## Machine Learning

## Inference and Learning in

GM
Eric Xing

Lecture 18, November 10, 2015


Reading: Chap. 8, C.B book

## Inference and Learning

- We now have compact representations of probability distributions: BN
- A BN M describes a unique probability distribution $P$
- Typical tasks:
- Task 1: How do we answer queries about $P$ ?
- We use inference as a name for the process of computing answers to such queries
- Task 2: How do we estimate a plausible model $M$ from data $D$ ?
i. We use learning as a name for the process of obtaining point estimate of $M$.
ii. But for Bayesian, they seek $p(M \mid D)$, which is actually an inference problem.
iii. When not all variables are observable, even computing point estimate of $M$ need to do inference to impute the missing data.


## Approaches to inference

- Exact inference algorithms
- The elimination algorithm
- Belief propagation
- The junction tree algorithms (but will not cover in detail here)
- Approximate inference techniques
- Variational algorithms
- Stochastic simulation / sampling methods
- Markov chain Monte Carlo methods


## Marginalization and Elimination

- A food web:


What is the probability that hawks are leaving given that the grass condition is poor?
Query: $P(h)$

$$
P(h)=\sum_{g} \sum_{f} \sum_{e} \sum_{d} \sum_{c} \sum_{b} \sum_{a} P(a, b, c, d, e, f, g, h)
$$



> | a naïve summation needs to |
| :--- |
| enumerate over an exponential |
| number of terms |

- By chain decomposition, we get

$$
=\sum_{g} \sum_{f} \sum_{e} \sum_{d} \sum_{c} \sum_{b} \sum_{a} P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) P(h \mid e, f)
$$

## Variable Elimination

- Query: $P(A \mid h)$
- Need to eliminate: $B, C, D, E, F, G, H$
- Initial factors:

$$
P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) P(h \mid e, f)
$$

- Choose an elimination order: $H, G, F, E, D, C, B$

- Step 1:
- Conditioning (fix the evidence node (i.e., $h$ ) on its observed value (i.e., $\widetilde{h}$ ):

$$
m_{h}(e, f)=p(h=\widetilde{h} \mid e, f)
$$

- This step is isomorphic to a marginalization step:

$$
m_{h}(e, f)=\sum_{\substack{h \\ \text { ©Eric Xing @ CMU, 2006-2015 }}} p(h \mid e, f) \delta(h=\widetilde{h})
$$



## Example: Variable Elimination

- Query: $P(B \mid h)$
- Need to eliminate: $B, C, D, E, F, G$
- Initial factors:

$$
\begin{aligned}
& P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) P(h \mid e, f) \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) m_{h}(e, f)
\end{aligned}
$$



- Step 2: Eliminate G

$$
\begin{aligned}
& \text { compute } \quad m_{g}(e)=\sum_{g} p(g \mid e)=1 \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) m_{g}(e) m_{h}(e, f) \\
& \left.=P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) \underline{m_{h}(e, f}\right)
\end{aligned}
$$



## Example: Variable Elimination

- Query: $P(B \mid h)$
- Need to eliminate: $B, C, D, E, F$
- Initial factors:

$$
\begin{aligned}
& P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) P(h \mid e, f) \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) m_{h}(e, f)
\end{aligned}
$$



- Step 3: Eliminate F
- compute

$$
m_{f}(e, a)=\sum_{f} p(f \mid a) m_{h}(e, f)
$$

$\Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) m_{f}(a, e)$


## Example: Variable Elimination

- Query: $P(B \mid h)$
- Need to eliminate: $B, C, D, E$
- Initial factors:

$$
\begin{aligned}
& P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) P(h \mid e, f) \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) m_{f}(a, e)
\end{aligned}
$$



