

6.01 Introduction to
Electrical Engineering and
Computer Science

Exam Written: Nov 25, 2011

Scott Guy

<u>Sections</u>	Score	Value
State Estimation	20	20
Friend or Foe	15	15
A Puzzle	16	20

$$\text{Total } \boxed{51} / \boxed{55} \\ = 92\%$$

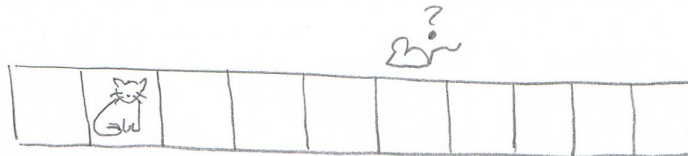
Assume 20 pts

State Estimation

Now, let's assume a particular environment for the cat and mouse, and that the cat doesn't know exactly where the mouse is. The cat and mouse live in a world that is a 1 by 10 grid. On each step, the mouse moves to one of the two neighboring (E, W) grid squares uniformly at random, unless it is currently at location 0 or 9, in which case it stays where it is with probability 0.5 and moves to the neighboring square with probability 0.5.

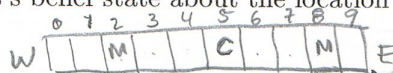
The cat always knows where it is in this world, but not where the mouse is.

The cat's actions are to move east and west, and they always have the intended effect (except when they would cause it to move off the edge of the world, in which case they have no effect). The observation is the result of listening; on each step, the cat will "hear" a distance between 0 and 10. If it is 0, then the mouse is on the same square as the cat. If it is 1, then it is one square away; if 2, two squares, away; etc. If, for example, the cat is on location 1 and it "hears" a distance of 5,



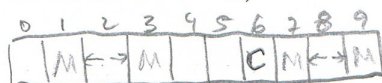
then the mouse *must* be on location 6 (because for it to be 5 squares away in the other direction, it would have to be off the end of the world.)

Question 19: Imagine that cat starts at location 5 and observes that the mouse is 3 squares away. What is the cat's belief state about the location of the mouse?



$$B = \begin{cases} MG2 = .50 \\ MG8 = .50 \end{cases} \quad \checkmark$$

Question 20: Now, let the cat move east (and the mouse moves as described above). Before making any observation, what is the state of the system? (List both the cat's location and the belief state about the mouse's location).



cat's position = 6 ✓

$$B = \begin{cases} MG1 = .25 \\ MG3 = .25 \\ MG7 = .25 \\ MG9 = .25 \end{cases} \quad \checkmark$$

* Note these results assume east is in positive direction, which is not clear from the stated problem

Question 21: If the cat now observes that the mouse is 3 squares away, what is the new belief state about the location of the mouse?

$$B' = \begin{cases} MG3 = .50 \\ MG9 = .50 \end{cases} \quad \checkmark$$

Question 22: If, instead, the cat observes that the mouse is 5 squares away (after the transition in question 20), what is the new belief state about the location of the mouse?

$$B = M@1 = 1.0$$



Question 23: Now, we'll construct a class that can do belief state estimation for this problem. The class definition and the first line of the initialization method are provided below.

```
class CatMouseState():
    def __init__(self, catLoc):
```

Write the rest of the initialization method, which takes as input the known location of the cat, and defines two instance variables: `self.catLocation`, which contains an integer from 0 to 9 indicating the cat's location; and `self.mouseBelief` which contains a list of 10 probabilities indicating the cat's belief about the mouse's location. The initial value of `self.mouseBelief` should be a representation of the distribution corresponding to having no information about the location of the mouse.

```
def __init__(self, catLoc):
    self.catLocation = catLoc
    self.mouseBelief = {0:0.1, 1:0.1, 2:0.1, 3:0.1,
                        4:0.1, 5:0.1, 6:0.1, 7:0.1,
                        8:0.1, 9:0.1}
```

Question 24: Provide the bodies of the following methods for the `CatMouseState` class. The procedures should modify the components of the object, and need not return any values. `actionUpdate` should update the state (cat's location and mouse belief state) given the cat's action (but remember that the mouse also moves), and `observationUpdate` should update the belief state based on its observation of the mouse.

