Artificial Intelligence Homework 2

In the book, you are responsible for the following sections:

- Chapter 5 (Adversarial search): 5.1 through 5.4, but ignore 5.4.3.
- Chapter 13 (Probability): 13.1 through 13.5
- Chapter 14 (Bayes nets): 14.1, 14.2, (not 14.3), 14.4 (only the exact inference algorithm, not variable elimination and clustering), 14.5 (only direct sampling, rejection sampling, and likelihood weighting)

1. Run minimax 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.

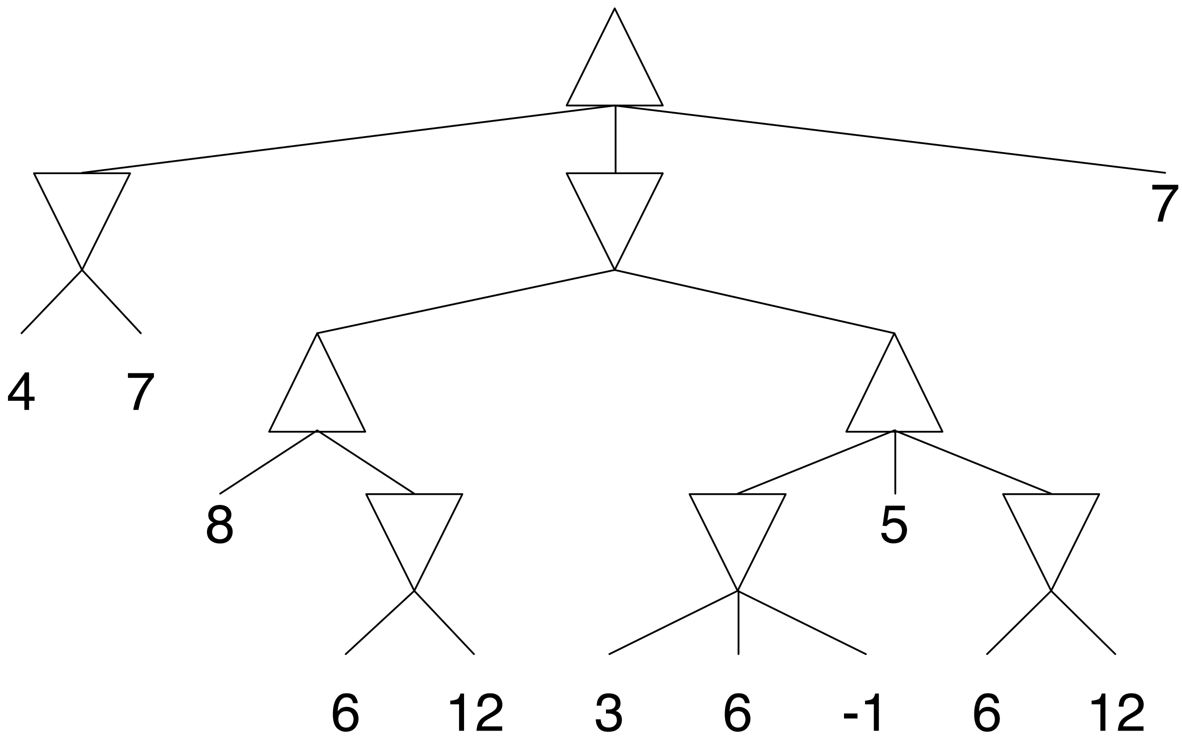

2. Run minimax with alpha-beta pruning on the same tree, with left-to-right node expansion (that is, consider the children of a node in left-to-right order, as we've done normally). Show the values of alpha and beta at each node (you may want to copy this tree onto another sheet of paper for more room),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

