

Classic HMM tutorial – see class website:

*L. R. Rabiner, "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition," Proc. of the IEEE, Vol.77, No.2, pp.257--286, 1989.

HMMs (cont.)

Machine Learning – 10701/15781

Carlos Guestrin

Carnegie Mellon University

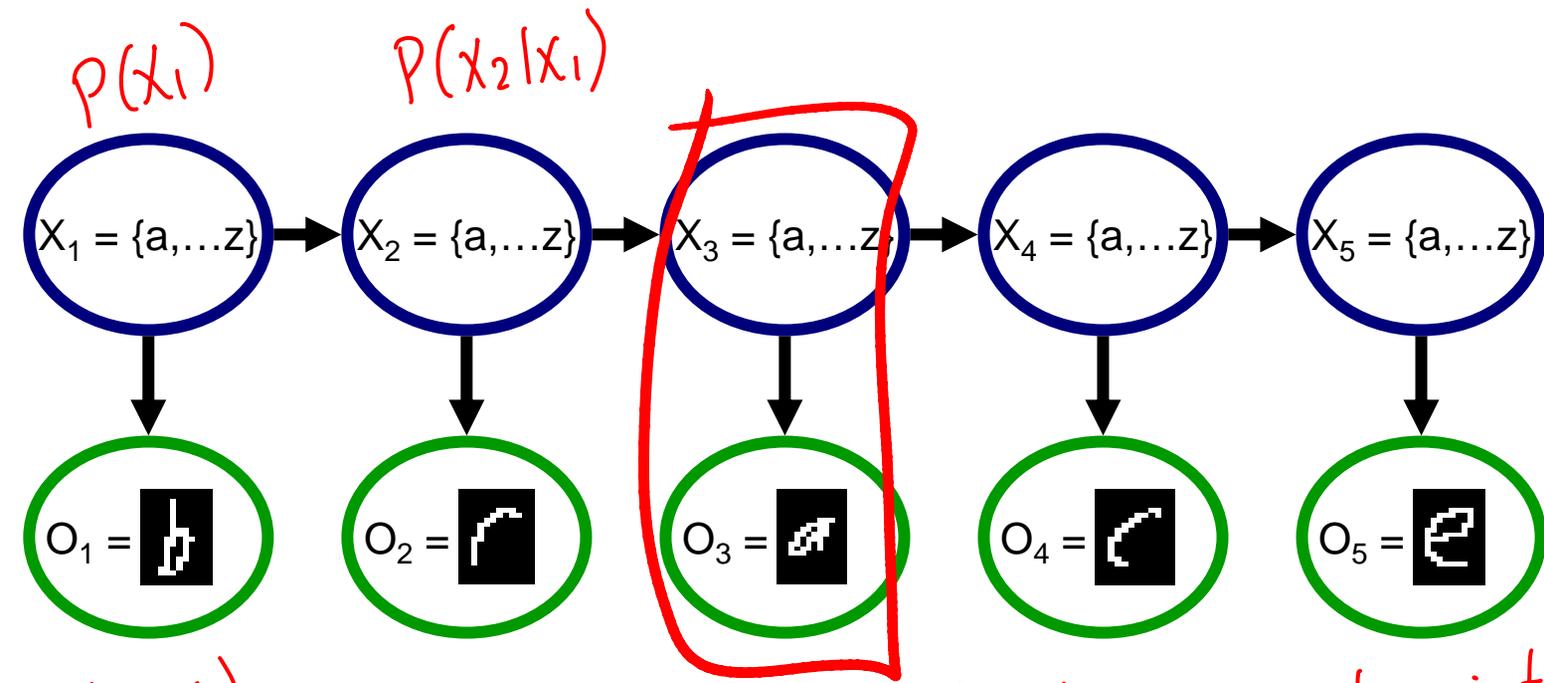
March 29th, 2006

Announcements



- This weeks recitation:
 - Go through several BNs topics, representation, inference, learning, in the context of an example → very useful for homework

Understanding the HMM Semantics



$P(O_i | X_i)$
 one option is
 Naive Bayes

$$P(O_i | X_i) = \prod_{\text{pixels}} P(O_i^j | X_i)$$

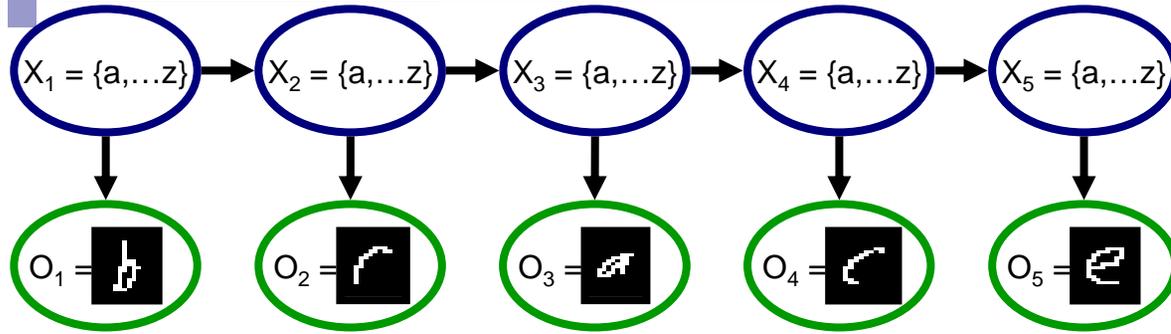
Markov assumption: the future is
 indep. of past given present

$$X_{1:t+1} \perp X_{t+1:n} | X_t$$

$$O_t \perp \text{everybody} | X_t$$

Music!

