

CS 3750 Machine Learning Lecture 4

Markov Random Fields II

Milos Hauskrecht
milos@cs.pitt.edu
5329 Sennott Square

CS 3750 Advanced Machine Learning

Markov random fields

- Probabilistic models with symmetric dependences.

- Typically models spatially varying quantities

$$P(x) \propto \prod_{c \in cl(x)} \phi_c(x_c)$$

$\phi_c(x_c)$ - A potential function (defined over factors)

- If $\phi_c(x_c)$ is strictly positive we can rewrite the definition in terms of a log-linear model :

$$P(x) = \frac{1}{Z} \exp \left(- \sum_{c \in cl(x)} E_c(x_c) \right) \quad \text{- Energy function}$$

- Gibbs (Boltzman) distribution

$$Z = \sum_{x \in \{x\}} \exp \left(- \sum_{c \in cl(x)} E_c(x_c) \right) \quad \text{- A partition function}$$

CS 3750 Advanced Machine Learning

Graphical representation of MRFs

An undirected network (also called independence graph)

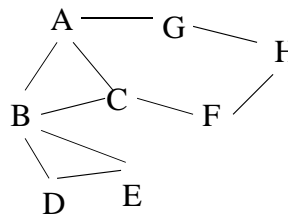
- $G = (S, E)$
 - $S = 1, 2, \dots, N$ correspond to random variables
 - $(i, j) \in E \Leftrightarrow \exists c : \{i, j\} \subset c$
or x_i and x_j appear within the same factor c

Example:

- variables A, B, \dots, H
- Assume the full joint of MRF

$$P(A, B, \dots, H) \sim$$

$$\phi_1(A, B, C) \phi_2(B, D, E) \phi_3(A, G) \\ \phi_4(C, F) \phi_5(G, H) \phi_6(F, H)$$

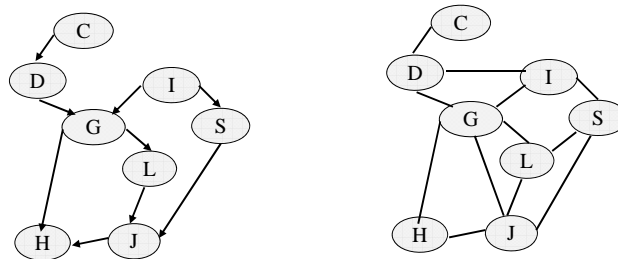


CS 3750 Advanced Machine Learning

Converting BBNs to MRFs

Moral-graph $H[G]$: of a Bayesian network over X is an undirected graph over X that contains an edge between x and y if:

- There exists a directed edge between them in G .
- They are both parents of the same node in G .

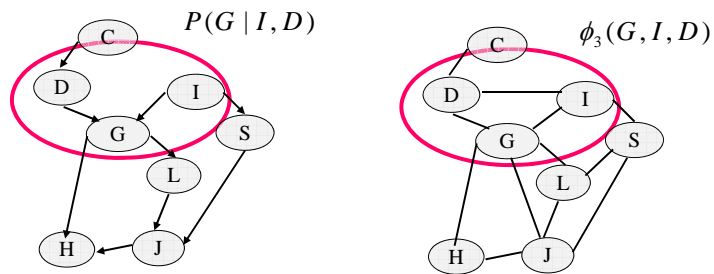


CS 3750 Advanced Machine Learning

Moral Graphs

Why moralization?

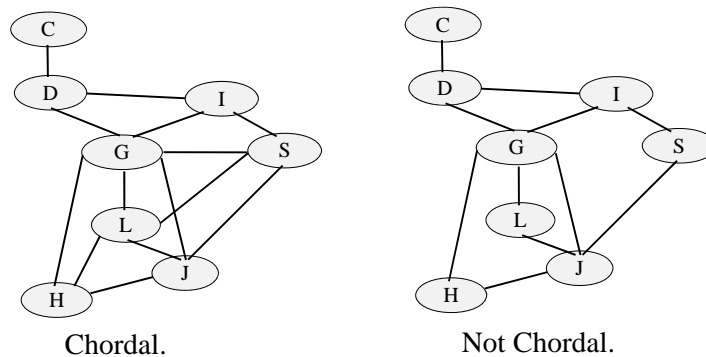
$$\begin{aligned}
 P(C, D, G, I, S, L, J, H) &= \\
 &= P(C)P(D|C)P(G|I, D)P(S|I)P(L|G)P(J|L, S)P(H|G, J) \\
 &= \phi_1(C)\phi_2(D, C)\phi_3(G, I, D)\phi_4(S, I)\phi_5(L, G)\phi_6(J, L, S)\phi_7(H, G, J)
 \end{aligned}$$



