CS221 Problem Workout

Week 8

Stanford University

Introduction

Jenn Grannen

General OH: Mondays 9:30-11:00 Bytes + Zoom
HW OH: Fridays 9:30-11:00 Huang Basement

Jeremy Kim

HW OH: Wednesdays 3:30-5:30 Huang
HW OH: Sunday 2:00-3:00 Online

Outline

- **HMM Review**
- Bayesian networks: Learning
  - Maximum likelihood
  - Smoothing
  - EM Algorithm
- Problem discussion

Hidden Markov Model (HMM) Review

Defining HMMs

Defining HMMs - Ice Cream Example

Famous problem by Jason Eisner (2002) where you want to predict if a day was COLD or HOT (your hidden states) based on records of the # of ice creams (your known evidence) Eisner ate that day.

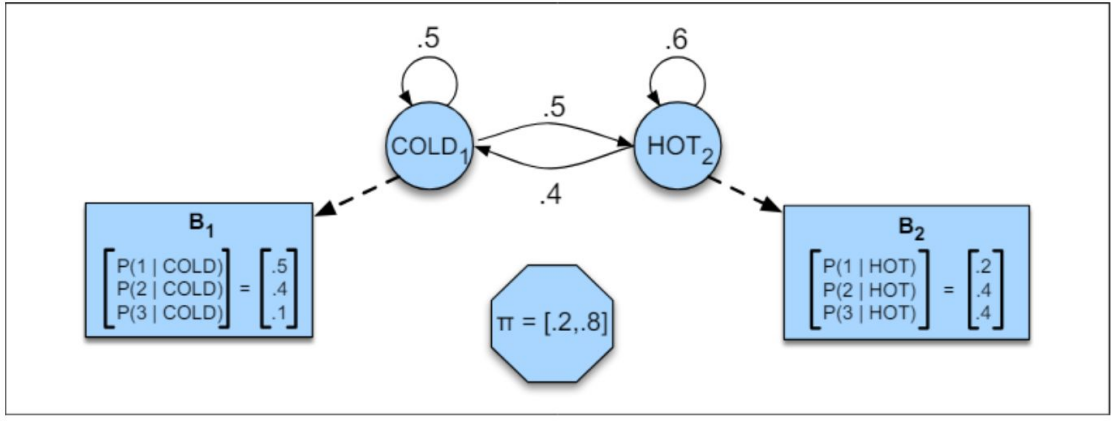

- $S = \{s_1 ... s_N\}$, N states (2 states: cold or hot)
- $A = a_{11} ... a_{ij} ... a_{NN}$, transition probabilities (e.g. cold $\rightarrow$ hot?)
- $B = b_i(o_t)$, emission probabilities (e.g. 3 ice creams $\rightarrow$ hot?)
- $\pi = \{\pi_1 ... \pi_N\}$, initial probabilities (e.g. start $\rightarrow$ hot?)

Defining HMMs - Ice Cream Example

Famous problem by Jason Eisner (2002) where you want to predict if a day was COLD or HOT (your hidden states) based on records of the # of ice creams (your known evidence) Eisner ate that day.

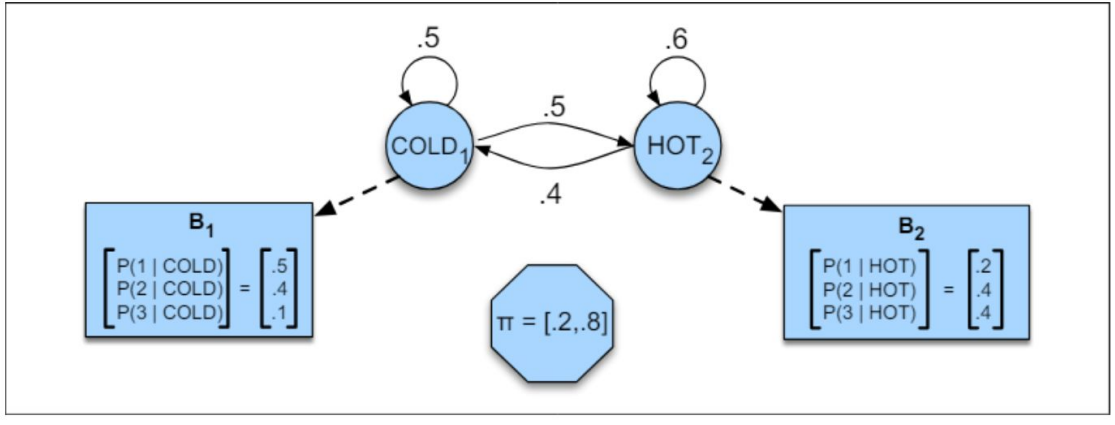

HMM problem to motivate the Forward Algorithm:

- Given HMM λ (like above), what is the probability $P(O|\lambda)$ of a specific observation sequence O (evidence e.g. 3 1 3)?

Hidden Markov Model (HMM) Review

The Forward Algorithm

The Forward Algorithm - Ice Cream Example

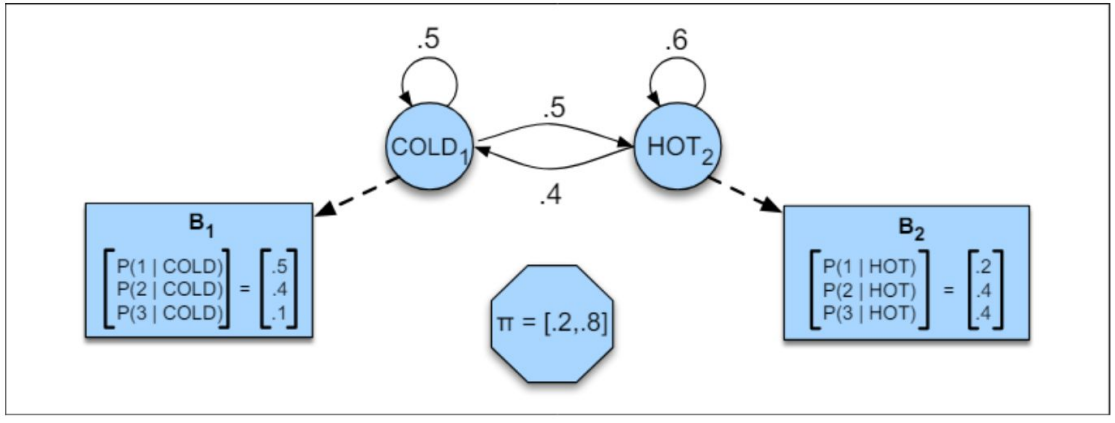

HMM problem to motivate the Forward Algorithm:

- Given HMM λ (like above), what is the probability $P(O|\lambda)$ of a specific observation sequence O (evidence e.g. 3 1 3)?

First, consider an easier problem: suppose our states are not hidden (we just have a "Markov model") and we have Q = (hot hot cold).

What is the probability (aka likelihood) of O = 3, 1, 3?

The Forward Algorithm - Ice Cream Example

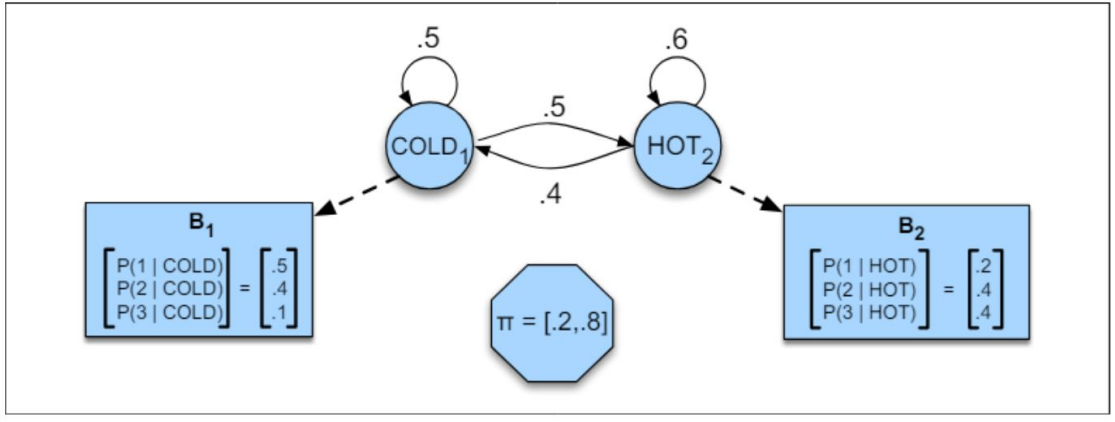

First, consider an easier problem: suppose our states are not hidden (we just have a "Markov model") and we have Q = (hot hot cold).

