Artificial Intelligence
Spring 2017
Homework 3

Show all your work for full credit!

1. Suppose a person goes for a heart check-up and the doctor tries to ascertain whether that person has a heart condition based on three characteristics: whether a patient is a *smoker*, the patient's *blood pressure*, and their *electrocardiograph (ECG) reading.* The characteristics can take on the following values:

   - Smoker: Yes or No
   - Blood pressure: Low or High
   - ECG reading: Normal or Abnormal

   We are given the following information:

   $$P(\text{having a heart condition}) = 0.1$$

   $$P(\text{patient is a smoker} \mid \text{a patient does not have a heart condition}) = 0.2$$
   $$P(\text{patient is a smoker} \mid \text{a patient has a heart condition}) = 0.4$$

   $$P(\text{high blood pressure} \mid \text{a patient does not have a heart condition}) = 0.6$$
   $$P(\text{high blood pressure} \mid \text{a patient has a heart condition}) = 0.7$$

   $$P(\text{abnormal ECG} \mid \text{a patient does not have a heart condition}) = 0.5$$
   $$P(\text{abnormal ECG} \mid \text{a patient has has a heart condition}) = 0.8$$

   You may assume the patient's smoker status, blood pressure, and ECG reading are all conditionally independent of each other given the presence or absence of a heart condition.

   (a) A patient who smokes has low blood pressure but an abnormal ECG. What are the maximum likelihood and MAP hypotheses regarding whether or not this patient has a heart condition?

   (b) What is the posterior probability of this patient having a heart condition?

   (c) A patient who smokes, has high blood pressure, and an abnormal ECG will still be predicted as *not* having a heart condition under the MAP hypothesis with the probabilities given here. What would the prior probability of a heart conditional have to be in order for this patient to be predicted to have a heart condition under the MAP hypothesis? Hint: Your answer should involve a greater-than or less-than sign.

2. The Rhodes College Math & CS Department is trying to do a better job at predicting who will become math and CS majors based on what classes they take. *For this problem, assume there are no double majors.* Suppose we know the following information:

   - There are 532 sophomore students at Rhodes this year.
   - 44 of those 532 students declared a CS major.
   - All of the 44 are currently taking a CS class.
   - 18 of the 44 are currently taking a math class.
   - 11 of those 532 students became Math majors.
   - 1 of the 11 is currently taking a CS class.
   - All of the 11 are currently taking a math class.
   - 477 of those 532 students decided to major in something else.
   - 28 of the 477 are currently taking a CS class.
   - 96 of the 477 are currently taking a math class.

   (a) Show the prior probabilities calculated for each of the hypotheses (CS major, math major, other major), as well as the probabilities for each of the features given each of the hypotheses.

   Two important points:

   Because many of the quantities that you are probabilities for have similar names (CS class, CS major, math class, math major, etc), I suggest the following naming scheme. Call your hypotheses H_{CS}, H_{M}, and H_{O}, for CS major, math major, and other major, respectively; and call your features F_{CS} and F_{M}, for being enrolled in a CS class and being enrolled in a math class, respectively. So for instance, $P(H_{\text{CS}})$ would be the (prior) probability that someone majors in CS, and $P(F_{\text{M}} \mid H_{\text{CS}})$ would be the probability that someone takes a math class given that they are a CS major.

   Remember that the probabilities for each feature given the hypothesis should be smoothed, but the priors for the hypotheses should not be smoothed. See your notes/slides for smoothing calculations.

   (b) If next year, a sophomore takes both a math class and a CS class in the spring, what major will they probably have? (show the calculations for the MAP hypothesis).

   (c) Give the probabilities of this sophomore (from part b) declaring each of (1) a CS major, (2) a math major, and (3) a different major.

3. *Note: This problem involves multiplying vectors and matrices. You are welcome to use software to do the multiplications itself, however, you should show all your work. Another way of thinking about this is that you should understand the math well enough to do it by hand if I asked you to. That is, you should know how to multiply 2-by-2 matrices and vectors by hand. (Hint, hint.)*

Professor Somnus is investigating whether the students in his class are getting enough sleep. He collects some data and deduces that the probability that a student will get enough sleep on a given night only depends on whether they got enough sleep the night before. If a student gets enough sleep the previous night, the probability they will get enough sleep tonight is 0.8. If they didn't get enough sleep the previous night, the probability they will get enough sleep tonight is only 0.3.

(a) Formulate this problem using a Markov chain.

(1) Draw the Markov chain diagram showing the probability of transitioning between "enough sleep" and "not enough sleep." (This is the not the Bayes net diagram; this is the diagram that shows up at the top of the page at en.wikipedia.org/wiki/Markov_chain.)

(2) Write down the transition matrix T.

(b) Assuming you got enough sleep on night 0, what's the probability you get enough sleep on night 3?

(1) Show the initial state vector, v_0.

(2) Calculate and show the probability distribution for nights 1, 2, and 3 by multiplying by T.

(c) In the far, far, future, what is the probability of getting enough sleep on some night?

4. Continuing with the situation in Problem 3, Professor Somnus has no way of directly observing whether or not his students are getting enough sleep. All he can observe is whether they are falling asleep in his class or not. He knows that if a student gets enough sleep on some night, the next day there is a 0.1 probability that they will fall asleep in class. If they don't get enough sleep the night before, there's a 0.3 probability of falling asleep in class.

For this question, assume that night x is followed by day x.

(a) Professor Somnus observes a certain student falling asleep in class on day 1, staying awake on day 2, but falling asleep again on day 3. Calculate the probability, using the forward algorithm, that the student got enough sleep on night 3, given that sequence of observations. (Remember, night 3 happens right before day 3). The professor assumes there's an equal prior probability of enough sleep/not enough sleep on night zero.

(1) Write out the appropriate matrices O_1, O_2, and O_3.

(2) Write out the initial vector f_0.

(3) Use the forward algorithm to calculate the probability distributions for enough sleep/not enough sleep on nights 1, 2, and 3. Show your work.

(b) Professor Somnus remembers that in Hidden Markov Models, the observations can affect our beliefs about the past, as well as the future. Use the backward algorithm to re-calculate Professor Somnus's belief of getting enough sleep on night 1, given the the sequence of observations. Show your work.