Artificial Intelligence: Assignment 6

Due date: April 12 (Wednesday)

Problem 1 (2 points)
The authors of the textbook argue that the human brain operates as a neural network, and its superior abilities are based on massive parallelism; however, most cognitive-science researchers believe otherwise. Empirical studies suggest that the human mind is a sequential computer with a very slow CPU, and its main advantage over artificial computers is superior algorithms.

Read Chapter 19 and argue that the functionality of the human mind is different from that of artificial neural networks. Give examples illustrating that the behavior of human problem solvers is closer to sequential symbolic algorithms than to neural networks.

Problem 2 (3 points)
Consider a multi-layer feed-forward network with a *linear* activation function; that is, the output of each unit is a linear function of the weighted sum of its inputs. Prove that we can replace this network with a perceptron that gives the same outputs.

Problem 3 (5 points)
Implement a neural network for learning basic boolean functions, with three input units and one output unit. Your network should be able to learn any of the following functions:

- **Conjunction:** Output 1 if all inputs are 1
- **Negated conjunction:** Output 1 if at least one input is 0
- **Disjunction:** Output 1 if at least one input is 1
- **Negated disjunction:** Output 1 if all inputs are 0
- **Majority:** Output 1 if at least two inputs are 1

Your program should read a file with training and test examples, use the training examples to adjust the weights of the network, and then use the resulting network to classify the test examples. The output of the program should include the training time (seconds) and the classification of the test examples. The input format is as follows:

```
<output>  <input-1> <input-2> <input-3>
   ...
<output>  <input-1> <input-2> <input-3>

<input-1> <input-2> <input-3>
   ...
<input-1> <input-2> <input-3>
```

The training examples are above the blank line, and the tests are below. The `<output>` value is 0 or 1, and each `<input>` is also 0 or 1; successive values are separated by one or more spaces. For instance, the following file includes six training examples and one test, for the majority function:

```
 1   1 1 0
 1   1 0 1
 1   0 1 1
 0   0 0 1
 0   0 1 0
 0   1 0 0

 0 0 0
```