

Readings:

K&F: 9.1, 9.2, 9.3, 9.4

K&F: 5.1, 5.2, 5.3, 5.4, 5.5, 5.6

Clique Trees 3

Let's get BP right

Undirected Graphical Models

Here the couples get to swing!

Graphical Models – 10708

Carlos Guestrin

Carnegie Mellon University

October 25th, 2006

1

Factor division

- Let X and Y be disjoint set of variables

- Consider two factors:
 $\phi_1(X, Y)$ and $\phi_2(Y)$

- Factor $\psi = \phi_1 / \phi_2$

□ $0/0=0$

$\phi_1(A, B)$

a^1	b^1	0.5
a^1	b^2	0.2
a^2	b^1	0
a^2	b^2	0
a^3	b^1	0.3
a^3	b^2	0.45

$\phi_2(A)$

a^1	0.8
a^2	0
a^3	0.6

$\psi(A, B)$

a^1	b^1	0.625
a^1	b^2	0.25
a^2	b^1	0
a^2	b^2	0
a^3	b^1	0.5
a^3	b^2	0.75

$$\psi(X=x, Y=y) = \frac{\phi_1(X=x, Y=y)}{\phi_2(Y=y)}$$

$$\psi(A=a^1, B=b^2) = \frac{0.2}{0.8} = 0.25$$

Introducing message passing with division

Variable elimination (message passing with multiplication)

message: $\sigma_{3 \rightarrow 4}(G, S) = \sum_d \pi_0(G, d, S) \cdot \sigma_{1 \rightarrow 3}(d) \cdot \sigma_{2 \rightarrow 3}(S)$


belief: $\pi_3(GDS) = \pi_0(GDS) \cdot \sigma_{1 \rightarrow 3} \cdot \sigma_{2 \rightarrow 3} \cdot \sigma_{4 \rightarrow 3}$

Message passing with division:

message: $\sigma_{3 \rightarrow 4}(GS) = \sum_d \pi_3(G, d, S) \cdot \sigma_{2 \rightarrow 3}(S) \cdot \sigma_{4 \rightarrow 3}(GS)$

belief update:

$$\pi_4(GDS) \leftarrow \pi_3(GDS) \cdot \sigma_{3 \rightarrow 4}(GS)$$



10-708 - ©Carlos Guestrin 2006

3

Lauritzen-Spiegelhalter Algorithm (a.k.a. belief propagation)

Separator potentials μ_{ij}

- one per edge (same both directions)
- holds "last message"
- initialized to 1

Message $i \rightarrow j$

- what does i think the separator potential should be?


$$\sigma_{i \rightarrow j}(S_{ij}) = \sum_{C_i / S_{ij}} \pi_i(C_i)$$

- update belief for j :

$$\pi_j \leftarrow \pi_j \cdot \sigma_{i \rightarrow j}$$

- replace separator potential:

$$\mu_{ij} \leftarrow \sigma_{i \rightarrow j}$$



10-708 - ©Carlos Guestrin 2006

4

Clique tree invariant

- **Clique tree potential:** $\pi_T(X) = \frac{\prod_i \pi_i(c_i)}{\prod_{ij} \mu_{ij}(s_{ij})}$
 - Product of clique potentials divided by separator potentials
- initialization: $\pi_T(X) = \frac{\prod_i \pi_i(c_i)}{\prod_{ij} 1} = \prod_j P(x_j | P_{x_j})$

$$\frac{\prod_i \pi_i(c_i)}{\prod_{ij} 1} = P(X)$$
- **Clique tree invariant:**
 - $P(X) = \pi_T(X)$
 - send message $i \rightarrow j$
 - use π_j, μ_{ij} $\pi_j' \leftarrow \pi_j - \frac{\sigma_{ij}}{\mu_{ij}}$
 - use π_j' $\mu_{ij}' \leftarrow \sigma_{ij}$
 - $\pi_T(X)$ before message
 - $\pi_T'(X)$ after message
 - $\left. \begin{aligned} \frac{\pi_j'}{\mu_{ij}'} &= \frac{\pi_j \cdot \frac{\sigma_{ij}}{\mu_{ij}}}{\frac{\sigma_{ij}}{\mu_{ij}}} \\ &= \frac{\pi_j}{\mu_{ij}} \end{aligned} \right\}$