HMMs semantics: Details



Just 3 distributions:

$P(X_1)$

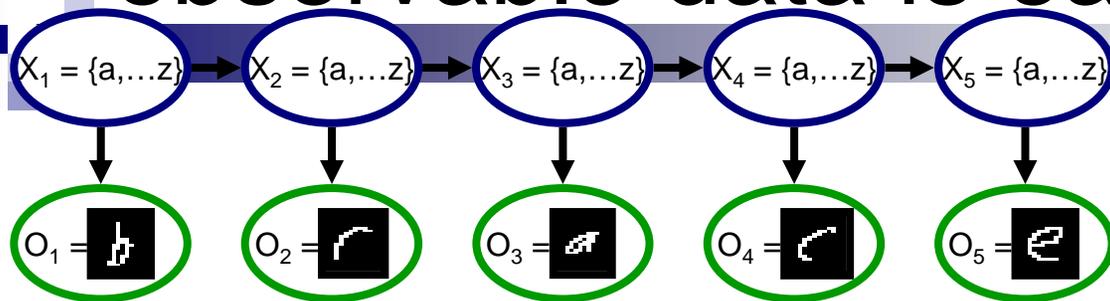
$P(X_i | X_{i-1})$

$P(O_i | X_i)$

observing a distribution $P(x_3 | x_2) \sim P(x_2 | x_1) \sim P(x_1 | x_4)$
 a after b has same prob. no matter where in word

prob. image given letter is same for all positions in word.

Learning HMMs from fully observable data is easy



Learn 3 distributions:

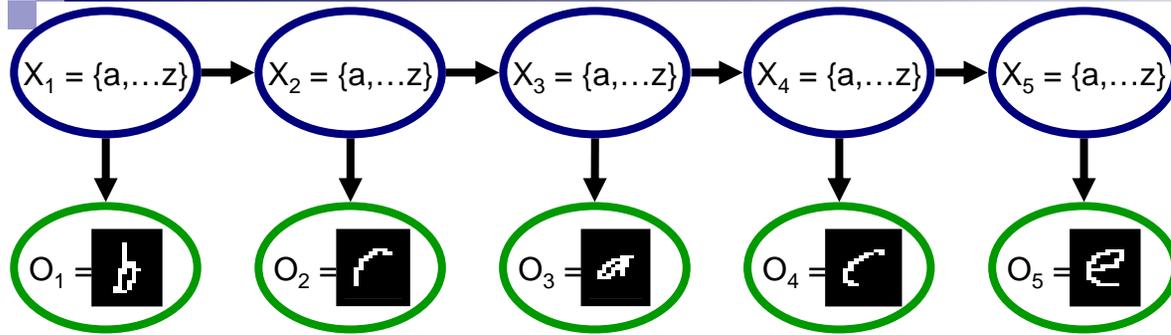
$$P(X_1^a) = \frac{\text{count}(\# \text{ first letter was } a)}{N = \text{dataset size}}$$

$$P(O_i^{\text{pixel } l \text{ is white}} | X_i^a) = \frac{\text{count}(\text{pixel } l \text{ was white, } X_i = a)}{\text{count}(X_i = a)}$$

$$P(X_i^a | X_{i-1}^b) = \frac{\text{count}(a \text{ appears after } b)}{\text{count}(\# \text{ of } b\text{'s that are not at the end of the word})}$$

select training data where letter was a

Possible inference tasks in an HMM



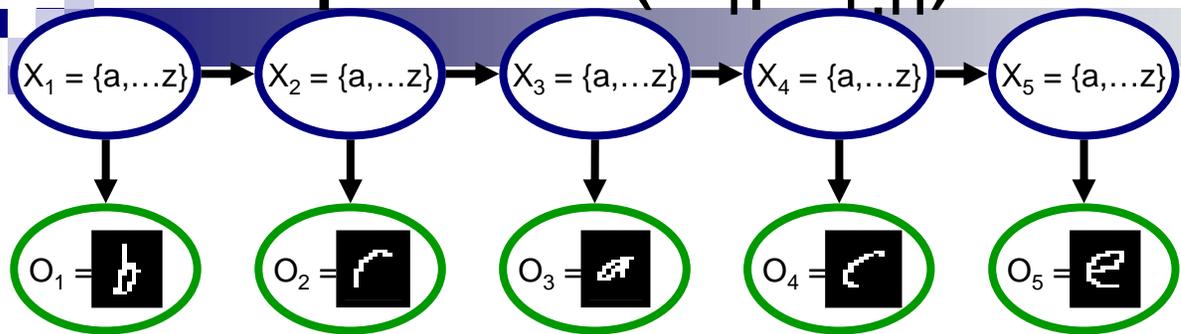
Marginal probability of a hidden variable:

$$\rightarrow P(X_i | O_1 = [1], O_2 = [7], O_3 = [4], \dots)$$

Viterbi decoding – most likely trajectory for hidden vars:

$$\underset{x_1 x_2 x_3 x_4 x_5}{\text{argmax}} P(x_1, x_2, x_3, x_4, x_5 | O_{1:5})$$

Using variable elimination to compute $P(X_i | o_{1:n})$



Compute:

$$P(X_i | o_{1..n})$$

Variable elimination order?

