CSci 5512
Midterm 2

Name: _____

Student ID: _____

Instructions: The time limit is 75 minutes. Please write your answers in the space below. The exam is open book and notes. You may use electronic devices to ONLY look at either an e-book version or electronic notes. You may not use the internet, program/run code or any other outside resources. For all questions you must **show work**.

Problem (1) [20 points]
Consider the Hidden Markov Model shown below:

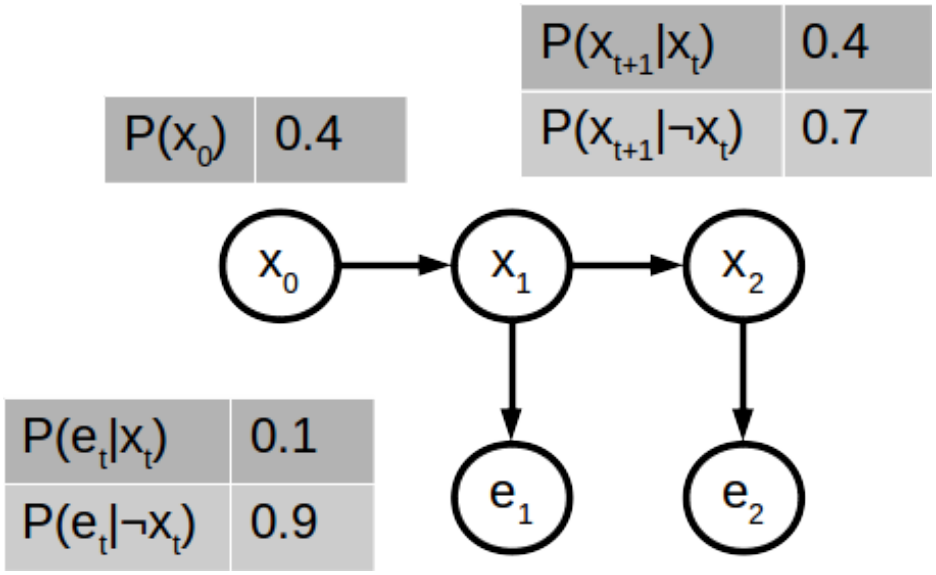

Find $P(x_1|e_1, \neg e_2)$.

Problem (2) [20 points]
Using the same HMM as the last problem (copied below):