CS 3750 Advanced Machine Learning

Chordal graphs

Chordal Graph: an undirected graph G whose minimum cycle contains 3 vertices.

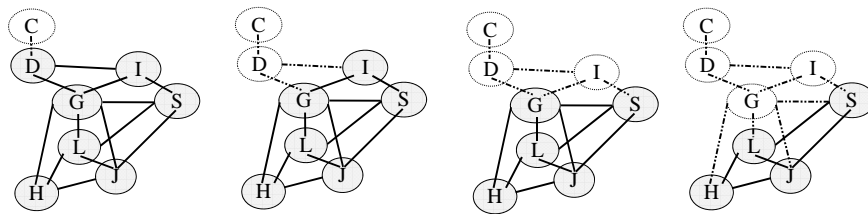


CS 3750 Advanced Machine Learning

Chordal Graphs

Properties:

- There exists an elimination ordering that adds no edges.
- The minimal induced treewidth of the graph is equal to the size of the largest clique - 1.



CS 3750 Advanced Machine Learning

Triangulation

The process of converting a graph G into a *chordal graph* is called Triangulation.

A new graph obtained via triangulation is:

- 1) Guaranteed to be chordal.
- 2) Not guaranteed to be (treewidth) optimal.

There exist exact algorithms for finding the *minimal chordal graphs*, and heuristic methods with a guaranteed upper bound.

CS 3750 Advanced Machine Learning

Chordal Graphs

- Given a minimum triangulation for a graph G , we can carry out the variable-elimination algorithm in the minimum possible time.
- **Complexity** of the optimal triangulation:
 - Finding the minimal triangulation is **NP-Hard**.
- **The inference limit:**
 - Inference time is exponential in terms of the largest clique (factor) in G .

CS 3750 Advanced Machine Learning

Inference: conclusions

- We cannot escape **exponential costs in the treewidth**.
- But in many graphs the treewidth is much smaller than the total number of variables
- Still a problem: Finding the optimal decomposition is hard
 - But, paying the cost up front may be worth it.
 - Triangulate once, query many times.
 - Real cost savings if not a bounded one.

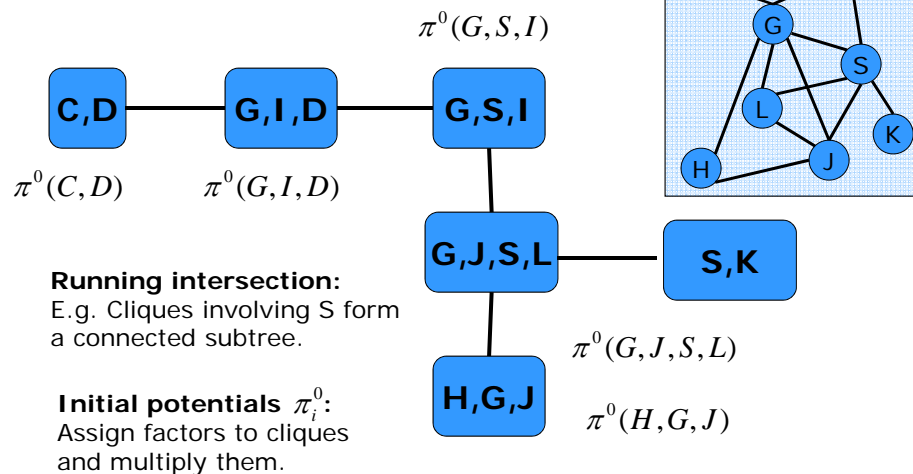
CS 3750 Advanced Machine Learning

Clique tree properties

- **A clique tree:**
 - a tree where nodes correspond to sets of variables
 - used for performing probabilistic inferences
- **Sepset** $S_{ij} = C_i \cap C_j$
 - **separation set:** Variables **X** on one side of sepset are separated from the variables **Y** on the other side in the factor graph given variables in **S**
- **Running intersection property**
 - if C_i and C_j both contain **X**, then all cliques on the unique path between them also contain **X**

