AIL 722: Reinforcement Learning

Lec 4: Hidden Markov Models (Part 3: Learning)

Raunak Bhattacharyya

ScAI | YARDI SCHOOL OF ARTIFICIAL INTELLIGENCE
INDIAN INSTITUTE OF TECHNOLOGY DELHI

Outline

- Observation Probability
- Backward Probability
- Transition Probability
- Baum-Welch Algorithm

Problem 3: Learning

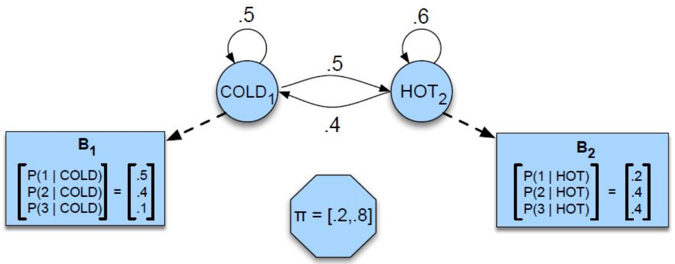

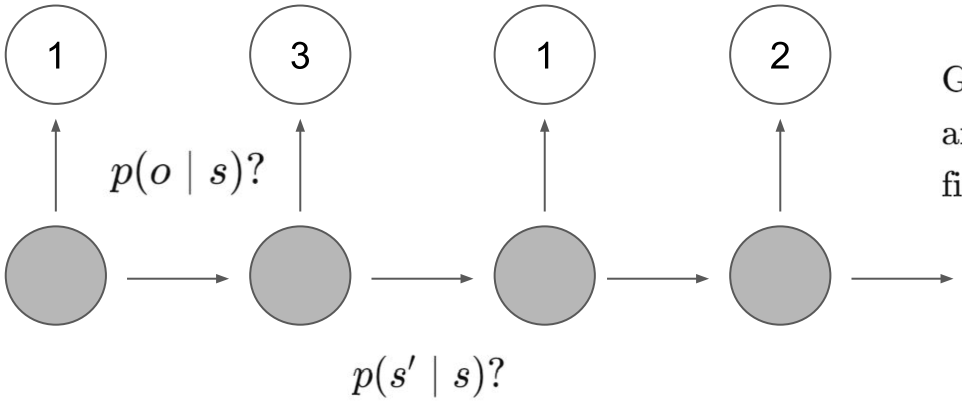

Given $O = \{1, 3, 1, 2, 3, 1, \ldots\}$,
and the set of possible hidden states $\{H, C\}$,
find the T and B matrices.

