CS 530 — Principles of AI
Written Exercises
Out: November 8, 2001
Due: November 20, 2001

Problem 1. The goal of this problem is to familiarize yourself with HMMs, and to get a sense of some of the strengths and weaknesses of HMMs as models of the world.

We model the following situation. We are observing an agent. The agent has two coins, possibly unfair; the agent starts with a specific coin c_1. At each step, the agent selects a coin—perhaps the same one as last time with some probability, perhaps the other one—and flips it. You get to observe the result, which is either heads h or tails t. You do not know which coin was used, so you do not know what state the agent winds up in.

Overall then we have an HMM which can be visualized as follows:

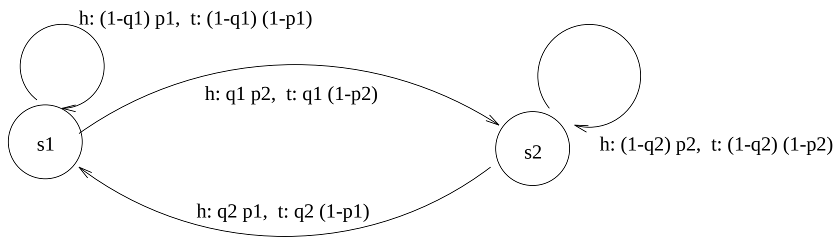

There are two states, s^1 and s^2, corresponding to the coin in use; there are two possible observations h and t. There are four basic probabilistic parameters: p_1 is the probability of seeing a head in transitions that end in state s_1; p_2 is the probability of seeing a head in transitions that end in state s_2; q_1 is the probability of moving to state s^2 when we start in state s^1, and q_2 is the probability of moving to state s^1 when we start in state s^2. Particular probabilities are assigned to complete transitions by multiplying as appropriate, as shown in the diagram. For example, $P(s^1 \xrightarrow{t} s^2) = q_1(1-p_2)$ where q_1 gives the probability of the state change and $(1-p_2)$ is the probability of observation t resulting.

1a. Suppose the coin in s^1 *always* comes up heads and that in state s^2 always comes up tails, but that each state changes to the other (or stays the same) with probability .5. Draw a specialized version of the HMM above to describe this situation, filling in specific numbers for the probabilities and omitting transitions with zero probability.

1b. For a given sequence of heads and tails $w_{1,n}$ of length n, how many paths through the HMM of problem 1a generate $w_{1,n}$ with nonzero probability?

1c. For a given sequence of heads and tails $w_{1,n}$ of length n, what is the probability of observing $w_{1,n}$ according to the HMM model of problem 1a? Ie., give the numerical value of $P(w_{1,n})$.

1d. Now suppose the coins in states s^1 and s^2 are both fair: they come up heads or tails with probability .5. Assume we switch from one state to the other with a common probability q. Again, draw a specialized version of the introductory HMM to describe this situation, filling in specific numbers for the probabilities and omitting transitions with zero probability.

1e. Use the following notation:

$$P(w_{1,k}, S_{k+1} = s^1) = a$$
$$P(w_{1,k}, S_{k+1} = s^2) = b$$

In other words, suppose the probability of seeing the first k symbols and being in state s^1 is a and the probability of seeing the first k symbols and being in state s^2 is b.

Use this notation to calculate

$$P(w_{1,k+1}, S_{k+2} = s^1)$$
$$P(w_{1,k+1}, S_{k+2} = s^2)$$

as a function of a and b, assuming the HMM model of problem 1d.

1f. Use your answer to the previous problem to write $P(w_{1,k+1})$—the probability of seeing the first $k+1$ observations in the HMM model of problem 1d—as a function of $P(w_{1,k})$, the probability of seeing the first k observations.

1g. Therefore, what probability does this second model assign to the sequence $w_{1,n}$?

1h. Jeff Siskind of NEC, in talks at RuCCS on using HMMs for action recognition, has complained that the forward-backward or Baum-Welch training algorithm for HMMs was not doing what he would have liked. It was learning "high frequency" information when learning a general HMM network from his data. In other words, it was assigning meaning to states spuriously to capture rapid changes of observations over time.

Use the examples of this problem to comment on Siskind's difficulty.

Problem 2. This problem offers a case study in the contrasts between using different models to keep track of a changing environment. Our agent is interested in estimating the value of a variable X at each point in time. You might think of this variable as the income or success that the agent is achieving in each cycle of deliberation and action. X takes on one of two values over time: there is a "normal mode" in which the value of X is high and a "failure mode" in which the value of X is low.

Our agent acts in this environment in trials of a fixed length of time. In each trial, the environment starts out normal, but there may be a failure at any time; once there is a failure, the failure persists until the end of the trial. At each stage, the agent makes an observation that gives noisy information about the value of X.

In all, then, within a trial, the agent might see a pattern of observations such as that shown below.

