Assignment4 – Markov Models

Given: May 22   Due: May 27

Problem 4.1 (The Value of Information)

Chef Giordana runs a kitchen that provides food for a large organisation. A salad is sold for €6 and costs €4 to prepare. So a sold salad is profit of €2, and an unsold salad a loss of €4. Actual demand will be 40 (with probability 0.5) or 60 (also with probability 0.5) each day.

Each day, Giordana must decide in advance between two options: prepare 40 or 60 salads.

1. In the absence of additional information, compute the expected utility of each decision option, and choose the best option.
2. She is considering a new ordering system, where she knows the demand perfectly in advance. So she can always choose the better of the two options. State the formula for computing the value of this perfect information, explain the components in Giordana's case, and compute the value.

Problem 4.2 (Expected Utility)

1. State the formal definition of *expected utility* of an action in the current state of an agent? Explain the meaning of every variable in the defining equation.
2. How do we use expected utility to make decisions?

Problem 4.3 (Decision Theory)

You are offered the following game: You pay x dollars to play. A fair coin is then tossed repeatedly until it comes up heads for the first time. Your payout is 2^n, where n is the number of tosses that occurred.

1. Assume your utility function is exactly the monetary value. How much should you, as a rational agent, be willing to pay to play? Use the formal definition of "expected utility" from the *lecture*.
2. Assume now, that your utility function for having k dollars is $U(k) = m \log_n k$ for some $m, n \in \mathbb{N}^+$. How does this change the result?
3. What is wrong with the result from the first exercise? Which implicit assumption leads to the apparently nonsensical result? How could it be fixed?

Hint: The series $\sum_{k=1}^{\infty} \frac{k}{2^k}$ is convergent with limit 2.

Problem 4.4 (Markov Mood Detection)

On any given day d, your roommate Moody is either happy or sad, so $M_d \in \{h, s\}$. Usually when he is sad, he stays sad for a while, and $P(M_{d+1} = s \mid M_d = s) = 0.7$. But aside from that he is a cheery guy, and $P(M_{d+1} = h \mid M_d = h) = 0.85$.

He either listens to jazz or metal music, so $L_d \in \{j, m\}$. On a happy day he usually listens to Jazz, and $P(L_d = j \mid M_d = h) = 0.7$. On a sad day, he slightly prefers metal, and $P(L_d = m \mid M_d = s) = 0.6$.

1. Model this situation as a Markov process. Explain what the state and evidence variables are. What order does the process have? Is the transition model stationary? Is the sensor model stationary? Does the transition model have the Markov property? Does the sensor model have the sensor Markov property? Explain all answers in one short sentence each.
2. State the formula for the full joint probability distribution. You know that he was happy at day d_0. What is the probability that he is happy and plays Jazz for the next two days?

Problem 4.5 (Moody HMM)

Consider the Markov process from before about the roommate Moody (which in particular gives the concrete probabilities needed below). We have already modeled it as an HMM with state variables M_d and evidence variables L_d.

Because the transition model is first-order and stationary, we can collect the conditional probabilities for the state transitions into a matrix $T_{ij} = P(M_d = x_j \mid M_{d-1} = x_i)$ where x_i, x_j are two states (i.e., two possible values of the state variable). We use k for the number of states, and T is an $k \times k$ matrix.

Because the sensor model is stationary and has the sensor Markov property, we can collect the conditional probabilities for the observations into a matrix $S_{ij} = P(L_d = y_j \mid M_d = x_i)$ where x_i is a state and the y_j are the possible observations. If there are l possible observations, this is an $k \times l$ matrix. For a fixed observation e, the diagonal $k \times k$ matrices O_e from the *lecture notes* are obtained from the columns of this matrix.

Clarify the modeling as an HMM. Concretely:

1. What is k? Give the transition matrix T.
2. What is l? Give the sensor matrix S.

Note that you need to choose and state orderings for the states and observations so that it is clear which state/observation corresponds to which row/column of T and S.

Now consider a fixed sequence $L_1 = e_1, L_2 = e_2$ of observations that we have made for two days. Concretely, you heard Moody play metal on day $d = 1$ and jazz on day $d = 2$.

3. Give the diagonal sensor matrices O_1 and O_2 corresponding to the observation at $d = 1$ and $d = 2$.
4. You are not sure what kind of mood your flatmate was in on day $d = 0$, but it was either good or bad with equal probability. The HMM algorithm for filtering and smoothing uses compact matrix/vector equation to compute f and b. Use those equation to determine the probability distribution of Moody's mood on day $d = 1$.
   Note that T^T in the filtering equation in the *lecture notes* denotes the transpose of T.
   Grading will focus on writing out the matrices with the correct probabilities in them and on formally stating the computations that need to applied to those matrices. You should also actually do those computations, but that is secondary.

Problem 4.6 (Prediction, Filtering, Smoothing)

Consider an HMM consisting of a stationary Markov process and a stationary sensor model. We restrict attention to HMMs with a single state-variable X_t and a single evidence variable E_t. (The sensor model does not necessarily have the Markov property.)

Explain the results of the *prediction*, *filtering*, *smoothing* algorithms. For each one, state the motivation, the expression for the conditional probability that is to be computed, and explain the components of the formula. You do not have to explain how the algorithms work.