Simple Case: Fully Visible Markov Model

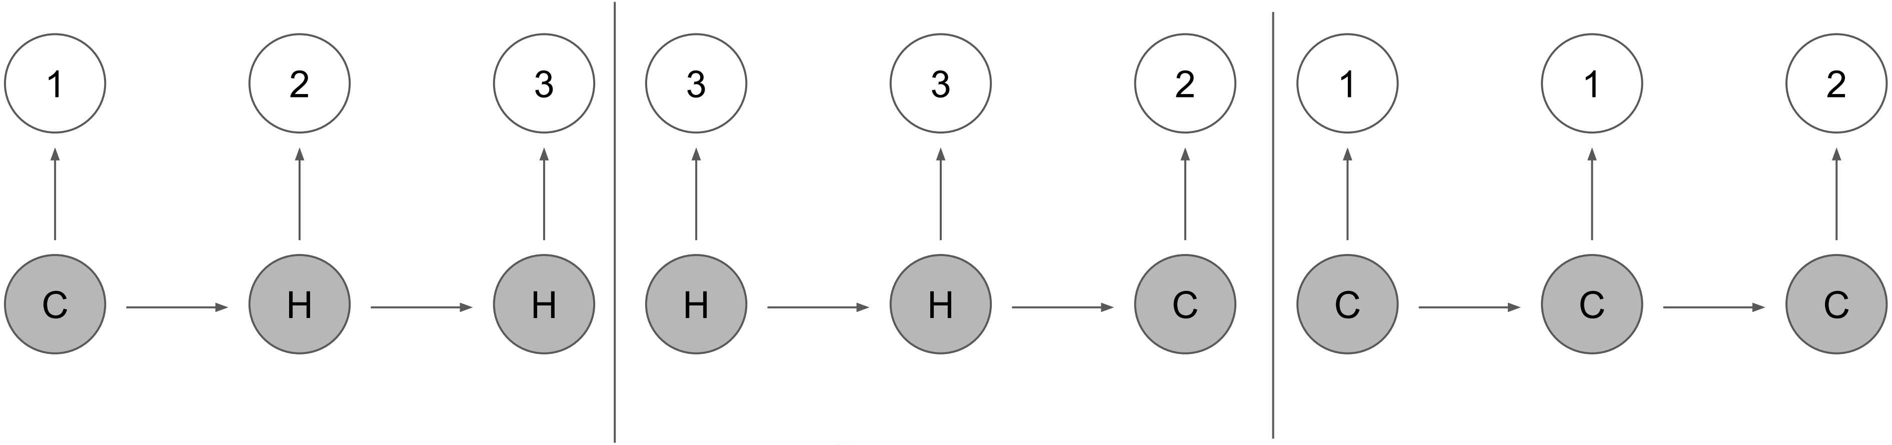

$$p(s_1 = H) = \frac{1}{3}$$

$$p(s_1 = C) = \frac{2}{3}$$

$$p(H \mid H) = \frac{2}{3} \qquad p(H \mid C) = \frac{1}{3}$$

$$p(C \mid H) = \frac{1}{3} \qquad p(C \mid C) = \frac{2}{3}$$

$$p(1 \mid H) = 0 \qquad p(1 \mid C) = 0.6$$
$$p(2 \mid H) = 0.25 \qquad p(2 \mid C) = 0.4$$
$$p(3 \mid H) = 0.75 \qquad p(3 \mid C) = 0$$

Why does this not work for an actual HMM?

Observation Probabilities

Observation Probs

Given $O = \{1, 3, 1, 2, 3, 1, \ldots\}$,
and the set of possible hidden states $\{H, C\}$,
find the T and B matrices.

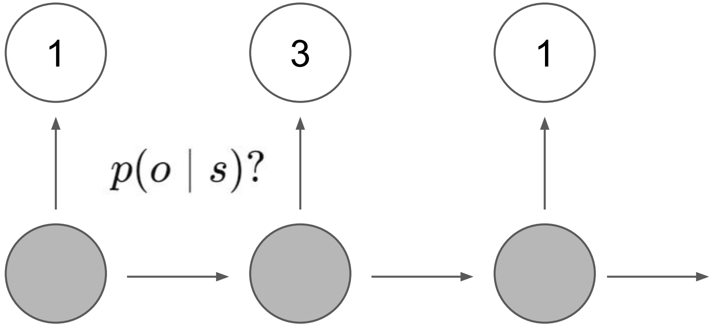

$$p(o^k \mid s^j) = \frac{\text{expected number of times in state } j \text{ and observing } o^k}{\text{expected number of times in state } j}$$

Key: Expected number of times in state j

A Side Problem

You are the admissions coordinator for your department. You are trialing an office assistant robot that your colleague has requested you to test. You have to send offer letters to candidates who recently took the admission test. You ask the robot to do the task.

Unfortunately, there was an error in the visual perception so the robot randomly inserted the letters into envelopes (instead of the correct letter into the correct envelope) and posted them.

Out of the N accepted candidates, how many got their offer letter?

Letter Envelope Matching Problem

Random Var?

$$I_i = \begin{cases} 1 & \text{if the } i\text{th letter is in the correct envelope} \\ 0 & \text{otherwise} \end{cases}$$

Number of correct placements?

$$S = \sum_{i=1}^{n} I_i$$

$$\mathbb{E}[S] = \mathbb{E}\left[\sum_{i=1}^{n} I_i\right] = \sum_{i=1}^{n} \mathbb{E}[I_i] = \sum_{i=1}^{n} \frac{1}{n} = 1$$

Expected number of correct letters sent

Learning the Observation Prob.