What is the probability (aka likelihood) of O = 3, 1, 3?

$$P(O|Q) = \prod_{t}^{T} P(o_t|q_t)$$

$$P(3\ 1\ 3|\text{hot hot cold}) = P(3|\text{hot})P(1|\text{hot})P(3|\text{cold})$$

The Forward Algorithm - Ice Cream Example

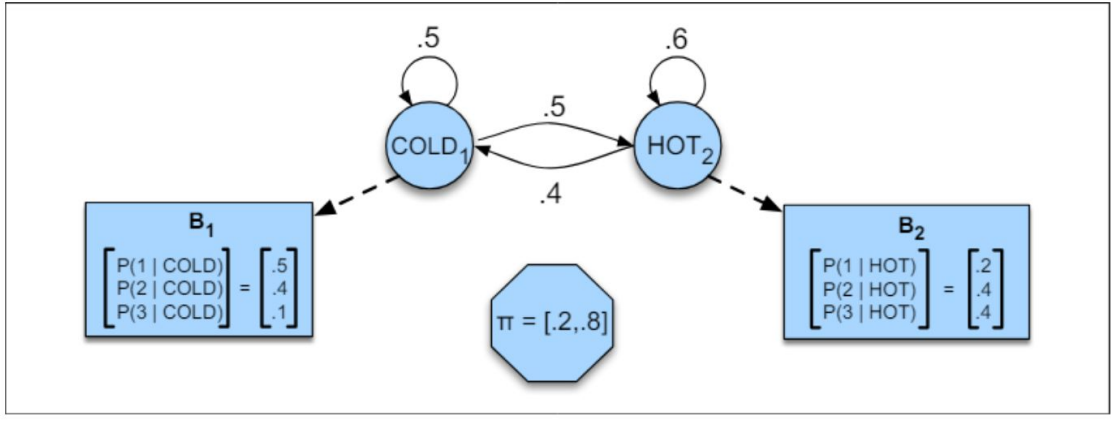

Simplification: probability $O = 3, 1, 3$ given $Q =$ (hot hot cold)?

$$P(3\ 1\ 3|\text{hot hot cold}) = P(3|\text{hot})P(1|\text{hot})P(3|\text{cold})$$

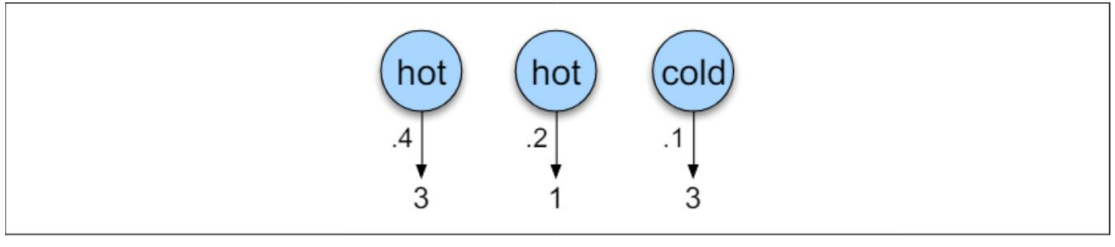

The Forward Algorithm - Ice Cream Example

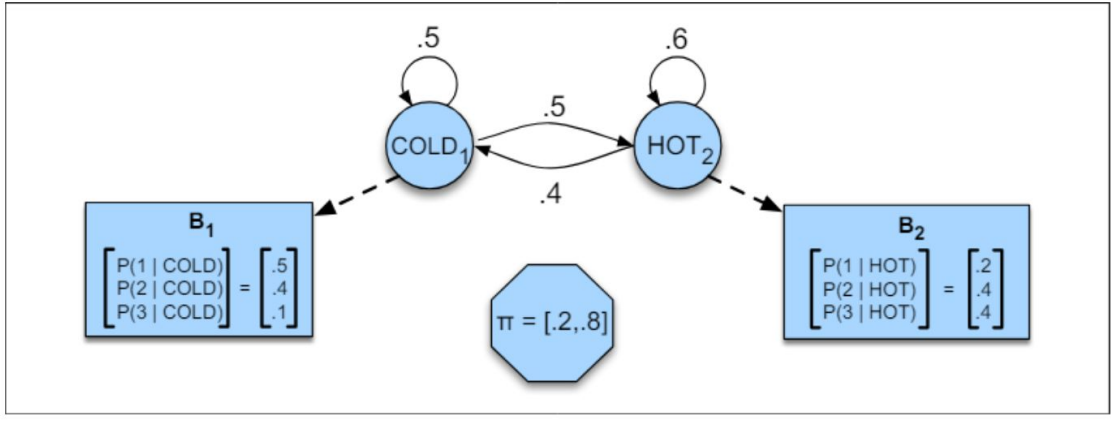

Simplification: probability $O = 3, 1, 3$ given $Q =$ (hot hot cold)?

$$P(3\ 1\ 3|\text{hot hot cold}) = P(3|\text{hot})P(1|\text{hot})P(3|\text{cold})$$

Back to the original problem: we don't know the actual weather sequence – it's a HIDDEN Markov model!

What is the probability of 3 1 3 given the HMM?

The Forward Algorithm - Ice Cream Example

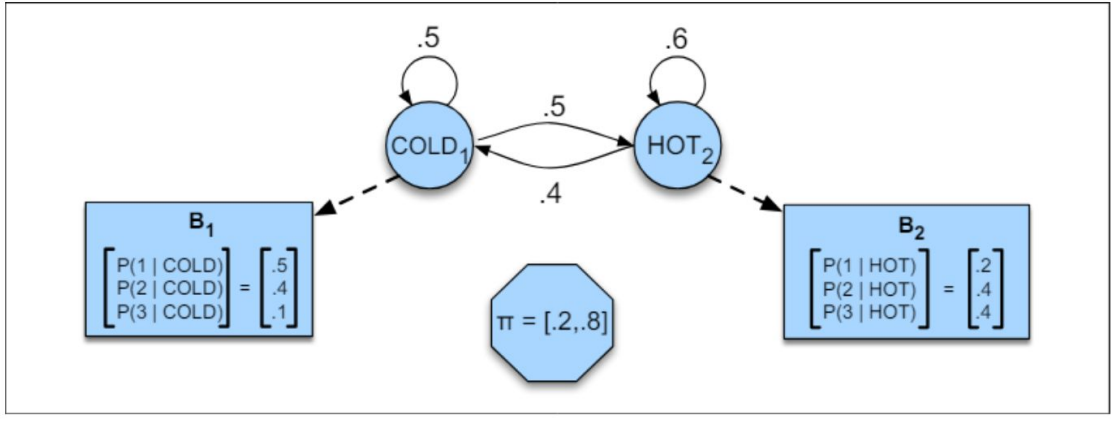

Back to the original problem: we don't know the actual weather sequence – it's a HIDDEN Markov model!

What is the probability of 3 1 3 given the HMM?

Brute Force: Sum over all possible weather sequences:

P(3 1 3, cold cold cold)? *P*(3 1 3, hot cold cold)?

P(3 1 3, hot hot cold)? etc...?

Then add them all together...

The Forward Algorithm - Ice Cream Example

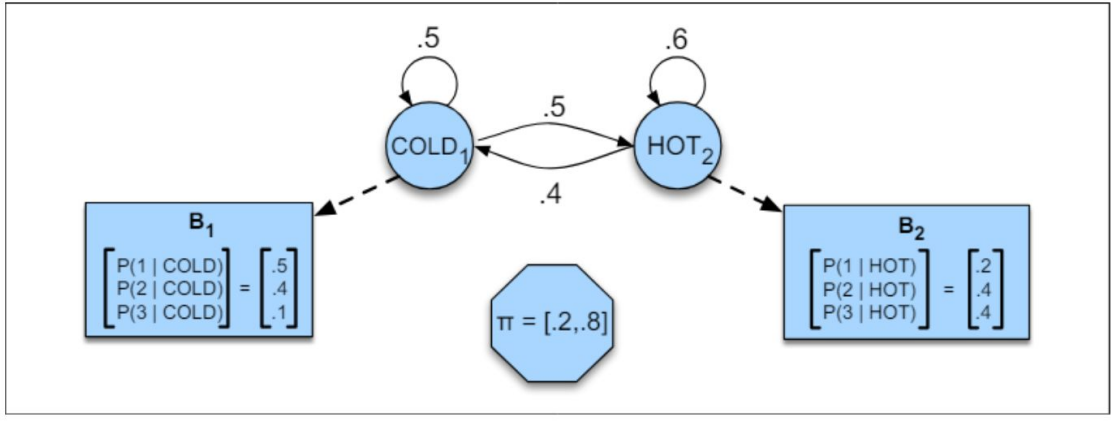

What is the probability of 3 1 3 given the HMM?

