Review

* Hidden Markov Model (HMMs)

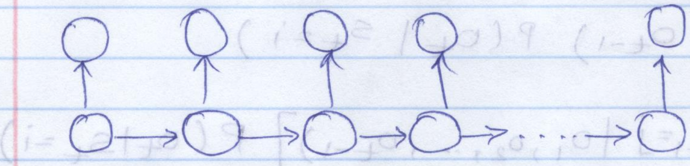

observations $o_t \in \{1 \ldots m\}$

states $s_t \in \{1, 2, \ldots, n\}$

* Parameters

$$a_{ij} = P(s_{t+1} = j \mid s_t = i)$$

$$b_{ik} = P(o_t = k \mid s_t = i)$$

$$\pi_i = p(s_1 = i)$$

* Key questions

1. How to compute likelihood $P(o_1, o_2, \ldots, o_T)$?
2. How to decode $\underset{s_1, \ldots, s_T}{\operatorname{argmax}} P(s_1, s_2, \ldots, s_T \mid o_1, o_2, \ldots, o_T)$?
3. How to monitor $P(s_t = i \mid o_1, o_2, \ldots, o_t)$?
4. How to estimate parameters $\{\pi_i, a_{ij}, b_{ik}\}$?

(3) Belief Updating

Define $q_{it} = P(s_t = i \mid o_1, o_2, \ldots, o_t)$

How to update from $q_{j,t-1}$ to q_{it}?

From product rule:

$$P(s_t = i \mid o_1, o_2, \ldots, o_t) = \frac{P(s_t = i, o_t \mid o_1, o_2, \ldots, o_{t-1})}{P(o_t \mid o_1, o_2, \ldots, o_{t-1})}$$

Numerator:

$$P(s_t = i, o_t \mid o_1, o_2, \ldots, o_{t-1})$$

Product Rule $= P(s_t = i \mid o_1, o_2, \ldots, o_{t-1})\, P(o_t \mid s_t = i, \cancel{o_1, o_2, \ldots, o_{t-1}})$

Conditional independance $= P(s_t = i \mid o_1, \ldots, o_{t-1})\, P(o_t \mid s_t = i)$

Marginalization $= \left[\sum_{j=1}^{n} P(s_t = i, s_{t-1} = j \mid o_1, o_2, \ldots, o_{t-1}) \right] P(o_t \mid s_t = i)$

Product rule $= \left[\sum_{j=1}^{n} P(s_{t-1} = j \mid o_1, o_2, \ldots, o_{t-1})\, P(s_t = i \mid s_{t-1} = j, \cancel{o_1, o_2, \ldots, o_{t-1}}) \right] \cdot P(o_t \mid s_t = i)$

$$= \sum_j q_{j,t-1} \cdot a_{ji} \cdot b_i(o_t)$$

Denominator:

$$P(o_t \mid o_1, o_2 \ldots, o_{t-1}) = \sum_{i=1}^{n} \underbrace{P(o_t, s_t = i \mid o_1, o_2 \ldots, o_{t-1})}_{\text{Numerator}}$$

Ratio (recursion):

$$q_{it} = \frac{\sum_j q_{j,t-1} \cdot a_{ji} \cdot b_i(o_t)}{\sum_{i', j'} q_{j',t-1} \cdot a_{j'i'} \cdot b_{i'}(o_t)}$$

for $i = 1 \ldots n$

$O(n^2)$ update from time $t-1$ to time t.

$\boxed{\text{Base case} \quad q_{i1} = \,? \quad \text{Exercise}}$

(4) Learning HMMs

Given: sequence of observations
$\{o_1, o_2, \ldots, o_T\}$
(just one sequence (for simplicity))

Goal: estimate (π_i, a_{ij}, b_{ik}) to maximize

$$P(o_1, o_2, \ldots, o_T)$$

Assume: # hidden states is known; n fixed

* EM algorithm in HMMs:

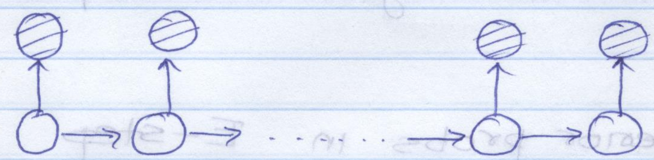

CPTs to estimate:
$\pi_i = P(s_1 = i)$
$a_{ij} = P(s_{t+1} = j \mid s_t = i)$
$b_{ik} = P(o_t = k \mid s_t = i)$

* E-step

Here are posterior prob. needed for EM

$$\begin{cases} P(s_1 = i \mid o_1, o_2, \ldots, o_T) \\ P(s_{t-1} = i, s_t = j \mid o_1, o_2, \ldots, o_T) \\ P(s_t = i, o_t = k \mid o_1, o_2, \ldots o_T) = I(o_t, k)\, P(s_t = i \mid o_1, o_2 \ldots o_T) \end{cases}$$

* How to compute these posterior probs?