- Step 4: Eliminate E
- compute

$$
m_{e}(a, c, d)=\sum_{e} p(e \mid c, d) m_{f}(a, e)
$$

$$
\Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) m_{e}(a, c, d)
$$



## Example: Variable Elimination

- Query: $P(B \mid h)$
- Need to eliminate: $B, C, D$
- Initial factors:

$$
\begin{aligned}
& P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) P(h \mid e, f) \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) m_{f}(a, e) \\
& \Rightarrow P(a) P(b) P(c \mid b) P(d \mid a) m_{e}(a, c, d)
\end{aligned}
$$



- Step 5: Eliminate D

$$
\begin{aligned}
& \text { compute } \quad m_{d}(a, c)=\sum_{d} p(d \mid a) m_{e}(a, c, d) \\
& \Rightarrow P(a) P(b) P(c \mid d) m_{d}(a, c)
\end{aligned}
$$



## Example: Variable Elimination

- Query: $P(B \mid h)$
- Need to eliminate: $B, C$
- Initial factors:

$$
\begin{aligned}
& P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) P(h \mid e, f) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) P(f \mid a) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) m_{f}(a, e) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) m_{e}(a, c, d) \\
& \Rightarrow P(a) P(b) P(c \mid d) m_{d}(a, c)
\end{aligned}
$$



- Step 6: Eliminate C

$$
\begin{aligned}
& \text { compute } \quad m_{c}(a, b)=\sum_{c} p(c \mid b) m_{d}(a, c) \\
& \Rightarrow P(a) P(b) P(c \mid d) m_{d}(a, c)
\end{aligned}
$$

## Example: Variable Elimination

- Query: $P(B \mid h)$
- Need to eliminate: $B$
- Initial factors:

$$
\begin{aligned}
& P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) P(h \mid e, f) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) P(f \mid a) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) m_{f}(a, e) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) m_{e}(a, c, d) \\
& \Rightarrow P(a) P(b) P(c \mid d) m_{d}(a, c) \\
& \Rightarrow P(a) P(b) m_{c}(a, b)
\end{aligned}
$$

- Step 7: Eliminate $B$


$$
\begin{array}{ll}
\text { - compute } & m_{b}(a)=\sum_{b} p(b) m_{c}(a, b) \\
\Rightarrow P(a) m_{b}(a) &
\end{array}
$$

## Example: Variable Elimination

- Query: $P(B \mid h)$
- Need to eliminate: $B$
- Initial factors:

$$
\begin{aligned}
& P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) P(h \mid e, f) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) P(f \mid a) m_{h}(e, f) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) P(e \mid c, d) m_{f}(a, e) \\
& \Rightarrow P(a) P(b) P(c \mid d) P(d \mid a) m_{e}(a, c, d) \\
& \Rightarrow P(a) P(b) P(c \mid d) m_{d}(a, c) \\
& \Rightarrow P(a) P(b) m_{c}(a, b) \\
& \Rightarrow P(a) m_{b}(a)
\end{aligned}
$$



- Step 8: Wrap-up $p(a, \widetilde{h})=p(a) m_{b}(a), \quad p(\widetilde{h})=\sum_{a} p(a) m_{b}(a)$

$$
\Rightarrow P(a \mid \tilde{h})=\frac{p(a) m_{b}(a)}{\sum_{\text {© Eric Xing @ ©MU, 2006-2015 }} p(a) m_{b}(a)}
$$

## Complexity of variable elimination

- Suppose in one elimination step we compute

$$
\begin{gathered}
m_{x}\left(y_{1}, \ldots, y_{k}\right)=\sum_{x} m_{x}^{\prime}\left(x, y_{1}, \ldots, y_{k}\right) \\
m_{x}^{\prime}\left(x, y_{1}, \ldots, y_{k}\right)=\prod_{i=1}^{k} m_{i}\left(x, \boldsymbol{y}_{c_{i}}\right)
\end{gathered}
$$

