

Homework #7: Chapters 18 - 20


The following exercises are due at the beginning of class on Friday, April 24. Note, this homework is continued on the reverse side of the paper.

1. [40 pts. total] In this problem we'll consider the following training set:

Example	Color	Legs	Tail	Fur	Goal Predicate
X ₁	Brown	4	Yes	Yes	Yes
X ₂	Brown	2	No	Yes	No
X ₃	Green	4	Yes	No	No
X ₄	Brown	0	Yes	No	No
X ₅	Black	4	Yes	Yes	Yes
X ₆	Black	4	No	Yes	No
X ₇	Gold	2	Yes	No	No
X ₈	Gold	4	Yes	Yes	Yes

- [8 pts.] Convert the training set into a set of first-order logic description and classification sentences. Use the predicates $Color(x,c)$, $Legs(x,n)$, $Tail(x,t)$ and $Fur(x,f)$ in your description sentences and $Q(x)$ for your goal predicate. Note, that by using binary predicates for $Tail$ and Fur we are able to restrict our hypothesis space to include only candidate definitions that are positive conjunctive sentences.
 - [32 pts.] Use current best-hypothesis search learning on the training set. As was demonstrated in class, assume the algorithm is implemented as a depth-first search. The initial hypothesis should either be "True" or "False", and each node in your search tree should show the current hypothesis and how it classifies the current example (i.e., true positive, true negative, false positive or false negative). When a node is expanded, its children are determined by the consistency of its hypothesis. If it is consistent, then it has a single child with the same hypothesis considering the next example. If it is inconsistent, then the children's hypotheses are the immediate specializations (or generalizations) that are consistent with all examples to that point. Assume that the examples are received in the order given in the table and that the hypothesis space only contains hypotheses whose candidate definitions are positive conjunctive sentences. When there is a choice of otherwise equivalent nodes to expand, always choose the one that adds the leftmost remaining condition from the attributes in the table.
- [10 pts.] Give an example training set for which logical-based learning techniques, such as current best hypothesis search and version space learning, will fail to find a consistent hypothesis.
 - [30 pts.] Use the perceptron learning algorithm to teach the perceptron shown below to recognize implications (i.e. $X_1 \Rightarrow X_2$). Assume that a **threshold activation function** is being used and that the threshold function returns 1 when its input is > 0 and returns 0 otherwise. For initial weights, use $W_0 = 0.2$, $W_1 = -0.5$, and $W_2 = 0.1$. For the learning rate, use $\alpha = 0.1$. Use only the examples in the table to the right of the network in your learning process.

Stop the training once the weights remain unchanged for one full pass through the examples. The examples must be used in the order given by the table below. Start again with the first example whenever you exhaust all of the examples but have not yet reached the stopping criteria. Show all of the intermediate calculations and values (not just the answer or the updated weights after each example).



Training Set		
X_1	X_2	out
0	1	1
1	0	0
0	0	1
1	1	1

4. [20 pts.] Consider the following neural network in which the hidden units and output units use a **threshold activation function** (as above, assume it returns 1 when its input is > 0 and returns 0 otherwise). The number of each node is written in bold above it. The $t=x$ notation means that a unit has threshold x (as opposed to 0). Recall, this is shorthand for an ordinary threshold node which has an additional bias weight of x on a fixed input of -1 . Given the activation levels written in the boxes for the input units on the left, compute the activation levels ($a_3, a_4, a_5, a_6, a_7, a_8$ and a_9) of the remaining nodes in the network. Show your work for each activation level.

