Extra Practice for Midterm

Heuristic/State Space Problems:

There exists a game (called The Initiative*) where you are searching for a specific piece of treasure on an n x n board. Each grid location has up to two places where you can search for treasure, and you have tokens that start in the top-left corner. The actions you can take each turn are: move up to 2 spaces on the grid, reveal all treasure in your current space, and reveal 1 of the treasures on any square in the board. You win the game once you find the treasure, navigate to the square containing it, and then travel to the bottom right square on the board.
*This description is slightly modified from the original game.

(a) What variables would you use to represent the game state for this problem? Remember that the game state should be represented by a minimal set of variables.

Player location (as a tuple/Location)
The locations of all the potential treasures and their individual states {unknown, not treasure, treasure} (you'd probably use a 2D array to keep track of this information)
Whether the player has collected the treasure (as a Boolean)

(b) Using your state representation from a, how would the below state be represented/encoded? (? Represents a potential unrevealed treasure, N represents no treasure, and T represents a treasure. * represents the player, and the contents of each || represent a grid cell. A 3x3 grid is being represented here.)

```
|*??  | N?  | ??  |
| ??  | ?   | N?  |
| ??  | NN  | ??  |
```

Player location: 0,0
Treasure locations:
[[?? , N? , ??],
[?? , ? , N?],
[?? , NN , ??]]
Has collected the treasure: False

(c) What are the possible actions at each game state? What is the maximum number of possible actions you could take on a turn?

Reveal all treasures on current square (1 option)
Reveal one item on any square (n*n*2 options)
Move up to two squares (13 options)
Max possible actions: 2*n^2*13 = 26*n^2

(d) What would the goal state look like, given your answer from a?

Player location: 2,2
-One of the locations has been revealed to have the treasure, example below
Treasure locations:
[[?? , N? , ??],
[?? , T   , N?],
[?? , NN , ??]]
Has collected the treasure: True

(e) What heuristic might you use to solve this problem using A* (assuming that the cost of each action is 1)? Is the heuristic admissible? Consistent?

Until the treasure is revealed,
h(n) = number of unknown treasures/(2*n^2) + number of moves to the nearest unknown + number of moves from the nearest unknown to the goal location

Once the treasure is revealed and treasure not collected,
h(n) = number of moves to the treasure + number of moves from the treasure to the goal location

Once the treasure is revealed and treasure collected,
h(n) = number of moves from the treasure to the goal location

This is admissible, but not necessarily consistent. (Consider the effect of checking a nearby item that isn't the treasure; artificially causes the heuristic to increase.)

ML/MAP Problem:

Suppose there is a bowl of chocolate candies at a party. Each candy is filled either with marshmallow filling or peanut-butter filling, but from the outside, the two types of candy are indistinguishable. Suppose you know that the friend of yours who prepared the bowl of candy made the bowl in one of three ways: the bowl either has:

25% marshmallow candy and 75% peanut butter (Option A), or
50% each kind (Option B), or **Hypotheses!**
75% marshmallow and 25% peanut butter (Option C)

Suppose you reach into the bowl and draw out three pieces of candy, one at a time. Assume the type of candy you get when you choose each piece is conditionally independent of all the other candy choices, given you know the proportions of the candy in the bowl (this is not true in general, because each time you take a piece of candy it changes the proportions in the bowl, but let's just make the math easier and assume they don't change.)

The <u>first piece you choose is peanut butter,</u> and <u>the next two are each marshmallow</u>. **Data!**

Suppose your goal for this problem is to figure out of the three different possible proportions of candy in the bowl, which one your friend used.

(a) What is the ML hypothesis after you choose the first piece of candy? After the second piece? After the third piece?

P(first is peanut butter): P(pb1)
P(second is marshmallow): P(m2)
P(third is marshmallow): P(m3)

P(b1|A) = .75 P(b1|B) = .5 P(b1|C) = .25
P(m2|A) = .25 P(m2|B) = .5 P(m2|C) = .75
P(m3|A) = .25 P(m3|B) = .5 P(m3|C) = .75

Maximum likelihood hypothesis = argmax(P(data|hypotheses))

First piece; data = pb1
P(pb1|A) = .75 vs P(pb1|B) = .5 vs P(pb1|C) = .25
Answer is A.