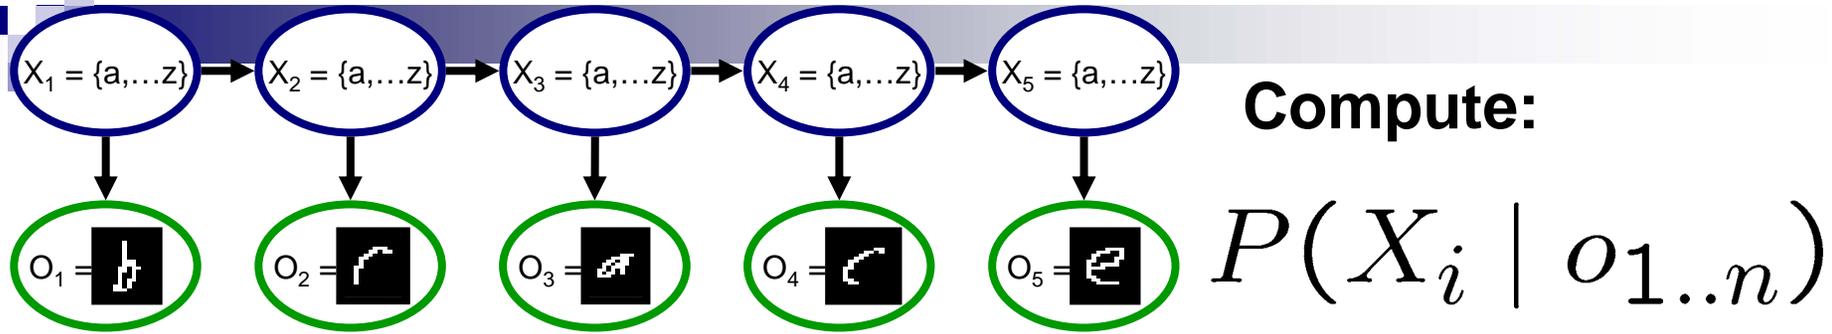
1, 2, 3, 4, ..., n

Example:
$$P(X_3 | o) = \sum_{x_1, x_2, x_4, x_5} P(x_1, x_2, x_3, x_4, x_5 | o)$$

$$\sum_{x_1, x_2, x_4, x_5} P(x_1) P(o_1 | x_1) P(x_2 | x_1) P(o_2 | x_2) P(x_3 | x_2) \dots$$

$$= \sum_{x_2, x_4, x_5} P(o_2 | x_2) P(x_3 | x_2) \dots \underbrace{\sum_{x_1} P(x_1) P(o_1 | x_1) P(x_2 | x_1)}_{g_1(x_2, o_1)}$$

What if I want to compute $P(X_i | o_{1:n})$ for each i ?



Variable elimination for each i ?

$P(X_1 | o)$, $P(X_2 | o)$, $P(X_3 | o)$...

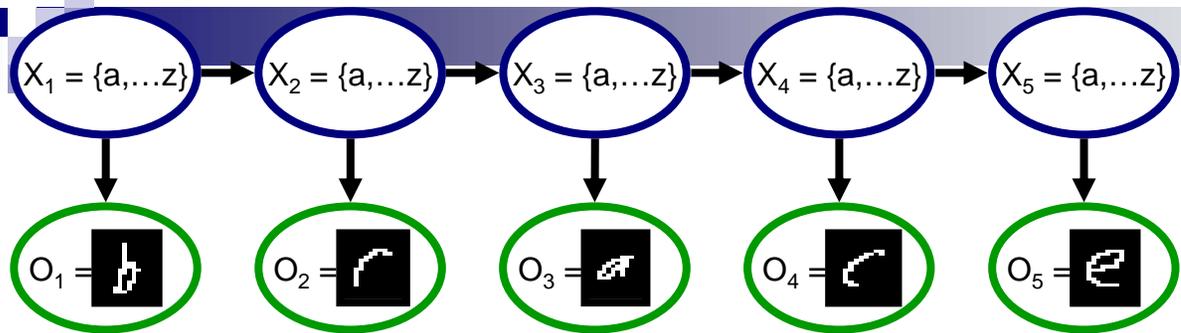
Variable elimination for each i , what's the complexity?

n letters $P(X_i | o) \rightarrow O(n)$

$O(n^2)$

[can solve in $O(n)$]

Reusing computation



Compute:

$$P(X_i | o_{1..n})$$

$$P(X_5 | O)$$

elimination
 \rightarrow order 1, 2, 3, 4

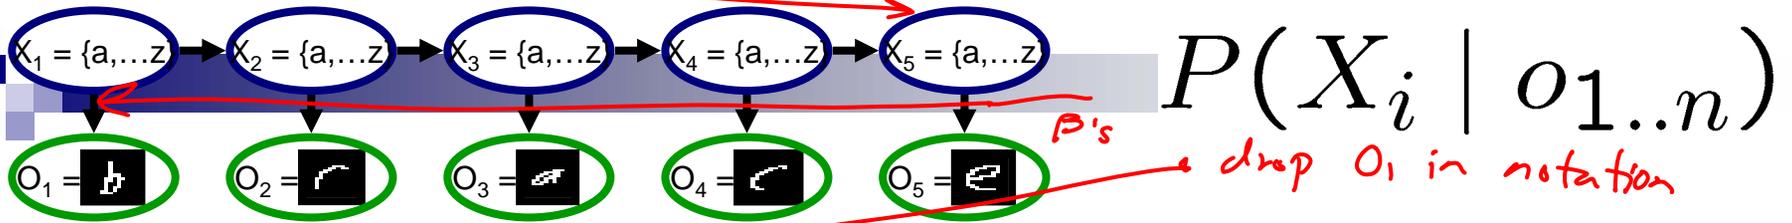
$$g_1(x_2, o_1), g_2(x_3, o_{1:2}), g_3(x_4, o_{1:3}), g_4(x_5, o_{1:4})$$