Brute Force: Sum over all possible weather sequences:
P(3 1 3, cold cold cold)? P(3 1 3, hot cold cold)? etc...?

$P(O, Q)$ is the joint probability:

$$P(O, Q) = P(O|Q)P(Q) = \prod_{t}^{T} P(o_t|q_t) \prod_{t}^{T} P(q_t|q_{t-1})$$

The Forward Algorithm - Ice Cream Example

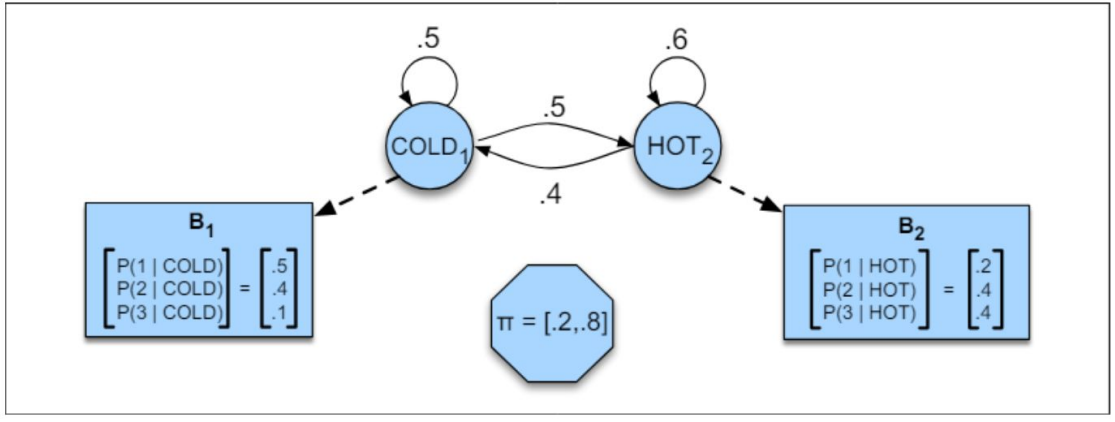

$P(O, Q)$ is the joint probability:

$$P(O, Q) = P(O|Q)P(Q) = \prod_t^T P(o_t|q_t) \prod_t^T P(q_t|q_{t-1})$$

Example: joint probability of $O = 3\ 1\ 3$ and $Q =$ hot hot cold

$P(O, Q) = P(3|\text{hot})P(1|\text{hot})P(3|\text{cold})P(\text{hot}|\text{start})P(\text{hot}|\text{hot})P(\text{cold}|\text{hot})$

The Forward Algorithm - Ice Cream Example

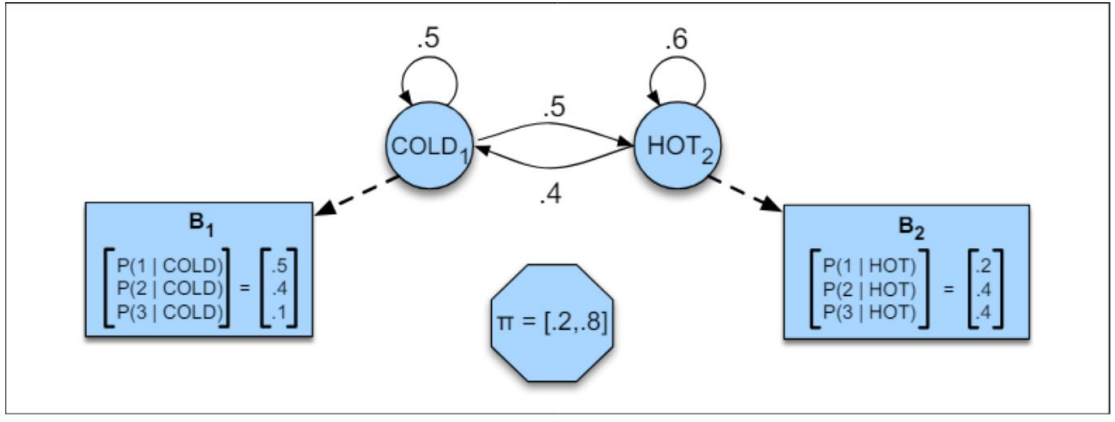

Example: joint probability of $O = 3\ 1\ 3$ and $Q =$ hot hot cold

$P(O, Q) = P(3|\text{hot})P(1|\text{hot})P(3|\text{cold})P(\text{hot}|\text{start})P(\text{hot}|\text{hot})P(\text{cold}|\text{hot})$

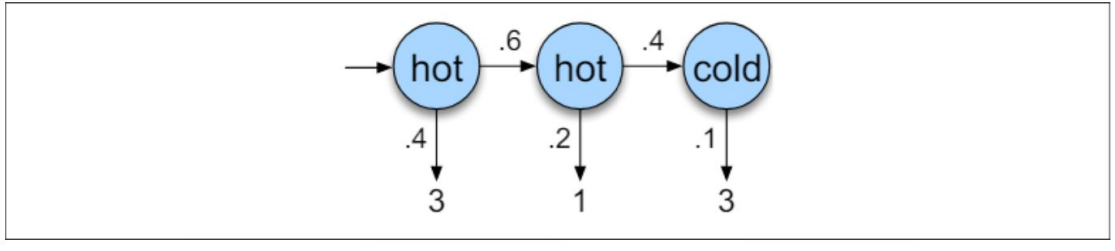

The Forward Algorithm - Ice Cream Example

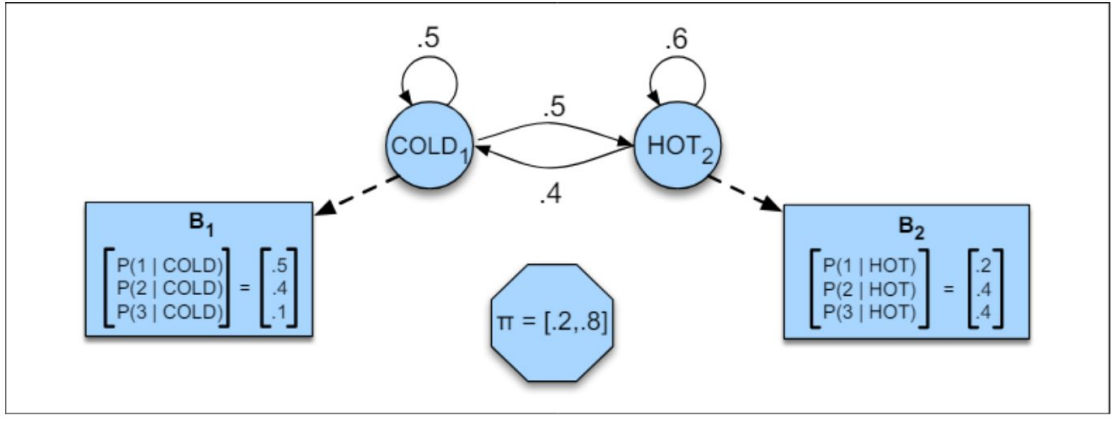

What is the probability of 3 1 3 given the HMM?

Brute Force: Sum over all possible weather sequences:

P(3 1 3, cold cold cold) + P(3 1 3, hot cold cold) + P(3 1 3, hot hot cold) + ...

This is a N^T operation with N states and T observations!

Not efficient for more complex problems!

The Forward Algorithm - Ice Cream Example

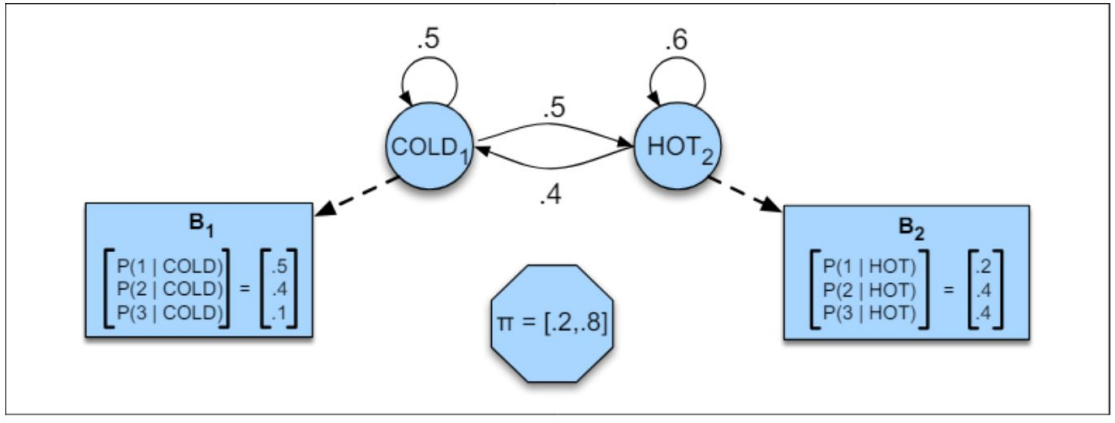

What is the probability of 3 1 3 given the HMM?