CS 3750 Advanced Machine Learning

Inference in clique trees

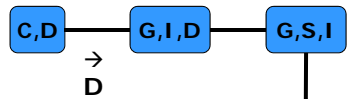


CS 3750 Advanced Machine Learning

Message Passing VE

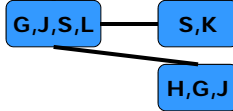
- Query for $P(J)$

– Eliminate C: $\tau_1(D) = \sum_C \pi_1^0[C, D]$

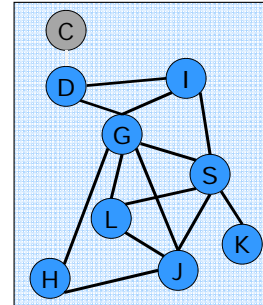


Message sent
from $[C, D]$
to $[G, I, D]$

Message received
at $[G, I, D]$ --
 $[G, I, D]$ updates:



$$\pi_2[G, I, D] = \tau_1(D) \times \pi_2^0[G, I, D]$$

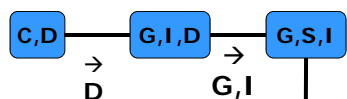


CS 3750 Advanced Machine Learning

Message Passing VE

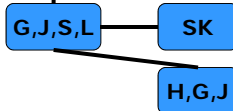
- Query for $P(J)$

– Eliminate D: $\tau_2(G, I) = \sum_D \pi_2[G, I, D]$

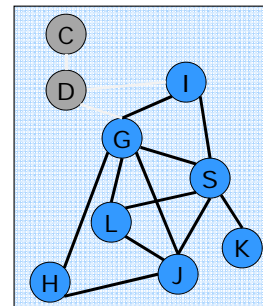


Message sent
from $[G, I, D]$
to $[G, S, I]$

Message received
at $[G, S, I]$ --
 $[G, S, I]$ updates:



$$\pi_3[G, S, I] = \tau_2(G, I) \times \pi_3^0[G, S, I]$$

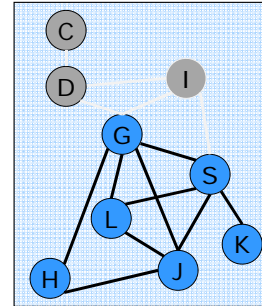
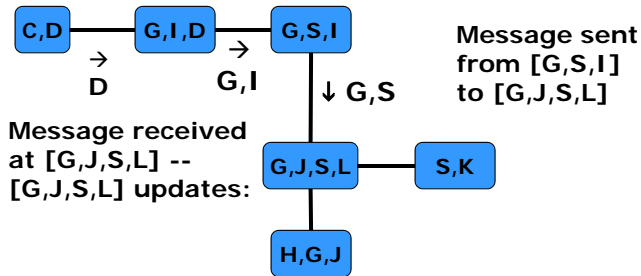


CS 3750 Advanced Machine Learning

Message Passing VE

- Query for $P(J)$

– Eliminate I: $\tau_3(G, S) = \sum_I \pi_3[G, S, I]$



$$\pi_4[G, J, S, L] = \tau_3(G, S) \times \pi_4^0[G, J, S, L] \quad !$$

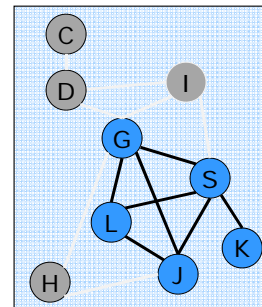
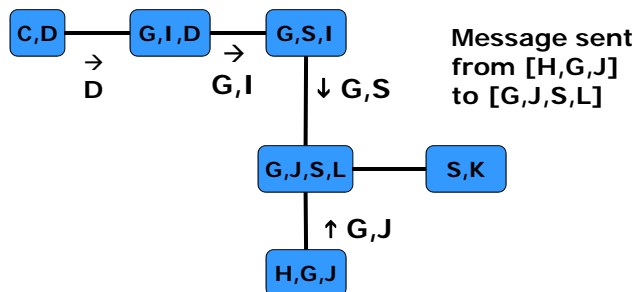
$[G, J, S, L]$ is not **ready**!

