Artificial Intelligence Homework 2

1. *(For this problem, you will probably want to print out these two pages and turn them in with the rest of your written work.)*

   Run **minimax** (no alpha-beta) on the following game tree, filling in values for each internal node. The first player is MAX (triangle pointing up). MIN nodes are downward-pointing triangles.

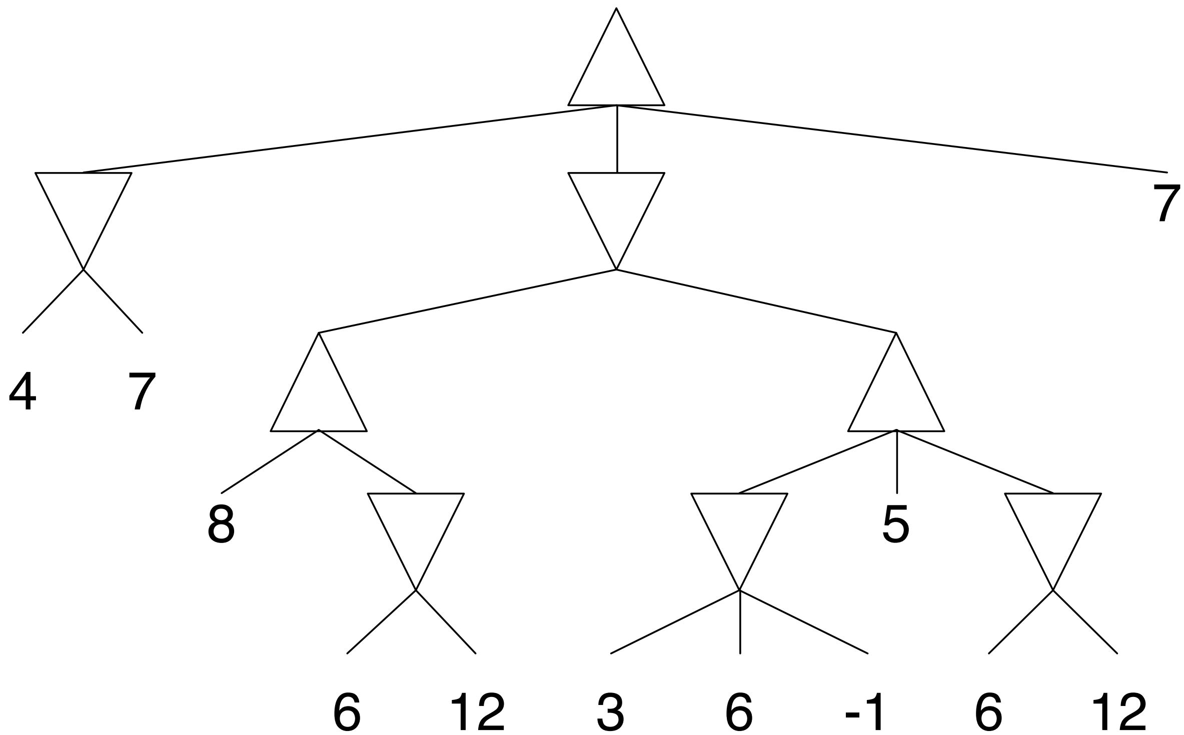

2. Run **minimax with alpha-beta pruning** on the same tree, with left-to-right node expansion (that is, consider the children of a node in left-to-right order, as we've done normally). Show the values of alpha and beta at each node (you may want to copy this tree onto another sheet of paper if you need more room), which values get passed up the tree by using arrows, and which nodes are not examined at all by crossing them out.