This requires

- $k \bullet|\operatorname{Val}(X)| \cdot \prod_{i}\left|\operatorname{Val}\left(\boldsymbol{Y}_{C_{i}}\right)\right|$ multiplications
- For each value of $x, y_{1}, \ldots, y_{k}$, we do $k$ multiplications
- $|\operatorname{Val}(X)| \cdot \prod_{i}\left|\operatorname{Val}\left(\boldsymbol{Y}_{C_{i}}\right)\right|$ additions
- For each value of $y_{1}, \ldots, y_{k}$, we do $|\operatorname{Val}(X)|$ additions

Complexity is exponential in number of variables in the intermediate factor

## Elimination Clique

- Induced dependency during marginalization is captured in elimination cliques
- Summation <-> elimination
- Intermediate term <-> elimination clique

$$
\begin{aligned}
& P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) P(h \mid e, f) \\
\Rightarrow & P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) P(g \mid e) \phi_{h}(e, f) \\
\Rightarrow & P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) P(f \mid a) \phi_{g}(e) \phi_{h}(e, f) \\
\Rightarrow & P(a) P(b) P(c \mid b) P(d \mid a) P(e \mid c, d) \phi_{f}(a, e) \\
\Rightarrow & P(a) P(b) P(c \mid b) P(d \mid a) \phi_{e}(a, c, d) \\
\Rightarrow & P(a) P(b) P(c \mid b) \phi_{d}(a, c) \\
\Rightarrow & P(a) P(b) \phi_{c}(a, b) \\
\Rightarrow & P(a) \phi_{b}(a) \\
\Rightarrow & \phi(a)
\end{aligned}
$$

- Can this lead to an generic inference algorithm?



## From Elimination to Message Passing

- Elimination $\equiv$ message passing on a clique tree

- Messages can be reused


## From Elimination to Message Passing

- Elimination $\equiv$ message passing on a clique tree
- Another query ...

- Messages $m_{f}$ and $m_{h}$ are reused, others need to be recomputed


## From elimination to message passing

- Recall ELIMINATION algorithm:
- Choose an ordering $z$ in which query node $f$ is the final node
- Place all potentials on an active list
- Eliminate node $i$ by removing all potentials containing $i$, take sum/product over $x_{i}$.
- Place the resultant factor back on the list
- For a TREE graph:
- Choose query node $f$ as the root of the tree
- View tree as a directed tree with edges pointing towards from $f$
- Elimination ordering based on depth-first traversal
- Elimination of each node can be considered as message-passing (or Belief Propagation) directly along tree branches, rather than on some transformed graphs
$\rightarrow$ thus, we can use the tree itself as a data-structure to do general inference!!



## Message passing for trees



- Elimination on trees is equivalent to message passing along



## The message passing protocol:

- A two-pass algorithm:



## Belief Propagation (SP-algorithm): Sequential implementation

$\operatorname{Sum}-\operatorname{Product}(\mathcal{T}, E)$
Evidence $(E)$
$f=$ ChooseRoot $(\mathcal{V})$
for $e \in \mathcal{N}(f)$
$\operatorname{Collect}(f, e)$
for $e \in \mathcal{N}(f)$
Distribute $(f, e)$
for $i \in \mathcal{V}$
Computemarginal $(i)$
Evidence( $E$ )
for $i \in E$
$\psi^{E}\left(x_{i}\right)=\psi\left(x_{i}\right) \delta\left(x_{i}, \bar{x}_{i}\right)$
for $i \notin E$
$\psi^{E}\left(x_{i}\right)=\psi\left(x_{i}\right)$
$\operatorname{Collect}(i, j)$
for $k \in \mathcal{N}(j) \backslash i$
$\operatorname{Collect}(j, k)$
SendMessage $(j, i)$
Distribute $(i, j)$
$\operatorname{SendMESSAGE}(i, j)$
for $k \in \mathcal{N}(j) \backslash i$
Distribute $(j, k)$
SendMessage $(j, i)$

$$
m_{j i}\left(x_{i}\right)=\sum_{x_{j}}\left(\psi^{E}\left(x_{j}\right) \psi\left(x_{i}, x_{j}\right) \prod_{k \in \mathcal{N}(j) \backslash i} m_{k j}\left(x_{j}\right)\right)
$$

