NTNU

Det medisinske fakultet
Institutt for nevromedisin

Exam NEVR3004
Neural Networks

Monday June 7th 2010, 9am-1pm

ECTS credits: 7,5
Supporting materials: Dictionary
Number of pages (front page included): 10

Contact person during the exam: Keith Downing (73 59 02 71)

Exam results: June 28th 2010
Exam results are announced on: http://studweb.ntnu.no/

NEVR3004 Exam (Spring, 2010)

May 28, 2010

Notation

In the questions below, the sum-of-weighted-inputs to neuron i is denoted V_i, while the activation value (also known as the *output value*) for neuron i is denoted X_i. V_i is computed as:

$$V_i = \sum_j X_j w_{i,j} \tag{1}$$

where j iterates over all pre-synaptic neurons to neuron i, and $w_{i,j}$ is the weight on the connection from neuron j to neuron i.

X_i is computed from V_i via an *activation function*.

The weights of all connections appear as labels on those connections, unless otherwise specified.

Answer Format

For all questions that involve calculations, your written answer should BEGIN with the numbers or table of numbers that are requested in the exercise. CLEARLY mark them as your answer. Below that, feel free to include your calculations, although they are optional. Partial credit will be available in situations where the sensors notice that you made a simple arithmetic error. However, in cases where the calculations were not set up properly, are difficult to follow or hard to read, little or no partial credit will be given. Feel free to include verbal descriptions as accompaniment to calculations in cases where you fear that they are not straightforward for a sensor to follow.

1 Feed-Forward Network (10 points)

In the artificial neural network (ANN) of Figure 1, send the output values of the upper (input) layer of neurons through the ANN from top to bottom and record V_i and X_i for the neurons in the middle and lower (output) layers. The output values of each input neuron are written inside the small circle that represents the neuron, while the contents of the middle- and output-layer neuron figures are their activation functions. The neurons are numbered 1-5 by the small octagons. Fill in the V_i and X_i values in Table 1.

V_1		X_1	
V_2		X_2	
V_3		X_3	
V_4		X_4	
V_5		X_5	

Table 1: Sums-of-weighted-inputs and activation values for the 5 neurons in Figure 1