Same g 's useful for $P(X_4 | O)$, order 1, 2, 3, 5

eliminate 5 \rightarrow $\sum_{x_5} P(O_5 | x_5) \cdot P(x_5 | x_4)$ (computing $P(x_4 | O)$):
 use g_1, g_2, g_3, g_5

$$g_5(x_4, O_5)$$

α's The forwards-backwards algorithm



■ Initialization: $\alpha_1(X_1) = P(X_1)P(o_1 | X_1)$

■ For $i = 2$ to n

□ Generate a forwards factor by eliminating X_{i-1}

sum out previous var prob obs

$$\alpha_i(X_i) = \sum_{x_{i-1}} P(o_i | X_i) P(X_i | X_{i-1} = x_{i-1}) \alpha_{i-1}(x_{i-1})$$

transition prob

■ Initialization: $\beta_n(X_n) = 1$

■ For $i = n-1$ to 1

□ Generate a backwards factor by eliminating X_{i+1}

$\alpha_5(a)$
 $\alpha_5(b)$
 \vdots
 $\alpha_5(z)$

$\alpha_n(X_n)$
 normalized
 $= P(X_n | O_{1:n})$

$\beta_1(X_1) \alpha_1(X_1)$
 normalized
 $= P(X_1 | O_{1:n})$

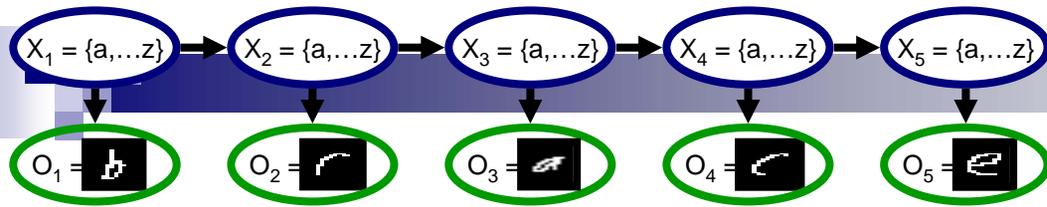
∗ xi

$$\beta_i(X_i) = \sum_{x_{i+1}} P(o_{i+1} | x_{i+1}) P(x_{i+1} | X_i) \beta_{i+1}(x_{i+1})$$

∗ xi

■ $\forall i$, probability is: $P(X_i | O_{1..n}) = \alpha_i(X_i) \beta_i(X_i)$

Most likely explanation



not equal to $\arg\max_{x_i} P(x_i | O_{1:5})$

Compute: $\arg\max_{x_1, \dots, x_5} P(x_1, \dots, x_5 | O_{1:5})$

Variable elimination order?

1, 2, 3, 4, 5

Example:

$$P(x_1, \dots, x_5 | O_{1:5})$$

$$\max_{x_1, x_2, x_3, x_4, x_5}$$

$$= \max_{x_2, \dots, x_5} P(x_3 | x_2) P(O_2 | x_2) P(O_3 | x_3) P(x_4 | x_3) \dots$$

$$\max_{x_1} P(x_1) P(O_1 | x_1) P(x_2 | x_1)$$

$$x_5^* = \arg\max_{x_5} \alpha_5(x_5)$$

; backwards

$$x_4^* = \arg\max_{x_4} \alpha_4(x_4)$$

$$P(O_5 | x_4^*)$$

A, B binary

$P(A, B) =$

	A	t	f
B	t	0.3	0
f	0.3	0.4	

$\arg\max_{a,b} P(a,b) = \begin{cases} a=f \\ b=f \end{cases}$

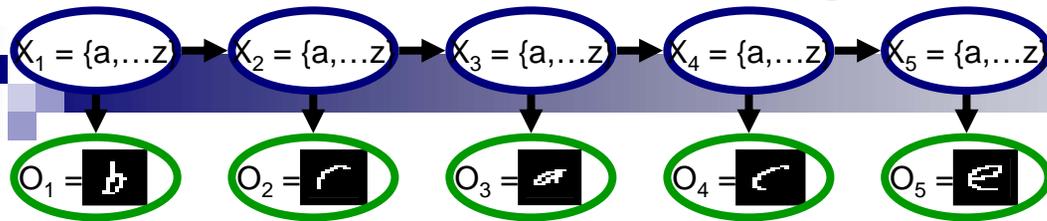
$P(a=t) = 0.3$

$P(a=f) = 0.4$

$\arg\max_a P(a) = (a=f)$

doesn't depend on x_4

The Viterbi algorithm



■ Initialization: $\alpha_1(X_1) = P(X_1)P(o_1 | X_1)$

■ For $i = 2$ to n

□ Generate a forwards factor by eliminating X_{i-1}

$$\alpha_i(X_i) = \max_{x_{i-1}} P(o_i | X_i) P(X_i | X_{i-1} = x_{i-1}) \alpha_{i-1}(x_{i-1})$$

Handwritten annotations:
 - $\alpha_i(X_i)$ is annotated with $\forall x_i$
 - $P(o_i | X_i)$ is annotated with o_i and "obs. prob."
 - $P(X_i | X_{i-1} = x_{i-1})$ is annotated with x_{i-1} and "transition"
 - $\alpha_{i-1}(x_{i-1})$ is annotated with "message"

■ Computing best explanation: $x_n^* = \operatorname{argmax}_{x_n} \alpha_n(x_n)$

■ For $i = n-1$ to 1

□ Use argmax to get explanation:

$$x_i^* = \operatorname{argmax}_{x_i} P(x_{i+1}^* | x_i) \alpha_i(x_i)$$

Handwritten annotations:
 - x_{i+1}^* is annotated with "best for next letter"

