

Probabilistic Graphical Models

Lecture 9 – Undirected Models


CS/CNS/EE 155
Andreas Krause

Announcements

- Homework 2 due next Wednesday (Nov 4) in class
 - Start early!!!
- Project milestones due Monday (Nov 9)
 - 4 pages of writeup, NIPS format
 - <http://nips.cc/PaperInformation/StyleFiles>

Best project award!!

Answering multiple queries



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

Can compute in $O(n)$ operations (+, *)


$P(x_2, x_5)$ costs $O(n)$ ops.

$P(x_i, x_5)$ — " $O(n)$ ops

If compute $P(x_i | x_j) \forall i \Rightarrow O(n^2)$

Can we reduce this to $O(n)$??

Reusing computation



$$P(x_1, x_5) = P(x_1) \sum_{x_2} P(x_2 | x_1) \underbrace{\sum_{x_3} P(x_3 | x_2) \underbrace{\sum_{x_4} P(x_4 | x_3) P(x_5 | x_4)}_{g_4(x_3, x_5)}}_{g_3(x_2, x_5)}$$

$$P(x_2, x_5) = \sum_{x_1} P(x_1) P(x_2 | x_1) \sum_{x_3} P(x_3 | x_2) \underbrace{\sum_{x_4} P(x_4 | x_3) P(x_5 | x_4)}_{g_4(x_3, x_5)}$$


↑
already computed!

Want to "cache" our computations!

Next

- Will learn about algorithm for efficiently computing all marginals $P(X_i \mid \mathbf{E}=\mathbf{e})$ given fixed evidence $\mathbf{E}=\mathbf{e}$
- Need appropriate data structure for storing the computation
 - ➔ Junction trees


Junction trees



A junction tree for a collection of factors:


- A tree, where each node is a cluster of variables
- Every factor contained in some cluster C_i
- **Running intersection property:** If $X \in C_i$ and $X \in C_j$, and C_m is on the path between C_i and C_j , then $X \in C_m$

VE constructs a junction tree




- One clique C_i for each factor f_i created in VE
- C_i connected to C_j if f_i used to generate f_j
- Every factor used only once \rightarrow Tree
- **Theorem:** resulting tree satisfies RIP

Constructing JT from chordal graph




1. Moralize
2. Triangulate (make chordal)
3. Identify max. cliques
4. Connect cliques into undirected graph
 $w(C_i, C_j) = |C_i \cap C_j|$
5. Find Max ST