$P(3\ 1\ 3, \text{cold cold cold}) + P(3\ 1\ 3, \text{hot cold cold}) + P(3\ 1\ 3, \text{hot hot cold}) + \ldots$

This is a N^T operation with N states and T observations!

Forward Algorithm does this in $O(N^2 T)$ via dynamic programming!

The Forward Algorithm - Ice Cream Example

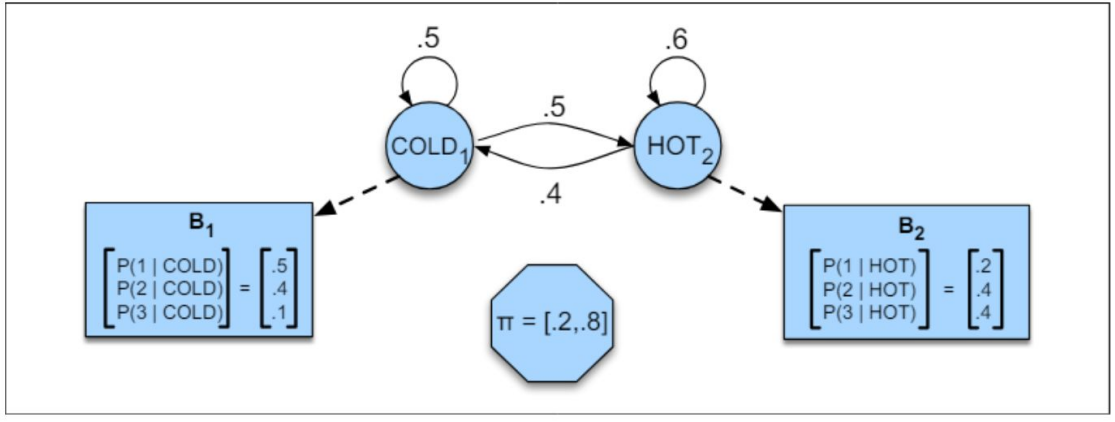

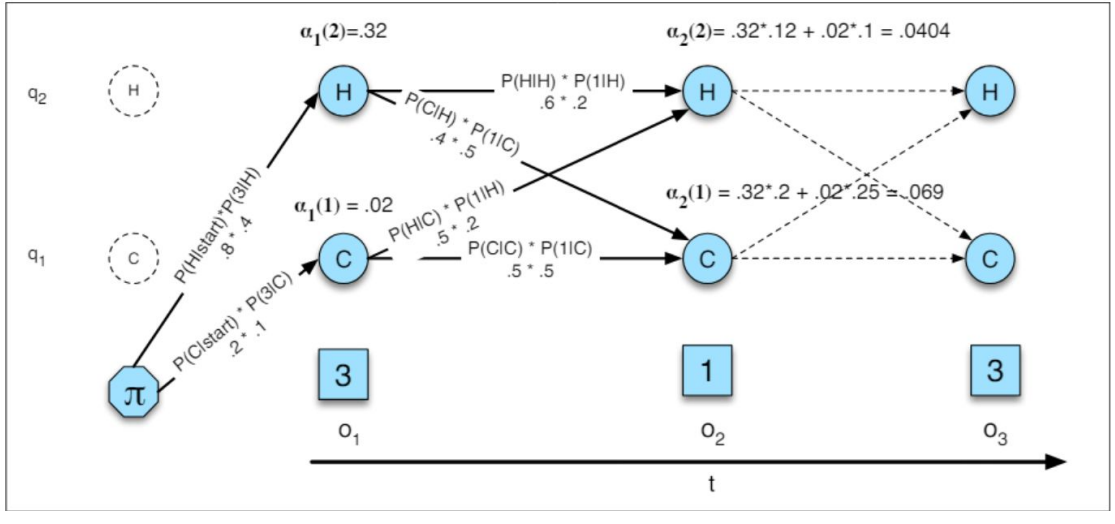

The Forward Algorithm - Ice Cream Example

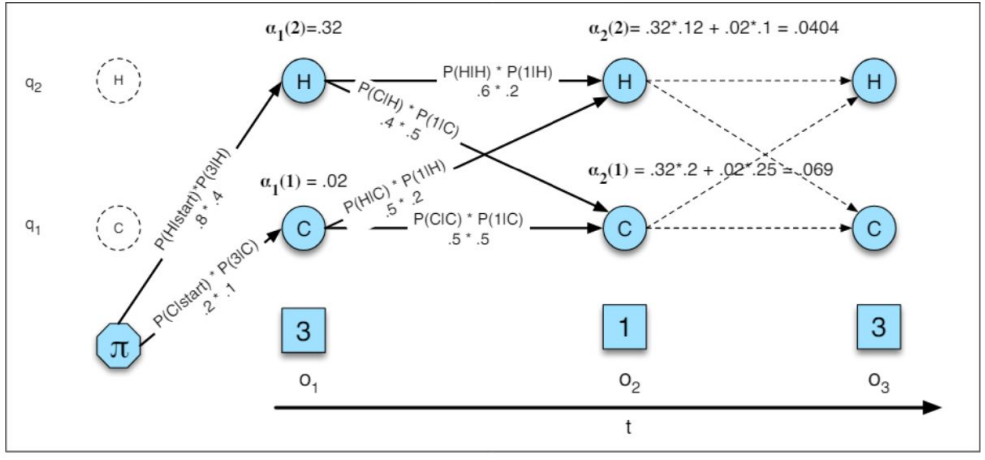

Formally, for each cell $\alpha_t(j)$ in our lattice structure, we compute

$$\alpha_t(j) = \sum_{i}^{N} \alpha_{t-1}(i) a_{ij} b_j(o_t)$$

and the probability of sequence 3 1 3 is at the end

$$P(O|\lambda) = \sum_{i}^{N} \alpha_T(i)$$

Hidden Markov Model (HMM) Review

Relating back to Lecture

The Problem of Filtering

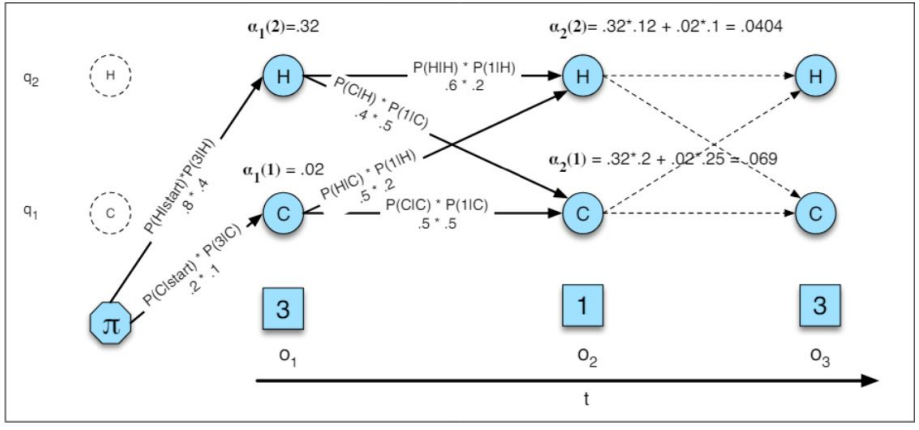

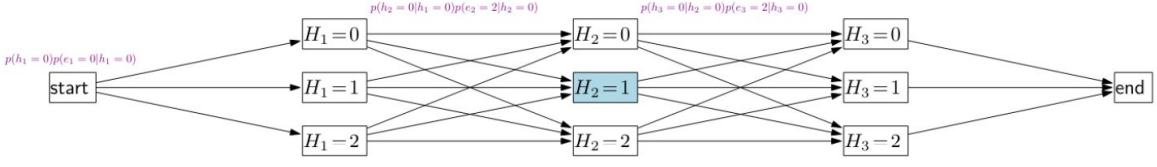

Problem of Filtering: what is the distribution of a hidden state H_i based on the observations aka evidence (E in lecture) so far?