What you'll implement 1: multiplication

$$\alpha_i(X_i) = \max_{x_{i-1}} P(o_i | X_i) P(X_i | X_{i-1} = x_{i-1}) \alpha_{i-1}(x_{i-1})$$

$$f_1(X_i) \cdot f_2(X_{i-1}, X_i) \quad (\text{factors})$$

def. new factor g , domain union $\text{dom}(f_1) \cup \text{dom}(f_2)$
 $\forall x_i, x_{i-1}$

$$g(x_{i-1} = x_{i-1}, x_i = x_i) = f_1(x_i = x_i) \cdot f_2(x_{i-1} = x_{i-1}, x_i = x_i)$$

What you'll implement 2: max & argmax

$$\alpha_i(X_i) = \max_{x_{i-1}} P(o_i | X_i) P(X_i | X_{i-1} = x_{i-1}) \alpha_{i-1}(x_{i-1})$$

$$\max_{x_{i-1}} g(x_{i-1}, x_i)$$

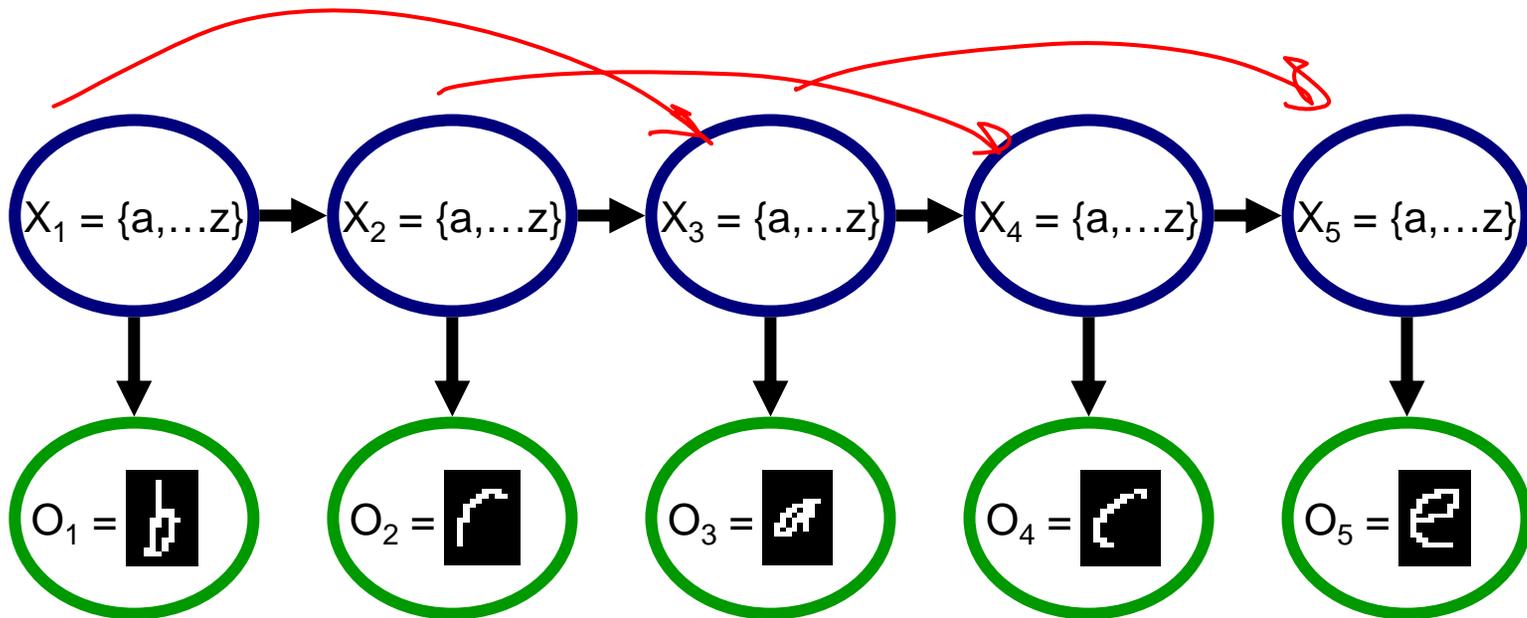
get function h , domain $\text{dom}(g) - \{x_{i-1}\}$

$\forall x_i$

$$h(x_i = x_i) = \max_{x_{i-1}} g(x_{i-1} = x_{i-1}, x_i = x_i)$$

Higher-order HMMs

2nd order \rightarrow depend on last 2 time steps



**Add dependencies further back in time \rightarrow
better representation, harder to learn**

What you need to know

- Hidden Markov models (HMMs)
 - Very useful, very powerful!
 - Speech, OCR,...
 - Parameter sharing, only learn 3 distributions
 - Trick reduces inference from $O(n^2)$ to $O(n)$
 - Special case of BN

Koller & Friedman Chapters (handed out):

Chapter 11 (short)

Chapter 12: 12.1, 12.2, 12.3 (covered in the beginning of semester)
12.4 (Learning parameters for BNs)

Chapter 13: 13.1, 13.3.1, 13.4.1, 13.4.3 (basic structure learning)

Learning BN tutorial (class website):

<ftp://ftp.research.microsoft.com/pub/tr/tr-95-06.pdf>

TAN paper (class website):