⇒ Results in valid junction tree


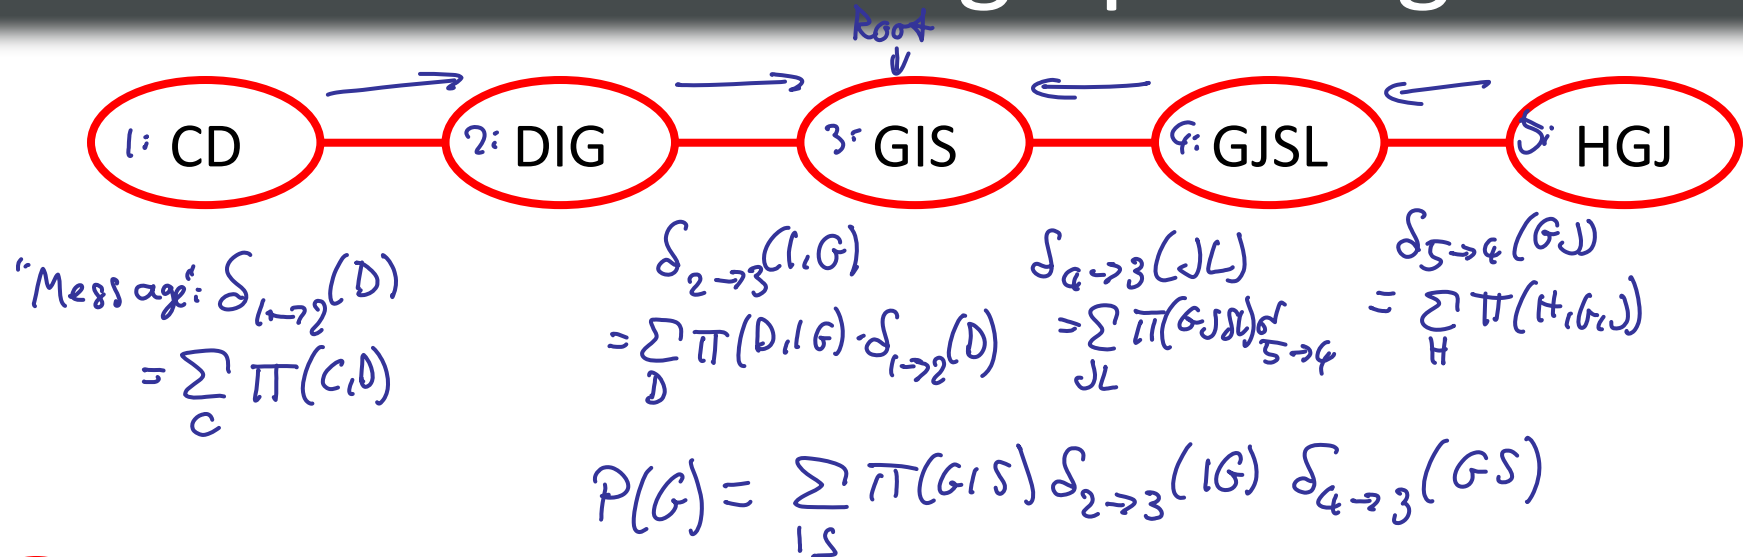
Junction trees and independence



Theorem:

- Suppose
 - T is a junction tree for graph G and factors F
 - Consider edge $C_i - C_j$ with separator $S_{i,j} = C_i \cap C_j$
 - Variables X and Y on opposite sites of separator
- Then $X \perp Y \mid S_{i,j}$
- Furthermore, $I(T) \subseteq I(G)$

VE as message passing



VE for computing X_i


- Pick root (any clique containing X_i)
- Don't eliminate, only send messages recursively from leaves to root
 - Multiply incoming messages with clique potential
 - Marginalize variables not in separator
- Root "ready" when received all messages

Correctness of message passing




- **Theorem:** When root ready (received all messages), all variables in root have correct potentials
 - Follows from correctness of VE
- So far, no gain in efficiency ☹️

Does the choice of root affect messages?



instead of $O(n^2) \rightarrow O(n)$
1 message per edge per direction

Shenoy-Shafer algorithm



- Clique i ready if received messages from all neighbors *but 1*
 - Leaves always ready
- While there exists a message $\delta_{i \rightarrow j}$ ready to transmit send message

Complexity? $O(n 2^{\text{treewidth}})$
1 msg per edge per direction
"Only" exp. in treewidth

Theorem: At convergence, every clique has correct beliefs

Inference using VE

- Want to incorporate evidence $E=e$
- Multiply all cliques containing evidence variables with indicator potential 1_e

$$\textcircled{AB} \quad I_{A=T}(a,b) = \begin{cases} 1 & \text{if } a=T \\ 0 & \text{if } a=F \end{cases}$$

- Perform variable elimination

Summary so far

- Junction trees represent distribution
 - Constructed using elimination order
 - Make complexity of inference explicitly visible
- Can implement variable elimination on junction trees to compute correct beliefs on all nodes
- Now:
 - **Belief propagation** – an important alternative to VE on junction trees.
 - Will later generalize to approximate inference!
 - Key difference: Messages obtained by division rather than multiplication

Message passing by factor division

- Variable elimination:

- Message \rightarrow Belief

$$\delta_{2 \rightarrow 3}(IG) = \sum_D \pi_2^{(0)}(DIG) \delta_{1 \rightarrow 2}(D)$$

$$\text{Belief at 3: } \pi_3^{(0)}(GIS) \cdot \delta_{2 \rightarrow 3}(IG)$$


- Factor division:

- Belief \rightarrow Message

$$\text{Belief at 2: } \pi_2^{(0)}(DIG) \cdot \delta_{1 \rightarrow 2}(D) \cdot \delta_{3 \rightarrow 2}(IG)$$

$$\text{Belief about Sep. IG: } \delta_{2 \rightarrow 3}^{(1)}(IG) = \sum_D \pi_2^{(1)}(DIG) \quad \leftarrow \text{Send as msg}$$

$$\text{Belief at 3: } \pi_3^{(k+1)}(GIS) = \pi_3^{(0)}(GIS) \cdot \frac{\delta_{2 \rightarrow 3}^{(1)}(IG)}{\delta_{3 \rightarrow 2}(IG)} = \frac{\sum_D \pi_3^{(0)}(GIS) \pi_2^{(0)}(DIG) \delta_{1 \rightarrow 2}(D) \delta_{3 \rightarrow 2}(IG)}{\delta_{3 \rightarrow 2}(IG)}$$



Factor division

$$f_1(A, B, C) \quad f_2(B, C)$$

$$f = \frac{f_1}{f_2} \quad f(A, B, C) = \frac{f_1(A, B, C)}{f_2(B, C)}$$

with the convention that $\frac{0}{0} = 0$

Clique and separator potentials

$A \rightarrow B \rightarrow C$



$$\pi_1^{(0)}(AB) = P(A) \cdot P(B|A)$$

$$\pi_2^{(0)}(BC) = P(C|B)$$

Belief
about
sep. B

$$\sigma_{1 \rightarrow 2}^{(0)}(B) = \sum_A \pi_1^{(0)}(AB) \not\equiv \sigma_{2 \rightarrow 1}^{(0)}(B) = \sum_C \pi_2^{(0)}(BC)$$


At convergence:

$$\pi_1^{(*)}(AB) = P(AB) \quad \pi_2^{(*)} = P(BC)$$

$$\sigma_{1 \rightarrow 2}^{(*)}(B) = \sum_A \pi_1^{(*)}(AB) = P(B) = \sigma_{2 \rightarrow 1}^{(*)} = \mu_{12}(B)$$

↑
Call such a JT with potentials $\pi_1^{(*)} \dots \pi_n^{(*)}$ that agree on separators "calibrated"

Lauritzen-Spiegelhalter algorithm (a.k.a. Belief Propagation)



- Initialize separator potentials μ_{ij}
 - One per edge, initialized to 1
stores last msg across edge $i-j$

- Messages $i \rightarrow j$

$$\sigma_{i \rightarrow j}(C_j \cap C_i) = \sum_{C_i \setminus C_j} \pi_i^{(t)}$$

$$\pi_j^{(t+1)} = \pi_j^{(t)} \cdot \frac{\sigma_{i \rightarrow j}}{\mu_{ij}}$$


$$\mu_{i\bar{j}} = \sigma_{i \rightarrow j}$$

Correctness of Belief propagation

- Complexity linear in #cliques
- **Theorem:**
At convergence, every clique has correct beliefs
(when using correct message order, i.e., leaves to root
and back)
- **Corollary:**
Junction tree is calibrated (cliques agree on separator)


Comparison of VE and BP messages

- Variable elimination



$$\delta_{3 \rightarrow 4} = \sum_{C_3 \setminus C_4} \pi_3^{(0)} \delta_{1 \rightarrow 3} \delta_{2 \rightarrow 3}$$

- Belief propagation



$$\begin{aligned} \phi_{3 \rightarrow 4}^{(k)} &= \sum_{C_3 \setminus C_4} \pi_3^{(k)} \\ \pi_4^{(k+1)} &= \pi_4^{(k)} \cdot \frac{\phi_{3 \rightarrow 4}^{(k)}}{\mu_{34}^{(k)}} \\ \mu_{34}^{(k+1)} &= \phi_{3 \rightarrow 4}^{(k)} \end{aligned}$$

Understanding BP

- Junction tree potential

$$\Pi_T(x) = \frac{\prod_i \pi_{C_i}(C_i)}{\prod_{ij} \mu_{ij}(C_i \cap C_j)}$$

- Junction tree invariant

$$\Pi_T(x) = P(x)$$

- Theorem: BP maintains Junction tree invariant
➔ BP reparametrizes clique and separator potentials

