

## Reinforcement Learning & HMM

### 1 Learning with Feature-based Representations

We would like to use a Q-learning agent for Pacman, but the state size for a large grid is too massive to hold in memory. To solve this, we will switch to feature-based representation of Pacmans state.

- Let's assume our two minimal features are the number of ghosts within 1 step of Pacman ( $F_g$ ) and the number of food pellets within 1 step of Pacman ( $F_p$ ). Calculate  $F_p$  and  $F_g$  for the following pacman board:




- With Q Learning, we train off of a few episodes, so our weights begin to take on values. Right now  $w_g = 70$  and  $w_p = 5$ . Calculate the Q value for the state above
  
  
- We receive an episode, so now we need to update our values. An episode consists of a start state  $s$ , an action  $a$ , an end state  $s'$ , and a reward  $R(s, a, s')$ . The start state of the episode is the state above (where you already calculated the feature values and the expected Q value). The next state has feature values  $F_g = 0$  and  $F_p = 4$  and the reward is 50. Assuming a discount of 0.4, calculate the new estimate of the Q value for  $s$  based on this episode.



## 2 Hidden Markov Model

Consider the following Hidden Markov Model:



Suppose that  $O_1 = A$  and  $O_2 = B$  is observed. Use the Forward algorithm to compute the probability distribution  $Pr(X_2, O_1 = A, O_2 = B)$ . Show your work.

$X_1$	$P(X_1)$
0	0.25
1	0.75

$X_t$	$X_{t+1}$	$Pr(X_{t+1} X_t)$
0	0	0.3
0	1	0.7
1	0	0.6
1	1	0.4

$X_t$	$O_t$	$Pr(O_t X_t)$
0	A	0.8
0	B	0.2
1	A	0.5
1	B	0.5

### 3 DNA Sequence

A DNA sequence is a series of components from  $\{A, C, G, T\}$ . Now let's assume there is one hidden variable  $S$  that controls the generation of DNA sequence.  $S$  takes 2 possible states  $\{S_1, S_2\}$ . Assume the following transition probabilities for HMM  $M$ :

$$P(S_1|S_1) = 0.75, P(S_2|S_1) = 0.25, P(S_1|S_2) = 0.25, P(S_2|S_2) = 0.75$$

emission probabilities as following:

$$P(A|S_1) = 0.4, P(C|S_1) = 0.1, P(G|S_1) = 0.4, P(T|S_1) = 0.1$$

$$P(A|S_2) = 0.1, P(C|S_2) = 0.4, P(G|S_2) = 0.1, P(T|S_2) = 0.4$$

and start probabilities as following

$$P(S_1) = 0.5, P(S_2) = 0.5$$

Assume the observed sequence is  $x = \text{CGTCAG}$ , calculate  $P(x|M)$  using the forward algorithm. Draw the HMM, and show your work to get full credit.