

2D1432

Artificial neural networks and other learning systems

Tentamen 2004-03-13 kl 14.00 – 19.00

No books or other tools allowed.

Non-native swedish persons may use a standard dictionary.

Question 1

(4p)

Which of the following statements are true? Correct the ones that are wrong.

1. Action potentials are positive in excitatory cells and negative in inhibitory.
2. The weights in artificial neural nets correspond to the synapses in real neural systems.
3. The cell membrane of a nerve cell is electrically conductive.
4. Ion pumps normally give rise to a negative potential inside nerve cells.

Question 2


(3p)

Which of the classification problems A, B and C below can be solved by


1. a single layer perceptron
2. a two layer perceptron
3. a RBF-net

Motivate your answers!


Problem A:



Problem B:



Problem C:



Question 3

(3p)

What happens if you use a large step length in the *perceptron learning rule*? How is convergence affected? How are the resulting weights affected? How is the final classification affected?

Question 4

(3p)

Show that all mappings that can be made by three cascaded (feedforward, connected in series), linear networks, can also be made by a single-layer linear network.

Question 5

(3p)

What is meant by the term *overfitting* (overtraining) in the context of learning from samples? Under what conditions is the risk for this phenomenon big? What can be done to detect the problem?

Question 6

(4p)

You have decided to use a multi-layer perceptron to predict the water level of a river. You have data for the amount of rain and water levels, on a day-by-day basis, covering several years. You figure that the water level primarily depends on these two variables from the previous week.

Describe how you should do to make a reasonable prediction of the water level of tomorrow with the help of the network. In particular, state what you use as input and output and how you go about training the network. Which learning algorithm do you choose?

Question 7

(3p)

In the *Boltzmann machine*, the term *temperature* plays an important role. How do you control this temperature of the system? What happens when you have a high or a low temperature? How should you control the temperature in order to find a globally optimal solution?

Question 8

(3p)

In a *Hopfield net* with asynchronous updating, this rule is normally used for the updates:

$$x_i \leftarrow \text{sign} \left(\sum_j w_{ij} x_j \right)$$

Show that these updates will always converge towards a fixpoint when $w_{ij} = w_{ji}$ and $w_{ii} = 0$.

Question 9

(3p)

What is meant by the term *spurious states* (also called *spurious attractors*)? In what kind of network does this phenomenon occur? Is it something good or bad?

Question 10

(3p)

Assume that you have an application where you receive measured values from 40 sensors. You know beforehand that the signal detected in reality only has approximately five independent components. The sensors are in fact measuring different weighted mixtures of the five underlying independent signals. Describe how you can use an artificial neural network to extract a five-dimensional signal which preserves most of the relevant information. To get full score you must supply the exact learning rule that should be used, and you must describe (e.g. draw) the topology of the network.

Question 11

(2p)

The *vigilance*-parameter has a central role in the ART networks. What are reasonable values for this parameter and how does the value influence the behavior of the network?

Question 12

(4p)

A multilayered feedforward net with threshold units can not be trained using a gradient following method, like e.g. back propagation. Why? Other optimization methods, e.g. genetic algorithms, can be used instead. Describe how this could be done. In particular, describe what should be used as the fitness function, and what kind of representation should be used in the chromosomes.

Question 13

(5p)

Suppose that a mail-order company has thousands of items for sale and wants to organize them in their catalogue so that similar items end up close to each other. To accomplish this, they intend to use a SOM (Self Organizing Feature Map) network.

Describe how you should go about to do this. How do you represent the input and output? What constitutes the training patterns? How is the actual training done? How do you get the final order?

Question 14

(4p)

In the theory on *reinforcement learning*, at least four different functions are used:

1. the Value function
2. the Policy function
3. the Reward function
4. the Q function

For each one, describe what it represents and, in particular, what constitutes its input and output.

Question 15

(3p)

Suppose that we are using *competitive learning* with four prototype vectors which, at a certain stage, have the following values (positions):

(0.7, 0.7, 0.0)
(0.7, -0.7, 0.0)
(-0.7, 0.0, 0.7)
(-0.7, 0.0, -0.7)

Describe in detail how these values are changed when a new input pattern (0.8, 0.3, 0.5) is presented to the net.