Advantages and disadvantages of junction tree inference

- Advantages

- Can answer multiple queries (for same evidence) efficiently
- Can perform incremental updates

- Disadvantages

- No factors are “deleted”
- Can't prune away unnecessary variables
- Slower for a single query


Summary so far

- Bayesian Networks
 - Representation
 - Learning (MLE / Bayesian) with fully observed data
 - Exact Inference
- Next
 - Undirected models
 - Approximate inference
 - Hidden variables

Representing the world using BNs



represent



True distribution P'
with cond. ind. $I(P')$

Bayes net (G, P)
with $I(P)$


- Want to make sure that $I(P) \subseteq I(P')$
- Ideally: $I(P) = I(P')$
- Want BN that **exactly** captures independencies in P' !

Perfect maps

- Minimal I-maps are easy to find, but can contain many unnecessary dependencies.
- A BN structure G is called **P-map** (perfect map) for distribution P if $I(G) = I(P)$
- Does every distribution P have a P-map?

Existence of perfect maps

$$X_1, \dots, X_4 \quad X_1 \perp X_3 \mid X_2, X_4 \\ X_2 \perp X_4 \mid X_1, X_3$$




← Undirected GM
is a P-map

but NO BN is P-map

- Will have undirected model as P-map

Undirected Parameterization



Specify factors over cliques in undirected graph

$$\psi_{34}(X_3, X_4) \geq 0$$

ψ_{24}

X_2	X_4	$\psi_{24}(X_2, X_4)$
0	0	3
0	1	17
1	0	1
1	1	31 & 15

$$P(X_1, \dots, X_4) = \frac{1}{Z} \psi_{12}(X_1, X_2) \psi_{13}(X_1, X_3) \psi_{24} \psi_{34}$$

$$Z = \sum_{X_1 \dots X_4} \prod_{ij} \psi_{ij}(X_i, X_j)$$


Markov Networks

(a.k.a., Markov Random Field, Gibbs Distribution, ...)

- A Markov Network consists of
 - An undirected graph, where each node represents a RV
 - A collection of factors defined over cliques in the graph

- Joint probability

$$P(x) = \frac{1}{Z} \prod_i \psi_i(C_i)$$



- A distribution factorizes over undirected graph G if


\exists factors $\psi_1 \dots \psi_k$ over cliques of G s.t.

$$P(x) = \frac{1}{Z} \prod_i \psi_i(C_i)$$

Example MN: Image denoising



Markov Network



X_i : noisy pixels

Y_i : "true" pixels

Computing Joint Probabilities

- Computing joint probabilities in BNs

$$P(X_1, \dots, X_n) = \prod_i P(X_i | Pa_i)$$

$P(X_i | X_n)$
actually comp. $P(X_i, X_n)$

- Computing joint probabilities in Markov Nets

$$P(X_1, \dots, X_n) = \frac{1}{Z} \prod_i \psi_i(C_i)$$

Need to know partition "function" Z

Can do $V_{\mathcal{F}}$


Can compute

$$\frac{P(x_1, \dots, x_n)}{P(x'_1, \dots, x'_n)} = \frac{\prod_i \psi_i(C_i)}{\prod_i \psi_i(C'_i)}$$

Independences in Markov Nets?


- In Bayes Nets (G,P)
 - Local Markov Assumption: $X \perp \text{NonDesc}(X) \mid \text{Pa}_X$
 - G is I-map for distribution P if Local Markov Assumption holds
 - Factorization Thm: P factorizes over G \Leftrightarrow G is an I-map
 - Global independences: d-separation
 - Completeness and soundness of d-separation
- How about Markov Nets?
 - What's the analog of the Local Markov Assumption?
 - Is there a factorization theorem for Markov Nets?
 - What replaces d-separation?

