

AIL 722: Reinforcement Learning

Lecture 5: Baum-Welch

Raunak Bhattacharyya



ScAI

YARDI SCHOOL OF ARTIFICIAL INTELLIGENCE
INDIAN INSTITUTE OF TECHNOLOGY DELHI

Outline

- Consolidation: Observation Prob. and Transition Prob.
- Baum-Welch Algorithm
- Introducing Decisions into the Markov Model: MDPs
- Example MDPs

Consolidation

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

$$\beta_t(i) = p(o_{t+1:T} \mid s_t = i)$$

Observation

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

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

Expected state occupancy count

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

Does this require knowing transition and observation?

Transition

$$\hat{p}(s^j \mid s^i) = \frac{\text{expected number of transitions from state } i \text{ to state } j}{\text{expected number of transitions from state } i}$$

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

Expected state transition count

$$p(s_t = i, s_{t+1} = j, o_{1:T}) = \alpha_t(i) \cdot a_{ij} \cdot b_j(o_{t+1}) \cdot \beta_{t+1}(j)$$

Can't we just use alpha?

Consolidation

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

$$\beta_t(i) = p(o_{t+1:T} \mid s_t = i)$$

Observation

Transition

$$p(o_{1:T}) = \sum_{i=1}^N \alpha_T(i)$$

$$\gamma_t(j) = \frac{\alpha_t(j) \cdot \beta_t(j)}{\sum_{i=1}^N \alpha_T(i)}$$

$$\xi_t(i, j) = \frac{\alpha_t(i) \cdot a_{ij} \cdot b_j(o_{t+1}) \cdot \beta_{t+1}(j)}{\sum_{i=1}^N \alpha_T(i)}$$

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

$$\hat{p}(s^j \mid s^i) = \frac{\sum_{t=1}^{T-1} \xi_t(i, j)}{\sum_{t=1}^{T-1} \sum_{k=1}^N \xi_t(i, k)}$$

Baum-Welch Algorithm

Baum-Welch Algorithm

function FORWARD-BACKWARD(*observations* of len T , *output vocabulary* V , *hidden state set* Q) **returns** $HMM=(A,B)$

initialize A and B

iterate until convergence

E-step

$$\gamma_t(j) = \frac{\alpha_t(j)\beta_t(j)}{\alpha_T(q_F)} \quad \forall t \text{ and } j$$
$$\xi_t(i, j) = \frac{\alpha_t(i)a_{ij}b_j(o_{t+1})\beta_{t+1}(j)}{\alpha_T(q_F)} \quad \forall t, i, \text{ and } j$$

M-step

$$\hat{a}_{ij} = \frac{\sum_{t=1}^{T-1} \xi_t(i, j)}{\sum_{t=1}^{T-1} \sum_{k=1}^N \xi_t(i, k)}$$
$$\hat{b}_j(v_k) = \frac{\sum_{t=1 \text{ s.t. } O_t=v_k}^T \gamma_t(j)}{\sum_{t=1}^T \gamma_t(j)}$$

return A, B

Alert: Use of q for state

Expected state occupancy and state transition counts

Re-estimate the transition and observation probabilities

Finding Forward and Backward Probs

Algorithm 3 Baum-Welch Algorithm

```
1: Input:  
2: Observation sequence:  $O = (O_1, O_2, \dots, O_T)$   
3: Number of states:  $N$   
4: Number of distinct observation symbols:  $M$   
5: Initial model parameters: transition probabilities  $A$ , emission probabilities  
    $B$ , and initial state probabilities  $\pi$   
6: Output: Updated model parameters:  $A, B, \pi$   
7: Initialize parameters  $A, B, \pi$   
8: repeat  
9:   Forward Procedure:  
10:  Initialize  $\alpha$ :  
11:  for  $i = 1$  to  $N$  do  
12:     $\alpha_1(i) = \pi_i b_i(O_1)$   
13:  end for  
14:  Recursively compute  $\alpha$ :  
15:  for  $t = 2$  to  $T$  do  
16:    for  $j = 1$  to  $N$  do  
17:       $\alpha_t(j) = \left( \sum_{i=1}^N \alpha_{t-1}(i) a_{ij} \right) b_j(O_t)$   
18:    end for  
19:  end for
```

