COMP.4510 and COMP.5490
Fall 2017

Problem Set 1: Introduction to Probability, Bayes Rule, and Bayes Nets

Out: Thursday, 7 September 2017
Due: Thursday, 14 September 2017 at start of class

If you have not recently studied probability, Bayes Rule, and Bayes Networks, you will need to be re-acquainted to it, because Chapter 2 of Probabilistic Robotics simply states the conclusions with the assumption that you have already understood them. Start out by working through the mini-lectures from last fall's Stanford AI online class. The lectures include little quizzes which test your understanding, which you should do. While you do not need to turn in anything for this portion of the assignment, this knowledge is the basis for your future success it the course.

Watch videos 1 through 34 of Unit 3 (through to the video “Unit 3 8d Answer”) at https://www.youtube.com/playlist?list=PLA9DA7AA5FBDBA44A. (Link on course web page.)

Important note: As you watch the videos, try to pause at the end of the questions so the playlist doesn't flow right into the answers. That will give you a chance to work on them.

Including doing problems along the way and viewing the answer lectures, this will take you two to three hours – but this is an essential basis for moving forward in the course.

Important note: Many of Thrun's examples are for a hypothetical case of cancer, given evidence from a cancer-test. To translate this into robotics, consider the “probability of cancer” to be analogous to some state of the robot and its world, e.g., “is the door open?” Then, the test-for-cancer translates into a sensor reading. So the question, “Given a positive test, what is the likelihood that I have cancer?” is mathematically equivalent to “Given that my sensor seems to see that the door is open, what is the likelihood that it's really open?” Keep this in mind as you work through the Stanford material.

After completing the online material:

- Read pages 3 through 10 of Chapter 1.
- From Chapter 2, “Recursive State Estimation,” read the introduction 2.1 and section 2.2, “Basic Concepts in Probability.”

I highly suggest that you complete this portion of the assignment before class on Tuesday, 12 September.

<Problem set continues on next page>

Probability Problems (turn in solutions to these problems)

Problem 1

Consider the Bayes Net below:

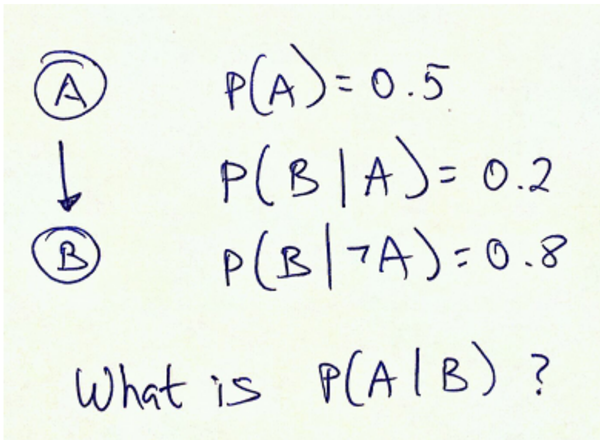

What is the value of P(A | B) ? Show your work.

Problem 2

Consider the representation of a robot and its sensor below. The robot is moving one-dimensionally along a wall. The robot's state is represented by the X variable, which may be in one of only two states: *at-wall* (as shown in the drawing), or *at-door*. (The position of the robot along the corridor is not being modeled.)

The robot has a sensor which gives a value represented by the Z variable, which is either *see-wall* or *see-door*.

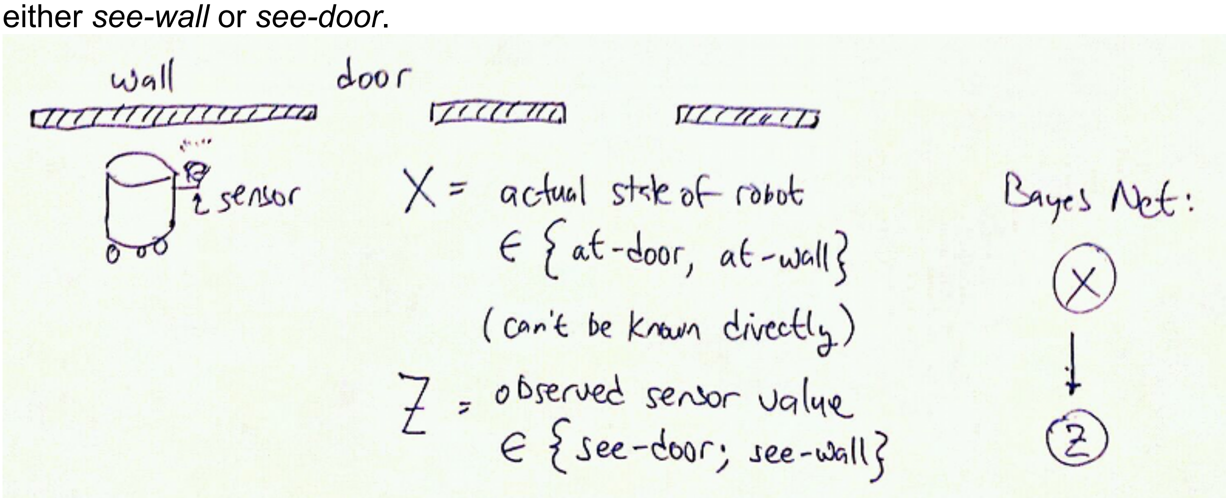

Remember that the X variable, representing the robot's “true” state, cannot be known directly. It only can be inferred probabilistically given the Z sensor values. (This is the fundamental modeling assumption of probabilistic robotics.)

This model is represented by the Bayes net on the right of the diagram, where the X state variable has influence on the Z sensor measurements.

The door sensor behaves as follows.

P(Z = see-door | X = at-door) = 0.95
P(Z = see-wall | X = at-wall) = 0.9

In other words, the sensor is slightly more reliable at seeing a door than seeing a wall. (If you imagine the door sensor as a robot sonar device, this is reasonable because sensing a door opening represents the lack of return ping, whereas sensing a wall requires the ping to be returned. Sometimes pings go out and are lost.)

Problem 2.1

Given the above, specify the values of:

P(Z = see-wall | X = at-door)

and

P(Z = see-door | X = at-wall)

Problem 2.2

Assume a prior probability P (X = at-door) = 0.5, and compute the posterior probability X = at-door given a sensor reading of Z = see-door. In other words, what is

P(X = at-door | Z = see-door)

Show your work.

Problem 2.3

Given the prior **P (X = *at-wall*) = 0.5**, compute the posterior probability X = *at-wall* given a sensor reading of Z = *see-wall*. In other words, what is

P(X = at-wall | Z = see-wall)

Problem 2.4

Which posterior gives more confidence of the robot's true state: being at the door and seeing a door, or being at the wall and seeing a wall? Explain why.

Problem 3

Complete Exercise 1 on page 36 of PR. Compute specific values for the posterior probability of a sensor fault for N= 1, 2, and 10 consecutive readings of less than 1m.