Computemarginal(i)

$$
p\left(x_{i}\right) \propto \psi^{E}\left(x_{i}\right) \prod_{j \in \mathcal{N}(i)} m_{j i}\left(x_{i}\right)
$$

## Inference on general GM

- Now, what if the GM is not a tree-like graph?
- Can we still directly run message message-passing protocol along its edges?
- For non-trees, we do not have the guarantee that message-passing will be consistent!
- Then what?
- Construct a graph data-structure from $P$ that has a tree structure, and run message-passing on it!
$\rightarrow$ Junction tree algorithm


## Summary: Exact Inference

- The simple Eliminate algorithm captures the key algorithmic Operation underlying probabilistic inference:
--- That of taking a sum over product of potential functions
- The computational complexity of the Eliminate algorithm can be reduced to purely graph-theoretic considerations.
- This graph interpretation will also provide hints about how to design improved inference algorithms
- What can we say about the overall computational complexity of the algorithm? In particular, how can we control the "size" of the summands that appear in the sequence of summation operation.


## Approaches to inference

- Exact inference algorithms
- The elimination algorithm
- Belief propagation
- The junction tree algorithms (but will not cover in detail here)
- Approximate inference techniques
- Variational algorithms
- Stochastic simulation / sampling methods
- Markov chain Monte Carlo methods


## Monte Carlo methods

- Draw random samples from the desired distribution
- Yield a stochastic representation of a complex distribution
- marginals and other expections can be approximated using sample-based averages

$$
E[f(x)]=\frac{1}{N} \sum_{t=1}^{N} f\left(x^{(t)}\right)
$$

- Asymptotically exact and easy to apply to arbitrary models
- Challenges:
- how to draw samples from a given dist. (not all distributions can be trivially sampled)?
- how to make better use of the samples (not all sample are useful, or eqally useful, see an example later)?
- how to know we've sampled enough?


## Example: naive sampling

- Construct samples according to probabilities given in a BN.


Alarm example: (Choose the right sampling sequence)

1) Sampling: $P(B)=<0.001,0.999>$ suppose it is false, B0. Same for E0. P(A|B0, E0) $=<0.001,0.999>$ suppose it is false...
2) Frequency counting: In the samples right,

| $E 0$ | $B 0$ | $A 0$ | M0 | J0 |
| :---: | :---: | :---: | :---: | :---: |
| $E 0$ | $B 0$ | $A 0$ | M0 | $J 0$ |
| $E 0$ | $B 0$ | $A 0$ | $M 0$ | $J 1$ |
| $E 0$ | $B 0$ | $A 0$ | $M 0$ | $J 0$ |
| $E 0$ | $B 0$ | $A 0$ | $M 0$ | $J 0$ |
| $E 0$ | $B 0$ | $A 0$ | $M 0$ | $J 0$ |
| $E 1$ | $B 0$ | $A 1$ | $M 1$ | $J 1$ |
| $E 0$ | $B 0$ | $A 0$ | $M 0$ | $J 0$ |
| $E 0$ | $B 0$ | $A 0$ | $M 0$ | $J 0$ |
| $E 0$ | $B 0$ | $A 0$ | $M 0$ | $J 0$ | $P(J \mid A 0)=P(J, A 0) / P(A 0)=<1 / 9,8 / 9>$.

## Example: naive sampling

- Construct samples according to probabilities given in a BN.

Alarm example: (Choose the right sampling sequence)
3) what if we want to compute $\mathrm{P}(\mathrm{J} \mid \mathrm{A} 1)$ ? we have only one sample ... $P(J \mid A 1)=P(J, A 1) / P(A 1)=<0,1>$.
4) what if we want to compute $\mathrm{P}(\mathrm{J} \mid \mathrm{B} 1)$ ?