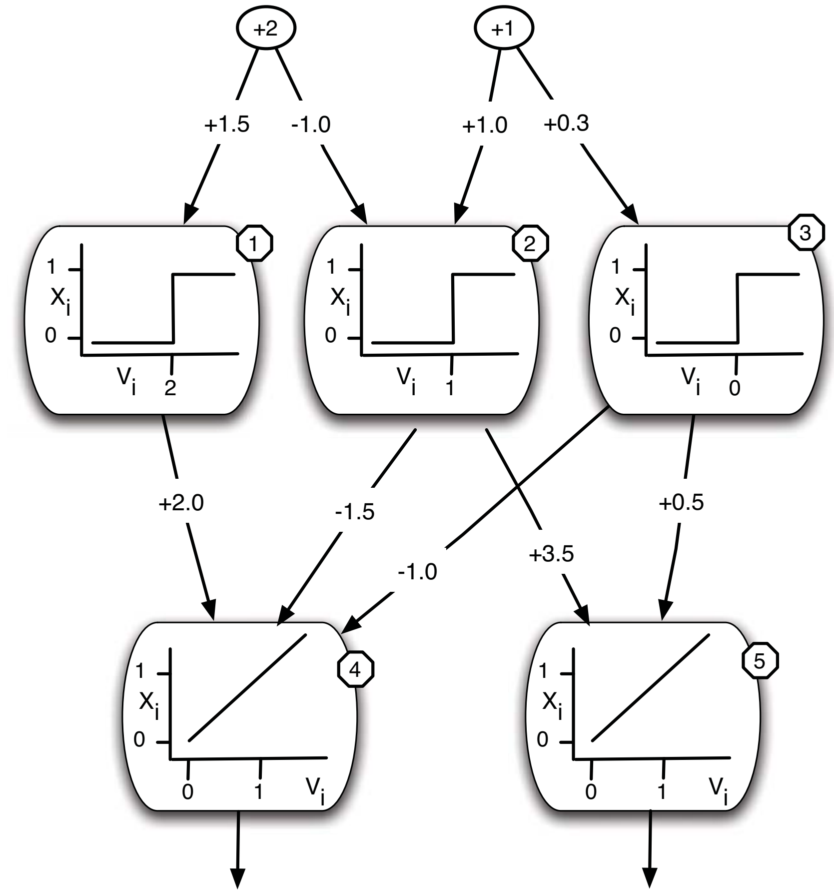

Figure 1: A 3-layered network

2 Recurrent Network (10 points)

In the recurrent network of Figure 2, the initial activation values (X_i) of the 5 neurons are given in Table 2. Complete the final 2 rows of the table by computing the output values (X_i) of each neuron for 2 complete timesteps. Use synchronous updating such that V_i and X_i at time t are based solely on the connection weights and the X values of all neurons at time t-1.

The activation function, shown on the right of Figure 2, is expressed mathematically as:

$$X_i = \begin{cases} 1 & \text{if } V_i \geq 0 \\ -1 & \text{otherwise} \end{cases} \tag{2}$$

	Activation Values				
Timestep	X_1	X_2	X_3	X_4	X_5
0	1	1	1	1	1
1					
2					

Table 2: Activation history for the recurrent network of Figure 2.

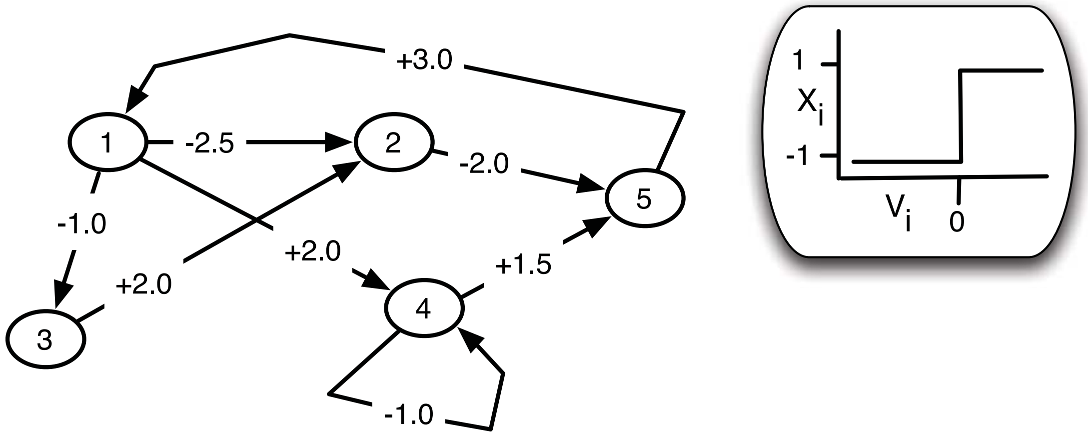

Figure 2: (Left) A simple recurrent network. (Right) The activation function for all 5 neurons in the network.

3 Oscillations in the Brain (20 points)

Explain why regions of the brain whose neurons oscillate at similar frequencies are more likely to exchange information than those that do not. Explain Laura Colgin's findings as to how CA1 accepts signals from EC and CA3, but using different gamma frequencies. Given what you understand about STDP, why should these two communication frequencies produce LTP during signalling from EC but not from CA3?

4 Willshaw Network (5 points)

Table 4 depicts the synaptic weights of a Willshaw weight that has been previously trained. The 6-bit patterns to be associated are entered along the left-hand side and along the bottom, as in the Willshaw networks used in class. So, for example, if the patterns 111000 and 001100 were to be associated, then the 3 1's of the first pattern would be entered along the upper half of the left side, with the three 0's along the bottom half of the left side, while 001100 would be entered from left to right along the bottom, as shown in Figure 3

Given the weights shown in Table 4, which of the two pairs of associations (shown below) have been learned? For each pair (denoted by 6 bits, $\leftrightarrow$, and 6 more bits), the leftmost 6 bits are entered along the left side of the network, while the rightmost 6 bits are entered along the bottom. Associations are then formed in the standard manner discussed in the lecture(s) and lecture notes.

