Machine Learning *CS 726, Fall 2011*

HW12: HMMs

Hand in at: `http://www.cs.utah.edu/~hal/handin.pl?course=cs726`. Remember that only PDF submissions are accepted. We encourage using LaTeX to produce your writeups. See `hw00.tex` for an example of how to do so. You can make a `.pdf` out of the `.tex` by running "`pdflatex hw00.tex`".

1. We all know the weather in College Park is wacky. We want to be able to figure out what the weather is like by touching the ground. The ground can either be Hot (H), Cold (C) or Wet (W). Let's say each day is either a good day (G) or bad day (B), and that the weather operates according to the semantics of a first order Markov model.

   More formally, each day is either G or B. Good days yield Hot grounds with probably 2/3, Cold grounds with probability 1/4 and Wet grounds with probability 1/6. Bad days yield Hot grounds with probability 1/4, Cold grounds with probability 1/4 and Wet grounds with probability 1/2. If today is Good, then there's a 2/3 chance tomorrow will be Good. If today is Bad, then there's a 1/2 chance tomorrow will be Bad. Suppose that at the beginning of time there's a 50-50 chance of it being Good or Bad weather.

   (a) If you models this as an HMM, what are the probability distributions like? You should write everything that's relevant in the form $p(\text{blah} \mid \text{blah}') = \text{something}$.

   (b) Suppose we have no idea what the weather has been like for the past few days, but we measure the ground and find out it's Cold. What is the probability distribution over weather given this observation.

   (c) Suppose that we observe the following sequence of ground observations: C W C H. What is the most likely sequence of states the world went through to produce these observation (that is, run the Viterbi algorithm by hand).