No such sample available!
$P(J \mid A 1)=P(J, B 1) / P(B 1)$ can not be defined.
For a model with hundreds or more variables, rare events will be very hard to garner evough samples even after a long time or sampling ...

| E0 | B0 | A0 | M0 | J0 |
| :---: | :---: | :---: | :---: | :---: |
| E0 | B0 | A0 | M0 | J0 |
| E0 | B0 | A0 | M0 | J1 |
| E0 | B0 | A0 | M0 | J0 |
| E0 | B0 | A0 | M0 | J0 |
| E0 | B0 | A0 | M0 | J0 |
| E1 | B0 | A1 | M1 | J1 |
| E0 | B0 | A0 | M0 | J0 |
| E0 | B0 | A0 | M0 | J0 |
| E0 | B0 | A0 | M0 | J0 |

## Markov chain Monte Carlo (MCMC)

- Construct a Markov chain whose stationary distribution is the target density $=P(X \mid e)$.
- Run for $T$ samples (burn-in time) until the chain converges/mixes/reaches stationary distribution.
- Then collect $M$ (correlated) samples $x_{m}$.
- Key issues:
- Designing proposals so that the chain mixes rapidly.
- Diagnosing convergence.


## Markov Chains

- Definition:
- Given an n-dimensional state space
- Random vector $\mathbf{X}=\left(x_{1}, \ldots, x_{n}\right)$
- $\mathbf{x}^{(t)}=\mathbf{x}$ at time-step t
- $\mathbf{x}^{(t)}$ transitions to $\mathbf{x}^{(t+1)}$ with prob

$$
\mathrm{P}\left(\mathbf{x}^{(t+1)} \mid \mathbf{x}^{(t)}, \ldots, \mathbf{x}^{(1)}\right)=\mathrm{T}\left(\mathbf{x}^{(t+1)} \mid \mathbf{x}^{(t)}\right)=\mathrm{T}\left(\mathbf{x}^{(t)} \rightarrow \mathbf{x}^{(t+1)}\right)
$$

- Homogenous: chain determined by state $\mathbf{x}^{(0)}$, fixed transition kernel Q (rows sum to 1)
- Equilibrium: $\pi(\mathbf{x})$ is a stationary (equilibrium) distribution if

$$
\pi\left(\mathbf{x}^{\prime}\right)=\Sigma_{\mathbf{x}} \pi(\mathbf{x}) \mathrm{Q}\left(\mathbf{x} \rightarrow \mathbf{x}^{\prime}\right)
$$

i.e., is a left eigenvector of the transition matrix $\pi^{T}=\pi^{T} \mathbf{Q}$.
$\left(\begin{array}{lll}0.2 & 0.5 & 0.3\end{array}\right)=\left(\begin{array}{lll}0.2 & 0.5 & 0.3\end{array}\right)\left(\begin{array}{ccc}0.25 & 0 & 0.75 \\ 0 & 0.7 & 0.3 \\ 0.5 & 0.5 & 0\end{array}\right)$
© Eric Xing @ CMU, 2006-2015


## Gibbs sampling

- The transition matrix updates each node one at a time using the following proposal:

$$
Q\left(\left(x_{i}, \mathbf{x}_{-i}\right) \rightarrow\left(x_{i}^{\prime}, \mathbf{x}_{-i}\right)\right)=p\left(x_{i}^{\prime} \mid \mathbf{x}_{-i}\right)
$$

- It is efficient since $p\left(x_{i}^{\prime} \mid \mathbf{x}_{-i}\right)$ only depends on the values in $X_{i}^{\prime}$ 's Markov blanket