1. 00000 $\leftrightarrow$ 111111 and 101010 $\leftrightarrow$ 010101
2. 110000 $\leftrightarrow$ 000011 and 101010 $\leftrightarrow$ 010101
3. 011000 $\leftrightarrow$ 001101 and 101010 $\leftrightarrow$ 001101
4. 001110 $\leftrightarrow$ 000110 and 111000 $\leftrightarrow$ 000111
5. 110000 $\leftrightarrow$ 000011 and 101010 $\leftrightarrow$ 010101
6. 011000 $\leftrightarrow$ 001101 and 001110 $\leftrightarrow$ 000110
7. 101101 $\leftrightarrow$ 001101 and 001110 $\leftrightarrow$ 010101
8. 101010 $\leftrightarrow$ 010101 and 001110 $\leftrightarrow$ 000110
9. None of the above
10. All pairs in choices 3 and 6 above.

1	○	○	○	○	○	○
1	○	○	○	○	○	○
1	○	○	○	○	○	○
0	○	○	○	○	○	○
0	○	○	○	○	○	○
0	○	○	○	○	○	○
	0	0	1	1	0	0

Table 3: Positioning of two patterns (111000 and 001100) in preparation for associative learning in a Willshaw network.

○	○	○	○	○	○
○	○	●	●	○	●
○	○	●	●	●	●
○	○	○	●	●	○
○	○	○	●	●	○
○	○	○	○	○	○

Table 4: The synaptic weights of a previously-trained, hetero-associative Willshaw network, where $\bullet = 1$ and $\circ = 0$.

5 Learning Paradigms (5 points)

Name the 3 major learning paradigms discussed in the context of artificial neural networks, and for each paradigm, give an example of one brain region that (at least abstractly) is believed to perform that type of learning.

6 Hebbian Learning (10 points)

Consider the network in Figure 3, a chain of 3 neurons, each providing input to the next in the chain. Neuron 1 (top) has a time-series activation pattern (i.e. series of X_1 values) as shown inside the large node at the top of the figure. Note that this neuron emits two high pulses, each of duration 2 timesteps, with a one-timestep resting period in between. The activation function for neurons 2 and 3 is $X_i = V_i$ (for all positive, zero and negative values of V_i), which is sketched on the right of the figure.

The weights, w_1 and w_2, are modified using the following version of the Hebb Rule:

$$\triangle w = \lambda X_{pre}(t-1)X_{post}(t) \tag{3}$$

where $\lambda = 0.2$ is the learning rate, $X_{pre}(t-1)$ is the activation value of the pre-synaptic neuron at time t-1, and $X_{post}(t)$ is the activation value of the post-synaptic neuron at time t. Hence, this rule modifies the weight based on the pre-synaptic value that comes one timestep before the post-synaptic value. Note that the weight will not appear as modified until the following timestep. So, for example, $X_{pre}(1)$ and $X_{post}(2)$ will be used to compute a weight change at the end of timestep 2 such that, in the table, the weight will not be shown to be modified until timestep 3.

Assume that it takes exactly one timestep for the output of neuron 1 to propagate through neuron 2 and on to the output of neuron 3. So on each timestep, you must update the output values of all 3 neurons. Fill in the empty cells of Table 5, which gives the values of the 5 main variables for 7 successive timesteps.

	Variable				
Timestep	X_1	w_1	X_2	w_2	X_3
0	0	1.0	0	-2.0	0
1	1	1.0	1	-2.0	-2.0
2	1	1.0	1	-2.0	-2.0
3	0	1.2	0		0
4	1	1.2			
5	1				
6	0		0		0