Check your understanding: what is the distribution of q_2 in the ice cream example given observations O: $o_1 = 3$ and $o_2 = 1$?

The Problem of Filtering

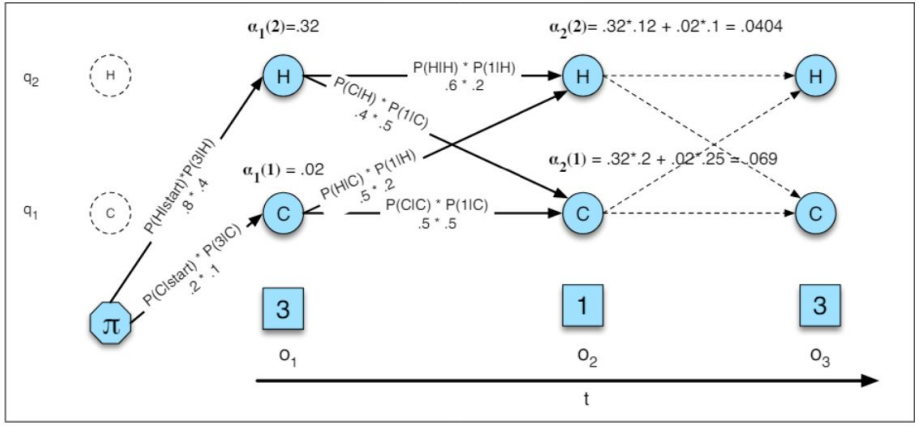

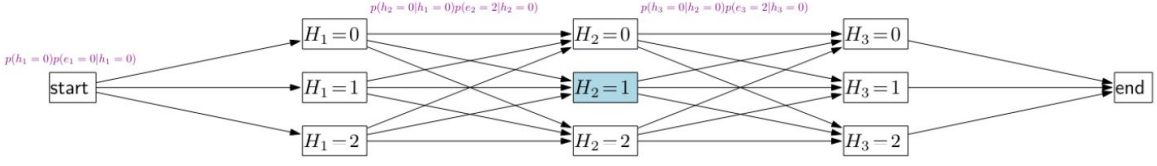

Problem of Filtering: what is the distribution of q_2 in the ice cream example given observations O: $o_1 = 3$ and $o_2 = 1$?

$$P(q_2 = \text{H} \mid o_1, o_2) = \frac{0.0404}{0.0404+0.069}, \; P(q_2 = \text{C} \mid o_1, o_2) = \frac{0.069}{0.0404+0.069}$$

The Problem of Smoothing

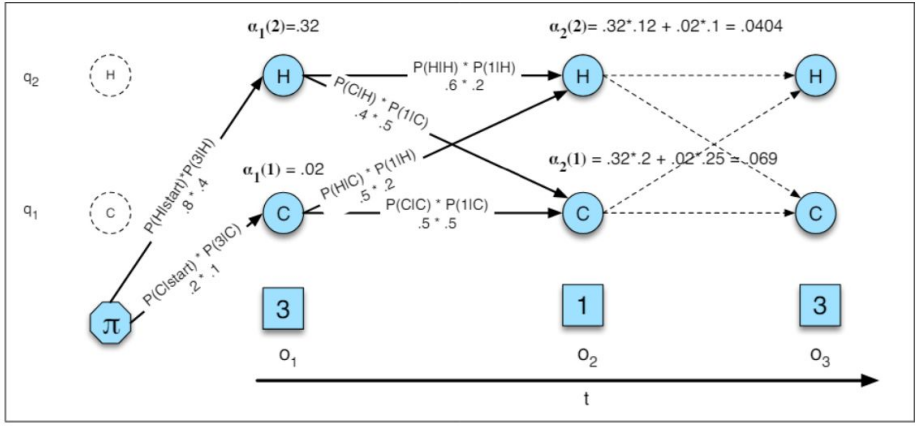

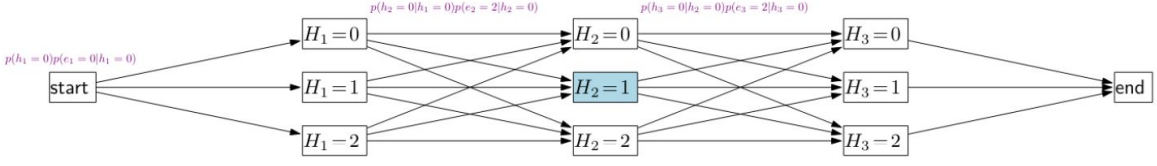

Problem of Smoothing: what is the distribution of a hidden state H_i based ALL observations aka evidence from start to end?

Forward Algorithm is not enough! What if hypothetically a later transition is 0?

The Problem of Smoothing

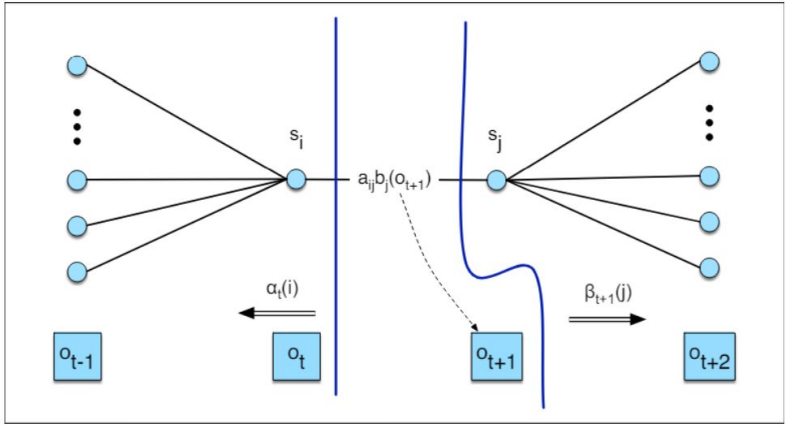

For Smoothing, need Forward AND Backward passes!

- Forward: compute $\alpha_t(i)$ or F from lecture.
- Backward: compute $\beta_t(i)$ or B from lecture.
- Define $S = FB$, that is for each cell in the lattice, multiply the forward and backward results together.

What happens now if there is a 0 along the backward pass?

The Problem of Smoothing

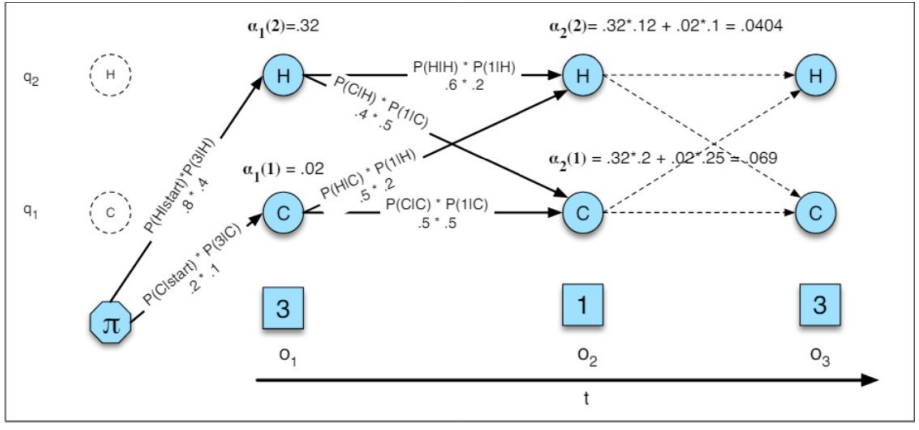

- Forward: compute $\alpha_t(i)$ or F from lecture.
- Backward: compute $\beta_t(i)$ or B from lecture.
- Define $S = FB$

Suppose $\beta_2(1) = 0.03, \beta_2(2) = 0.02$ (made up numbers).
What is the distribution of q_2 given all observations O?

The Problem of Smoothing

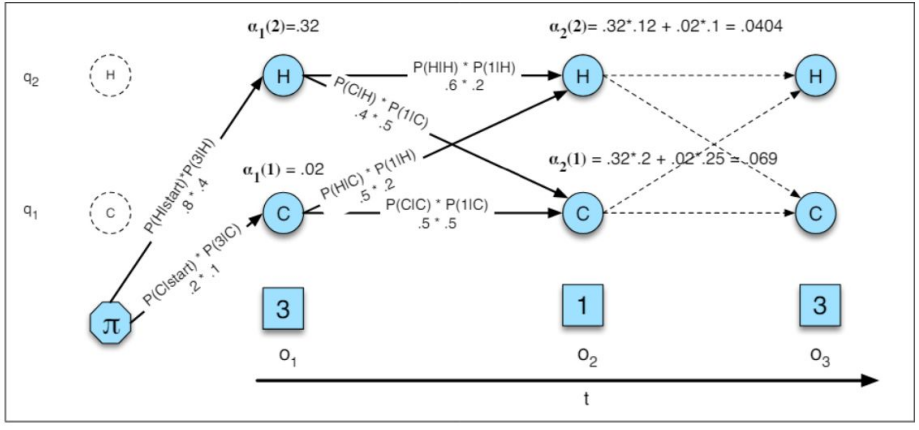

Suppose $\beta_2(1) = 0.03, \beta_2(2) = 0.02$ (made up numbers).