Analogous to $\alpha_{it} = P(o_1, \ldots, o_t, s_t = i)$

define $\beta_{it} = P(o_{t+1}, o_{t+2}, \ldots, o_T \mid s_t = i)$
↑ start at t+1 ↑ conditioning on $s_t = i$

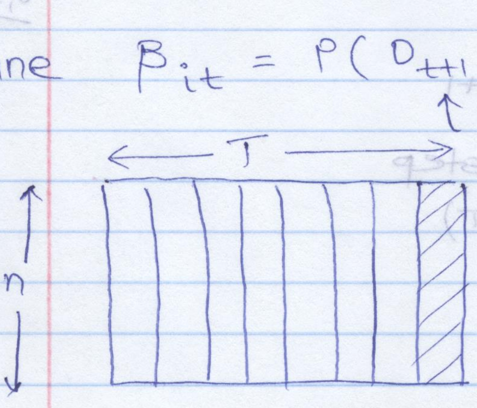

* Recursion
1) base case $\beta_{iT} = 1$ for all i

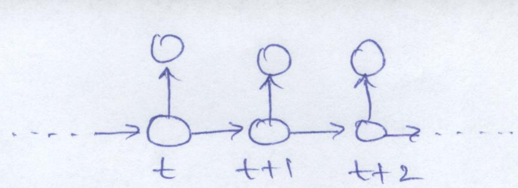

(2) $\beta_{it} = P(o_{t+1}, \cdots, o_T \mid s_t = i)$

Marginalization $= \sum_{j=1}^{n} P(s_{t+1} = j, o_{t+1}, \ldots, o_T \mid s_t = i)$

Product rule + Conditional Independance $= \sum_{j=1}^{n} P(s_{t+1} = j \mid s_t = i) \cdot P(o_{t+1} \mid s_{t+1} = j, \cancel{s_t = i}) \cdot P(o_{t+2}, \ldots, o_T \mid s_{t+1} = j, \cancel{s_t = i, o_{t+1}})$

$$= \boxed{\sum_{j=1}^{n} a_{ij} \, b_j(o_{t+1}) \, \beta_{j,t+1}}$$

" Forward - backward algorithm computes α_{it}, β_{it} in HMMs

* Return to posterior probs in E-step

Product rule $P(s_{t+1} = j, s_t = i \mid o_1, o_2, \ldots, o_T) = \dfrac{P(s_t = i, s_{t+1} = j, o_1, o_2, \ldots, o_T)}{P(o_1, o_2, \ldots, o_T)}$

Denominator $= P(o_1, o_2, \ldots, o_T) = \sum_{i=1}^{n} \alpha_{iT}$

Numerator :

$P(o_1, o_2, \ldots, o_t, s_t = i, s_{t+1} = j, o_{t+1}, o_{t+2}, \ldots, o_T)$

Product Rule $= P(o_1, o_2, \ldots, o_t, s_t = i) \, P(s_{t+1} = j \mid s_t = i, \cancel{o_1, \ldots, o_t}) \, P(o_{t+1} \mid s_{t+1} = j, \cancel{s_t = i, o_1, \ldots, o_t}) \, P(o_{t+2}, \ldots, o_T \mid s_{t+1} = j, \cancel{s_t = i, o_1, \ldots, o_{t+1}})$

$= \alpha_{it} \cdot a_{ij} \, b_j(o_{t+1}) \, \beta_{j,t+1}$

Posterior probs for E-step

$P(s_t = i, s_{t+1} = j \mid o_1, o_2 \ldots o_T)$

$P(s_t = i \mid o_1, o_2, \ldots, o_T)$

$P(s_1 \mid o_1, o_2, \ldots, o_T)$

* last two are special cases of first.

Marginalization $P(s_t = i \mid o_1, \ldots, o_T) = \sum_{j=1}^{n} P(s_t = i, s_{t+1} = j \mid o_1, o_2 \cdots, o_T)$

$P(s_1 = i \mid o_1, \ldots, o_T)$ set $t = 1$ (above).

$\boxed{\text{EM algorithm for HMMs}}$

- Initialize $\{\pi_i, a_{ij}, b_{ik}\}$ with intelligent guesses
- Until convergence:
  - E-Step ⇒ – compute $P(o_1, o_2, \ldots, o_T)$
  - – compute posterior probs $P(s_t = i, s_{t+1} = j \mid o_1, o_2, \ldots, o_T)$ for all t, i, j
  - M-Step ⇒ – update CPTs

M-step updates:

$\pi_i \quad P(s_1 = i) \longleftarrow P(s_1 = i \mid o_1, o_2 \cdots, o_T)$

$$a_{ij} \quad P(s_{t+1} = j \mid s_t = i) \longleftarrow \frac{\sum_t P(s_t = i, s_{t+1} = j \mid o_1, \ldots, o_T)}{\sum_t P(s_t = i \mid o_1, o_2 \cdots, o_T)}$$

$$b_{ik} \quad P(o_t = k \mid s_t = i) \longleftarrow \frac{\sum_t P(s_t = i \mid o_1 \cdots, o_T)\, I(o_t, k)}{\sum_t P(s_t = i \mid o_1, \ldots, o_T)}$$

* Complexity of inference | learning in HMMs

1> compute $P(o_1, \ldots, o_T)$

2> decode $S^* = \operatorname{argmax} P(s_1, \ldots, s_T \mid o_1, \ldots o_T)$

4> performing one iteration of EM algorithm

} all have complexity $O(Tn^2)$
No dependence on $m = \#$ observat^{n}s

Quiz 2

$\boxed{\text{Reinforcement learning}}$

Q: How should [embodied | embedded | situated | decision-making | autonomous] agent act and learn from experience in the world?

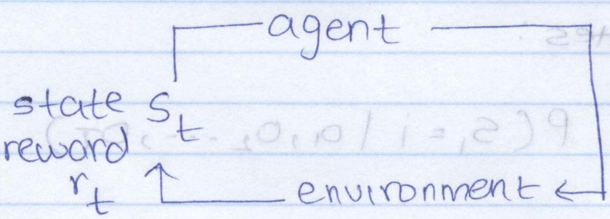

Ex: robot navigation
game-playing