10-708 - ©Carlos Guestrin 2006

5

Belief propagation and clique tree invariant

- **Theorem:** Invariant is maintained by BP algorithm!
 - $\pi_T(X) = P(X)$
 - at convergence: $\pi_i(c_i) = P(c_i)$
 - $\mu_{ij}(s_{ij}) = P(s_{ij})$
- BP reparameterizes clique potentials and separator potentials
 - At convergence, potentials and messages are marginal distributions

$$P(X) = \frac{\prod_i P(c_i)}{\prod_{ij} P(s_{ij})}$$

10-708 - ©Carlos Guestrin 2006

6

Subtree correctness

- **Informed message** from i to j , if all messages into i (other than from j) are informed
 - Recursive definition (leaves always send informed messages)
- **Informed subtree:**
 - All incoming messages informed
- **Theorem:**
 - Potential of connected informed subtree T' is marginal over $\text{scope}[T']$
- **Corollary:**
 - At convergence, clique tree is *calibrated*
 - $\pi_i = P(\text{scope}[\pi_i])$
 - $\mu_{ij} = P(\text{scope}[\mu_{ij}])$

10-708 – ©Carlos Guestrin 2006

7

Clique trees versus VE

- Clique tree advantages
 - Multi-query settings
 - Incremental updates
 - Pre-computation makes complexity explicit
- Clique tree disadvantages
 - Space requirements – no factors are “deleted”
 - Slower for single query
 - Local structure in factors may be lost when they are multiplied together into initial clique potential

10-708 – ©Carlos Guestrin 2006

8

Clique tree summary

- Solve marginal queries for all variables in only twice the cost of query for one variable
- Cliques correspond to maximal cliques in induced graph
- Two message passing approaches
 - VE (the one that multiplies messages)
 - BP (the one that divides by old message)
- Clique tree invariant
 - Clique tree potential is always the same
 - We are only reparameterizing clique potentials
- Constructing clique tree for a BN
 - from elimination order
 - from triangulated (chordal) graph
- Running time (only) exponential in size of largest clique
 - Solve **exactly** problems with thousands (or millions, or more) of variables, and cliques with tens of nodes (or less)

10-708 – ©Carlos Guestrin 2006

9

Announcements

- Recitation tomorrow, don't miss it!!!
 - Khalid on Undirected Models

10-708 – ©Carlos Guestrin 2006


10

Swinging Couples revisited

- This is no perfect map in BNs
- But, an undirected model will be a perfect map

$$w_1 \perp w_2 \mid M_1, M_2$$


$$M_1 \perp M_2 \mid w_1, w_2$$



10-708 - ©Carlos Guestrin 2006

11

Potentials (or Factors) in Swinging Couples



$$\pi_1(A, B) \neq 0 \quad \forall A, B$$

$\pi_1^0[A, B]$			$\pi_2^0[B, C]$			$\pi_3^0[C, D]$			$\pi_4^0[D, A]$		
a^0	b^0	30	b^0	c^0	100	c^0	d^0	1	d^0	a^0	100
a^0	b^1	5	b^0	c^1	1	c^0	d^1	100	d^0	a^1	1
a^1	b^0	1	b^1	c^0	1	c^1	d^0	100	d^1	a^0	1
a^1	b^1	10	b^1	c^1	100	c^1	d^1	1	d^1	a^1	100

$$P(A, B, C, D) = \frac{1}{Z} \pi_1(A, B) \cdot \pi_2(B, C) \cdot \pi_3(C, D) \cdot \pi_4(D, A)$$

normalization

$$Z = \sum_{a,b,c,d} \pi_1(a,b) \cdot \pi_2(b,c) \cdot \pi_3(c,d) \cdot \pi_4(d,a)$$

