Neural Networks and Biological Modeling

Professor Wulfram Gerstner
Laboratory of Computational Neuroscience

Question set 6

Exercise 1: Hopfield network with probabilistic update

So far we have studied Hopfield networks with deterministic activity dynamics. That is, for the same input potential h a neuron always takes the same state:

$$S_i(t+1) = sign(h_i(t)) \tag{1}$$

In this exercise we model stochastic neurons by replacing that equation with a probabilistic state update:

$$P\{S_i(t+1) = 1 | h_i(t)\} = g(h_i(t)) \tag{2}$$

Let's say we have stored M patterns p^μ in a network of N neurons. We then set the network to an initial state $S(t_0)$ that has significant overlap with the third pattern and no overlap with other patterns: $m^{\mu \neq 3}(t_0) = 0$. For the deterministic update (eq. 1) we know (either from the textbook or from the proof done last week) we would retrieve pattern p^3 in a single update: $m^3(t_0 + 1) = g(m^3(t_0)) = 1$.

We now study how that result changes in the presence of noisy neurons (eq. 2). Look at figure 1 to get an intuition about the stochastic update.

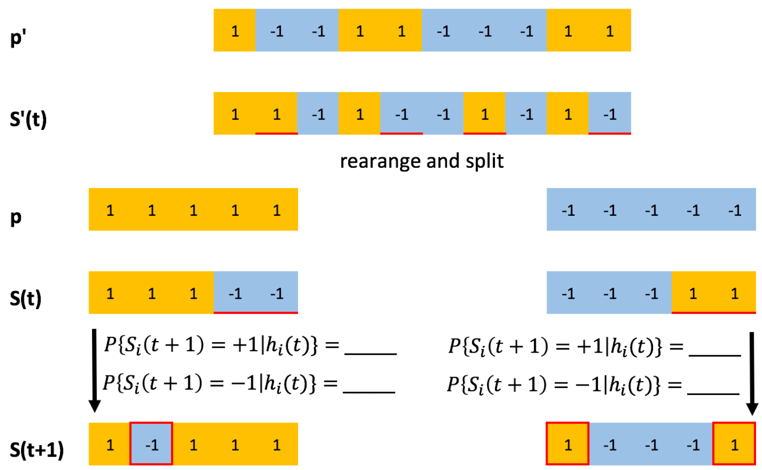

Figure 1: For the analysis of the overlap $m^3(t+1)$ it helps to rearrange pattern p and state S such that we can identify four sub-populations in the last row. We first split the neurons $S_i(t)$ into those that *should* be active and those that *should not* be active. All neurons in the same sub-population share the same probabilistic activity dynamics. In the last row, we see four groups of neurons which we label $\{p_i/S_i(t+1)\}$: {on/on}, {on/off}, {off/on}, {off/off}.

1.1 Derive the overlap $m^3(t_0 + 1)$ (eq. 3) under the state dynamics of eq. 2. Assume that there's only overlap with pattern p^3, and that for each pixel of the pattern 3, the probability to be on is $P\{p_i^3 = 1\} = 0.5$

$$m^3(t_0+1) = g(m^3(t_0)) - g(-m^3(t_0)) \tag{3}$$

Hints:

1. Use a result we derived earlier: $h_i(t_0) = p_i^3 m^3(t_0)$.
2. For each of the four groups (see figure 1) find the probabilities for $P\{S_i(t+1)|h_i(t_0)\}$
3. Recall the definition of the overlap m: $m^3(t_0+1) = \frac{1}{N}\sum_{i=1}^{N} p_i^3 S_i(t_0+1)$
4. For large N we can use the expected number of neurons in each of the four sub populations to express (the expected) overlap $m^3(t_0+1)$.

1.2

(a) In equation 2, what properties should the transfer function g have?

(b) Use $g(h) = \frac{1}{2}(tanh(\beta h)+1)$ in equation 3. Simplify it, plot the function graph and discuss it.

Exercise 2: Hopfield, asynchronous update and the energy picture

Consider a Hopfield network of N neurons with an **asynchronous** update regime. That is, only *one* randomly selected neuron k is updated at each step according to equation 4:

$$\begin{cases} S_k(t+1) = g(h_k(t)) = sign\left(\sum_j^N w_{kj} S_j(t)\right) & \text{for exactly one randomly chosen neuron k} \\ S_i(t+1) = S_i(t) & \text{for all other neurons,} \quad i \neq k \end{cases} \tag{4}$$

For each state S of a Hopfield network, we can compute a scalar value, known as the **energy E** of the network:

$$E := -\sum_i^N \sum_j^N w_{ij} S_i S_j. \tag{5}$$

The evolution of the network state and the change of energy are related in an interesting way:

> When a network is updated asynchronously then the energy function $E(S(t))$ does either decrease or stays at a (local) minimum.

We will now proof this property:

In the trivial case of $S_k(t+1) = S_k(t)\ \forall k$ the network has reached a stable state and therefore the energy function is stable too: $\Delta E = E(t+1) - E(t) = 0$.

Now consider the case of one neuron k changing its state and proof, in steps 4.1 to 4.3, that the energy decreases:

2.1 The energy $E(t)$ in eq. 5 is summed over all pre- and post- synaptic neurons i and j. Rewrite that sum such that the contribution of neuron k to the total energy E appears explicitly.

Hint: To simplify the resulting expression, remember that in a Hopfield network, the weight are symmetric: $w_{ij} = w_{ji}$ and there are no self recurrent connections: $w_{kk} = 0$

2.2 Write the change in energy $\Delta E = E(t+1) - E(t)$ when exactly one neuron k does changes its state.

2.3 Proof that $\Delta E < 0$ when exactly one neuron k does changes its state under the dynamics of eq. 4

Exercise 3: Binary codes and spikes

A Hopfield model is specified by a binary variable $S_i \in \{-1, +1\}$, the weights (eq. 6) and the update dynamics (eq. 7).

$$w_{ij} = c \sum_{\mu=1}^{M} p_i^\mu p_j^\mu \quad \text{with } c = \frac{1}{N} \tag{6}$$

$$S_i(t+1) = sign\left(\sum_{j=1}^{N} w_{ij} S_j(t)\right) \tag{7}$$

For an interpretation in terms of spikes it is, however, more appealing to work with a binary variable σ_i which is zero or 1.

3.1 Rewrite the Hopfield model in terms of $\sigma_i \in \{0, 1\}$, $S_i = 2\sigma_i - 1$.

3.2 Assume that the patterns have the property $\sum_{i=1}^{N} p_i^\mu = 0 \quad \forall \mu$.
Discuss that condition and use it to simplify the update dynamics found in the previous question.

3.3 Assume low-activity patterns $w_{ij} = \sum_\mu (\xi_i^\mu - b)(\xi_j^\mu - a)$, where the random variables $\xi_i^\mu \in \{0, 1\}$ have mean $\langle \xi_i^\mu \rangle = a$. For $b = 0$ can you restrict the weights to excitation only and move negative interaction into a group of inhibitory neurons?