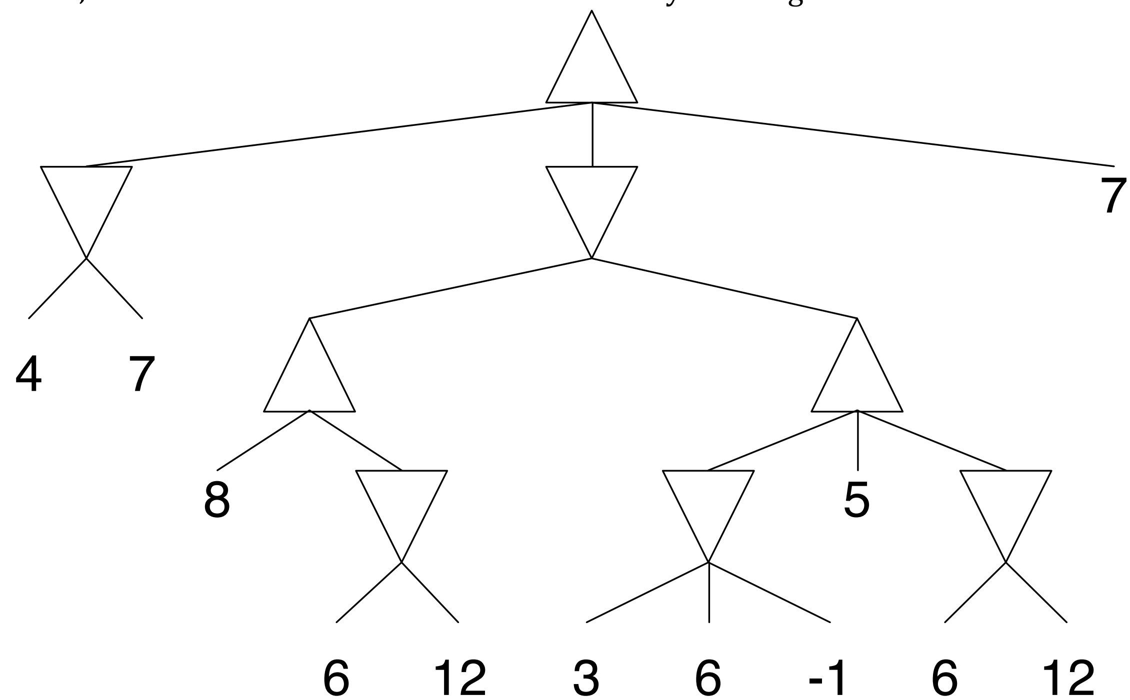

3. Run minimax with alpha-beta pruning, but now consider the children of a node in right-to-left order.

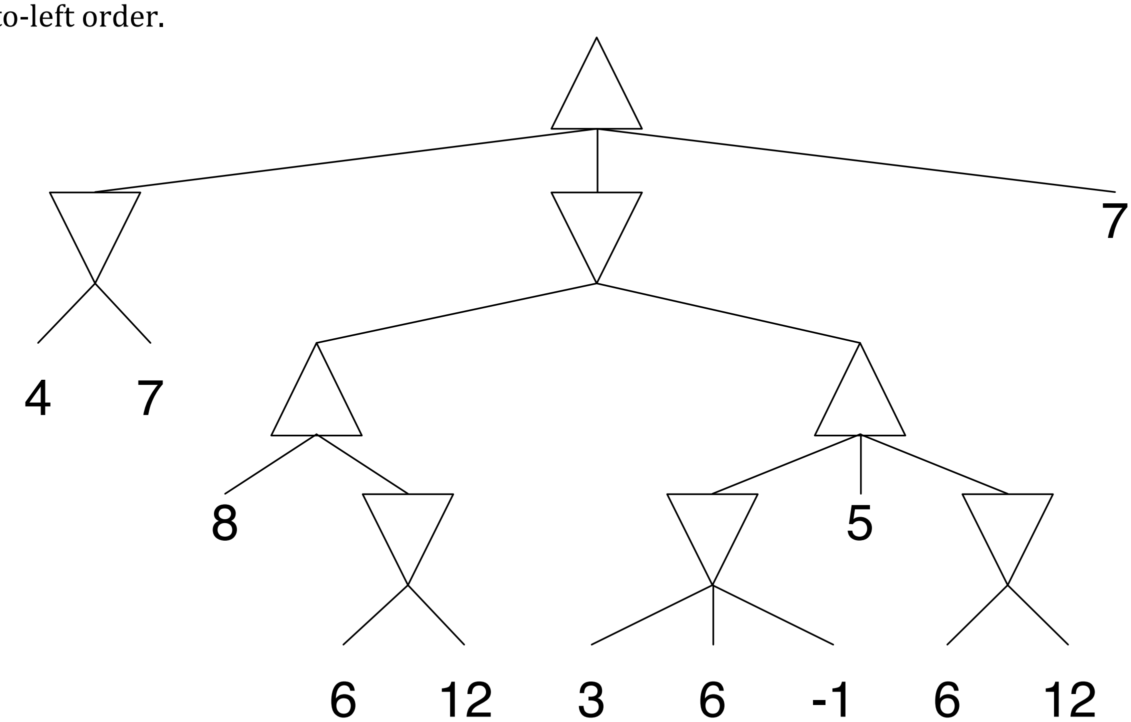

4. Consider the following Bayes network (note there are no conditional probabilities given yet):

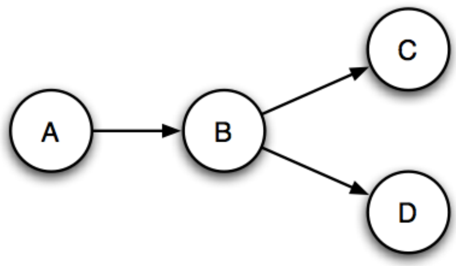

   a. Which of these statements are implied by the network structure? Explain why or why not, for each statement.
      i. $P(B \mid A, C, D) = P(B \mid A)$
      ii. $P(D \mid B) = P(D \mid B, C)$
      iii. $P(B \mid A) = P(B)$

   b. Suppose we now add the following conditional probability tables to the 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.)

      $P(a) = 0.4$

      $P(b \mid a) = 0.7$　　　$P(b \mid \sim a) = 0.3$

      $P(c \mid b) = 0.2$　　　$P(c \mid \sim b) = 0.6$

      $P(d \mid b) = 0.9$　　　$P(d \mid \sim b) = 0.5$

      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 $P(A \mid c, \sim d)$. (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.

   c. Suppose we know the value of random variable B is false, and we wish to calculate $P(A \mid \sim b)$. 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. In preparation for Rhodes College playing Humans vs. Zombies, you decide to investigate how exactly people become zombies and what happens when they do. Recent zombology research has revealed that
   - If someone becomes a zombie (Z), they have a chance of developing a hunger for brains (H), moaning a lot (M), and showing lack of coordination (C).
   - Zombification happens due to the interaction of two factors: carrying the zombie gene (G), which makes someone more susceptible to be zombified, and being infected (I) somehow.
   - The zombie gene is recessive, and there's a chance you carry it if your mother carried it (T).
   - To be infected, someone can be bitten (B) by another zombie or come into contact with zombie blood (L).
   - Zombologists only just learned that the presence of the full moon (F) makes zombie blood and zombie bites more infectious, and also increases the power of carrying the zombie gene.
   - Additionally, identifying zombified people is complicated because when a person is bitten, he or she may start moaning immediately due to the pain even if they don't turn into a zombie.

   a. Draw the Bayesian network that corresponds to the situation above, showing all the variables (Z, H, M, C, G, I, T, B, L, F). You do not need to show the CPTs.

   b. Based on your network, the probability of being bitten by a zombie (B) is conditionally independent of *all* other variables in the network given which set of variables?

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, but the coin-chosen 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.