10-708 - ©Carlos Guestrin 2006

12

Computing probabilities in Markov networks v. BNs

- In a BN, can compute prob. of an instantiation by multiplying CPTs

$P(A=a, B=b, C=c, D=d)$
 $= P(A=a) \cdot P(B=b|A=a) \cdot P(C=c|A=a) \cdot P(D=d|B=b, C=c)$

can't do this directly in a MN, you need to know Z

- In an Markov networks, can only compute ratio of probabilities directly

$\pi_1[A, B]$	$\pi_2[B, C]$	$\pi_3[C, D]$	$\pi_4[D, A]$
$a^0 b^0$ 30	$b^0 c^0$ 100	$c^0 d^0$ 1	$d^0 a^0$ 100
$a^0 b^1$ 5	$b^0 c^1$ 1	$c^0 d^1$ 100	$d^0 a^1$ 1
$a^1 b^0$ 1	$b^1 c^0$ 1	$c^1 d^0$ 100	$d^1 a^0$ 1
$a^1 b^1$ 10	$b^1 c^1$ 100	$c^1 d^1$ 1	$d^1 a^1$ 100

$$\frac{P(A=t, B=f, C=t, D=t)}{P(A=f, B=f, C=f, D=f)} = \frac{\frac{1}{Z} \pi_1(t, t) \pi_2(t, t) \pi_3(t, t) \pi_4(t, t)}{\frac{1}{Z} \pi_1(ff) \cdot \pi_2(ff) \cdot \pi_3(ff) \pi_4(ff)}$$

$$= \frac{1}{3}$$

10-708 - ©Carlos Guestrin 2006

13

Normalization for computing probabilities


- To compute actual probabilities, must compute normalization constant (also called partition function)

$$P(ABCD) = \frac{1}{Z} \pi_1(AB) \pi_2(BC) \pi_3(CD) \pi_4(DA)$$

Assignment	Unnormalized	Normalized
$a^0 b^0 c^0 d^0$	300000	0.03
$a^0 b^0 c^0 d^1$	300000	0.04
$a^0 b^0 c^1 d^0$	300000	0.03
$a^0 b^0 c^1 d^1$	30	$1.1 \cdot 10^{-5}$
$a^0 b^1 c^0 d^0$	500	$6.9 \cdot 10^{-5}$
$a^0 b^1 c^0 d^1$	500	$6.9 \cdot 10^{-5}$
$a^0 b^1 c^1 d^0$	5000000	0.60
$a^0 b^1 c^1 d^1$	500	$6.9 \cdot 10^{-5}$
$a^1 b^0 c^0 d^0$	100	$1.4 \cdot 10^{-5}$
$a^1 b^0 c^0 d^1$	100000	0.14
$a^1 b^0 c^1 d^0$	100	$1.4 \cdot 10^{-5}$
$a^1 b^0 c^1 d^1$	10	$1.4 \cdot 10^{-6}$
$a^1 b^1 c^0 d^0$	100000	0.014
$a^1 b^1 c^0 d^1$	100000	0.014
$a^1 b^1 c^1 d^0$	100000	0.014

- Computing partition function is hard! → Must sum over all possible assignments

$$Z = \sum_{abcd} \pi_1(ab) \pi_2(bc) \cdot \pi_3(cd) \pi_4(da)$$



10-708 - ©Carlos Guestrin 2006



14

Factorization in Markov networks



- Given an undirected graph H over variables $\mathbf{X} = \{X_1, \dots, X_n\}$
- A distribution P **factorizes** over H if $\exists$
 - subsets of variables $\mathbf{D}_1 \subseteq \mathbf{X}, \dots, \mathbf{D}_m \subseteq \mathbf{X}$, such that the $\mathbf{D}_i$ are fully connected in H
 - non-negative potentials (or factors) $\pi_1(\mathbf{D}_1), \dots, \pi_m(\mathbf{D}_m)$
 - also known as clique potentials
 - such that

$$P(\mathbf{X}) = \frac{1}{Z} \prod_{i=1}^m \pi_i(\mathbf{D}_i)$$



- Also called Markov random field H , or Gibbs distribution over H


10-708 - ©Carlos Guestrin 2006

15

Global Markov assumption in Markov networks



- A path $X_1 - \dots - X_k$ is **active** when set of variables $\mathbf{Z}$ are observed if none of $X_i \in \{X_1, \dots, X_k\}$ are observed (are part of $\mathbf{Z}$)
- Variables $\mathbf{X}$ are **separated** from $\mathbf{Y}$ given $\mathbf{Z}$ in graph H , $\text{sep}_H(\mathbf{X}; \mathbf{Y} | \mathbf{Z})$, if there is no active path between any $X \in \mathbf{X}$ and any $Y \in \mathbf{Y}$ given $\mathbf{Z}$




- The **global Markov assumption** for a Markov network H is $\mathcal{I}(H) = \{ X_1, X_2 \mid \text{sep}_H(X_1, X_2 | Z) \}$

active if
I don't observe
35, 1, 9...