CS 3750 Advanced Machine Learning

Message Passing VE

- Query for $P(J)$

– Eliminate H: $\tau_4(G, J) = \sum_H \pi_5[H, G, J]$



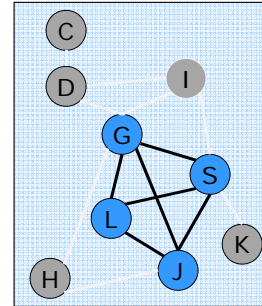
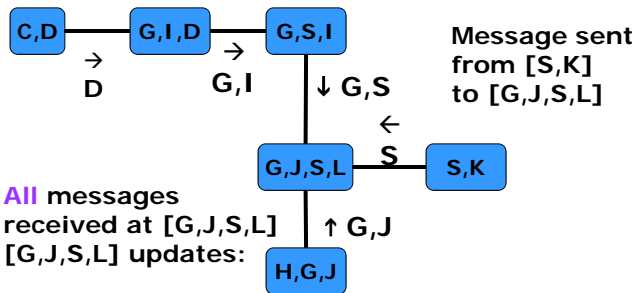
$$\pi_4[G, J, S, L] = \tau_3(G, S) \times \tau_4(G, J) \times \pi_4^0[G, J, S, L]$$

And ...

CS 3750 Advanced Machine Learning

Message Passing VE

- Query for $P(J)$
 - Eliminate K : $\tau_6(S) = \sum_K \pi^0[S, K]$



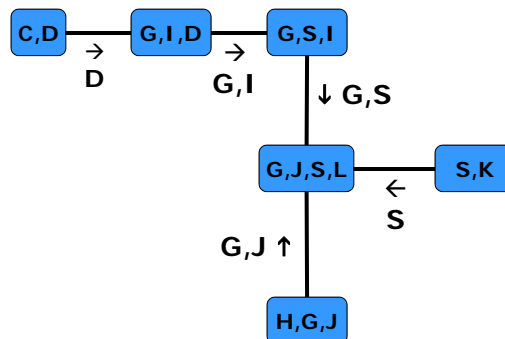
$$\pi_4[G, J, S, L] = \tau_3(G, S) \times \tau_4(G, J) \times \tau_6(S) \times \pi_4^0[G, J, S, L]$$

And calculate $P(J)$ from it by summing out G, S, L

CS 3750 Advanced Machine Learning

Message Passing VE

- $[G, J, S, L]$ clique potential
- ... is used to finish the inference



CS 3750 Advanced Machine Learning

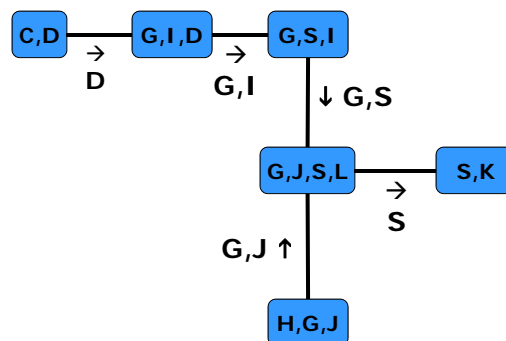
Message passing VE

- Often, **many marginals are desired**
 - Inefficient to re-run each inference from scratch
 - One distinct message per edge & direction
- **Methods :**
 - Compute (**unnormalized**) marginals for any vertex (clique) of the tree
 - Results in a **calibrated clique tree** $\sum_{C_i - S_{ij}} \pi_i = \sum_{C_j - S_{ij}} \pi_j$
- Recap: three kinds of factor objects
 - Initial potentials, final potentials and messages