## Gibbs sampling

- Gibbs sampling is an MCMC algorithm that is especially appropriate for inference in graphical models.
- The procedue
- we have variable set $X=\left\{x_{1}, x_{2}, x_{3}, \ldots x_{N}\right\}$ for a GM

- at each step one of the variables $X_{i}$ is selected (at random or according to some fixed sequences), denote the remaining variables as $\boldsymbol{X}_{-i}$, and its current value as $X_{-i}{ }^{(t-1)}$
- Using the "alarm network" as an example, say at time $t$ we choose $X_{E}$, and we denote the current value assignments of the remaining variables, $\boldsymbol{X}_{-E}$, obtained from previous samples, as $x_{-E}^{(t-1)}=\left\{x_{B}^{(t-1)}, x_{A}^{(t-1)}, x_{J}^{(t-1)}, x_{M}^{(t-1)}\right\}$
- the conditonal distribution $p\left(X_{i} \mid x_{-i}{ }^{(t-1)}\right)$ is computed
- a value $\boldsymbol{x}_{i}^{(t)}$ is sampled from this distribution
- the sample $\boldsymbol{X}_{i}^{(t)}$ replaces the previous sampled value of $X_{i}$ in $X$.
- i.e., $x^{(t)}=x_{-E}^{(t-1)} \cup x_{E}^{(t)}$


## Markov Blanket

## - Markov Blanket in BN

- A variable is independent from others, given its parents, children and children's parents (dseparation).
- MB in MRF
- A variable is independent all its non-neighbors, given all its direct
 neighbors.
$\Rightarrow p\left(X_{i} \mid X_{-i}\right)=p\left(X_{i} \mid \operatorname{MB}\left(X_{i}\right)\right)$
- Gibbs sampling
- Every step, choose one variable and sample it by $P(X \mid M B(X))$ based on previous sample.


## Gibbs sampling of the alarm network

- To calculate P(J|B1,M1)

- Choose (B1,E0,A1,M1,J1) as a start
- Evidences are B1, M1, variables are A, E, J.
- Choose next variable as $\mathbf{A}$
- Sample A by $P(A \mid M B(A))=P(A \mid B 1, E 0, M 1, J 1)$ suppose to be false.
- (B1, E0, A0, M1, J1)
- Choose next random variable as E, sample E~P(E|B1,A0)
$M B(A)=\{B, E, J, M\}$
$M B(E)=\{A, B\}$


## Example

First 100 iterations of sample3


## Example:


© Eric Xing @ CMU, 2006-2015

## Example


© Eric Xing @ CMU, 2006-2015

## Example

$P(J 1 \mid B 1, M 1)=0.90$
$P(J 1 \mid E 1, M 0)=0.14$
$P(E 1 \mid J 1)=0.01$
$P(E 1 \mid M 1)=0.04$
$P(E 1 \mid M 1, J 1)=0.17$

Gibbs sampling of alarm network


## The

## of simulation

- Run several chains
- Start at over-dispersed points
- Monitor the log lik.
- Monitor the serial correlations
- Monitor acceptance ratios
- Re-parameterize (to get approx. indep.)
- Re-block (Gibbs)
- Collapse (int. over other pars.)
- Run with troubled pars. fixed at reasonable vals.


## Learning Graphical Models

## The goal:

Given set of independent samples (assignments of random variables), find the best (the most likely?) Bayesian Network (both DAG and CPDs)


## Learning Graphical Models (cont.)

- Scenarios:
- completely observed GMs
- directed
- undirected
- partially observed GMs
- directed
- undirected (an open research topic)
- Estimation principles:
- Maximal likelihood estimation (MLE)
- Bayesian estimation
- Maximal conditional likelihood
- Maximal "Margin"
- We use learning as a name for the process of estimating the parameters, and in some cases, the topology of the network, from data.