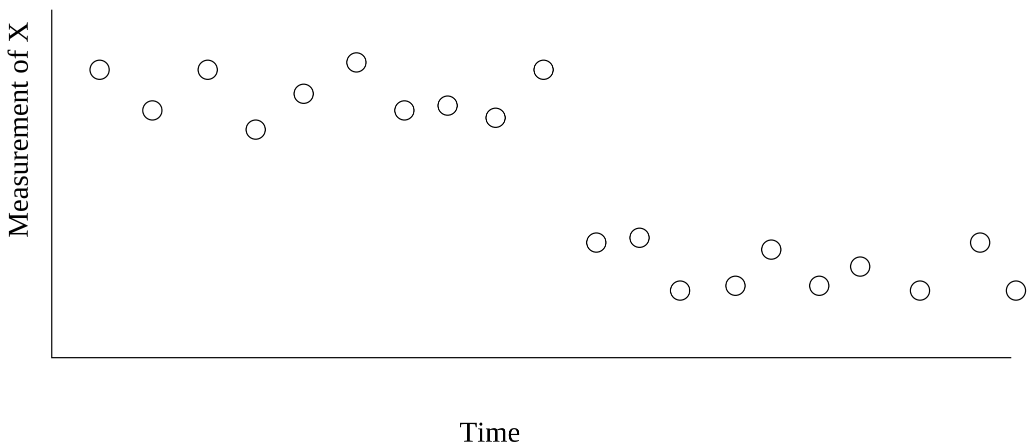

2a. Suppose our agent is designed to keep track of the value of X by just updating a running estimate. This corresponds to a coarse, general model in which the agent assumes that the actual value of X is more similar to the values seen in more recent observations. One algorithm for computing this estimate (related to so-called *reinforcement learning* algorithms) works as follows. Initially the system's estimate of X, $\hat{X}$ is the first data point w_1. Thereafter, on making an observation X_i, the system updates $\hat{X}$ by

$$\hat{X} \leftarrow w_i/3 + 2\hat{X}/3$$

Thus, we write w_i for the reward observed on step i, and write $w_{i,j}$ for the sequence of rewards observed on steps i through j, as we have commonly done for sequences of observations. Use a graph such as that above to describe how the agent's estimate changes over time according to this algorithm.

2b. Now we consider an alternative design for our agent: it has a model of failure, and reasons by finding the best estimates for this model given the evidence it has. We design the model to interpret a set of data like those graphed above. There may be a failure after any such step t—refer to this event as $F = t$. Another possible kind of hypothesis is that there is no failure before step t—refer to this event as $F \geq t$. So after n steps, the agent must consider alternative hypotheses $F = 1$, $F = 2$, $\ldots F = n-1$, $F \geq n$.

Note that we stop at $n-1$ because we will only be able to detect whether there was a failure *after* observation n once we have obtained observation $n+1$! Note also that the first try is always a successful case: we can only hypothesize a failure that starts afterwards! By the way, these assumptions mirror the assumptions that we have been making for HMMs, too. For the rest of the problem, however, we will diverge from HMM theory.

We assume that the model assigns a prior probability to each of the following hypotheses:

$$P(F = 1), \ldots, P(F = n-1), P(F \geq n)$$

Regard this as a prior distribution $P(F)$ for variable F.

We also formalize a probabilistic model of our observations as follows. For any i and j, w_i is conditionally independent of w_j given F, r_s and r_f. If there is no failure before step i, the reward on step i is normally distributed with mean r_s (for success) and variance $\sigma^2 = 1$. on the other hand, if there is a failure after any step up to step $i-1$, the reward on step i is normally distributed with mean r_f (for failure) and (again) variance $\sigma^2 = 1$.

We can summarize this mathematically as follows:

$$F \geq i \Rightarrow w_i \sim N(r_s, 1)$$
$$F < i \Rightarrow w_i \sim N(r_f, 1)$$

Specifically, this gives us the second ingredient of the model:

$$p(w_i|F, r_s, r_f) = \begin{cases} ce^{-\frac{(w_i - r_f)^2}{2}} & \text{if } F = t \text{ and } t < i \\ ce^{-\frac{(w_i - r_s)^2}{2}} & \text{otherwise} \end{cases}$$

We assume priors for r_s and r_f that are uniformly distributed on some interval, so $p(r_s) = p(r_f) = h$.

The rest of problem 2b walks you through the derivation of the maximum a posteriori joint estimate for F, r_s and r_f given a sequence of observations $w_{1,n}$. Before proceeding, note that the formalization developed so far just spells out some details that are implicit in the problem statement, and that you probably used to interpret the graphs above. Your intuitions provide a good sense of what kind of answers to expect here.

b1. Use Bayes's theorem to find an expression for the quantity $p(F, r_s, r_f|w_{1,n})$ in terms of $p(w_{1,n}|F, r_s, r_f)$ (and other things).

b2. Use the independence assumptions to rewrite your answer to b1 in terms of the parameters of the model just outlined, expanding the conditional definition of $p(w_i|F, r_s, r_f)$. (Use two product operators.)

b3. Rewrite your answer to b2, dropping all constant factors in anticipation of maximization, and take the logarithm.

b4. What value of r_s maximizes the quantity described in your answer to b3? This is the best estimate of r_s.

b5. Find the best estimate for r_f the analogous way—noting that for $F \geq n$, you actually have no evidence about r_f.

b6. Since F is a discrete set of hypotheses, you have to maximize by enumerating the alternatives. What estimation procedure results?

2c. Suppose the failure occurs after observation $n-1$. The question we want to address is how different the nth observation has to be from previous observations in order to recognize the failure immediately using the model. In other words, when do we prefer the hypothesis $F = n - 1$ to the hypothesis $F \geq n$ (using our best estimates for r_s and r_f) based on the nth data point? Provide an answer, using the rough assumption that we have enough data that the average of the first $n-1$ observations is about the same as the average of the first n observations.

2d. Use a graph such as that sketched initally to describe how the agent's estimate changes over time according to the algorithm derived in problem 2b; assume the criterion in problem 2c is met at the moment of failure when the value of X changes. Contrast this with the coarse model of problem 2a.