CS 3750 Advanced Machine Learning

Two-pass message passing VE

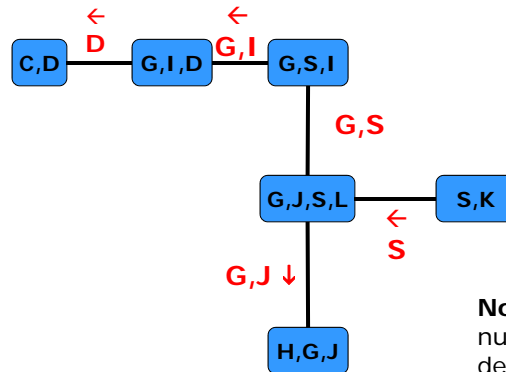
- Chose the root clique, e.g. [S,K]
- Propagate messages to the root



CS 3750 Advanced Machine Learning

Two-pass message passing VE

- Send messages back from the root



Notation:
number the cliques and
denote the messages

$$\delta_{i \rightarrow j}$$

CS 3750 Advanced Machine Learning

Message Passing: BP

- Graphical model of a **distribution**
 - More edges = larger expressive power
 - Clique tree also a model of distribution
 - Message passing preserves the model but changes the parameterization
- Different but equivalent algorithm

CS 3750 Advanced Machine Learning

Factor division

A=1	B=1	0.5
A=1	B=2	0.4
A=2	B=1	0.8
A=2	B=2	0.2
A=3	B=1	0.6
A=3	B=2	0.5

A=1	0.4
A=2	0.4
A=3	0.5

A=1	B=1	0.5/0.4=1.25
A=1	B=2	0.4/0.4=1.0
A=2	B=1	0.8/0.4=2.0
A=2	B=2	0.2/0.4=2.0
A=3	B=1	0.6/0.5=1.2
A=3	B=2	0.5/0.5=1.0

Inverse of factor product

CS 3750 Advanced Machine Learning

Message Passing: BP

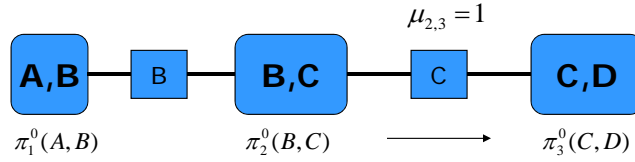
- Each node: multiply all the messages and divide by the one that is coming from node we are sending **the message to**
 - Clearly the same as VE

$$\delta_{i \rightarrow j} = \frac{\sum_{C_i - S_{ij}} \pi_i}{\delta_{j \rightarrow i}} = \frac{\sum_{C_i - S_{ij}} \prod_{k \in N(i)} \delta_{k \rightarrow i}}{\delta_{j \rightarrow i}} = \sum_{C_i - S_{ij}} \prod_{k \in N(i) \setminus j} \delta_{k \rightarrow i}$$

- Initialize the messages on the edges to 1

CS 3750 Advanced Machine Learning

Message Passing: BP



Store the last message on the edge and divide each passing message by the last stored.

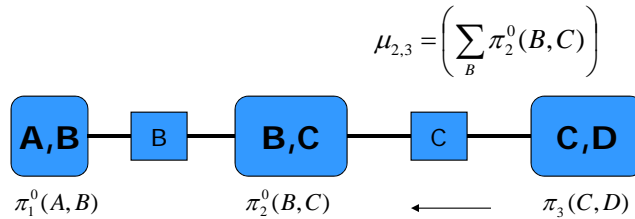
$$\delta_{2 \rightarrow 3} = \left(\sum_B \pi_2^0(B, C) \right)$$

$$\pi_3(C, D) = \pi_3^0(C, D) \frac{\delta_{2 \rightarrow 3}}{\mu_{2,3}} = \pi_3^0(C, D) \sum_B \pi_2^0(B, C)$$

$$\mu_{2,3} = \delta_{2 \rightarrow 3} = \left(\sum_B \pi_2^0(B, C) \right) \quad \text{New message}$$

CS 3750 Advanced Machine Learning

Message Passing: BP



Store the last message on the edge and divide each passing message by the last stored.

