
CAP 5636 – Midterm 2

Date: November 14, 2017

Name:

Instructions:

- This exam is open book and open notes. Textbooks and notes on tablet devices are acceptable but they must be put into airplane mode. No device with a keyboard is acceptable.
- Allotted time is 75 minutes.
- In your answer, please follow the number of sentences as specified.
- Note that the points add up to 100 + 20 bonus points.

Problem 1 (40 pts)

Rain influences whether the sprinkler is activated, and both rain and the sprinkler influence whether the grass is wet. The following table describes the joint probability of the rain, sprinkler and wet grass:

Rain	Sprinkler	GrassWet	Joint probability
-rain	-sprinkler	-grasswet	0.475200
-rain	-sprinkler	+grasswet	0.004800
-rain	+sprinkler	-grasswet	0.032000
-rain	+sprinkler	+grasswet	0.288000
+rain	-sprinkler	-grasswet	0.039600
+rain	-sprinkler	+grasswet	0.158400
+rain	+sprinkler	-grasswet	0.000020
+rain	+sprinkler	+grasswet	0.001980

Calculate the following probabilities and explain the calculations:

a) $P(+rain)$

b) $P(+grasswet)$

c) $P(+grasswet \mid -rain)$

d) $P(+grasswet \mid -rain, -sprinkler)$

Problem 2 (40 pts)

Your neighbors in apartment 13B like to party. Parties are usually noisy, but the students in apartment 13B might be noisy even when there is no party. Their behavior is described by the following Hidden Markov Model:

Party(n)	Party (n+1)	P(Party(n+1) Party(n))	Party(n)	Noise (n)	P(Noise(n) Party(n))
+party	+party	0.2	+party	+noise	0.8
+party	-party	0.8	+party	-noise	0.2
-party	+party	0.3	-party	+noise	0.5
-party	-party	0.7	-party	-noise	0.5


We know that in Day 0, there was no party. We have the following observations: on Day 0 there is no noise, while on Days 1 and 2, there was noise.

- a) What algorithm can you use to calculate the probability of apartment 13B having a party on Day 2?

- b) Calculate the probability of apartment 13B having a party on Day 2 using the algorithm in your answer to a). Show the calculations as well as the result.

- c) What is your prediction about apartment 13B having a party on Day 3? Show how you reached this conclusion

Problem 3 (40 pts)



BlackDisk lives in the world of four squares shown above. Every timestep it must make an action of moving Up, Down, Left or Right. There is no uncertainty in the movements: BlackDisk will always end up where it is headed, except when it hits the wall, in which case it stays in place. The figure shows its current Q values. Every time BlackDisk finishes an action in a gray state, it will get a reward of +30. Every time it ends up in a white state, it gets a reward of 0.

BlackDisk is trying to improve its life using Q learning with a learning rate $\alpha=0.3$ and discount factor $\gamma = 0.9$.

a) Trace two steps of Q-learning. Show the location of the agent after these steps. Show the changes of the Q-values at each step

b) What do you think the optimal policy is for BlackDisk?