First *and* second piece; data = pb1, m2
P(pb1|A) * P(m2|A) vs P(pb1|B)) * P(m2|B) vs P(pb1|C)) * P(m2|C)
.75 * .25 vs .5*.5 vs .25*.75
Answer is B.
Note: We can do this because each piece is conditionally independent, given the hypotheses

First, second, and third piece; data = pb1, m2, m3
P(pb1|A) * P(m2|A)) * P(m3|A) vs
P(pb1|B)) * P(m2|B) * P(m3|B)) vs
P(pb1|C)) * P(m2|C)) * P(m3|C)

.75 * .25 *.25 vs .5*.5*.5 vs .25*.75*.75
Answer is C.
Note: We can do this because each piece is conditionally independent, given the hypotheses

(b) Suppose that you know your friend really likes peanut butter candy, so you estimate that there is a 2/3 chance the bowl is actually from Option A, and a 1/6 chance (each) from Option B and Option C. What is the MAP hypothesis after the first, second, and third pieces of candy are chosen from the bowl?

To find this, we multiply each conditional probability by the likelihood; same as above with that extra add-in:

P(A) = 2/3
P(B) = 1/6
P(C) = 1/6

First piece; data = pb1
P(pb1|A)*P(A) vs P(pb1|B)*P(B) vs P(pb1|C) *P(C)
.75 *2/3 vs .5*1/6 vs .25*1/6

Answer is A.

First *and* second piece; data = pb1, m2
P(pb1|A) * P(m2|A))*P(A) vs P(pb1|B)) * P(m2|B) *P(B) vs P(pb1|C)) * P(m2|C) *P(C)
.75 * .25*2/3 vs .5*.5*1/6 vs .25*.75*1/6
Answer is A.
Note: We can do this because each piece is conditionally independent, given the hypotheses

First, second, and third piece; data = pb1, m2, m3
P(pb1|A) * P(m2|A)) * P(m3|A))*P(A) vs
P(pb1|B)) * P(m2|B) * P(m3|B)) *P(B) vs
P(pb1|C)) * P(m2|C)) * P(m3|C) *P(C)

.75 * .25 *.25*2/3 vs .5*.5*.5*1/6 vs .25*.75*.75*1/6
Answer is A.
Note: We can do this because each piece is conditionally independent, given the hypotheses

