CS475 Spring 2006
Homework 3

Due Date: Friday March 31 2006

Problem 1

1. Exercise 15.1 from the AIMA book.
2. Augment the Rain-Umbrella model from the book (page 540, figure 15.2) to include information about whether or not the umbrella is actually wet as an observable (assume that sometimes people take umbrellas to work "just in case", even if it is not actually raining) and humidity and pressure as state variables partially influencing whether it rains or not. Draw the resulting Dynamic Bayesian Network and explain why you included the edges you did, and excluded others.
3. Now turn the DBN from above into a Hidden Markov Model. Make all necessary assumptions (in particular, do you need to discretize the state values from the DBN?). Describe your states (and observables), and compute how many there are in your model. Would all transitions (from any state to any other) be allowed in the transition model of your HMM?

Problem 2

Exercise 15.2 from the AIMA book. It might help you to first write a little program that computes the probabilities and look at how it behaves (optional for part (a)). Then try to show the answers to the questions rigorously (the matrix interpretation from page 549 might help you).

Problem 3

Exercise 15.4 from the AIMA book. Construct a model that can serve as a counter-example to the flawed "smoothing-one-by-one" algorithm. Then compute the most-likely sequence given some observations and show that such sequence is not possible in the model you constructed.

Problem 4

Consider the Rain-Umbrella world again. Assume that $P(R_0) =< 0, 1 >$ (e.g. it is not raining when you go to the underworld). Otherwise use the values given in the book (page 540, figure 15.2). In the following, u_i indicates whether the umbrella was observed on day i. Compute:

- $P(R_3|u_1 = true, u_2 = false, u_3 = true)$ (filtering)
- $P(u_1 = true, u_2 = false, u_3 = true)$ (likelihood of evidence sequence)
- $P(R_5|u_1 = true, u_2 = false, u_3 = true)$ (prediction)
- $P(R_2|u_1 = true, u_2 = false, u_3 = true)$ (smoothing)
- find the most likely sequence of rain/no-rain for days $1 \ldots 4$ ($\{r_i\}_{i=1}^{4}$) given that $u_1 = true, u_2 = false, u_3 = true, u_4 = true$. Use the Viterbi algorithm.

You may find applet at
`http://i13pc1.ira.uka.de/speechCourse/slides/hmm/applet/applet.html`
useful for checking your answers. The model they have should (after appropriate mapping of states) correspond to our example. You can change the values in the transition model by clicking at the blue numbers, and values of the sensor model by clicking at the little sun/cloud icon bellow each state ("sun" and "cloud" are the two observable values).