# CSC 380/480 – Foundations of Artificial Intelligence
**Winter 2007**
## Assignment 4 Selection Solutions

**1. Minimax and Alpha-Beta Pruning:**

**Solution:**

(a) The annotated tree below shows the final backed up values resulting from the Minimax procedure:

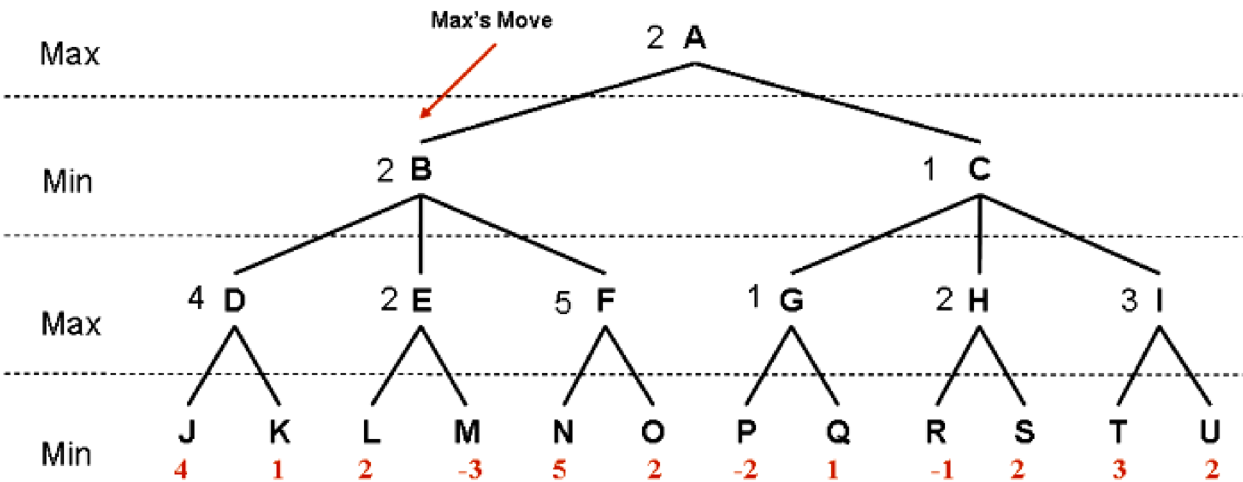


(b) The annotated tree below shows the progress of Minimax procedure with alpha-beta pruning:

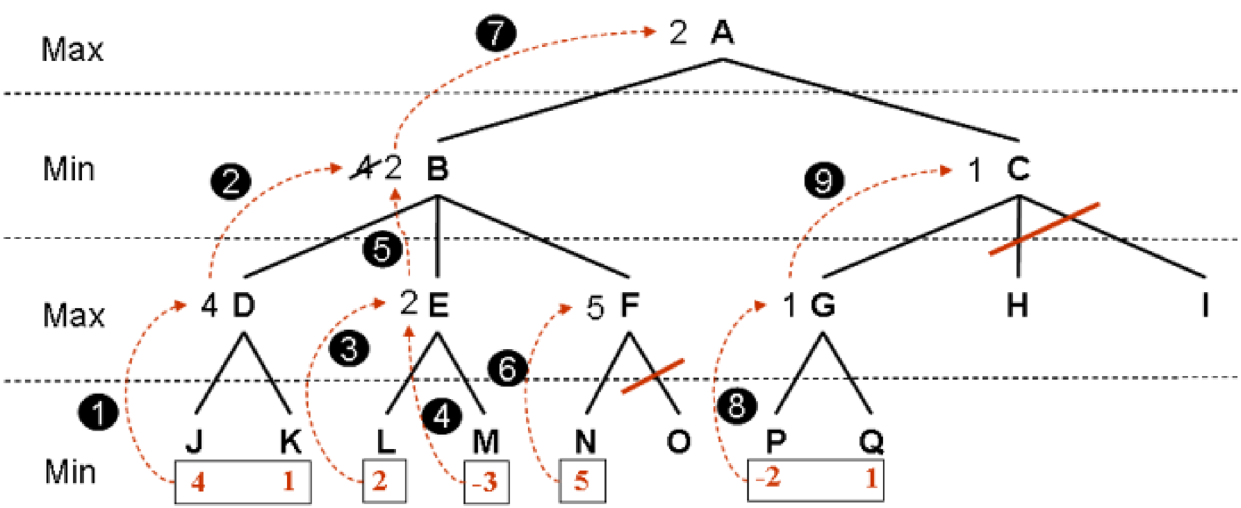


The step numbers in black circles and the red dotted lines indicate the step-by-step depth-first search and the corresponding backed up values. Note that the search below the Max node F is discontinued (pruned) because the alpha value of F (5) is greater than the current beta value its Min ancestor B (2). Similarly, the right two branches of C are prunes since C's beta value (1) is less than the alpha value of it Max ancestor A (which at the time of step 9 is 2 based on the previous backed up value from the left subtree of A).

**2.** Three prisoners, A, B, C are in their cells. They are told that one of them will be executed the next day and the others will be pardoned. Only the governor knows who will be executed. Prisoner A asks the guard a favor. “Please ask the governor who will be executed, and then tell either prisoner B or C that they will be pardoned.” The guard does as was asked and then comes back and tells prisoner A that he has told prisoner B that he (B) will be pardoned. What are prisoner A’s chances of being executed, given this message? Is there more information than before his request to the guard?

**Solution:**

Prior probability of each of A, B, C get executed is 1/3

| A | B | C | P |
|---|---|---|---|
| T | F | F | 1/3 |
| F | T | F | 1/3 |
| F | F | T | 1/3 |

Now, using the definition of conditional probability, we have:

$P(A \mid \neg B) = P(A \wedge \neg B) / P(\neg B)$

From the above table, we observe that:

$P(A \wedge \neg B) = 1/3$ [this is the same as P(A) since A begin executed means B will not be executed]
$P(\neg B) = 1/3 + 1/3$ [the sum of the two cases where B is false]

Therefore, $P(A \mid \neg B) = (1/3) / (1/3 + 1/3) = ½ = 0.5$

Since this conditional probability is greater than the prior probability, P(A), there is now more information than before A made the request to the guard.

**3.** Suppose an automobile insurance company classifies a driver as good, average, or bad. Of all their insured drivers, 25% are classified good, 50% are average, and 25% are bad. Suppose for the coming year, a good driver has a 5% chance of having an accident, and average driver has 15% chance of having an accident, and a bad driver has a 25% chance. If you had an accident in the past year what is the probability that you are a good driver?

**<u>Solution:</u>**

First let's list all of the know information:

p(good) = 0.25
p(avg) = 0.50
p(bad) = 0.25
p(acc | good) = 0.05
p(acc | avg) = 0.15
p(acc | bad) = 0.25

Now, we need to use Bayes' Rule to obtain the probability that you are a good driver given that you have an accident: p(good | acc).

$$p(\text{good} \mid \text{acc}) = \frac{p(\text{acc} \mid \text{good}) * p(\text{good})}{p(\text{acc})}$$

But, we don't know p(acc) directly from our prior information. We need to use normalization to figure this out:

Note that p(good | acc), p(bad | acc), and p(avg | acc) must add up to 1. Inverting each of these using Bayes' Rule (as in the case of p(good | acc), as above) and taking their sum, we have:

$$\frac{p(\text{acc} \mid \text{good}) * p(\text{good})}{p(\text{acc})} + \frac{p(\text{acc} \mid \text{bad}) * p(\text{bad})}{p(\text{acc})} + \frac{p(\text{acc} \mid \text{avg}) * p(\text{avg})}{p(\text{acc})} = 1$$

Therefore, p(acc) = p(acc | good) * p(good) + p(acc | bad) * p(bad) + p(acc | avg) * p(avg)

In other words:

$$p(\text{good} \mid \text{acc}) = \frac{p(\text{acc} \mid \text{good}) * p(\text{good})}{p(\text{acc} \mid \text{good}) * p(\text{good}) + p(\text{acc} \mid \text{bad}) * p(\text{bad}) + p(\text{acc} \mid \text{avg}) * p(\text{avg})}$$

$$= (0.05)(0.25) / [\ (0.05)(0.25)+ (0.25)(0.25) +(0.15)(0.50)\ ]$$

$$= \mathbf{0.083}$$