## MLE for general BN parameters

- If we assume the parameters for each CPD are globally independent, and all nodes are fully observed, then the loglikelihood function decomposes into a sum of local terms, one per node:

$$
\ell(\theta ; D)=\log p(D \mid \theta)=\log \prod_{x_{B}}\left(\prod_{i} p\left(x_{n, i} \mid \mathbf{x}_{n, \pi_{i}}, \theta_{i}\right)\right)=\sum_{i}\left(\sum_{n} \log p\left(x_{n, i} \mid \mathbf{x}_{n, \pi_{i}}, \theta_{i}\right)\right)
$$



## Example: decomposable likelihood of a directed model

- Consider the distribution defined by the directed acyclic GM:

$$
p(x \mid \theta)=p\left(x_{1} \mid \theta_{1}\right) p\left(x_{2} \mid x_{1}, \theta_{2}\right) p\left(x_{3} \mid x_{1}, \theta_{3}\right) p\left(x_{4} \mid x_{2}, x_{3}, \theta_{4}\right)
$$

- This is exactly like learning four separate small BNs, each of which consists of a node and its parents.



## E.g.: MLE for BNs with tabular CPDs

- Assume each CPD is represented as a table (multinomial) where

$$
\theta_{i j k} \stackrel{\text { def }}{=} p\left(X_{i}=j \mid X_{\pi_{i}}=k\right)
$$

- Note that in case of multiple parents, $\mathbf{X}_{\pi_{i}}$ will have a composite state, and the CPD will be a high-dimensional table
- The sufficient statistics are counts of family configurations


$$
n_{i j k} \stackrel{\operatorname{def}}{=} \sum_{n} x_{n, i}^{j} x_{n, \pi_{i}}^{k}
$$

- The log-likelihood is $\ell(\theta ; D)=\log \prod_{i, j, k} \theta_{i j k}^{n_{j k}}=\sum_{i, j, k} n_{i j k} \log \theta_{i j k}$
- Using a Lagrange multiplier to enforce $\sum_{j} \theta_{j j k}=1$, we get: $\quad \theta_{i j k}^{M L}=\frac{n_{i j k}}{\sum_{i, j^{\prime}, k} n_{i j \prime}}$


## Definition of HMM

- Transition probabilities between any two states

$$
p\left(y_{t}^{j}=1 \mid y_{t-1}^{i}=1\right)=a_{i, j},
$$


or

$$
p\left(y_{t} \mid y_{t-1}^{i}=1\right) \sim \operatorname{Multinomial}\left(a_{i, 1}, a_{i, 2}, \ldots, a_{i, M}\right), \forall i \in \mathrm{I} .
$$

- Start probabilities

$$
p\left(y_{1}\right) \sim \operatorname{Multinomial}\left(\pi_{1}, \pi_{2}, \ldots, \pi_{M}\right) .
$$

- Emission probabilities associated with each state

$$
p\left(x_{t} \mid y_{t}^{i}=1\right) \sim \operatorname{Multinomial}\left(b_{i, 1}, b_{i, 2}, \ldots, b_{i, K}\right), \forall i \in \mathrm{I} .
$$

or in general: $\quad p\left(x_{t} \mid y_{t}^{i}=1\right) \sim \mathrm{f}\left(\cdot \mid \theta_{i}\right), \forall i \in \mathbb{I}$.

## Supervised ML estimation

- Given $x=x_{1} \ldots x_{N}$ for which the true state path $y=y_{1} \ldots y_{N}$ is known,

$$
\begin{aligned}
& \left.\begin{array}{c}
\begin{array}{c}
\ell(\boldsymbol{\theta} ; \mathbf{x}, \mathbf{y}) \\
\text { Define: }
\end{array}
\end{array}\right) \log p(\mathbf{x}, \mathbf{y})=\log \prod_{n}\left(p\left(y_{n, 1}\right) \prod_{t=2}^{T} p\left(y_{n, t} \mid y_{n, t-1} \prod_{t=1}^{T} p\left(x_{n, t} \mid x_{n, t}\right)\right)\right. \\
& \text { - Define: }
\end{aligned}
$$

$$
\begin{array}{ll}
A_{i j} & =\# \text { times state transition } i \rightarrow j \text { occurs in } \mathbf{y} \\
B_{i k} & =\# \text { times state } i \text { in } \mathbf{y} \text { emits } k \text { in } \mathbf{x}
\end{array}
$$