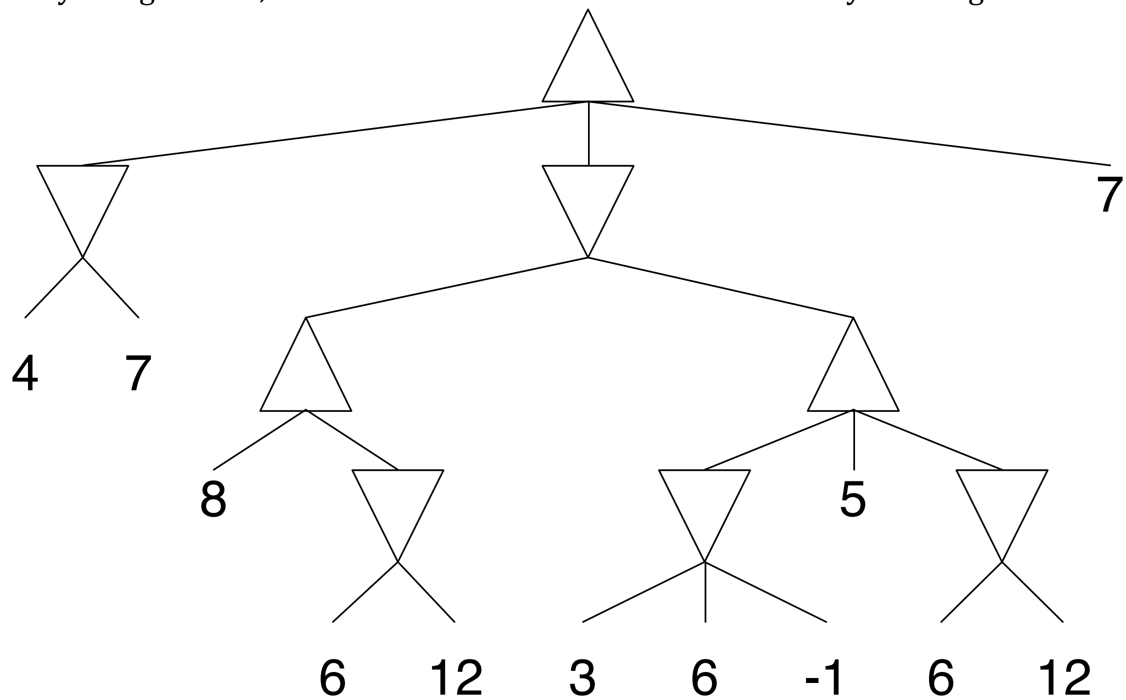

3. Run **minimax with alpha-beta pruning**, but now consider the children of a node in right-to-left order.

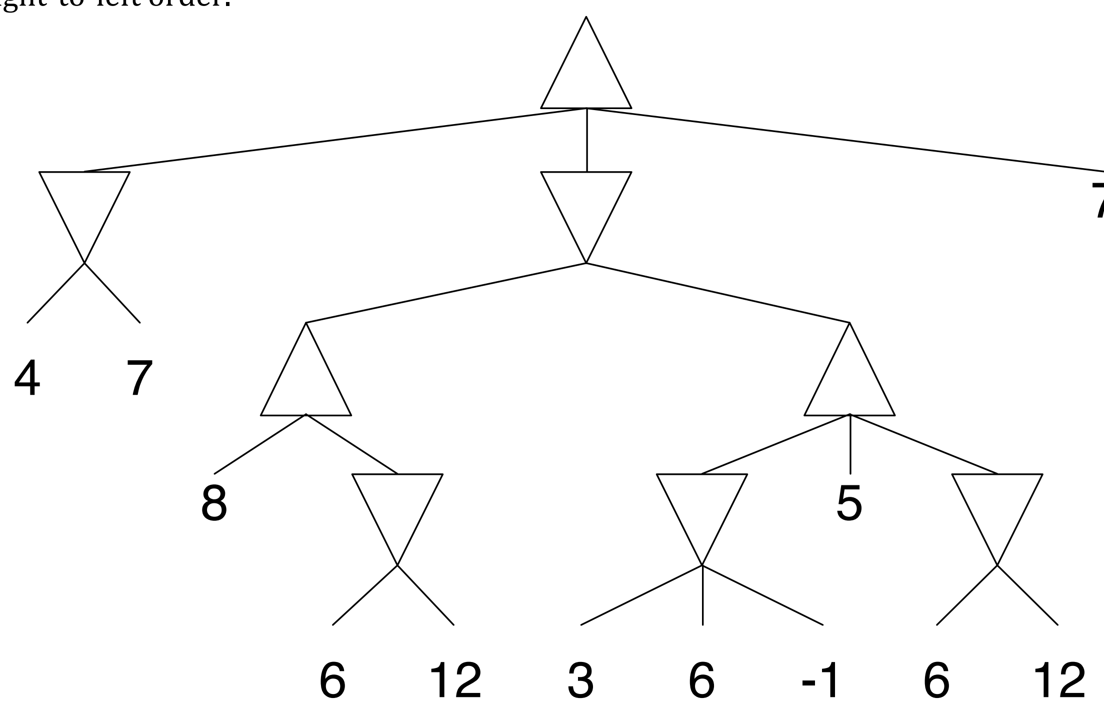

4. Consider the following Bayes network:

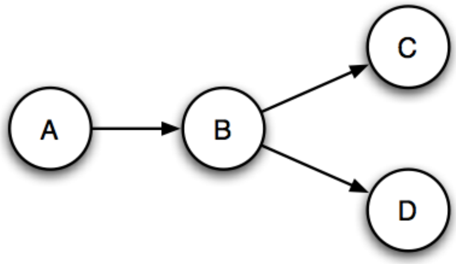

Here are the CPTs (conditional probability tables) for this network:

(I follow the book's convention of using uppercase letters to stand for a random variable, and lowercase letters to for a specific assignment of a value to the random variable. For instance "A" is a random variable, but "a" is the specific setting of "A = true" and "~a" means "A = false.")

P(a) = 0.4

P(b | a) = 0.7 P(b | ~a) = 0.3

P(c | b) = 0.2 P(c | ~b) = 0.6

P(d | b) = 0.9 P(d | ~b) = 0.5

a. Suppose we know the value of random variables C and D; specifically, assume C is true and D is false. Use the Bayes net exact inference algorithm to calculate P(A | c, ~d). (This means calculate the probability of A being true [and then being false] given the values of C and D). Show all of your work, including the steps involving the definition of conditional probability, where you introduce the normalization constant, the marginalization step, the re-arrangement of the summations to make the calculation as efficient as possible, drawing the tree to show your calculations, and the normalization step at the end.

b. Suppose we know the value of random variable B is false, and we wish to calculate P(A | ~b). Similar to part (b) above, illustrate how the exact inference algorithm works, **but only up through re-arranging the summations**. After you have re-arranged the summations, there will be an extra mathematical step you can do to make your calculation much easier. What is this step, and what *general* conclusions can you draw (about any Bayes net) that tell you when you will have such a step?

   Hint: What is $\sum_d \mathrm{P}(d \mid \sim b)$?

   Hint 2: read the last paragraph before the start of section 14.4.3 on page 528.

