

RECITATION 8

HIDDEN MARKOV MODELS AND BAYES NET

10-601: INTRODUCTION TO MACHINE LEARNING

4/1/2022

1 HMMs

You are given the following training data:

win_C league_C Liverpool_D

win_C Liverpool_D league_C

Liverpool_D win_C

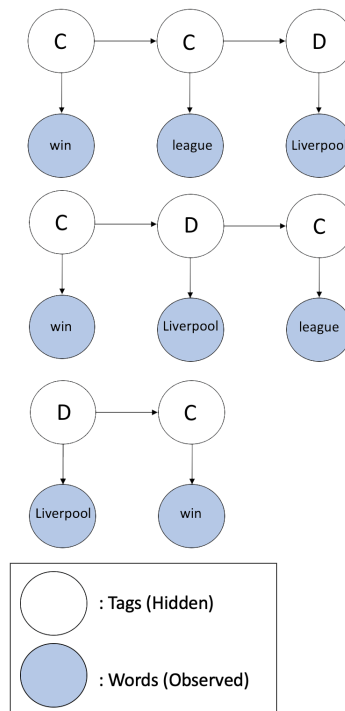


Figure 1: Visualization of Sequences

You are also given the following observed (validation) data: Liverpool win league

1.1 Initial, Emission, and Transition Matrices

Let each observed state $x_t \in \{1, 2, 3\}$, where 1 corresponds to **win**, 2 corresponds to **league**, and 3 corresponds to **Liverpool**. Let each hidden state $Y_t \in \{C, D\}$, where $s_1 = C$ and $s_2 = D$.

First, we need to estimate the HMM parameters - the initial probabilities: π , the transition probability matrix: \mathbf{B} , and the emission probability matrix: \mathbf{A} . Remember that we use MLE estimation to do so:

$$\hat{C}_k = \frac{\#(y_i^{(i)} = k)}{N} \quad \forall i, k$$

$$\hat{B}_{jk} = \frac{\#(y_t^{(i)} = k \text{ and } y_{t-1}^{(i)} = j)}{\#(y_{t-1}^{(i)} = j)} \quad \forall i, t > 1, j, k$$

$$\hat{A}_{jk} = \frac{\#(x_t^{(i)} = k \text{ and } y_t^{(i)} = j)}{\#(y_t^{(i)} = j)} \quad \forall i, t, j, k$$

Note: When learning an HMM, we add 1 to each count to make a pseudocount. This improves performance when evaluating unseen cases in the validation set or test set.

1. Find the initial matrix π . Recall that $\pi_j = P(Y_1 = s_j)$.

- Find count matrix and pseudocount matrix:

C	<i>Count</i>	$\xrightarrow{\text{Pseudocount}}$	C
D			D

- Normalize:

$$\pi = \begin{matrix} C \\ D \end{matrix}$$

2. Find the transition matrix \mathbf{B} . Recall that $B_{jk} = P(Y_t = s_k | Y_{t-1} = s_j)$.

- Find count matrix and pseudocount matrix:

C	C	D	$\xrightarrow{\text{Pseudocount}}$	C	D
D				D	

- Normalize:

$$\mathbf{B} = \begin{matrix} & C & D \\ C & & \\ D & & \end{matrix}$$

3. Find the emission matrix \mathbf{A} . Recall that $A_{jk} = P(X_t = k \mid Y_t = s_j)$.

- Find count matrix and pseudocount matrix:

$$\begin{matrix} & \text{win} & \text{league} & \text{Liverpool} & & & \text{win} & \text{league} & \text{Liverpool} \\ C & & & & \xrightarrow{\text{Pseudocount}} & C & & & \\ D & & & & & D & & & \end{matrix}$$

- Normalize:

$$\mathbf{A} = \begin{matrix} & C & \text{win} & \text{league} & \text{Liverpool} \\ C & & & & \\ D & & & & \end{matrix}$$

1.2 The Forward Algorithm

One type of inference problem that can be answered by an HMM is **Evaluation** - computing the probability of a sequence of observations. We calculate the likelihood of observing the validation sequence:


Liverpool win league

To do so, we calculate the forward probability matrix α . Recall that

$$\alpha_t(k) = P(x_{1:t}, Y_t = s_k)$$

We have the following bottom-up dynamic programming algorithm to calculate the forward probabilities:

1. $\alpha_1(j) = \pi_j A_{jx_1}$
2. For $t > 1$, $\alpha_t(j) = A_{jx_t} \sum_{k=1}^J \alpha_{t-1}(k) B_{kj}$



$$\bullet \alpha_1 = \begin{bmatrix} \alpha_1(C) \\ \alpha_1(D) \end{bmatrix} = \begin{bmatrix} \pi_C * A_{C,x_1} \\ \pi_D * A_{D,x_1} \end{bmatrix} = \begin{bmatrix} \pi_C * A_{C,Liverpool} \\ \pi_D * A_{D,Liverpool} \end{bmatrix} =$$

Now, we have our α and β matrices:

$$\alpha = \begin{array}{c} 1 \\ 2 \\ 3 \end{array} \begin{array}{cc} C & D \\ \left[\begin{array}{cc} 0.0750 & 0.2667 \\ 0.1150 & 0.0186 \\ 0.0225 & 0.0123 \end{array} \right] \end{array}$$

$$\beta = \begin{array}{c} 1 \\ 2 \\ 3 \end{array} \begin{array}{cc} C & D \\ \left[\begin{array}{cc} 0.0823 & 0.1072 \\ 0.2500 & 0.3229 \\ 1.0000 & 1.0000 \end{array} \right] \end{array}$$

1. What is $P(Y_2 = C|\vec{x})$?

2. What is $P(Y_2 = D|\vec{x})$?

3. What is $P(Y_3 = C|\vec{x})$?

4. What is the minimum Bayes risk (MBR) decoder prediction for Y_2 ?

1.4 The Viterbi Algorithm

Instead of finding the most likely hidden state at some time t , we may instead want to find the most likely sequence of hidden states. This is known as **Viterbi Decoding** - computing the most probable assignment of hidden states, given a sequence of observations.

The sequence of words you observe is again the same: `Liverpool win league`

However, you are only given the tag of the last word: `league_C`


1. Recall that:

$$\omega_t(s_k) = \max_{y_{1:t-1}} P(x_{1:t}, y_{1:t-1}, y_t = s_k)$$

$$b_t(s_k) = \arg \max_{y_{1:t-1}} P(x_{1:t}, y_{1:t-1}, y_t = s_k)$$

Using the formulae above and the first order Markov assumption, derive a recursive definition for $\omega_t(s_k)$ and $b_t(s_k)$ that will let you employ bottom-up dynamic programming.

Below is the trellis corresponding to the given data:



2. Annotate the trellis at the nodes that correspond to:

- (a) $\omega_1(C)$
- (b) $\omega_1(D)$
- (c) $\omega_2(C)$
- (d) $\omega_2(D)$
- (e) $\omega_3(C)$
- (f) $\omega_3(D)$

3. Find the most likely sequence of tags given the observed data:

(a) Set up the matrices ω and b

$$\omega = \begin{matrix} & & \text{C} & \text{D} & \text{START} \\ \omega_0 & & \begin{bmatrix} 0 & 0 & 1 \end{bmatrix} \\ \omega_1 & & \begin{bmatrix} - & - & - \end{bmatrix} \\ \omega_2 & & \begin{bmatrix} - & - & - \end{bmatrix} \\ \omega_3 & & \begin{bmatrix} - & - & - \end{bmatrix} \end{matrix}$$

and

$$b = \begin{matrix} & & \text{C} & \text{D} & \text{END} \\ b_1 & & \begin{bmatrix} - & - & - \end{bmatrix} \\ b_2 & & \begin{bmatrix} - & - & - \end{bmatrix} \\ b_3 & & \begin{bmatrix} - & - & - \end{bmatrix} \\ b_4 & & \begin{bmatrix} - & - & - \end{bmatrix} \end{matrix}$$