Table 5: History of the activation levels and weights in the network of Figure 3

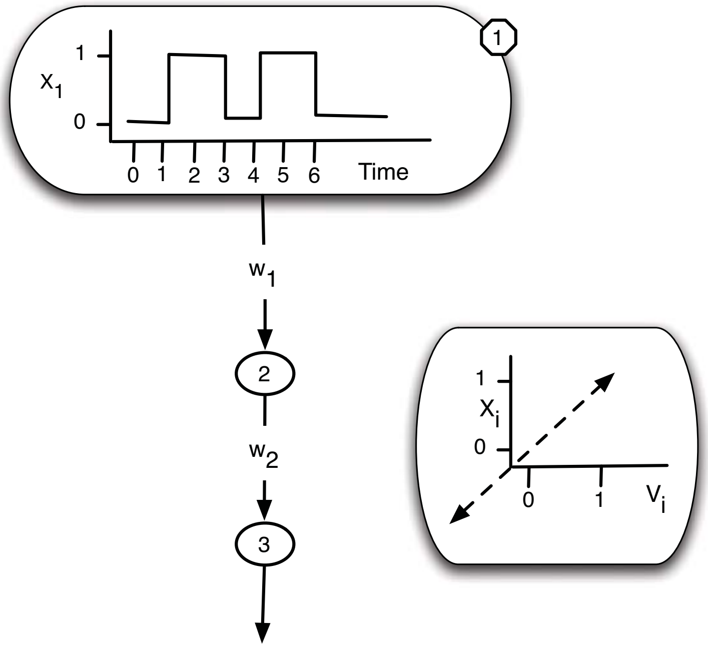

Figure 3: (Left) A 3-neuron feed-forward network with 2 modifiable synaptic weights., w_1 and w_2. The graph inside neuron 1 represents its time series of activation states (X_1).(Right) The activation function for neurons 2 and 3.

7 Spike-Time Dependent Plasticity (10 points)

- (4 points) Using one or more diagrams and a small amount of text, explain the concept of spike-time dependent plasticity
- (6 points) Consider the situation in Figure 4. A neural network has been trained to recognize different simple patterns such that each of the neurons fires when the pattern (A,B,C or D) on its label is present. Each of these four neurons is thus a *detector* for one pattern. This ability to detect these 4 patterns is implemented by a set of afferent connections (to the four neurons) which are not shown in the figure and are of no relevance for this exercise. We are concerned with the learning of a **series** of patterns and are thus interested in how the detector neuron for one pattern can stimulate another detector.

  At the beginning of the experiment, the strengths of the connections between these detector neurons are approximately equal, as shown by the equal width of all inter-neuron arrows in the network at the top of Figure 4.

  The network is then exposed to the following pattern sequence: ABCDABCDABCD, with a very short time span (of approximately 10 msec) between the presentation of each pattern.

  Assuming that the network weights are modified using a version of STDP, choose one of the four networks at the bottom of Figure 4 that most accurately illustrates the status of the network after repeated presentations of the pattern sequence (where thicker connections indicate stronger weights and thinner denote weaker). Explain your choice.

  Note: Do not confuse the issue by considering the real time for the brain's perceptual system to process information (such as visual or acoustic). Simply assume that the neurons A,B,C and D are activated in series, with a small delay between each activation time. The patterns that these neurons detect could simply be firing patterns (of other neurons) in the brain.