- We can show that the maximum likelihood parameters $\theta$ are:

$$
\begin{aligned}
& a_{i j}^{M L}=\frac{\#(i \rightarrow j)}{\#(i \rightarrow \bullet)}=\frac{\sum_{n} \sum_{t=2}^{T} y_{n, t-1}^{i} y_{n, t}^{j}}{\sum_{n} \sum_{t=2}^{\top} y_{n, t-1}^{i}}=\frac{A_{i j}}{\sum_{j^{\prime}} A_{j^{\prime}}} \\
& b_{i^{\prime}}^{M L}=\frac{\#(i \rightarrow k)}{\#(i \rightarrow \bullet)}=\frac{\sum_{n} \sum_{t=1}^{\top} y_{n, t}^{i} x_{n, t}^{k}}{\sum_{n} \sum_{t=1}^{\top} y_{n, t}^{i}}=\frac{B_{i k}}{\sum_{k^{\prime}} B_{i k^{\prime}}}
\end{aligned}
$$

- If $\mathbf{y}$ is continuous, we can treat $\left\{\left(x_{n, t}, y_{n, t}\right): t=1: T, n=1: N\right\}$ as $N \times T$ observations of, e.g., a Gaussian, and apply learning rules for Gaussian ...


## What if some nodes are not observed?

- Consider the distribution defined by the directed acyclic GM:

$$
p(x \mid \theta)=p\left(x_{1} \mid \theta_{1}\right) p\left(x_{2} \mid x_{1}, \theta_{1}\right) p\left(x_{3} \mid x_{1}, \theta_{3}\right) p\left(x_{4} \mid x_{2}, x_{3}, \theta_{1}\right)
$$



- Need to compute $\mathrm{p}\left(\mathrm{x}_{\mathrm{H}} \mid \mathrm{x}_{\mathrm{V}}\right) \rightarrow$ inference


## MLE for BNs with tabular CPDs

- Assume each CPD is represented as a table (multinomial) where

$$
\theta_{i j k} \stackrel{\text { def }}{=} p\left(X_{i}=j \mid X_{\pi_{i}}=k\right)
$$

- Note that in case of multiple parents, $\mathbf{X}_{\pi_{i}}$ will have a composite state, and the CPD will be a high-dimensional table
- The sufficient statistics are counts of family configurations


$$
n_{i j k} \stackrel{\text { def }}{=} \sum_{n}\left\langle x_{n, i}^{j} x_{n, \pi_{i}}^{k}\right\rangle
$$

- The log-likelihood is $\ell(\theta ; D)=\log \prod_{i, j, k} \theta_{i j k}^{n_{j k}}=\sum_{i, j, k} n_{i j k} \log \theta_{i j k}$
- Using a Lagrange multiplier to enforce $\sum_{j} \theta_{i j k}=1$, we get:

$$
\theta_{i j k}^{M L}=\frac{n_{i j k}}{\sum_{i, j^{\prime}, k} n_{i j^{\prime} k}}
$$

## Summary

- A GMM describes a unique probability distribution $P$
- Typical tasks:
- Task 1: How do we answer queries about $P$ ?
- We use inference as a name for the process of computing answers to such queries
- Task 2: How do we estimate a plausible model $M$ from data $D$ ?
i. We use learning as a name for the process of obtaining point estimate of $M$.
ii. But for Bayesian, they seek $p(M \mid D)$, which is actually an inference problem.
iii. When not all variables are observable, even computing point estimate of $M$ need to do inference to impute the missing data.