5. We continue with the Bayes network from problem 4. Suppose we generate twenty samples from the network (each sample lists the T/F values for A, B, C, D):

```
(False, True, False, False)
(False, False, False, False)
(False, False, True, True)
(False, False, False, False)
(False, False, True, False)
(False, False, True, False)
(False, True, False, False)
(False, False, True, False)
(True, False, True, True)
(True, False, False, True)
(True, True, False, True)
(False, True, False, False)
(False, False, True, True)
(False, False, False, True)
(False, False, True, True)
(False, False, False, False)
(False, False, False, True)
(False, True, False, False)
(True, False, False, False)
(False, False, False, True)
```

   a. Using direct sampling, use the twenty samples above to estimate the probability P(~a, ~b, ~c). Show your work.

   b. Using rejection sampling, use the twenty samples above to estimate the probability P(~a | ~b, ~c). Show your work, including stating how many of the 20 samples are rejected.

6. You have a bag containing three biased coins, called coin a, coin b, and coin c, with probabilities of coming up heads of 20%, 60%, and 80% respectively. You reach in and pick a coin randomly from the bag, but you can't tell which coin you picked (they all look the same to you). You flip that same coin three times and observe whether you got heads or tails each time.

   a. Define a complete Bayesian network for this situation, showing the structure of the network and the CPTs.

      Hint: you will need four random variables, one for which coin you chose, and three for the flips. The three coin flips are Boolean random variables (two-valued), but the coin-choice random variable is three-valued.

   b. Calculate which coin was most likely to have been drawn from the bag if the observed flips were heads, heads, and tails. Show all of your work. (Note that this last problem is not a ML/MAP problem, it's a direct probability calculation using the Bayes net.)