Initial State Distribution?

```
18: Backward Procedure:  
19: Initialize  $\beta$ :  
20: for  $i = 1$  to  $N$  do  
21:    $\beta_T(i) = 1$   
22: end for  
23: Recursively compute  $\beta$ :  
24: for  $t = T - 1$  to  $1$  do  
25:   for  $i = 1$  to  $N$  do  
26:      $\beta_t(i) = \sum_{j=1}^N a_{ij} b_j(O_{t+1}) \beta_{t+1}(j)$   
27:   end for  
28: end for
```

E-Step and M-step: Gamma and Obs Prob

```
27:   Compute  $\gamma$ :  
28:   for  $t = 1$  to  $T$  do  
29:     for  $i = 1$  to  $N$  do
```

$$\gamma_t(i) = \frac{\alpha_t(i)\beta_t(i)}{\sum_{j=1}^N \alpha_t(j)\beta_t(j)}$$

```
30:     end for  
31:   end for
```

```
49:   Re-estimate  $B$ :  
50:   for  $j = 1$  to  $N$  do  
51:     for  $k = 1$  to  $M$  do
```

$$b_j(k) = \frac{\sum_{t=1}^T \delta(O_t = k)\gamma_t(j)}{\sum_{t=1}^T \gamma_t(j)}$$

where $\delta(\cdot)$ is 1 if the argument is true, 0 otherwise

```
52:     end for  
53:   end for
```


Baum-Welch Properties

- Each iteration changes the parameters in a way that is guaranteed to increase the likelihood of the data
- Anytime algorithm: Can stop at any time prior to convergence to get an approximate solution
- Converges to a local maximum

HMM: Three Fundamental Problems

Problem 1 (Likelihood):

Given an HMM $\lambda = (\mathcal{T}, \mathcal{B})$ and an observation sequence O , determine the likelihood $P(O \mid \lambda)$.

Problem 2 (Decoding):

Given an observation sequence O and an HMM $\lambda = (\mathcal{T}, \mathcal{B})$, discover the best hidden state sequence.

Problem 3 (Learning):

Given an observation sequence O and the set of states in the HMM, learn the HMM parameters \mathcal{T} and \mathcal{B} .

Review

Likelihood


$$p(3, 1, 3)?$$

$$\alpha_t(j) = p(o_1, o_2, \dots, o_t, s_t = j)$$

$$\alpha_t(j) = \sum_{i=1}^N \alpha_{t-1}(i) \cdot p(s_t = j \mid s_{t-1} = i) \cdot p(o_t \mid s_t = j)$$

$$p(O) = \sum_{i=1}^N \alpha_T(i)$$

Decoding



$$\arg \max_{s_1, s_2, s_3} p(s_1, s_2, s_3 \mid o_1 = 3, o_2 = 1, o_3 = 3)?$$




$$\arg \max_{s_1, s_2, s_3} p(o_1, o_2, o_3, s_1, s_2, s_3)$$

$$v_t(j) = \max_{s_{1:t-1}} p(s_{1:t-1}, o_{1:t}, s_t = j)$$

$$v_t(j) = \max_{i=1}^N v_{t-1}(i) \cdot t_{ij} \cdot b_j(o_t)$$

$$\max_{i=1}^N v_T(i)$$

Trellis: Forward vs. Viterbi





Why HMMs?

- Uncertainty: Zone of probabilistic reasoning
- Constructs: Sequences of states, a.k.a. trajectories
- Algorithms: Dynamic Programming, Expectation Maximisation
- States, Transitions and Observations

Markov Decision Processes

HMM: State Evolution



Source: [Sutton & Barto](#)

MDP

MDP : Tuple $\langle \mathcal{S}, \mathcal{A}, T, R, \rho \rangle$

\mathcal{S} : State Space

\mathcal{A} : Action Space

$T : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$: Probabilistic Transition Function

$R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$: Reward Function

ρ : Initial State Distribution


Difference from HMM?

$$s_t, a_t, r(s_t, a_t)$$



Richard Bellman, Source: [Wikipedia](#)

$$x_t, u_t, c(x_t, u_t)$$



Lev Pontryagin, Source: [Wikipedia](#)

MDP: State Evolution