Cat's action can be 'E' or 'W'

`def actionUpdate(self, action):`

```
    if action == 'E' and self.catLocation < 9:
        self.catLocation += 1
    elif action == 'W' and self.catLocation > 0:
        self.catLocation -= 1
```

Update the mouse

`newBelief = {0:0, 1:0, 2:0, 3:0, 4:0, 5:0, 6:0, 7:0, 8:0, 9:0}`

`for possibility in self.mouseBelief.keys():`

`if possibility > 0:`

`newBelief[possibility-1] += self.mouseBelief[possibility]*0`

`if possibility < 9:`

`newBelief[possibility+1] += self.mouseBelief[possibility]*1`

`self if possibility == 9 or possibility == 1:`

`newBelief[possibility] += self.mouseBelief[possibility]*0.5`

`self.mouseBelief = newBelief`

Observation: distance to mouse, as described above.

`def observationUpdate(self, obs):`

first compute $P(A \cap B)$

`for p in self.mouseBelief.keys():`

`if abs(self.catLocation - p) != obs:`

`self.mouseBelief[p] = 0`

then compute total probability

`total = 0`
`for p in self.mouseBelief.keys():`

`total += self.mouseBelief[p]`

`for p in self.mouseBelief.keys():`

`self.mouseBelief[p] /= total`

5 Friend or Foe (15 points)

Imagine that you are defending a city and there is an aircraft flying toward you. It may be important to know whether that aircraft is a 'friend' or a 'foe' (enemy). Assume you have a radar sensor that can give you noisy information about the type of the aircraft that is approaching.

We will model this situation as an HMM, in which:

- The state space is described by two components d and a , where d is the number of miles (0, 1, ..., 10) away the target is and a is its *attitude* to you, which is either 'friend' or 'foe'. The values of d correspond to ranges of distance, so that, in fact, $d == 0$ means the aircraft is somewhere between 0 and 1 miles away, etc.
- The observation space is {'oFriend', 'oFoe'}, which stands for *observed friend* and *observed foe*.
- There are no actions (or, if you prefer, a single action, which just waits a time step).
- The transition model is specified in the following Python method, which takes a state s as input and returns a `DDist` over possible next states:

```
def transitionModel(s, i):
    # note that the i (the input action) is ignored
    (d, a) = s
    if d == 0:
        return DDist({(0, a): 1.0}) # if you're next to me, must stay there
    else:
        return DDist({(d, a): 0.5, (d-1, a): 0.5}) # otherwise 50% chance of moving closer, 50% chance of moving further
```

- The observation model is as described in the following Python method, which takes a state s as input and returns a `DDist` over possible observations:

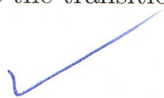
```
def observationModel(s):
    (d, a) = s
    if a == 'friend':
        return DDist({'oFriend': 1 - d / 20.0, 'oFoe' : d / 20.0})
    else:
        return DDist({'oFoe': 1 - d / 20.0, 'oFriend' : d / 20.0})
```

Questions:

1. Assume that the aircraft are approaching with a constant velocity. What velocity, in miles per time step, would generate the transition model over distance intervals given in the transition model?

Answer:

0.5 miles/time step



2. Provide an alternative transition model in which friendly aircraft move 1.5 miles per time step (and the foes move at the same rate as the given model.)

```
def transitionModel(s, i):
    (d, a) = s
    if d == 0:
        return DDist({(0, a): 1.0})
    else:
        if a == 'foe':
            return DDist({(d, a): 0.5, (d-1, a): 0.5})
        else:
            return DDist({(d-1, a): 0.5, (d-2, a): 0.5})
```

3. At what distance is our sensor most useful?

Answer:

0

How does it behave at that distance?

Answer:

if the plane is friendly, obs = 100% of being friendly
... if foe, obs = 100% of foe

4. Using the original transition and observation models, if the initial belief state, b_0 , is

$\text{DDist}(\{(10, \text{'friend'}) : 0.5, (10, \text{'foe'}) : 0.5\})$

then what would b'_0 be, after receiving observation $o_0 = \text{'oFriend'}$.

Answer:

$\text{DDist}(\{(10, \text{'friend'}) : 0.5, (10, \text{'foe'}) : 0.5\})$

5. After that, what would the belief state b_1 be?

Answer:

$\text{DDist}(\{(10, \text{'friend'}) : 0.25, (9, \text{'friend'}) : 0.25, (10, \text{'foe'}) : 0.25, (9, \text{'foe'}) : 0.25\})$

6. How many non-zero entries are there in b_4 (the answer will be the same for any sequence of observations)?

Answer:

10

$b_0 \rightarrow 2$

$b_3 \rightarrow 8$

$b_1 \rightarrow 4$

$b_4 \rightarrow 10$

$b_2 \rightarrow 6$

$0-4 = 5$ possible dist

$\times 2$ possible sta

10 possible stat

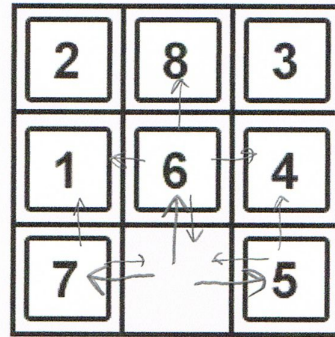
7. Starting again from the initial belief state, and imagining the following observation sequence: $[\text{'oFriend'}, \text{'oFoe'}, \text{'oFriend'}, \text{'oFoe'}]$, is it more likely that the aircraft is a friend or a foe?