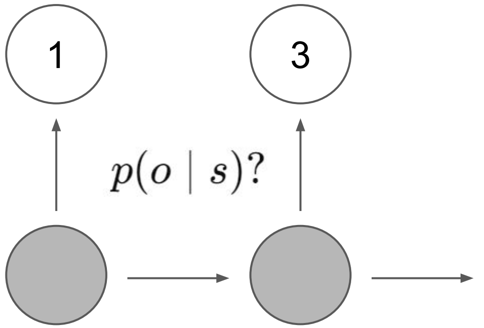

$$p(o^k \mid s^j) = \frac{\text{expected number of times in state } j \text{ and observing } o^k}{\text{expected number of times in state } j}$$

$$I_t = \begin{cases} 1 & \text{if } s_t = j \\ 0 & \text{otherwise} \end{cases}$$

$$S = \sum_{i=1}^{T} I_t$$

Number of times in state j over the sequence of T steps

$$\mathbb{E}[S] = \mathbb{E}\left[\sum_{i=1}^{T} I_t\right] = \sum_{i=1}^{T} \mathbb{E}\left[I_t\right] = \sum_{i=1}^{T} p(s_t = j)$$

Learning the Observation Prob.

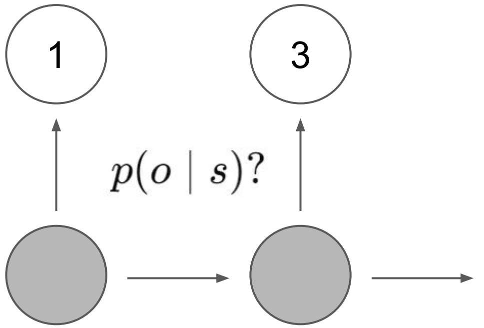

$$p(o^k \mid s^j) = \frac{\text{expected number of times in state } j \text{ and observing } o^k}{\text{expected number of times in state } j}$$

$$\gamma_t(j) = p(s_t = j \mid o_{1:T})$$

$$p(o^k \mid s^j) = \frac{\sum_{t=1}^{T} \gamma_t(j) 1(o_t = o_k)}{\sum_{t=1}^{T} \gamma_t(j)}$$

How do we find gamma?

Finding Gamma

$$p(o_k \mid s_j) = \frac{\sum_{t=1}^{T} \gamma_t(j) 1(o_t = o_k)}{\sum_{t=1}^{T} \gamma_t(j)}$$

$$\gamma_t(j) = p(s_t = j \mid o_{1:T})$$

$$= \frac{p(s_t = j, o_{1:T})}{p(o_{1:T})}$$

$$p(o_{1:T}, s_t = j)$$

Numerator

Look familiar?

$$\alpha_t(j) = p(o_{1:t}, s_t = j)$$

Backward Probability

$$p(o_{1:T}, s_t = j)$$

$$p(o_{1:T}, s_t = j) = p(o_{1:t}, s_t = j, o_{t+1:T})$$

$$\beta_t(i) = p(o_{t+1}, o_{t+2}, \ldots, o_T \mid s_t = i)$$

$$p(o_{1:T}, s_t = j) = \alpha_t(j) \cdot \beta_t(j)$$

Why is this true?

Finding Backward Probability

$$\beta_t(i) = p(o_{t+1}, o_{t+2}, \ldots, o_T \mid s_t = i)$$

1. **Initialization:**

$$\beta_T(i) \; = \; 1, \quad 1 \leq i \leq N$$

2. **Recursion**

$$\beta_t(i) = \sum_{j=1}^{N} a_{ij}\, b_j(o_{t+1})\, \beta_{t+1}(j), \quad 1 \leq i \leq N, 1 \leq t < T$$

3. **Termination:**

$$P(O|\lambda) = \sum_{j=1}^{N} \pi_j\, b_j(o_1)\, \beta_1(j)$$

Why is this true?