<http://www.cs.huji.ac.il/~nir/Abstracts/FrGG1.html>

Bayesian Networks – (Structure) Learning

Machine Learning – 10701/15781

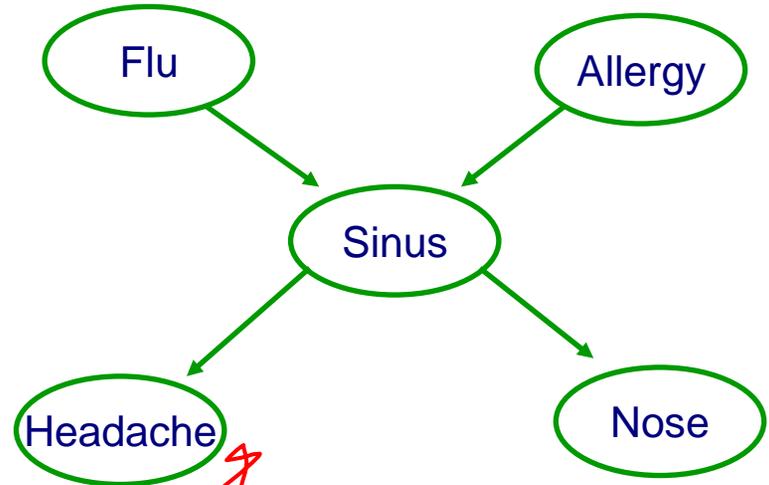
Carlos Guestrin

Carnegie Mellon University

March 29th, 2006

Review

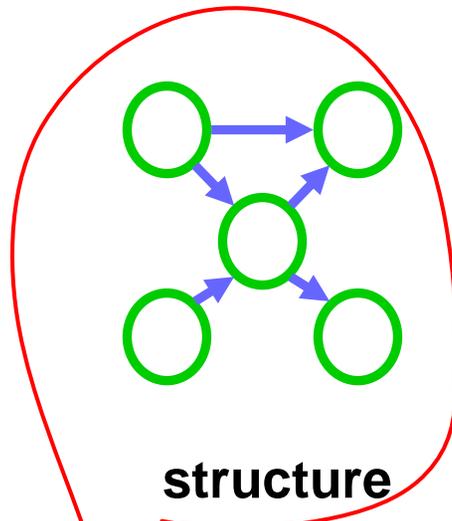
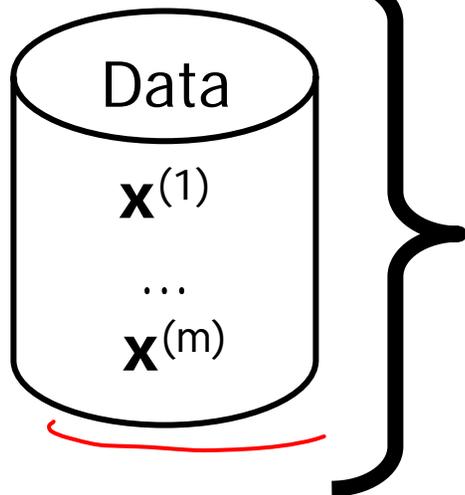
- Bayesian Networks
 - Compact representation for probability distributions
 - Exponential reduction in number of parameters
- Fast probabilistic inference using variable elimination
 - Compute $P(X|e)$
 - Time exponential in tree-width, not number of variables
- Today
 - Learn BN structure



Learning Bayes nets

	Known structure	Unknown structure
<u>Fully observable data</u>	easy !! 😊	NP-hard (not always) we'll see next week
Missing data	hard, in 2 weeks	really hard, next semester

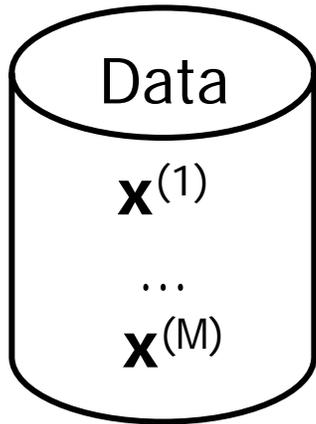
A, B, C
 $\langle A=a, B=?, C=c \rangle$



+

CPTs –
 $P(X_i | \mathbf{Pa}_{X_i})$
 parameters

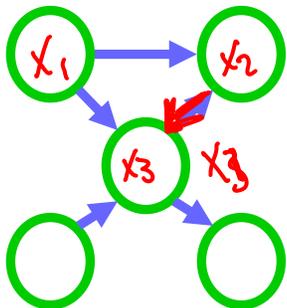
Learning the CPTs



For each discrete variable X_i

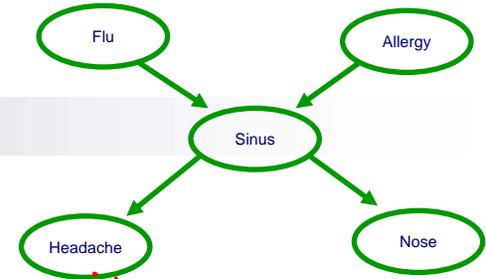
$$P(x_3 | x_1, x_2)$$
$$P(x_3 = x_3 | x_1 = x_1, x_2 = x_2) = \frac{\text{Count}(x_3 = x_3, x_1 = x_1, x_2 = x_2)}{\text{Count}(x_1 = x_1, x_2 = x_2)}$$

WHY????????????



