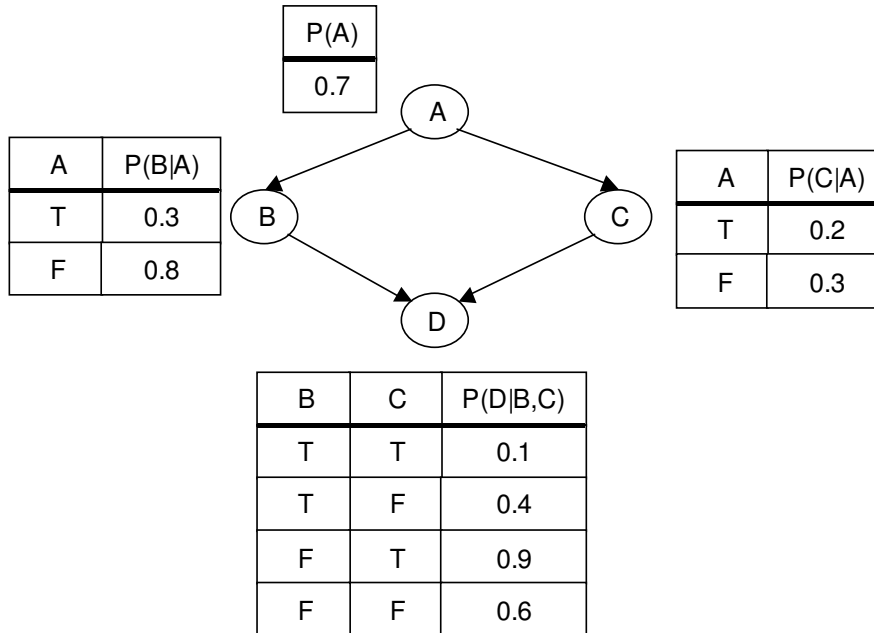


Homework #5: Chapters 14, 16, 18, 19, and 20

The following exercises are due at the beginning of class on April 27.

1. [20 pts.] Consider the Bayesian network below, where A,B,C and D are all Boolean random variables.



Compute the following probabilities and probability distribution.

- $P(a, \neg b, c, d)$
 - $P(\neg a, b, \neg c, \neg d)$
 - $\mathbf{P}(A|b, \neg d)$
2. [20 pts.] Consider the following variation of the Wumpus World agent. The agent has a choice of eight actions: move in one of the four directions (north, south, east, or west) or shoot its only arrow in one of those directions. If the agent moves into a square with the Wumpus or a pit in it, the agent dies; it can move into any other square safely. A square cannot contain both a Wumpus and a pit. The arrow will only travel one square, but if the Wumpus is in the square the arrow will kill it. The agent's utility is calculated as follows:

- 1000 for an action that results in death
- 100 for wasting the arrow by shooting at a square that doesn't contain the Wumpus
- 100 for moving into a safe square
- 200 for killing the Wumpus with an arrow

For each of the four adjacent squares (North, South, East, and West), the agent has determined the probability that they contain the Wumpus and the probability that they contain a pit. These probabilities are given in the table on the other side of this sheet:

Square	P(Wumpus)	P(Pit)
North	0.4	0.5
South	0.6	0.2
East	0.0	0.0
West	0.0	0.3

What is the expected utility of each action? To maximize the chance for partial credit, be sure to show your work. If the agent follows the principle of maximum expected utility and only considers single actions (as opposed to action sequences), which action will it choose?

3. [35 pts.] For this problem, assume that the hypothesis space only contains hypotheses whose candidate definitions are positive conjunctive sentences (i.e., a set of unnegated atoms separated by and (\wedge) symbols). Thus, the immediate generalization or specialization of a sentence should differ by only a single conjunct. Use the training set given below.

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	Gold	2	Yes	No	No
X ₇	Black	4	No	No	No
X ₈	Gold	4	Yes	Yes	Yes

- 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.
 - Which of these examples is the candidate definition $Color(x,Brown) \wedge Fur(x,Yes)$ consistent with? Which ones result in false positives and which ones in false negatives?
 - Give all the immediate specializations of $Tail(x,Yes) \wedge Legs(x,4)$ that are consistent with examples X₁ to X₃.
 - Give all the immediate generalizations of $Legs(x,4) \wedge Color(x,Gold)$ that are consistent with example X₅.
 - Use version space learning on the training set. Assume that the examples are received in the order given. Show your G-set and S-set after each new example is received. When done, list all hypotheses that are consistent with the training data.
4. [10 pts.] Construct by hand a perceptron that can calculate the logic function implies (\Rightarrow). Assume that 1 = true and 0 = false for all inputs and outputs. Be sure to say which activation function you are using.
5. [15 pts.] Construct by hand a feed-forward neural network of threshold units that computes the exclusive or (XOR) function of two inputs.