10-708 - ©Carlos Guestrin 2006

16

The BN Representation Theorem



Important because:
Independencies are sufficient to obtain BN structure G




Important because:
Read independencies of P from BN structure G

10-708 - ©Carlos Guestrin 2006

17

Markov networks representation Theorem 1



P factorizes according to H

- If you can write distribution as a normalized product of factors $\Rightarrow$ Can read independencies from graph

10-708 - ©Carlos Guestrin 2006

18

What about the other direction for Markov networks?

If H is an I-map for P

$I(H) \subseteq I(P)$

Then

joint probability distribution P :

$$P(X_1, \dots, X_n) = \frac{1}{Z} \prod_{i=1}^m \pi_i(D_i)$$

- Counter-example: $X_1, \dots, X_4$ are binary, and only eight assignments have positive probability:

(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)

$P(0,1,0,1) = 0$

$P(X_1=0 | X_2=0, X_4=0) = 1/2$

indep of value of X_3
- For example, $X_1 \perp X_3 | X_2, X_4$:

$P(X_1=0 | X_2=0, X_4=0) = 1/2$

indep of value of X_3
- But distribution doesn't factorize!!!

Homework-!!

Markov networks representation Theorem 2 (Hammersley-Clifford Theorem)

If H is an I-map for P
and
 P is a positive distribution

$\forall x \ P(x) > 0$

Then

joint probability distribution P :

$$P(X_1, \dots, X_n) = \frac{1}{Z} \prod_{i=1}^m \pi_i(D_i)$$

- Positive distribution and independencies $\Rightarrow$ P factorizes over graph

Representation Theorem for Markov Networks

If joint probability distribution P :

$$P(X_1, \dots, X_n) = \frac{1}{Z} \prod_{i=1}^m \pi_i(D_i)$$

Then

H is an I-map for P

If H is an I-map for P and P is a positive distribution

Then

joint probability distribution P :

$$P(X_1, \dots, X_n) = \frac{1}{Z} \prod_{i=1}^m \pi_i(D_i)$$

10-708 - ©Carlos Guestrin 2006

21

Completeness of separation in Markov networks

■ Theorem: Completeness of separation

- For “almost all” distributions that P factorize over Markov network H , we have that $I(H) = I(P)$ you know that $I(H) \subseteq I(P)$
- “almost all” distributions: except for a set of measure zero of parameterizations of the Potentials (assuming no finite set of parameterizations has positive measure)

■ Analogous to BNs

10-708 - ©Carlos Guestrin 2006

22

What are the “local” independence assumptions for a Markov network?

In a BN G :

- local Markov assumption: variable independent of non-descendants given parents $I_L(G)$
- d-separation defines global independence $I(G)$
- Soundness: For all distributions: $P \models I_L(G) \Leftrightarrow P \models I(G)$

In a Markov net H :

- **Separation** defines global independencies
- What are the notions of local independencies?

10-708 – ©Carlos Guestrin 2006

23

Local independence assumptions for a Markov network

Separation defines global independencies $I(H)$

Pairwise Markov Independence: $I_{pw}(H)$


- Pairs of non-adjacent variables are independent given all others

$$A \perp B \mid \mathcal{X} - \{A, B\}$$

Markov Blanket: $I_{MB}(H)$

- Variable independent of rest given its neighbors $N(A)$

$$A \perp \mathcal{X} - \{N(A), A\} \mid N(A)$$



$$T_7 \perp T_2 \mid T_1, T_3 = 6, T_8, T_9$$

$$T_7 \perp \{T_1, T_2, T_3, T_4, T_6, T_8, T_9\} \mid T_5$$

10-708 – ©Carlos Guestrin 2006

24

Equivalence of independencies in Markov networks

- **Soundness Theorem:** For all positive distributions P , the following three statements are equivalent:
 - P entails the global Markov assumptions
 $P \models I(H)$
 - P entails the pairwise Markov assumptions
 $P \models I_{pw}(H)$
 - P entails the local Markov assumptions (Markov blanket)
 $P \models I_{NB}(H)$

10-708 – ©Carlos Guestrin 2006

25

Minimal I-maps and Markov Networks

- A fully connected graph is an I-map
- Remember minimal I-maps?
 - A “simplest” I-map → Deleting an edge makes it no longer an I-map
- In a BN, there is no unique minimal I-map
- Theorem: **In a Markov network, minimal I-map is unique!!**
- Many ways to find minimal I-map, e.g.,
 - Take pairwise Markov assumption:
 - If P doesn't entail it, add edge:

10-708 – ©Carlos Guestrin 2006

26

How about a perfect map?