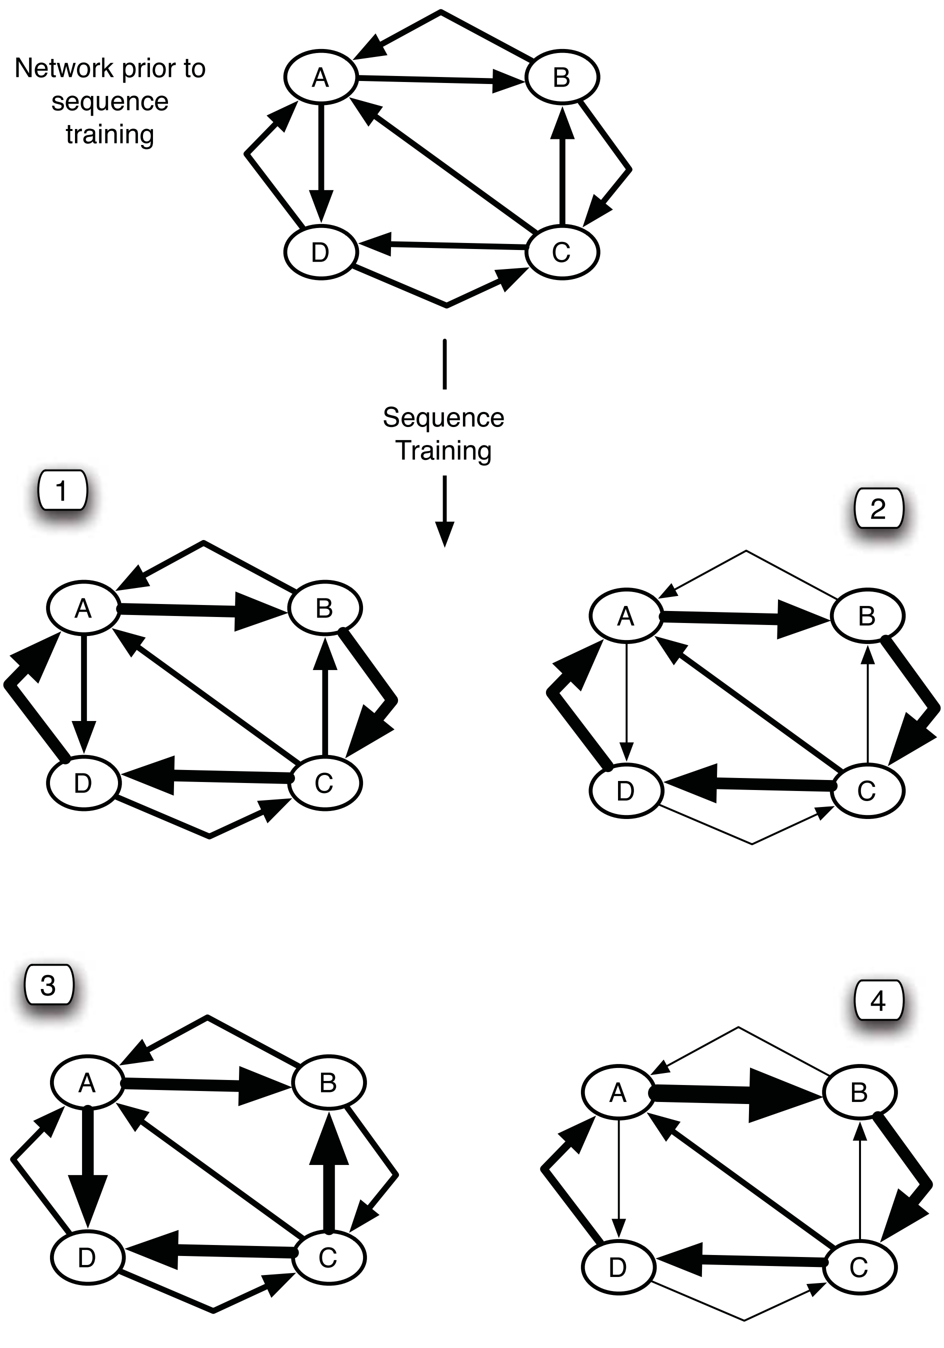

Figure 4: An abstract model of a sequence-learning neural network prior to pattern-sequence training (top), and four options (bottom) for the state of the network after training on the sequence ABCDABCDABCD.

8 Spiking Neurons (10 points)

Answer ONLY ONE of the two questions below. Each is worth 10 points on its own.

8.1 Spike Coding Concepts

Explain the difference between a firing-rate code and a spike-timing code in a neural network. Explain how oscillatory inhibition can be used to correlate the two codes in the sense that for a given firing rate, r, there would be a corresponding spike time, s, and vice versa.

8.2 Spiking Neuron Calculations

The following equation is often used to model the effects of pre-synaptic spikes upon the post-synaptic neuron, s:

$$G_s(t) = \sum_j \sum_k g_s(t - t_j^k) \tag{4}$$

where t_j^k is the time of the kth spike of the jth pre-synaptic neuron and $g_s(\triangle t)$ is a function that decreases with increasing $\triangle t$. In effect, $g_s(\triangle t)$ captures the basic concept that the further in the past a spike occurred, the less will be its effect upon the post-synaptic neuron.

Figure 5 provides the spike trains of 3 pre-synaptic neurons (left), which all send signals to the post-synaptic neuron, s, on the right. In this example, $g_s(\triangle t)$ is defined as follows:

$$g_s(\triangle t) = \begin{cases} 5 - \triangle t & \text{if } 0 \leq \triangle t \leq 5 \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

Use this information to calculate $G_s(4)$ and $G_s(7)$ for the post-synaptic neuron.