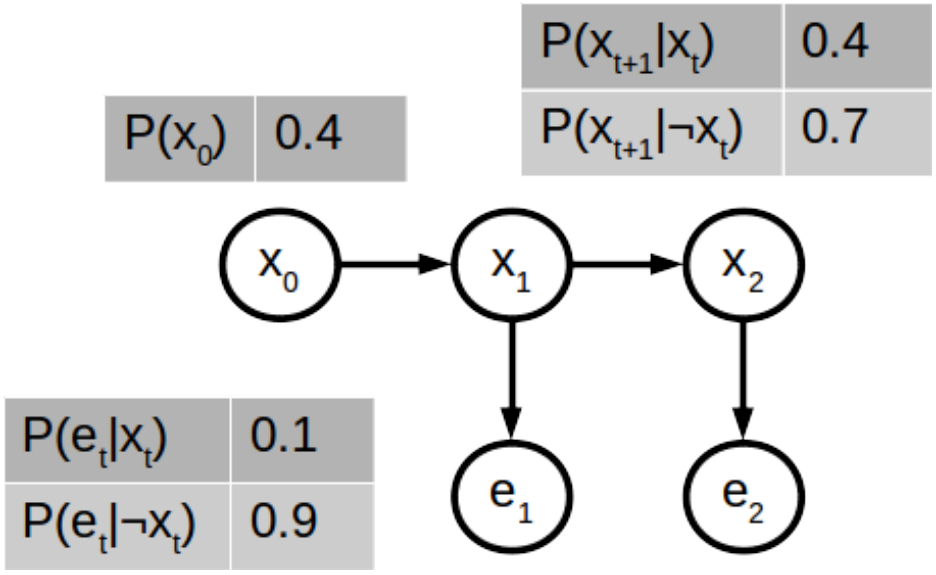

What is the result after doing particle filtering for two iterations (so up until "resampling" on x_2) with evidence: $\neg e_1, e_2$ (opposite of last time). Additionally, please use the following sequence of random numbers in order for your calculations, so the first two should be used to determine the initial distribution (e.g. if the probability you want to use is 50%, you should consider "random < 0.5" as a "true" outcome): (Note: there are more random numbers than you need.)
0.46, 0.82, 0.37, 0.79, 0.18, 0.07, 0.22, 0.39, 0.66, 0.43, 0.33, 0.90, 0.01, 0.42, 0.75
Obviously using only two particles are far too few. When working through the problem above, where does this stand out the most?

Problem (3) [20 points]
Assume we are using a Kalman filter which is initially distributed with $N(0, 7)$ ($\sigma^2_{t=0} = 7$) and the transition has variance/σ^2 of 5.
(1)How accurate does the evidence observation need to be for the estimate after one transition to have a variance/σ^2 of exactly 3?
(2) If the evidence shows a value 8 after one transition, where is mean for the actual location?

Problem (4) [20 points]
Shown below is a Markov Decision Process with each cell's reward. Assume there are only two actions: moving up or down. The results of each action are 50% moving in the indicated direction and 50% not moving anywhere. Use $\gamma = 0.5$.

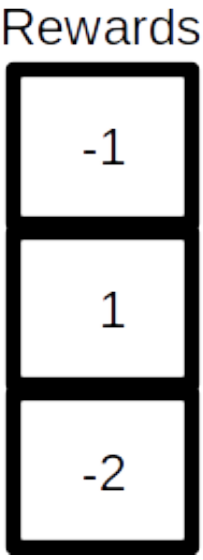

(1) If you initialize all states utility as zero, what will be their estimated utility after one step/iteration of value iteration?
(1) What are the estimated utilities for each cell after a second step/iteration of value iteration?
(3) If the actions had you move in the indicated direction 100% of the time, would this cause any issues for value iteration? How about policy iteration?

Problem (5) [20 points]
This is the same MDP from the previous problem (on the left), and again assume the movement is 50% in the indicated direction and 50% not moving at all. Both the initial distribution and probability of seeing evidence is shown as well.

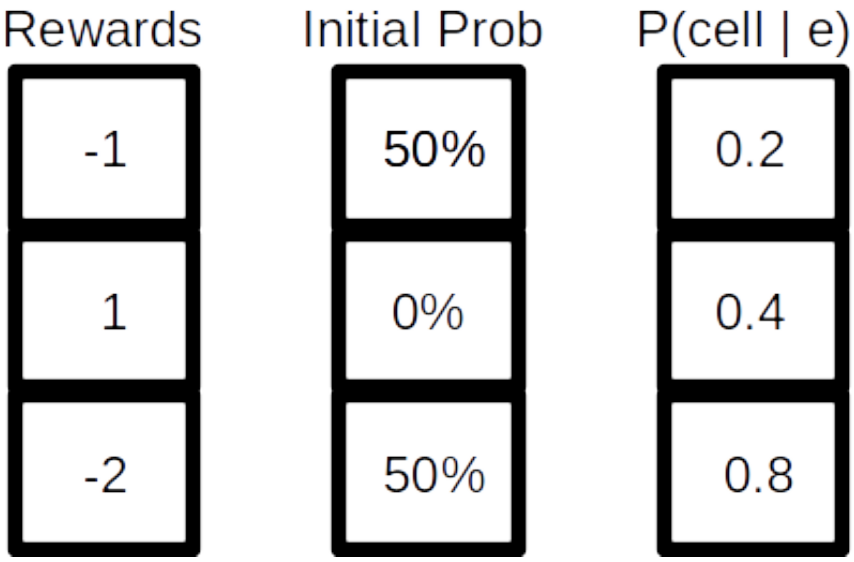

(1) If you move "up" and see evidence "e_1" (true), what is your probability of being in each cell now?
(2) What is the probability that you would see "e_1" (true) after moving up? (As opposed to $\neg e_1$.)