Initialize $w_0(\text{START}) = 1$

(b) Solve for matrix entries using Dynamic Programming:

$$\begin{aligned} \omega_1(\text{C}) &= \max_{s_j \in \text{C,D,START}} P(x_1 = \text{Liverpool} | Y_1 = \text{C}) \omega_0(s_j) P(Y_1 = \text{C}) \\ &= \end{aligned}$$

$$b_1(\text{C}) =$$

$$\begin{aligned} \omega_1(\text{D}) &= \max_{s_j \in \text{C,D,START}} P(x_1 = \text{Liverpool} | Y_1 = \text{D}) \omega_0(s_j) P(Y_1 = \text{D}) \\ &= \\ &= \end{aligned}$$

$$b_1(\text{D}) =$$

$$\begin{aligned} \omega_2(\text{C}) &= \max_{s_j \in \text{C,D}} P(x_2 = \text{win} | Y_2 = \text{C}) \omega_1(s_j) P(Y_2 = \text{C} | Y_1 = s_j) \\ &= \max \left(\right) \\ &= \end{aligned}$$

$$b_2(\text{C}) =$$

$$\begin{aligned}\omega_2(D) &= \max_{s_j \in \mathbf{C}, \mathbf{D}} P(x_2 = \text{win} | Y_2 = \mathbf{D}) \omega_1(s_j) P(Y_2 = \mathbf{D} | Y_1 = s_j) \\ &= \max \left(\right) \\ &= \end{aligned}$$

$$b_2(\mathbf{D}) =$$

$$\begin{aligned}\omega_3(\mathbf{C}) &= \max_{s_j \in \mathbf{C}, \mathbf{D}} P(x_3 = \text{league} | Y_3 = \mathbf{C}) \omega_2(s_j) P(Y_3 = \mathbf{C} | Y_2 = s_j) \\ &= \max \left(\right) \\ &= \end{aligned}$$

$$b_3(\mathbf{C}) =$$

$$\begin{aligned}\omega_3(D) &= \max_{s_j \in \mathbf{C}, \mathbf{D}} P(x_3 = \text{league} | Y_3 = \mathbf{D}) \omega_2(s_j) P(Y_3 = \mathbf{D} | Y_2 = s_j) \\ &= \max \left(\right) \\ &= \end{aligned}$$

$$b_3(\mathbf{D}) =$$

Now, to figure out the order, we set $\hat{y}_t = b_{t+1}(\hat{y}_{t+1})$

$$\hat{y}_{T+1} = \text{END}$$

$$\begin{aligned}\hat{y}_3 &= b_4(\text{END}) \\ &= \end{aligned}$$

$$\begin{aligned}\hat{y}_2 &= b_3(\quad) \\ &= \end{aligned}$$

$$\begin{aligned}\hat{y}_1 &= b_2(\quad) \\ &= \end{aligned}$$

$$\begin{aligned}\hat{y}_0 &= b_1(\quad) \\ &= \end{aligned}$$

So, the most likely sequence is:

2 Working in Log-space

2.1 Motivation

Some of the probabilities we work with in Homework 7 about are tiny and some of them are much larger. We tend to work with the tiny ones in log-space and only get back probabilities if we really need them for some other purpose. Throughout HW7 you will keep your probabilities in log-space. In this section we will motivate why we use log-space for small values.

Given the following series of probability values:

$P(x_1 = 1)$	$P(x_2 = 1 \mid x_1 = 1)$	$P(x_3 = 1 \mid x_2 = 1, x_1 = 1)$
0.002	0.004	0.003

We want to find $P(x_1 = 1, x_2 = 1, x_3 = 1)$. Suppose we have a calculator which only has 4 decimal places of precision, so it can only store values of format X.XXXX