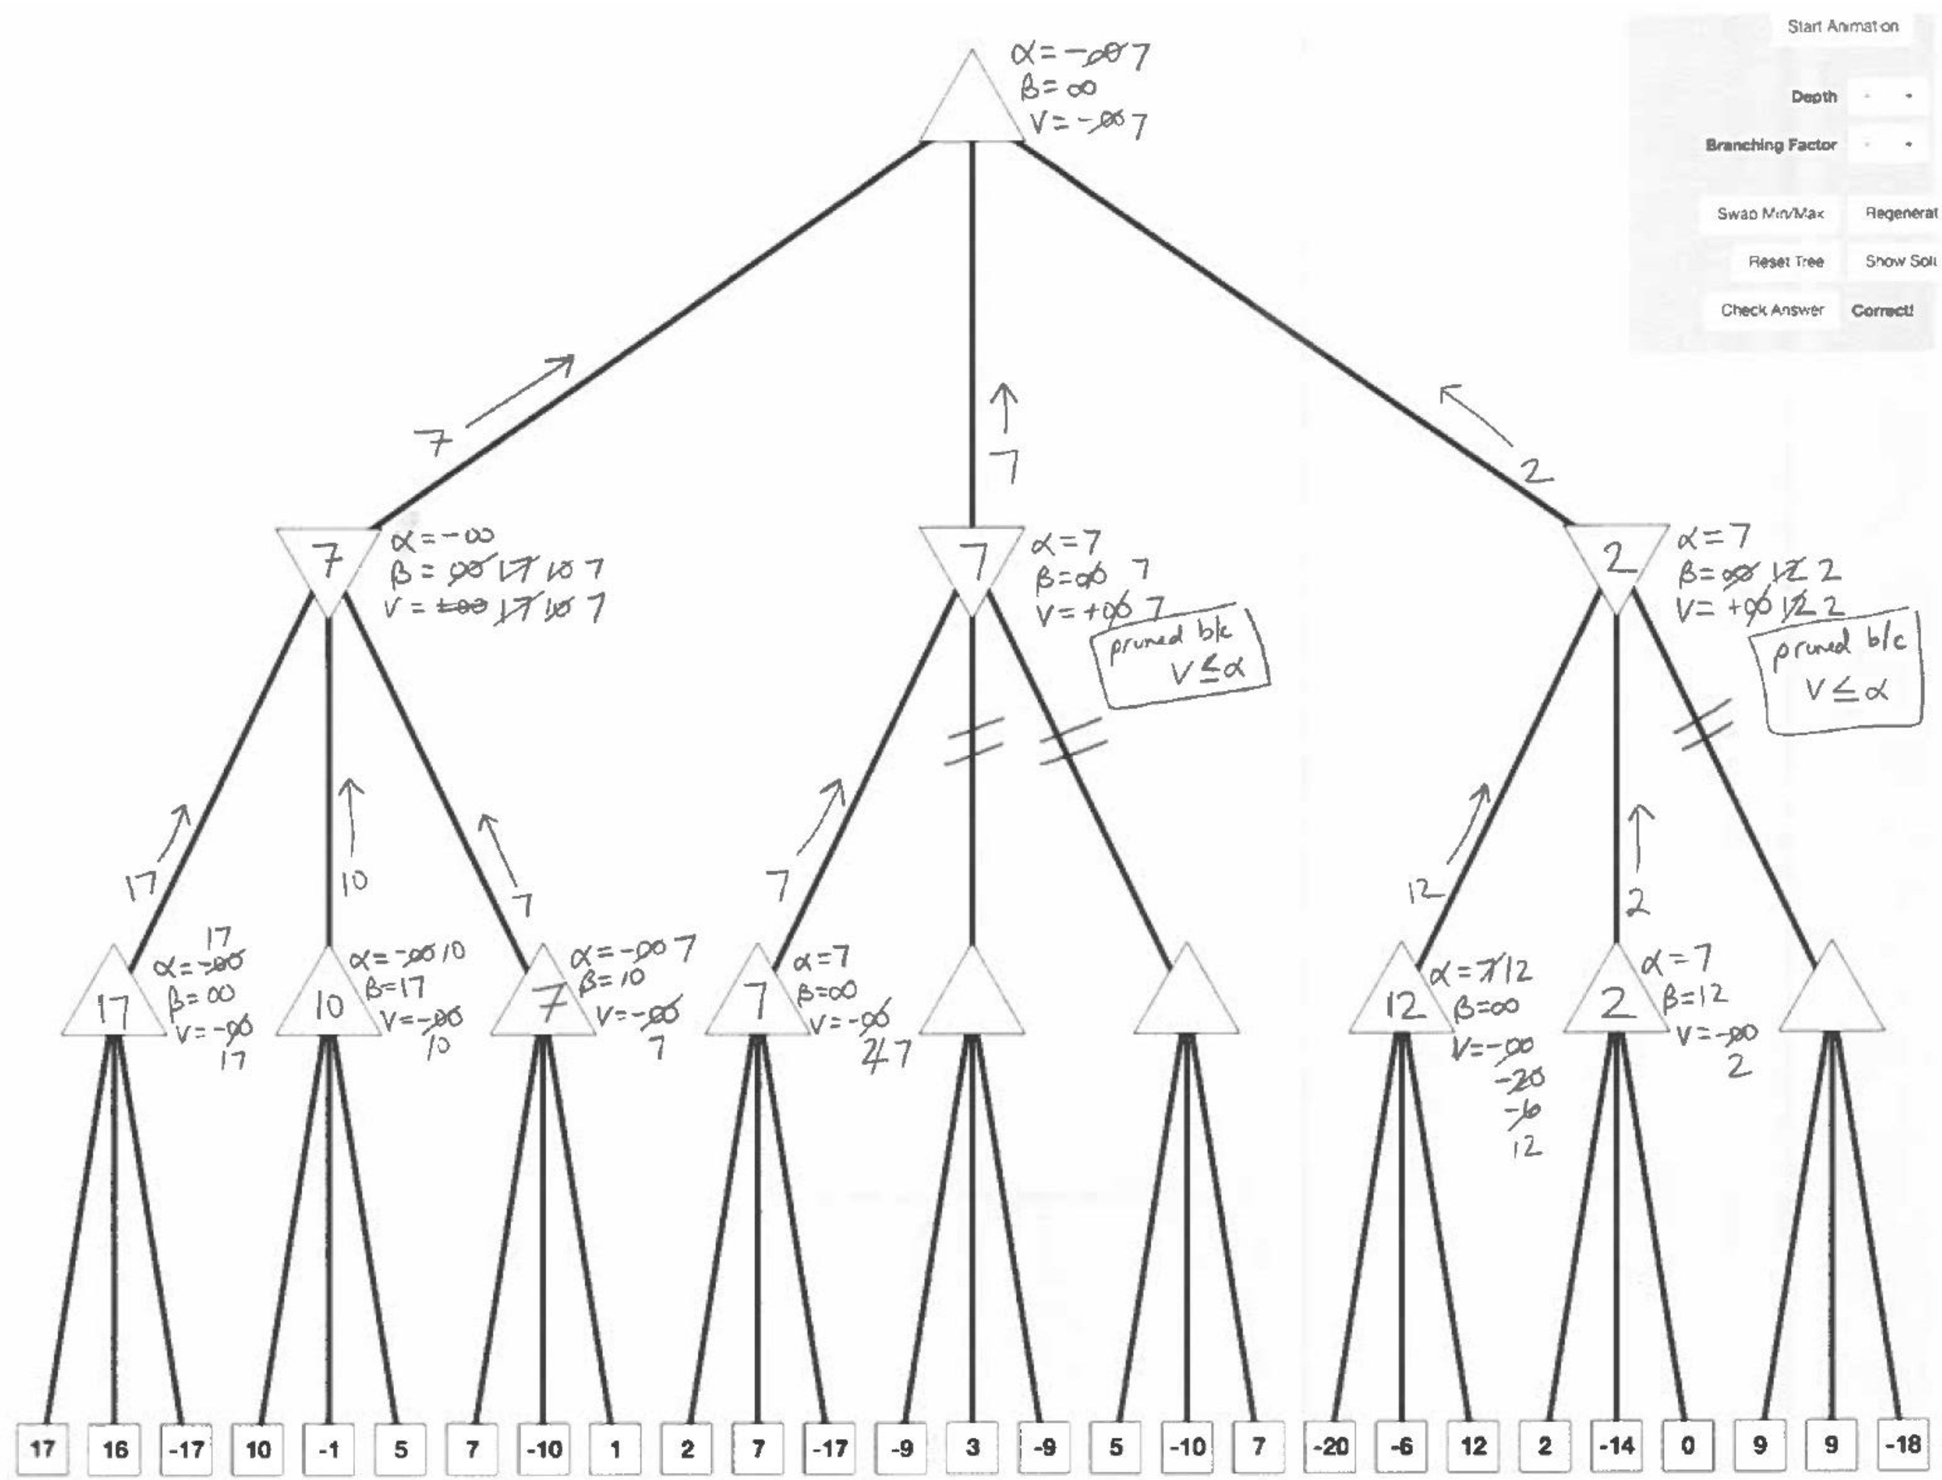
Start Animation
Depth
Branching Factor
Swap Min/Max
Regenerat
Reset Tree
Show Solt
Check Answer
Correct!
17
16
-17
10
-1
5
7
-10
1
2
7
-17
-9
3
-9
5
-10
7
-20
-6
12
2
-14
0
9
9
-18

You should be prepared for problems relating to:

- Bayes nets and probabilistic reasoning
- Djikstra's algorithm, A* algorithm, greedy best-first search
- Minimax (with both alpha/beta pruning and with heuristics)
- Which approach that we've talked about so far is best for a specific context and why
- Heuristic design (including consistent/admissible), problem setup, and discussion about the state space.