What is the distribution of q_2 given all observations O?

$$P(q_2 = \text{H} \mid O) = \frac{0.0404 * 0.02}{0.0404 * 0.02 + 0.069 * 0.03}$$

$$P(q_2 = \text{C} \mid O) = \frac{0.069 * 0.03}{0.0404 * 0.02 + 0.069 * 0.03}$$

Check back on the lecture slides to make sure you see the parallel!

Hidden Markov Model (HMM) Review

Particle Filtering

Motivation for Particle Filtering

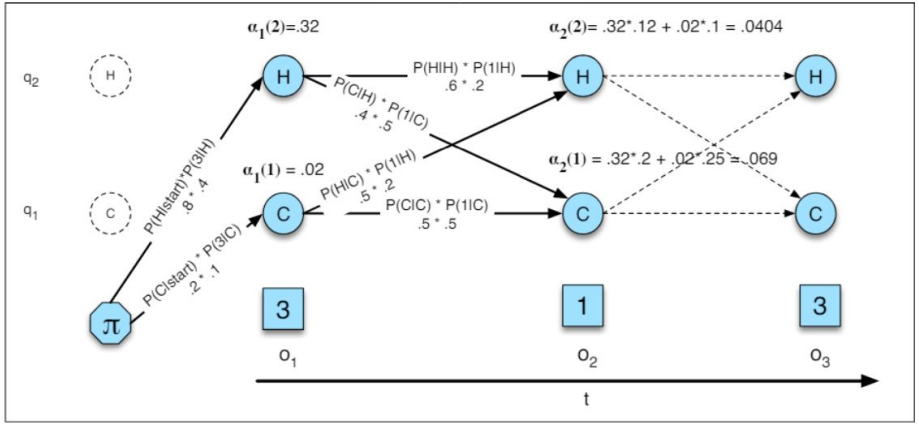

For T observations and N possible states (i.e. $|\text{domain}| = N$), the Forward-Backward Algorithm is $O(2 * N^2 T) \rightarrow O(N^2 T)$.

This can still be slow if N is large! Or consider if the domain is based on a continuous function, e.g. instead of just hot or cold, we have to consider a spectrum of floating point temperatures $[0, 100]$.

Particle Filtering

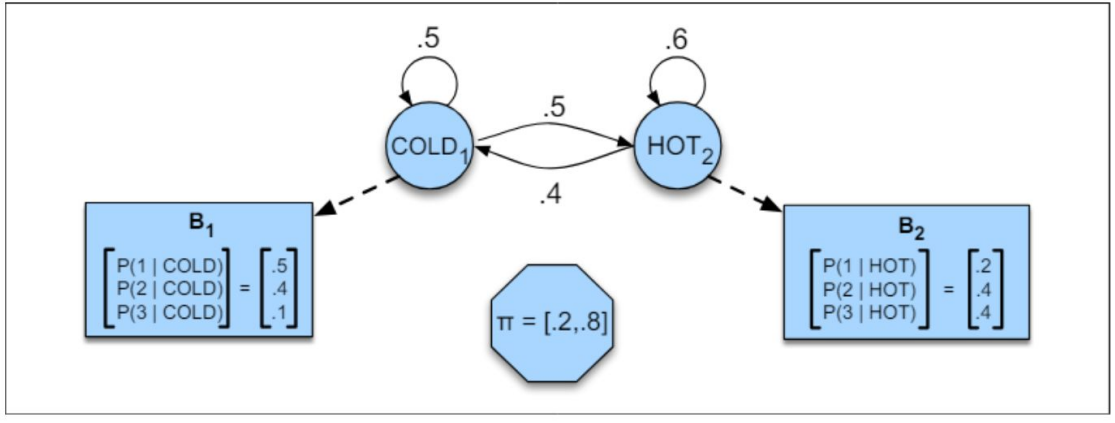

Big idea of Particle Filtering: introduce sampling!

1. First, we propose assignments aka particles to each hidden state by sampling from the transition probabilities.

Example: proposing a value for q_1 involves sampling from $P(H|\text{start}) = 0.8$ and $P(C|\text{start}) = 0.2$, i.e. we have an 80% chance to pick hot, 20% chance to pick cold.

Particle Filtering

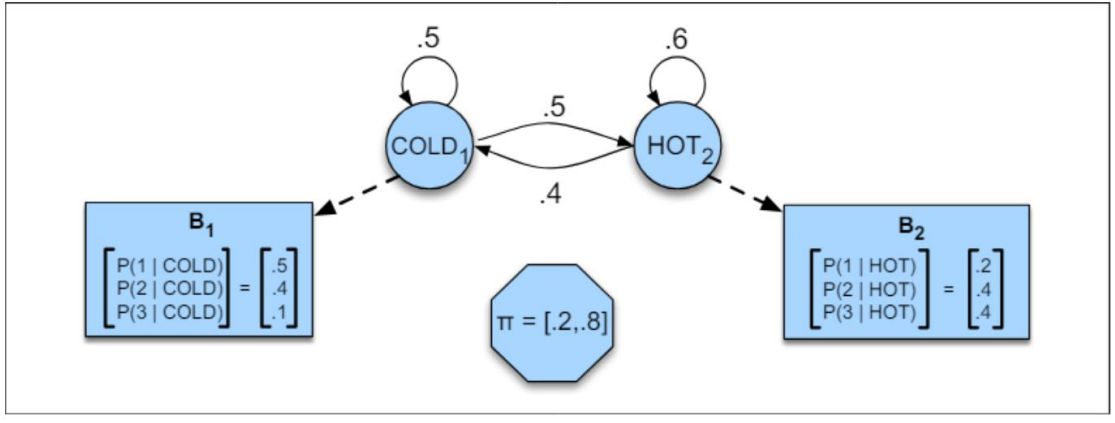

2 Second, we weight each assignment by the emission probabilities.

Example: suppose we have 3 particles of $q_1 = H, q_1 = H, q_1 = C$, and we have the observation $o_1 = 1$.
Then the weights of our particles are $P(1|H) = 0.2$, $P(1|H) = 0.2$, $P(1|C) = 0.5$ respectively.

Particle Filtering

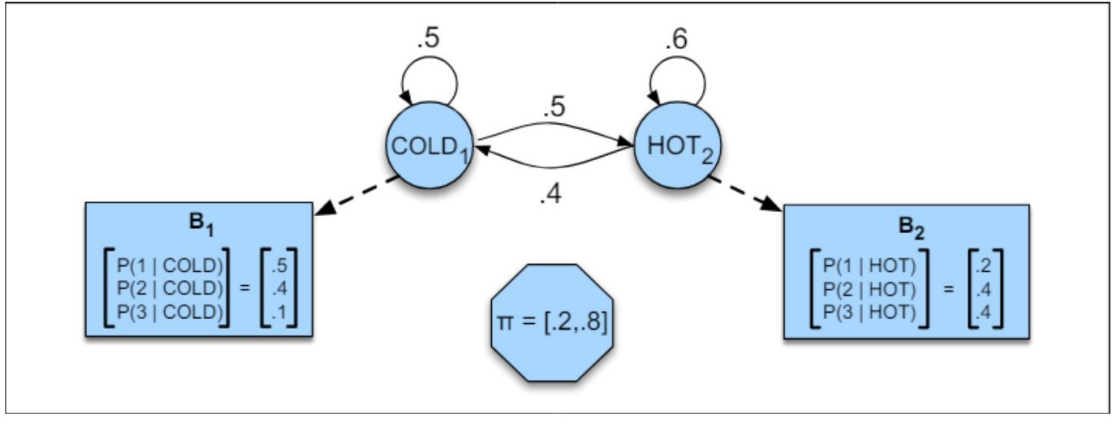

1. First, we propose assignments aka particles to each hidden state by sampling from the transition probabilities.
2. Second, we weight each assignment by the emission probabilities.
3. Third, we resample new assignments from the particles based on the weight distributions.