Answer:

foe, closer distances are more accurate and there is a higher likelihood of proximity in each step, meaning every observation is slightly more accurate than the preceding one

7 A Puzzle (20 points)

Consider the standard *Eight Puzzle*. It has 8 tiles arranged on a 3 by 3 grid. Here is one possible arrangement of the tiles:



The tiles neighboring an empty space can be slid into the space. We can think of the operations on the puzzle as *moving the space up, down, left, or right*. It obviously cannot be moved beyond the bounds of the puzzle. In the example above, if we were to slide the space **up**, then the **6** tile would be in the bottom row and the space would be in the middle.

We can solve this problem using search. A state of the search would be an array of numbers (and None for the space) showing the location of each tile on grid. A goal state would be some specified arrangement of the tiles.

Assuming we apply pruning rule 1, but not the others, how many descendants are there for the state shown above:

- at level 1 (immediate children)?

Answer:

3



- at level 2 (children of level 1)?

Answer:

8



7.1 Search

We would like to identify the advantages and disadvantages of each of the following search methods **for this problem**. We assume, as always, that we use pruning rules 1 and 2.

Enter T or F in the boxes provided if the following statement is true or false, respectively. You can assume that it is possible to reach the goal state from the initial state.

- **depth-first, no dynamic programming:**

- ☒ This method is guaranteed to find a path.
- ☒ The path that is found is guaranteed to be short.
- ☒ The same board state may be visited multiple times.
- ☒ The agenda is likely to remain short (relative to the other methods).

- **depth-first, with dynamic programming:**

- ☒ This method is guaranteed to find a path.
- ☒ The path that is found is guaranteed to be short.
- ☒ The same board state may be visited multiple times.
- ☒ The agenda is likely to remain short (relative to the other methods).

- **breadth-first, no dynamic programming:**

- ☒ This method is guaranteed to find a path.
- ☒ The path that is found is guaranteed to be short.
- ☒ The same board state may be visited multiple times.
- ☒ The agenda is likely to remain short (relative to the other methods).

- **breadth-first, with dynamic programming:**

- ☒ This method is guaranteed to find a path.
- ☒ The path that is found is guaranteed to be short.
- ☒ The same board state may be visited multiple times.
- ☒ The agenda is likely to remain short (relative to the other methods).

7.2 State machine

We can describe the puzzle and its evolution using a state machine, in much the same way we described the wolf-goat-cabbage problem. We will specify the state with two components:

- A pair of *row*, *column* indices, indicating where the space is; and
- A list of three lists of items; each item is a digit between 1 and 8, or **None**. This describes the state of the board.

So, for example, we could represent the puzzle state above as:

```
((2, 1), [[2, 8, 3], [1, 6, 4], [7, None, 5]])
```

Technically speaking, we don't need the first component of the state (it could always be computed from the list of three lists), but it will simplify our coding if we maintain both representations.

A skeleton of the state machine defining the eight puzzle is shown on the next page. Fill in the definition of `getNextValues`. If a move would produce an illegal state, then the new state should be the same as the current state. The output of the machine should just be the next state.

[2, 8, 3]
[1, 6, 4]
[7, None, 5]

```
class EightPuzzle(SM):
    startState = ((2, 1), [[2, 8, 3], [1, 6, 4], [7, None, 5]])
    size = 3
    offsets = {'up': (-1, 0), 'down': (1, 0), 'left': (0, -1), 'right': (0, 1)}
    legalInputs = ['up', 'down', 'left', 'right']

    def getNextValues(self, state, inp):
        ((p1, p2), board) = state
        bcopy = copy.deepcopy(board)
        (rOff, cOff) = offsets[inp]
        (r, c) = (p1 + rOff, p2 + cOff)
        if r <= 2 and r >= 0 and c <= 2 and c >= 0:
            square = board[r][c]
            board[r][c] = None
            board[p1][p2] = square
            (p1, p2) = (r, c)
        return (((p1, p2), board), ((p1, p2), board)))
```