$$\pi_3(C, D) = \pi_3^0(C, D) \sum_B \pi_2^0(B, C) = \pi_3^0(C, D) \mu_{2,3}$$

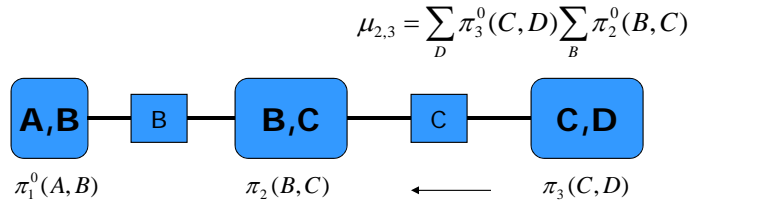
$$\delta_{3 \rightarrow 2} = \left(\sum_D \pi_3(C, D) \right)$$

$$\pi_2(B, C) = \pi_2^0(B, C) \frac{\delta_{3 \rightarrow 2}}{\mu_{2,3}(C)} = \frac{\pi_2^0(B, C)}{\mu_{2,3}(C)} \times \sum_D \pi_3^0(C, D) \times \mu_{2,3}(C) = \pi_2^0(B, C) \times \sum_D \pi_3^0(C, D)$$

$$\mu_{2,3} = \delta_{3 \rightarrow 2} = \left(\sum_D \pi_3(C, D) \right) = \sum_D \pi_3^0(C, D) \sum_B \pi_2^0(B, C) \quad \text{New message}$$

CS 3750 Advanced Machine Learning

Message Passing: BP



Store the last message on the edge and divide each passing message by the last stored.

$$\pi_3(C, D) = \pi_3^0(C, D) \sum_B \pi_2^0(B, C)$$

$$\delta_{3 \rightarrow 2} = \left(\sum_D \pi_3(C, D) \right)$$

$$\pi_2(B, C) = \pi_2^0(B, C) \times \sum_D \pi_3^0(C, D)$$

The same as before

$$\pi_2(B, C) = \pi_2(B, C) \frac{\delta_{3 \rightarrow 2}}{\mu_{2,3}(C)} = \pi_2(B, C) \times \frac{\sum_D \pi_3^0(C, D) \times \sum_B \pi_2^0(B, C)}{\sum_D \pi_3^0(C, D) \times \sum_B \pi_2^0(B, C)} = \pi_2(B, C)$$

CS 3750 Advanced Machine Learning

Message Propagation: BP

- **Lauritzen-Spiegelhalter algorithm**
- Two kinds of objects: clique and sepset potentials
 - Initial potentials not kept
- Improved “stability” of asynchronous algorithm (repeated messages cancel out)
- **New distribution representation**
 - clique tree potential

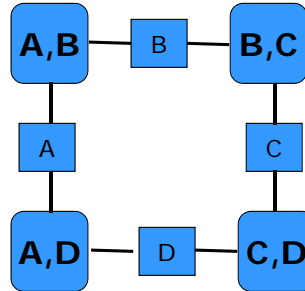
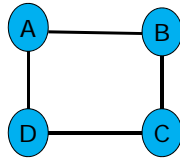
$$\pi_T = \frac{\prod_{C_i \in T} \pi_i(C_i)}{\prod_{(C_i \leftrightarrow C_j) \in T} \mu_{ij}(S_{ij})} = P_F(X)$$

- Clique tree invariant = P_F

CS 3750 Advanced Machine Learning

Loopy belief propagation

- The asynchronous BP algorithm works on clique trees
- What if we run the belief propagation algorithm on a non-tree structure?



- Sometimes converges
- If it converges it leads to an approximate solution
- **Advantage:** tractable for large graphs

CS 3750 Advanced Machine Learning

Loopy belief propagation

- If the BP algorithm converges, it converges to an optimum of the Bethe free energy

See papers:

- Yedidia J.S., Freeman W.T. and Weiss Y. Generalized Belief Propagation, 2000
- Yedidia J.S., Freeman W.T. and Weiss Y. Understanding Belief Propagation and Its Generalizations, 2001

CS 3750 Advanced Machine Learning