Particle Filtering

3 Third, we resample new assignments from the particles based on the weight distributions.

Example: suppose we have 3 particles of $q_1 = H, q_1 = H, q_1 = C$, and we have the observation $o_1 = 1$.
Then the weights of our particles are $P(1|H) = 0.2$, $P(1|H) = 0.2$, $P(1|C) = 0.5$ respectively.

Now to resample, we have the distribution:

- $P(q_1 \to H) = \frac{0.2}{0.2+0.2+0.5}$
- $P(q_1 \to H) = \frac{0.2}{0.2+0.2+0.5}$
- $P(q_1 \to C) = \frac{0.5}{0.2+0.2+0.5}$

Notice how even though our initial proposal had a higher chance to pick $q_1 = H$, we now have a higher chance to get $q_1 = C$!
The resampling takes into account the observations!

Particle Filtering

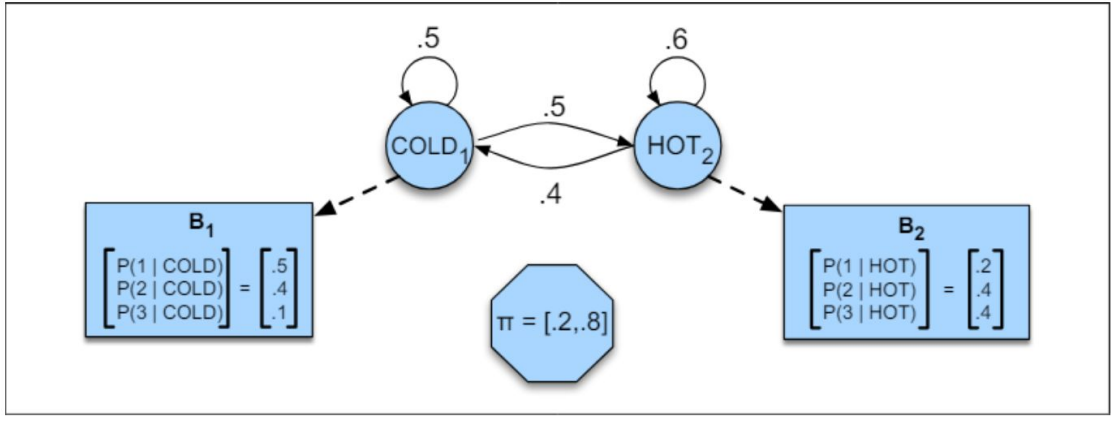

Suppose after all that, we have new assignments for our 3 particles: $q_1 = C, q_1 = C, q_1 = H...$

And repeated the propose process for q_2 to get: $(q_1, q_2) = (C, H)$; $(q_1, q_2) = (C, C)$; $(q_1, q_2) = (H, C)$ with $o_2 = 3...$

Particle Filtering

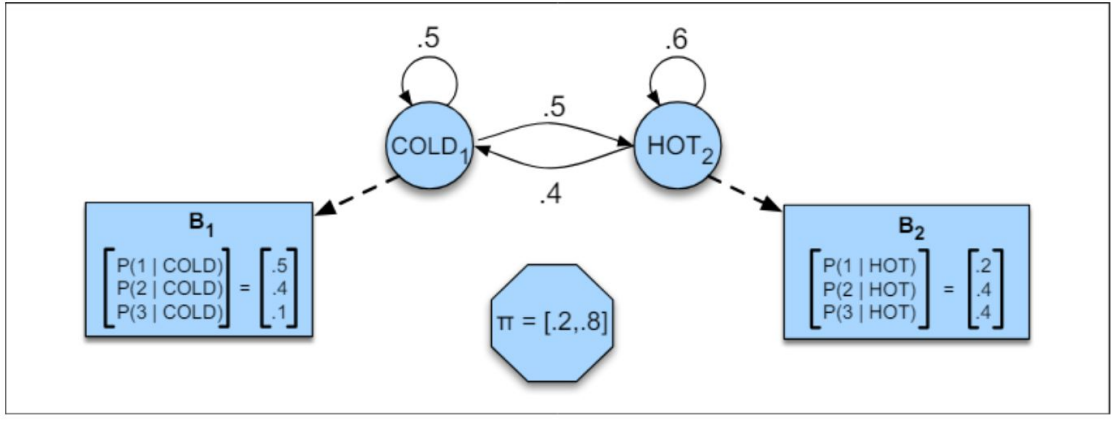

And repeated the propose process for q_2 to get: $(q_1, q_2) = (C, H)$; $(q_1, q_2) = (C, C)$; $(q_1, q_2) = (H, C)$ with $o_2 = 3$...

The weight process then assigns the particles:

- $(q_1, q_2) = (C, H)$: $P(3|H) = 0.4$
- $(q_1, q_2) = (C, C)$: $P(3|C) = 0.1$
- $(q_1, q_2) = (H, C)$: $P(3|C) = 0.1$

Particle Filtering

The weight process then assigns the particles:

- $(q_1, q_2) = (C, H)$: $P(3|H) = 0.4$
- $(q_1, q_2) = (C, C)$: $P(3|C) = 0.1$
- $(q_1, q_2) = (H, C)$: $P(3|C) = 0.1$

And the resample process then samples from the above 3 options, that is:

- $(q_1, q_2) = (C, H)$ has a $4/6$ chance of being picked.
- The other two each have a $1/6$ chance of being picked.

And so a possible resampling result might yield the particles: $(q_1, q_2) = (C, H), (C, H)$, and (C, C).
And you'd repeat the process with q_3...

Outline

- HMM Review
- **Bayesian networks: Learning**
  - Maximum likelihood
  - Smoothing
  - EM Algorithm
- Problem discussion

Bayesian networks: Learning

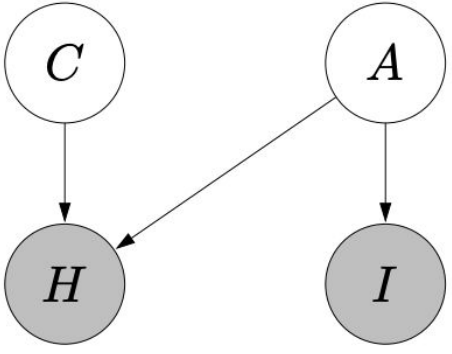

Given local probability distributions, i.e. P(x | parents(x))

Find conditional P(Q | E=e)

Inference

Given observations / samples

Find the local distributions, i.e. P(x | Parents(x))

Learning

Example

Variables:

- Genre $G \in \{\text{drama}, \text{comedy}\}$
- Rating $R \in \{1, 2, 3, 4, 5\}$

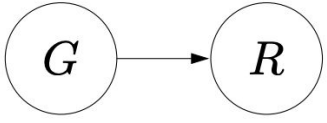

$$\mathbb{P}(G = g, R = r) = p_G(g) p_R(r \mid g)$$

$$\mathcal{D}_{\text{train}} = \{(\text{d}, 4), (\text{d}, 4), (\text{d}, 5), (\text{c}, 1), (\text{c}, 5)\}$$

Parameters: $\theta = (p_G, p_R)$

Example borrowed from lecture slides

Outline

- Bayesian networks: Learning
  - **Maximum likelihood**
  - Smoothing
  - EM Algorithm
- Problem discussion

Maximum likelihood

Input: training examples $\mathcal{D}_{\text{train}}$ of full assignments

Output: parameters $\theta = \{p_d : d \in D\}$

Algorithm: count and normalize

Count:

For each $x \in \mathcal{D}_{\text{train}}$:

For each variable x_i:

Increment $\text{count}_{d_i}(x_{\text{Parents}(i)}, x_i)$

Normalize:

For each d and local assignment $x_{\text{Parents}(i)}$:

Set $p_d(x_i \mid x_{\text{Parents}(i)}) \propto \text{count}_d(x_{\text{Parents}(i)}, x_i)$

Slide borrowed from lecture slides

Outline

- Bayesian networks: Learning
  - Maximum likelihood
  - **Smoothing**
  - EM Algorithm
- Problem discussion

Smoothing

Why?

- What if count is 0? Should P be 0?

How to solve?