Local Markov Assumption for MN




- The **Markov Blanket** $\text{MB}(X)$ of a node X is the set of neighbors of X
- Local Markov Assumption: $X \perp \text{EverythingElse} \mid \text{MB}(X)$
- $I_{\text{loc}}(G)$ = set of all local independences
- G is called an I-map of distribution P if $I_{\text{loc}}(G) \subseteq I(P)$

Factorization Theorem for Markov Nets “→”



True distribution P
can be represented exactly as
a Markov net (G, P)


$$P(X_1, \dots, X_n) = \frac{1}{Z} \prod_i \phi_i(\mathbf{C}_i)$$



$$I_{\text{loc}}(G) \subseteq I(P)$$

G is an **I-map** of P
(independence map)


Factorization Theorem for Markov Nets “←”



True distribution P
can be represented exactly as

$$P(X_1, \dots, X_n) = \prod_i P(X_i \mid \mathbf{Pa}_{X_i})$$

$I_{\text{loc}}(G) \subseteq I(P)$
 G is an **I-map** of P
(independence map)



*Not true
in general*


i.e., P can be represented as
a Markov net (G, P)

Counterexample

- G an I-map for P does not imply that P factorizes
- Binary variables X_1, \dots, X_4 .
- Only positive states

$(0,0,0,0), (1,0,0,0), (1,1,0,0), (1,1,1,0)$
 $(0,0,0,1), (0,0,1,1), (0,1,1,1), (1,1,1,1)$

each happens with prob $\frac{1}{8}$




G is I-map for P


$$X_1 \perp X_3 \mid X_2, X_4$$

Eg.: $X_2=1, X_4=1$

But to represent P , need fully connected graph



Factorization Theorem for Markov Nets “←” Hammersley-Clifford Theorem



True distribution P
can be represented exactly as

$$P(X_1, \dots, X_n) = \prod_i P(X_i \mid \mathbf{Pa}_{X_i})$$


i.e., P can be represented as
a Markov net (G, P)

$$I_{\text{loc}}(G) \subseteq I(P)$$



G is an **I-map** of P
(independence map)
and $P > 0$

Global independencies



- A trail $X - \underline{X_1 - \dots - X_m} - Y$ is called active for evidence E , if none of $X_1, \dots, X_m \in E$
- Variables X and Y are called **separated** by E if there is no active trail for E connecting X, Y . Write $\text{sep}(X, Y \mid E)$
- $I(G) = \{X \perp Y \mid E: \text{sep}(X, Y \mid E)\}$

Soundness of separation

- Know: For positive distributions $P > 0$

$$I_{\text{loc}}(G) \subseteq I(P) \Leftrightarrow P \text{ factorizes over } G$$

- **Theorem:** Soundness of separation

For positive distributions $P > 0$

$$I_{\text{loc}}(G) \subseteq I(P) \Leftrightarrow I(G) \subseteq I(P)$$

- Hence, separation captures only true independences
- How about $I(G) = I(P)$?

Completeness of separation

Theorem: Completeness of separation

$$I(G) = I(P)$$

for “almost all” distributions P that factorize over G


“almost all”: Except for of potential parameterizations of measure 0 (assuming no finite set have positive measure)

Independences in Markov Nets?

- In Bayes Nets (G,P)
 - Local Markov Assumption: $X \perp \text{NonDesc}(X) \mid \text{Pa}_X$
 - G is I-map for distribution P if Local Markov Assumption holds
 - Factorization Thm: P factorizes over G \Leftrightarrow G is an I-map
 - Global independences: d-separation
 - Completeness and soundness of d-separation
- How about Markov Nets?
 - Local Markov Assumption: $X \perp \text{EverythingElse} \mid \text{MB}(X)$
 - Factorization Thm: For positive P, P factorizes \Leftrightarrow G is an I-map
 - Global independences: separation
 - For positive P: separation is complete and sound
- How about minimal I-maps and P-maps??

Minimal I-maps


- For BNs: Minimal I-map not unique




- For MNs: For positive P, minimal I-map is unique!!

P-maps

- Do P-maps always exist?
- For BNs: no



- How about Markov Nets?




does not have
MN P-map!

Exact inference in MNs

- Variable elimination and junction tree inference work exactly the same way!
 - Need to construct junction trees by obtaining chordal graph through triangulation

Pairwise MNs

- A pairwise MN is a MN where all factors are defined over single variables or pairs of variables
- Can reduce any MN to pairwise MN!



Tasks

- Read Koller & Friedman Chapters 10 and 4.1-4.5