Source: SLP, Dan Jurafsky

Back to the Big Picture

$$p(o_{1:T}, s_t = j) = \alpha_t(j) \cdot \beta_t(j)$$

$$\gamma_t(j) = p(s_t = j \mid o_{1:T})$$

$$= \frac{p(s_t = j, o_{1:T})}{p(o_{1:T})}$$

How do we find this?

$$p(o^k \mid s^j) = \frac{\sum_{t=1}^{T} \gamma_t(j) 1(o_t = o_k)}{\sum_{t=1}^{T} \gamma_t(j)}$$

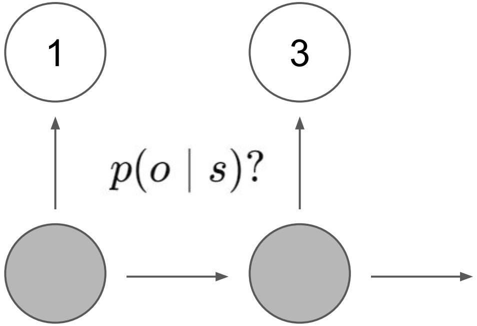

Transition Probabilities

Transition Probability

Given $O = \{1, 3, 1, 2, 3, 1, \ldots\}$,
and the set of possible hidden states $\{H, C\}$,
find the T and B matrices.

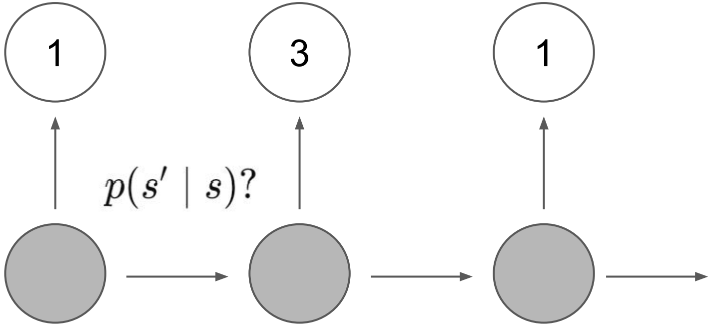

$$p(s_{t+1} = j \mid s_t = i) = \frac{\text{expected number of transitions from state } i \text{ to state } j}{\text{expected number of transitions from state } i}$$

Learning the Transition Prob.

$$\frac{\text{expected number of transitions from state } i \text{ to state } j}{\text{expected number of transitions from state } i}$$

$$\xi_t(i,j) = p(s_t = i, s_{t+1} = j \mid o_{1:T})$$

$$\text{almost-}\xi_t(i,j) = p(s_t = i, s_{t+1} = j, o_{1:T})$$

$$\text{almost-}\xi_t(i,j) = \alpha_t(i) \cdot a_{ij} \cdot b_j(o_{t+1}) \cdot \beta_{t+1}(j)$$

Why is this true?

$$\xi_t(i,j) = \frac{\text{almost-}\xi_t(i,j)}{p(o_{1:T})}$$

Learning the Transition Prob.

$$\frac{\text{expected number of transitions from state } i \text{ to state } j}{\text{expected number of transitions from state } i}$$

$$\xi_t(i,j) = p(s_t = i, s_{t+1} = j \mid o_{1:T})$$

$$\hat{a}_{ij} = \frac{\sum_{t=1}^{T-1} \xi_t(i,j)}{\sum_{t=1}^{T-1} \sum_{k=1}^{N} \xi_t(i,k)}$$

Why is this true?

Learning the Transition Prob.

$$\frac{\text{expected number of transitions from state } i \text{ to state } j}{\text{expected number of transitions from state } i}$$

Overall flow?

$$\text{almost-}\xi_t(i,j) = \alpha_t(i) \cdot a_{ij} \cdot b_j(o_{t+1}) \cdot \beta_{t+1}(j)$$

$$\xi_t(i,j) = \frac{\text{almost-}\xi_t(i,j)}{p(o_{1:T})}$$

$$\hat{a}_{ij} = \frac{\sum_{t=1}^{T-1} \xi_t(i,j)}{\sum_{t=1}^{T-1} \sum_{k=1}^{N} \xi_t(i,k)}$$