- Remember perfect maps?
 - independencies in the graph are exactly the same as those in P
- For BNs, doesn't always exist
 - counter example: Swinging Couples
- How about for Markov networks?

10-708 – ©Carlos Guestrin 2006

27

Unifying properties of BNs and MNs

- BNs:
 - give you: V-structures, CPTs are conditional probabilities, can directly compute probability of full instantiation
 - but: require acyclicity, and thus no perfect map for swinging couples
- MNs:
 - give you: cycles, and perfect maps for swinging couples
 - but: don't have V-structures, cannot interpret potentials as probabilities, requires partition function
- Remember PDAGS???
- skeleton + immoralities
 - provides a (somewhat) unified representation
 - see book for details

10-708 – ©Carlos Guestrin 2006

28

What you need to know so far about Markov networks

- Markov network representation:
 - undirected graph
 - potentials over cliques (or sub-cliques)
 - normalize to obtain probabilities
 - need partition function
- Representation Theorem for Markov networks
 - if P factorizes, then it's an I-map
 - if P is an I-map, only factorizes for positive distributions
- Independence in Markov nets:
 - active paths and separation
 - pairwise Markov and Markov blanket assumptions
 - equivalence for positive distributions
- Minimal I-maps in MNs are unique
- Perfect maps don't always exist

10-708 – ©Carlos Guestrin 2006

29

Some common Markov networks and generalizations

- Pairwise Markov networks
- A very simple application in computer vision
- Logarithmic representation
- Log-linear models
- Factor graphs

10-708 – ©Carlos Guestrin 2006


30

Pairwise Markov Networks

- All factors are over single variables or pairs of variables:

- ☐ Node potentials
- ☐ Edge potentials

- Factorization:



- Note that there may be bigger cliques in the graph, but only consider pairwise potentials

10-708 - ©Carlos Guestrin 2006

31

A very simple vision application

- Image segmentation: separate foreground from background

- Graph structure:

- ☐ pairwise Markov net
- ☐ grid with one node per pixel



- Node potential:

- ☐ "background color" v. "foreground color"

- Edge potential:

- ☐ neighbors like to be of the same class

10-708 - ©Carlos Guestrin 2006

32

Logarithmic representation

- Standard model: $P(X_1, \dots, X_n) = \frac{1}{Z} \prod_{i=1}^m \pi_i(D_i)$
- Log representation of potential (assuming positive potential):
 - also called the energy function
- Log representation of Markov net:

10-708 – ©Carlos Guestrin 2006

33

Log-linear Markov network (most common representation)


- **Feature** is some function $\phi[D]$ for some subset of variables D
 - e.g., indicator function
- **Log-linear model** over a Markov network H :
 - a set of features $\phi_1[D_1], \dots, \phi_k[D_k]$
 - each D_i is a subset of a clique in H
 - two ϕ 's can be over the same variables
 - a set of weights $w_1, \dots, w_k$
 - usually learned from data
 - $P(X_1, \dots, X_n) = \frac{1}{Z} \exp \left[\sum_{i=1}^k w_i \phi_i(D_i) \right]$

10-708 – ©Carlos Guestrin 2006

34

Structure in cliques

- Possible potentials for this graph:




10-708 – ©Carlos Guestrin 2006

35

Factor graphs

- Very useful for approximate inference
 - Make factor dependency explicit
- Bipartite graph:
 - variable nodes (ovals) for $X_1, \dots, X_n$
 - factor nodes (squares) for $\phi_1, \dots, \phi_m$
 - edge $X_i - \phi_j$ if $X_i \in \text{Scope}[\phi_j]$



10-708 – ©Carlos Guestrin 2006

36

Summary of types of Markov nets

- Pairwise Markov networks
 - very common
 - potentials over nodes and edges
- Log-linear models
 - log representation of potentials
 - linear coefficients learned from data
 - most common for learning MNs
- Factor graphs
 - explicit representation of factors
 - you know exactly what factors you have
 - very useful for approximate inference