MLE: $P(X_i = x_i | X_j = x_j) = \frac{\text{Count}(X_i = x_i, X_j = x_j)}{\text{Count}(X_j = x_j)}$

Information-theoretic interpretation of maximum likelihood



Given structure, log likelihood of data:

$$\log P(D | \theta_G, G) = \log \prod_j P(x^{(j)} | \theta, G)$$

$x^{(j)} = \langle F=t, S=f, A=f, H=t, N=f \rangle$

$$= \sum_j \log P(x^{(j)} | \theta, G)$$

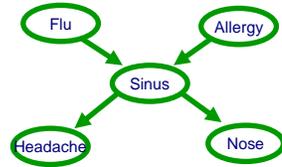
$$= \sum_j \log P(f^{(j)}) P(a^{(j)}) P(s^{(j)} | f^{(j)}, a^{(j)}) P(h^{(j)} | s^{(j)}) P(n^{(j)} | s^{(j)})$$

$$= \underbrace{\left[\sum_j \log P(f^{(j)}) \right]}_{\text{learn } P(F)} + \underbrace{\left[\sum_j \log P(a^{(j)}) \right]}_{\text{learn } P(A)} + \underbrace{\left[\sum_j \log P(s^{(j)} | f^{(j)}, a^{(j)}) \right]}_{\text{learn } P(S|F,A)}$$

log likelihood decomposes with BN structure

Maximum likelihood (ML) for learning BN structure

Possible structures



Learn parameters using ML

Score structure
 $\log P(D|\theta_G, G)$

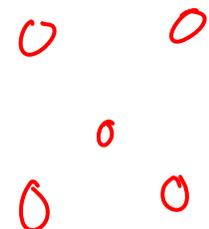
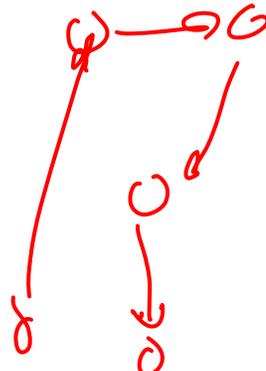
- 10000

wins

- 12000



$\langle X_1^{(1)}, \dots, X_n^{(1)} \rangle$
 ...
 $\langle X_1^{(M)}, \dots, X_n^{(M)} \rangle$

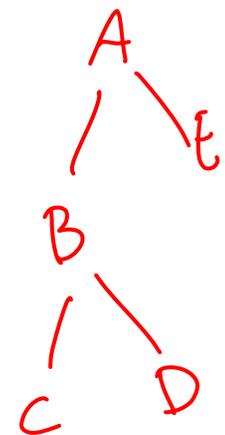
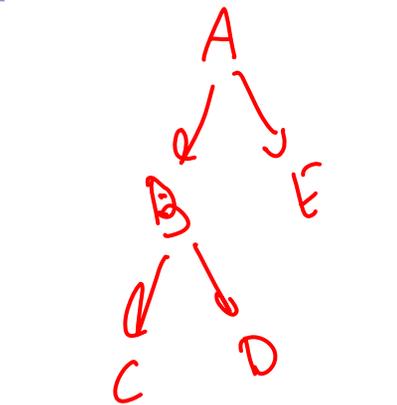


- 50000

How many trees are there?

with n variables?

$$\binom{n}{k} \geq \binom{n}{k}^k$$



how many undirected trees
 n vers $\rightarrow \frac{n}{2}(n-1)$ possible edges
 choosing $n-1$

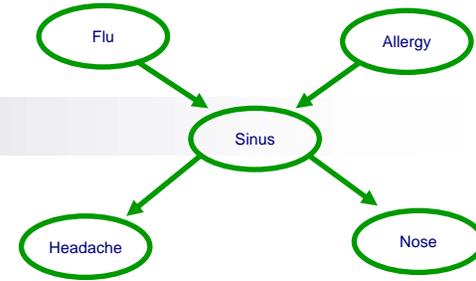
$$\binom{\frac{n}{2}(n-1)}{n-1} \geq \binom{\frac{n}{2}n-1}{n-1} = \binom{n}{2}^{n-1}$$

sorry... I counted A
 B-C
 B-E

go back to
 d tree slide

Nonetheless – Efficient optimal algorithm finds best tree

Information-theoretic interpretation of maximum likelihood 2



- Given structure, log likelihood of data:

$$\log P(D | \theta_G, G) = \sum_{i=1}^n \sum_{j=1}^M \log P(x_i^{(j)} | \text{Pa}_{x_i, G}^{(j)})$$

learning for each cPT

$$\sum_{j=1}^M \log P(x_i^{(j)} | \text{Pa}_{x_i, G}^{(j)}) = M \sum_{x_i, \text{Pa}_{x_i, G}} \frac{\text{count}(x_i, \text{Pa}_{x_i, G})}{\log \hat{P}(x_i | \text{Pa}_{x_i, G})}$$

$$= M \sum_{x_i, \text{Pa}_{x_i, G}} \hat{P}(x_i, \text{Pa}_{x_i, G}) \log \hat{P}(x_i | \text{Pa}_{x_i, G})$$

$$= -M \hat{H}(X_i | \text{Pa}_{x_i, G})$$

$s = t, f = t, d = t$
 ≈ 50 times

$$\frac{\text{count}(a, b)}{M} = \hat{P}(a, b)$$

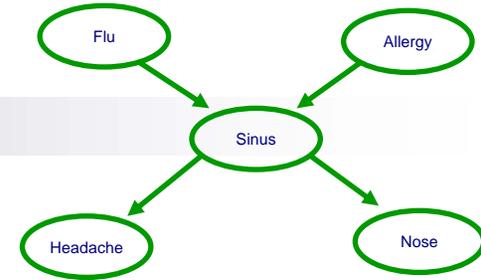
Conditional entropy
 $H(A|B)$

$$= - \sum_{a,b} P(a,b) \log P(a|b)$$

Information-theoretic interpretation of maximum likelihood 3

- Given structure, log likelihood of data:

$$\log P(\mathcal{D} | \theta_{\mathcal{G}}, \mathcal{G}) = M \sum_{i=1}^n \sum_{x_i, \mathbf{Pa}_{x_i, \mathcal{G}}} \hat{P}(x_i, \mathbf{Pa}_{x_i, \mathcal{G}}) \log \hat{P}(x_i | \mathbf{Pa}_{x_i, \mathcal{G}})$$