9 Anticipatory Shifts of Place Fields (20 points)

Explain the difference between symmetric and asymmetric place fields (of a place cell). How do asymmetric place cells relate to anticipation or prediction of future events? Using text and diagrams, explain how LTP can lead to anticipatory shifting of place fields in CA1.

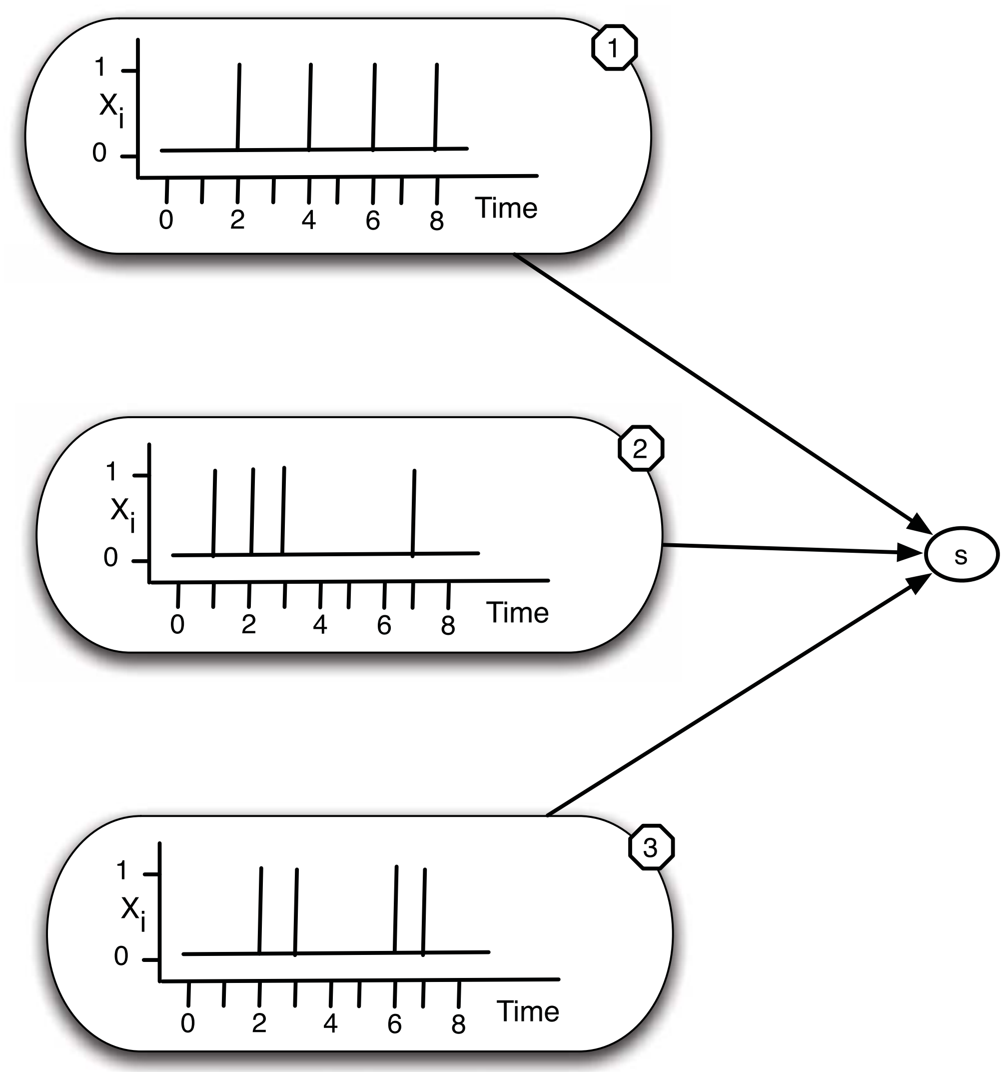

Figure 5: The spike trains of 3 pre-synaptic neurons whose contributions are to be calculated for the post-synaptic neuron, s, assuming no significant delays along the connections.