)$ of the corresponding neuron.

This model seems to be biologically plausible.

Network of Spiking Neurons

Let I be a set of neurons, and assume that the subsets $I_{input} \subseteq I$ and $I_{output} \subseteq I$ of neurons work as an interface with the external world. The I_{input} subset will also receive spikes from outside the network, while the I_{output} subset's spikes constitute the network output.

So a network of spiking neurons computes a function that maps a vector of time series $\langle \mathcal{F}_i \rangle_{i \in I_{input}}$ into another vector of time series $\langle \mathcal{F}_i \rangle_{i \in I_{output}}$.

Using a logarithmic encoding of the external sensory input could be a good way to treat data over a large range of values.

Remember that, for numerical reasons, one of the following two things has to be considered:

- weight normalization;
- bias. $\Leftarrow$

Learning

It is not clear which are the parameters that get changed to learn, and it is even not clear what the goal of learning is: learning a *function* (input vectors –> output vectors), or learning a *functional* (input functions –> output functions).

Synapses' parameters change depend on a series of factors, and this change happens in various forms depending on the time scale we consider.

Phenomenon	Duration	Locus of Induction
Short-term Enhancement		
Paired-Pulse Facilitation (PPF)	100 ms	Pre
Augmentation	10 s	Pre
Post-tetanic potentiation	1 min	Pre
Long-term Enhancement		
Short-Term Potentiation (STP)	15 min	Post
Long-Term Potentiation (LTP)	>30 min	Pre and Post
Depression		
Paired-Pulse Depression (PPD)	100 ms	Pre
Depletion	10 s	Pre
Long-term Depression (LTD)	>30 min	Pre and Post

The terms "Paired-pulse facilitation" and "Paired-pulse depression" are relative to the case where the stimulus consists of just two spikes, and the second spike causes a larger (smaller) postsynaptic response. The locus of induction indicates which places (presynaptic, postsynaptic, or both) have an influence on the synapse's learning.

Hebbian Learning

Hebbian learning is an unsupervised correlation-based adaptation mechanism:

"When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth process or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased."

LTP obeis Hebb rule. In the context of spiking neurons, a synapse is strengthened if pre- and postsynaptic neurons are "simultaneously" active, i.e. a presynaptic spike arrival and a postsynaptic firing occur in a certain time window.

Considering a synapse from neuron j to neuron i, let's define the spike trains in the following way:

presynaptic spike train: $S_j(t) = \sum_{t_j^{(f)} \in \mathcal{F}_j} \delta(t - t_j^{(f)})$

postsynaptic spike train: $S_i(t) = \sum_{t_i^{(f)} \in \mathcal{F}_i} \delta(t - t_i^{(f)})$

To be general, assume an arbitrary learning window $W(s)$ which depends on the time difference $s = t_j^{(f)} - t_i^{(f)}$. The change of the synaptic efficacy after a time $\mathcal{T}$ is:

$$\Delta w_{ij} = \int_0^{\mathcal{T}} \int_0^{\mathcal{T}} W(t'-t) S_i(t) S_j(t') dt\, dt' + \tilde{b}_i \int_0^{\mathcal{T}} S_i(t) dt + \tilde{b}_j \int_0^{\mathcal{T}} S_j(t) dt + \mathcal{T} \tilde{a}$$

The first term of the right-hand side forces simultaneity between pre- and postsynaptic firing through the learning window $W(s)$. The second and third terms are just the number of pre- and postsynaptic spikes that occur in the time window $0 ... \mathcal{T}$, without any correlation between them.

From experimental evidence, the following is a good approximation of the learning window $W(s)$:

$$W(s) = \begin{cases} [A_+ - A_-] \exp\left[-(s^* - s)/\tau^{syn}\right] & \text{for } s < s^* \\ A_+ \exp\left[-(s - s^*)/\tau_+\right] - A_- \exp\left[-(s - s^*)/\tau_-\right] & \text{for } s > s^* \end{cases}$$

The maximum value of $W(s)$ occurs for $s = s^*$, and to get stable learning s^* must be negative, i.e. the output spike must occur after the presynaptic spike arrival.

The parameters used so far depend on the kind of neuron we are considering. For neurons in the auditory pathway they will have the following values:

$A_+ = 0.5 \qquad A_- = 0.2$
$\tau^{syn} = 0.5\,\mathrm{ms} \qquad \tau_+ = 0.5\,\mathrm{ms} \qquad \tau_- = 5\,\mathrm{ms}$
$s^* = -0.05\,\mathrm{ms}$
$\tilde{a} = 0.0 \qquad \tilde{b}_i = 0.0 \qquad \tilde{b}_j = 0.1$

In the cortex and in the hippocampus the learning window is much wider, probably 50-200 ms.

The meaning of $W(s)$ is the following one: if a presynaptic spike arrives slightly *before* the postsynaptic firing, the synaptic efficacy is increased ($W(s) > 0$); if a presynaptic spike arrives a few milliseconds after the output spike, the synapse is weakened ($W(s) < 0$).

Since, using the previous equations, the value of the synaptic weights could grow to very high values, it is realistic to impose some upper and lower bounds, e.g. $(0, w^{max})$.

This model does not consider (all?) the dynamic effects of synapses.

Axonal transmission delays Δ_{ij} tuning by Hebbian learning

...

Stability-Plasticity dilemma

Learning and memory are based on long-lasting changes of synaptic efficacy. On the other hand, synaptic efficacy is a dynamical variable that changes in real time, depending on pre- and postsynaptic activity.

One way to implement this behaviour is to consider a synapse as a *bistable* system; i.e. a synapse dynamics is characterized by two stable and fixed states w_0 and w_1. In the absence of input spikes, the synaptic weight will settle down either at w_0 or at w_1. If pre- or postsynaptic activity has the necessary duration and amplitude, a transition can occur from one state to the other (long-lasting change). Otherwise, the synaptic weight will change to a value $\tilde{w}$, and it will immediately "decay" to its previous stable state (short term change).

Personal considerations

Spike train implementation

A possible way to represent a spike train for neuron i: using the sequence of bits contained, e.g., in a variable of type char or int or long, ... (depending on how much information about the past we want to keep trace of).

Every bit represents a “time slot”, and it contains a one if a spike occurs in that time frame, or a zero otherwise.

The most significant bit represents the most recent time slot.

At the next time step, the spike train of each neuron should be “translated” to the right (i.e. divided by 2), so that the oldest timeslot gets lost, and the space for a new timeslot is available in the most significant bit.

“Sender-oriented” vs. “receiver-oriented” connectivity ...

Appendix A - Synonims

Action potential, spike.

Facilitation, enhancement.

Synaptic weight, synaptic efficacy.

References

[1] Gerstner, W. (1999). Spiking Neurons.

[2] Gerstner, W., et al (1999). Hebbian Learning of Pulse Timing in the Barn Owl Auditory System.

[3] Gerstner, W., Kistler, W., (2002). Spiking Neuron Models. *Cambridge University Press.*

[4] Maass, W. (1999). Computing with Spiking Neurons.

[5] Maass, W., Zador, A.M. (1999). Computing and Learning with Dynamic Synapses.

[6] Marian, I., et al (2002). Efficient event-driven simulation of spiking neural networks.

[7] Gerstner, W., Kistler, W., (2002). Mathematical Formulations of Hebbian Learning.