Mutual information → Independence tests

- Statistically difficult task!
- Intuitive approach: **Mutual information**

$$I(X_i, X_j) = \sum_{x_i, x_j} P(x_i, x_j) \log \frac{P(x_i, x_j)}{P(x_i)P(x_j)}$$

- Mutual information and independence:
 - X_i and X_j independent if and only if $I(X_i, X_j) = 0$
- Conditional mutual information:

Decomposable score

- Log data likelihood

$$\log P(\mathcal{D} \mid \theta_{\mathcal{G}}, \mathcal{G}) = M \sum_i \hat{I}(X_i, \text{Pa}_{X_i, \mathcal{G}}) - M \sum_i \hat{H}(X_i)$$

Scoring a tree 1: equivalent trees

$$\log \hat{P}(\mathcal{D} \mid \theta, \mathcal{G}) = M \sum_i \hat{I}(x_i, \text{Pa}_{x_i, \mathcal{G}}) - M \sum_i \hat{H}(X_i)$$

Scoring a tree 2: similar trees

$$\log \hat{P}(\mathcal{D} \mid \theta, \mathcal{G}) = M \sum_i \hat{I}(x_i, \text{Pa}_{x_i, \mathcal{G}}) - M \sum_i \hat{H}(X_i)$$

Chow-Liu tree learning algorithm 1

- For each pair of variables X_i, X_j

- Compute empirical distribution:

$$\hat{P}(x_i, x_j) = \frac{\text{Count}(x_i, x_j)}{M}$$

- Compute mutual information:

$$\hat{I}(X_i, X_j) = \sum_{x_i, x_j} \hat{P}(x_i, x_j) \log \frac{\hat{P}(x_i, x_j)}{\hat{P}(x_i)\hat{P}(x_j)}$$

- Define a graph

- Nodes X_1, \dots, X_n
- Edge (i, j) gets weight $\hat{I}(X_i, X_j)$

Chow-Liu tree learning algorithm 2

$$\log \hat{P}(\mathcal{D} | \theta, \mathcal{G}) = M \sum_i \hat{I}(x_i, \text{Pa}_{x_i, \mathcal{G}}) - M \sum_i \hat{H}(X_i)$$

- Optimal tree BN
 - Compute maximum weight spanning tree
 - Directions in BN: pick any node as root, breadth-first-search defines directions

Can we extend Chow-Liu 1

■ Tree augmented naïve Bayes (TAN)

[Friedman et al. '97]

□ Naïve Bayes model overcounts, because correlation between features not considered

□ Same as Chow-Liu, but score edges with:

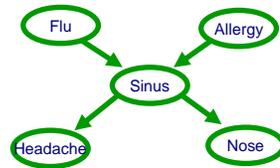
$$\hat{I}(X_i, X_j | C) = \sum_{c, x_i, x_j} \hat{P}(c, x_i, x_j) \log \frac{\hat{P}(x_i, x_j | c)}{\hat{P}(x_i | c) \hat{P}(x_j | c)}$$

Can we extend Chow-Liu 2

- (Approximately learning) models with tree-width up to k
 - [Narasimhan & Bilmes '04]
 - But, $O(n^{k+1})\dots$

Scoring general graphical models – Model selection problem

What's the best structure?



$\langle x_1^{(1)}, \dots, x_n^{(1)} \rangle$

...

$\langle x_1^{(m)}, \dots, x_n^{(m)} \rangle$

The more edges, the fewer independence assumptions, the higher the likelihood of the data, but will overfit...

Bayesian score avoids overfitting

- Given a structure, distribution over parameters

$$\log P(D | \mathcal{G}) = \log \int_{\theta_{\mathcal{G}}} P(D | \mathcal{G}, \theta_{\mathcal{G}}) P(\theta_{\mathcal{G}} | \mathcal{G}) d\theta_{\mathcal{G}}$$

- Difficult integral: use Bayes information criterion (BIC) approximation (equivalent as $M \rightarrow \infty$)

$$\log P(D | \mathcal{G}) \approx \log P(D | \mathcal{G}, \theta_{\mathcal{G}}) - \frac{\text{NumberParams}(\mathcal{G})}{2} \log M + \mathcal{O}(1)$$

- Note: regularize with MDL score
- Best BN under BIC still NP-hard

How many graphs are there?


$$\sum_{k=1}^n \binom{n}{k} = 2^n - 1$$

Structure learning for general graphs

- In a tree, a node only has one parent
- **Theorem:**
 - The problem of learning a BN structure with at most d parents is **NP-hard for any (fixed) $d \geq 2$**
- Most structure learning approaches use heuristics
 - Exploit score decomposition
 - (Quickly) Describe two heuristics that exploit decomposition in different ways

Learn BN structure using local search

Starting from
Chow-Liu tree

Local search,
possible moves:

- Add edge
- Delete edge
- Invert edge

Score using BIC

What you need to know about learning BNs

- Learning BNs
 - Maximum likelihood or MAP learns parameters
 - Decomposable score
 - Best tree (Chow-Liu)
 - Best TAN
 - Other BNs, usually local search with BIC score