Neural Nets and Symbolic Reasoning

Exercises

1. Read *A Brief History of Connectionism* (David A. Medler). The relevant link is given on the course website.

2. Learning:

   a. Consider the sigmoid-function in the working space [-1,+1]: $\sigma_{[-1,+1]}(net) = 2/(1+\exp(-net/T)) - 1$. Show that $\sigma(-net) = -\sigma(net)$.

   b. Use the generalized Hebbian rule and teach a binary threshold union first with the AND function and then with the OR function. Take a small learning parameter (η=0.1). Discuss the results in terms of *plasticity* and *stability*

3. Exercise practising the backpropagation algorithm:

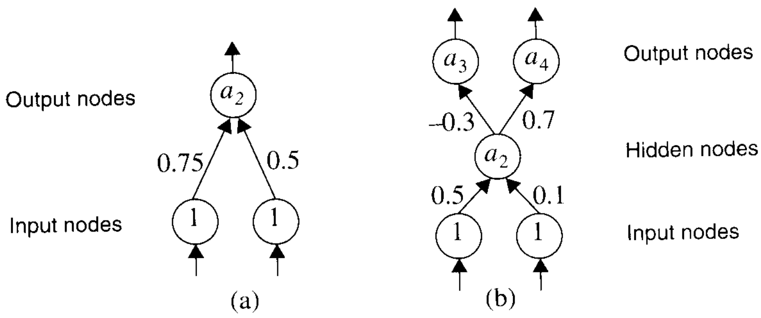

   a. Imagine we are training the single-layered network shown above on the left. The network is shown with a set of activation values for the input nodes and weights connecting the input nodes to the output node. Assume that the output node has a sigmoid activation function (assume T=1, working space [0,1]), that the desired output is 1.0, and that the learning rate is η=0.1. Calculate the changes that will be made to the two weights of the network.

   b. Repeat the calculation for the multi-layered network with hidden node using the same learning rate parameter and target activations of 1.0 on both output nodes. Calculate the relevant four weight changes.

4. Hopfield networks.

   a. Consider a discrete Hopfield network with 3 neurons. Teach the system with the input vectors (1 -1 1) and (-1 1 -1). What is the weight matrix? Why is the weight matrix not different from the matrix that results from either teaching with (1 -1 1) or (-1 1 -1).

b. Consider a discrete Hopfield network with possible activations ±1 and a binary threshold function. Prove that if **s** is an equilibrium state, then so is **–s** (there are monsters in state space!). Is the same valid for the sigmoid function with $T\neq0$?

c. Consider a discrete Hopfield network with possible activations ±1 and a binary threshold function. Teach the network with one training pattern **s**. Update then the network starting in state **s**. Does it change state? Do you get the same result if you use the sigmoid function with $T\neq0$?

5. Distributed processing and auto-association

   a. Explain how auto-association can be used to convert a *local* representation into a *distributed* representation. Give a concrete example.

   b. Do you think that auto-association can be useful to detect redundancies in a set of input patterns? Give an example and check your assumptions with the help of the neural network simulator tlearn.

   c. Do you think that a simple transformation network with three layers – where the input- and the output-layer represent strings of words and the hidden layer represents internal configurations – is adequate to learn simple transformations like the active-passive transformation: *Peter loves Maria* ➔ *Maria is loved by Peter*? For what reason did Chalmers use the RAAM? Why so complicated?

6. Structure in Time: Recurrent networks

   Why are the hidden nodes in Elman's simple recurrent network fully connected in one direction but not in the other? Why do you think the context nodes are set to be linear?

7. Past tense acquisition

   a. Show that for the words *algalgal* and *algal* the same set of *wickelphones* is generated.
   b. What is the difference between *wickelphones* and *wickelfeatures*? And why did Rumelhart and McClelland use *wickelfeatures* instead of *wickelphones*?
   c. Use your favoured model of default logic and develop a formal description of a simplified *past tense* morphology that is able to describe the general rule *stem* → *stem*◡*ed* (with the concatenation function "◡") and the two exceptions *go*→*went*, *hit*→*hit*.