ENGN 2520 Homework 4
Due Friday April 6 by 4pm

Problem 1

Suppose we have 4 binary random variables CLOUDY, RAIN, SPRINKLER, WETGRASS.

CLOUDY=True if the sky is cloudy today.

RAIN=True if it rains today.

SPRINKLER=True if the sprinkler comes on today.

WETGRASS=True if the grass is wet today.

(a) Suppose the sprinkler controller has no sensors. Draw a Bayesian network that models the natural relationships between these random variables.

(b) How many total parameters do you need to specify to define the necessary probability distributions associated with the bayesian network?

(c) How many parameters would you need to specify an arbitrary joint distribution over the random variables?

(d) What is the Markov blanket for SPRINKLER?

Problem 2

Suppose we have an HMM with states $S = \{s_1, \ldots, s_K\}$, observations $O = \{o_1, \ldots, o_M\}$, initia state distribution π, transition matrix B and observation matrix A.

Let $Y = \{y_1, \ldots, y_T\}$ be a sequence of observations. Let a, b and c be three states. Show how we can efficient compute the conditional probability $P(x_t = a, x_{t+1} = b, x_{t+2} = c|Y)$ using forward and backward weights.

Problem 3

In this assignment you will implement the Viterbi algorithm for correcting corrupted english sentences. For this assignment you should assume each english sentence is a sequence of symbols, where each symbol is one of the 26 lowercase letters or the space character.

We will assume that each correct english sentence is generated by a first order Markov model. We observe each sentence through a noisy channel that corrupts each symbol with probability ϵ. If a symbol is corrupted it gets replaced by an arbitrary symbol with equal probability (a symbol can be replaced by itself). This process leads to a HMM. The class website contains examples of uncorrupted english sentences that you will use to train the parameters of the model.

(a) What is the state space for the hidden variables? What is the observation space?

(b) What is the observation matrix B in terms of ϵ?

(c) Use the uncorrupted sentences available on the class website to estimate the initial state distribution π and the transition matrix A.

(d) Use your model to generate 10 independent samples from the joint distribution $p(x_1, \ldots, x_T, y_1, \ldots, y_T)$ when $\epsilon = 0.1$ and $\epsilon = 0.2$.

(e) Implement the Viterbi algorithm for computing the most likely sequence of hidden states $(x_1, \ldots, x_T)$ given a sequence of observations $(y_1, \ldots, y_T)$. Use the algorithm to "correct" the corrupted sentences available on the class website using the model from part (d) when $\epsilon = 0.1$.