- Initialize all counts with a non-zero constant **λ**

Observations

- Larger **λ** -> more uniform distributions, less influenced by data
- Smaller **λ** -> more influenced by data
- Infinite data -> effect of **λ** vanishes

Final algorithm

Input: training examples $\mathcal{D}_{\text{train}}$ of full assignments

Output: parameters $\theta = \{p_d : d \in D\}$

Algorithm: count and normalize

Count:

For each $x \in \mathcal{D}_{\text{train}}$:

For each variable x_i:

Increment $\text{count}_{d_i}(x_{\text{Parents}(i)}, x_i)$

Normalize:

For each d and local assignment $x_{\text{Parents}(i)}$:

Set $p_d(x_i \mid x_{\text{Parents}(i)}) \propto \text{count}_d(x_{\text{Parents}(i)}, x_i)$ + λ

Slide modified from lecture slides

Outline

- Bayesian networks: Learning
  - Maximum likelihood
  - Smoothing
  - **EM Algorithm**
- Problem discussion

EM Algorithm

Variables: H is hidden, $E = e$ is observed

Example:

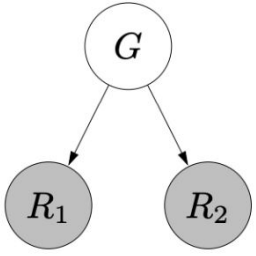

$H = G \qquad E = (R_1, R_2) \qquad e = (1, 2)$

$\theta = (p_G, p_R)$

Maximum marginal likelihood objective:

$$\max_\theta \prod_{e \in \mathcal{D}_{\text{train}}} \mathbb{P}(E = e; \theta)$$

$$= \max_\theta \prod_{e \in \mathcal{D}_{\text{train}}} \sum_h \mathbb{P}(H = h, E = e; \theta)$$

Slide borrowed from lecture slides

EM Algorithm

Initialize θ randomly

Until convergence:

 # E Step

for each e in Data:

for each h:

q(h; e) = P(H = h | E = e; θ) … inference

Update table from {e, count(e)} to {(h,e), (q(h; e) x count(e)}

Now no variables are hidden

 # M step

update(θ) using Table {(h,e), (q(h; e) x count(e)} … MLE

Summary

- Given data learn the parameters of bayesian net

MLE
$p \propto \text{count}(x_i; \text{parents}(x_i))$

Smoothing
$p \propto \text{count}(x_i; \text{parents}(x_i)) + \boldsymbol{\lambda}$

EM
Data is *incomplete*
E step: compute counts
M step: MLE

Outline

- HMM Review
- Bayesian networks: Learning
  - Maximum likelihood
  - Smoothing
  - EM Algorithm
- **Problem discussion**

Problem: P2, Winter 2021 Exam 2

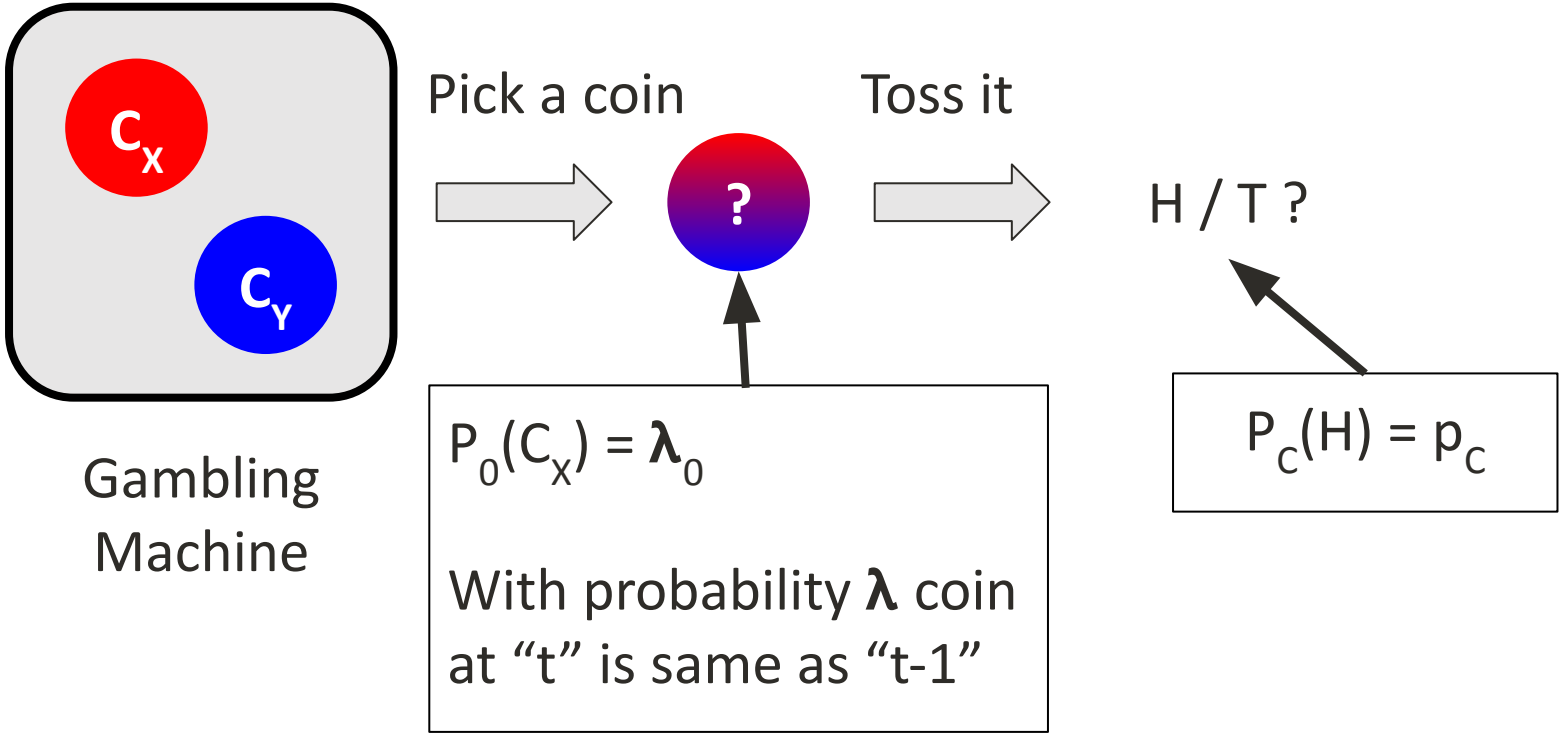

How does the bayesian net look?

$P(C_1) = \lambda_0$ if $C_1 = X$, $1-\lambda_0$ else

$P(C_{i+1} \mid C_i) = \lambda$ if $C_{i+1} = C_i$, $1-\lambda$ else

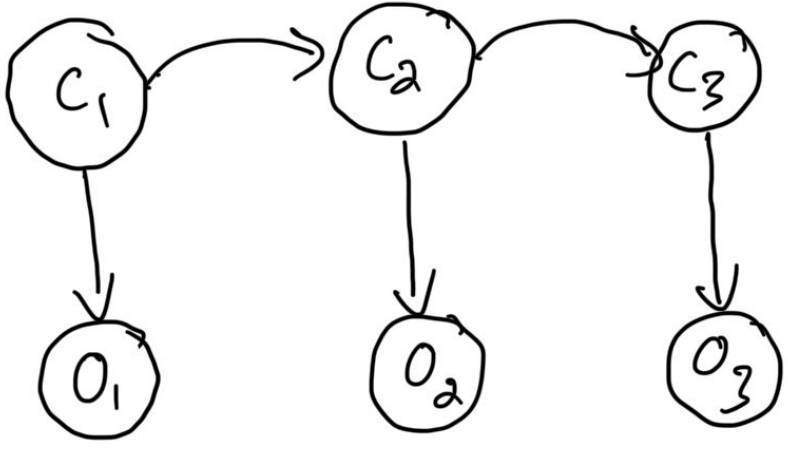

$P(O_i \mid C_i) = p_{C_i}$ if $O_i = H$, $1-p_{C_i}$ else

Learning using EM algorithm

- Data = {H, H, T}
- $\boldsymbol{\lambda}_0$ and $\boldsymbol{\lambda}$ are given. To find: p_X and p_Y

- Why do we need EM?
  - What is not observed?
  - C_i is not observed

- How do we use EM?
  - Compute $q(c_i)$ using p'_X and p'_Y
  - Use ML to update p'_X and p'_Y

Given q's compute update

	T=1	T=2	T=3
X	0.1	0.5	0.3
Y	0.9	0.5	0.7

Data = {H, H, T}

Compute:

C_i	O_i	Count
?	?	?

Given q's compute update

	T=1	T=2	T=3
X	0.1	0.5	0.3
Y	0.9	0.5	0.7

Data = {H, H, T}

C_i	O_i	Count
X	H	0.1
X	H	0.5
X	T	0.3

$P(H \mid X) = p'_X$

$= (0.1 + 0.5) / (0.1 + 0.5 + 0.3)$ … MLE

Thank You