1. What is the correct value of $P(x_1 = 1, x_2 = 1, x_3 = 1)$ without any precision limits?

$$P(x_1 = 1, x_2 = 1, x_3 = 1) = P(x_3 = 1 \mid x_2 = 1, x_1 = 1) * P(x_2 = 1 \mid x_1 = 1) * P(x_1 = 1)$$

$$=$$

2. What is the value of $P(x_1 = 1, x_2 = 1, x_3 = 1)$ using our faulty calculator?

$$P(x_1 = 1, x_2 = 1)$$

$$= P(x_2 = 1 \mid x_1 = 1)P(x_1 = 1) =$$

$$P(x_1 = 1, x_2 = 1, x_3 = 1) =$$

3. How do the values of $P(x_1 = 1, x_2 = 1, x_3 = 1)$ from part (1) and (2) compare?

$$\text{No precision limits: } P(x_1 = 1, x_2 = 1, x_3 = 1) =$$

$$\text{Faulty calculator: } P(x_1 = 1, x_2 = 1, x_3 = 1) =$$

4. What is the value of $P(x_1 = 1, x_2 = 1, x_3 = 1)$ if we perform the same computation but in log space?

$$\log(P(x_1 = 1, x_2 = 1, x_3 = 1))$$

$$= \log(x_1 = 1) + \log(P(x_2 = 1 \mid x_1 = 1)) + \log(P(x_3 = 1 \mid x_2 = 1, x_1 = 1))$$

$$=$$

$$\text{If we were to recover our value of } P(x_1 = 1, x_2 = 1, x_3 = 1) = e^{\log(P(x_1=1, x_2=1, x_3=1))} =$$

This is good! But we can use the log sum exp trick to extend its use to even smaller scales.

2.2 Forward and Backward Algorithm in Log Space

In the forward algorithm, recall that the entries in α can be computed using the bottom-up dynamic programming algorithm:

- $\alpha_1(j) = \pi_j A_{jx_1}$
 - For $t > 1$, $\alpha_t(j) = A_{jx_t} \sum_{k=1}^J \alpha_{t-1}(k) B_{kj}$
1. Derive $\log(\alpha_1(j))$ in terms of $\log(\pi_j)$ and $\log(A_{jx_1})$

$$\log(\alpha_1(j)) = \log(\pi_j A_{jx_1}) =$$
 2. Derive $\log(\alpha_t(j))$ in terms of $\log(\alpha_{t-1}(k))$ and $\log A_{kj}$

$$\begin{aligned} \log(\alpha_t(j)) &= \log(A_{jx_t} \sum_{k=1}^J \alpha_{t-1}(k) B_{kj}) \\ &= \log(A_{jx_t}) + \log\left(\sum_{k=1}^J \alpha_{t-1}(k) B_{kj}\right) \\ &= \end{aligned}$$

In the backward algorithm, we also have a similar bottom-up dynamic programming algorithm:

- $\beta_T(j) = 1$
 - For $1 \leq t \leq T - 1$, $\beta_t(j) = \sum_{k=1}^J A_{kx_{t+1}} \beta_{t+1}(k) B_{jk}$
1. Derive $\log(\beta_T(j))$

$$\log(\beta_T(j)) = \log(1) = 0$$
 2. Derive $\log(\beta_t(j))$ in terms of $\log(A_{kx_{t+1}})$, $\log(\beta_{t+1}(k))$, and $\log(B_{jk})$

$$\begin{aligned} \log(\beta_t(j)) &= \log\left(\sum_{k=1}^J A_{kx_{t+1}} \beta_{t+1}(k) B_{jk}\right) \\ &= \log\left(\sum_{k=1}^J e^{\log(A_{kx_{t+1}} \beta_{t+1}(k) B_{jk})}\right) \\ &= \end{aligned}$$

After transforming the equations into log form, you may discover calculations of the following type:

$$\log \sum_i \exp(v_i)$$

This may be programmed as is, but $\exp(v_i)$ may cause underflow when v_i is large and negative. One way to avoid this is to use the [log-sum-exp trick](#).

The log-sum-exp trick simply adds the maximum value in the vector to the log probabilities as follows:

$$\log \sum_i \exp(v_i - m) + \max_i(v_i)$$

3 Dynamic Programming

[DP Notebook](#)

To access this Colab notebook you will need to be logged into Google Drive with your Andrew email.