CS 4758/6758: Robot Learning: Homework 5 part 2

1 POMDP [20 pts]

This problem is known as the tiger problem. A cute little robot is facing two doors. Behind one is a terminator robot, behind the other a reward of +10. The robot can either listen or open one of the doors. When opening the door with the terminator robot, the robot will be annihilated, which has an associated cost of -20. Listening costs -1. When listening, the robot will hear a crushing noise that indicates the presence of the terminator robot, but only with 0.85 probability will the robot be able to localize the noise correctly. With 0.15 probability, the noise will appear as if it came from the door hiding the reward.

Your questions:

(a) Provide the formal model of POMDP, in which you define the state, action, and measurement spaces, the cost function, and the associated probability functions.

(b) What is the cumulative payoff/cost of the open-loop action sequence: "Listen, listen, open door 1"? Explain your calculation.

(c) What is the expected cumulative payoff/cost of the open-loop action sequence: "Listen, then open the door for which we did not hear a noise"? Again, explain your calculation.

2 Particle filters [10 pts]

Consider the following proposition: There is no point in using a particle filter if the number of distinct states at any time is no more than the number of particles one plans to allocate to the particle filter. Justify or refute this statement.