HW1 due tonight
Quiz posted soon
HW 2 posted soon, due ~Wed 10/16

COMPSCI 688: Probabilistic Graphical Models

Lecture 8: Undirected Graphical Models: Inference

Dan Sheldon

Manning College of Information and Computer Sciences
University of Massachusetts Amherst

Partially based on materials by Benjamin M. Marlin (marlin@cs.umass.edu) and Justin Domke (domke@cs.umass.edu)

Review

Markov Random Fields

$p(x_1,x_2,x_3,x_4) = \phi_{123}(x_1,x_2,x_3)\,\phi_{234}(x_2,x_3,x_4)\cdot\frac{1}{Z}$

- Markov random field

$$p(\mathbf{x}) = \frac{1}{Z}\prod_{c\in\mathcal{C}}\phi_c(\mathbf{x}_c)$$

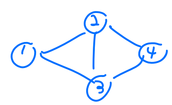

- *Dependence graph* $\mathcal{G}$
  - nodes i and j connected by an edge if they appear together in some factor
  - $\mathbf{X}_A \perp \mathbf{X}_B | \mathbf{X}_S$ if S separates A from B in $\mathcal{G}$

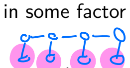

- Examples: Ising model, conditional random fields (images and text), models for differential privacy, Bayes nets
- Today: **inference** = answering probability queries

Inference: Conditioning

Inference in Markov Networks

$p(x) = p(x_Q, x_U, x_E) = \frac{1}{Z}\boxed{\prod_c \phi_c(x_c)}$

- Given a Markov network, the main task is *probabilistic inference*, which means answering probability queries of the form

$$p(\mathbf{x}_Q|\mathbf{x}_E) = \sum_{\mathbf{x}_U} p(\mathbf{x}_Q, \mathbf{x}_U|\mathbf{x}_E)$$

  - condition on *evidence variables* $\mathbf{x}_E$
  - marginalize *unobserved variables* $\mathbf{x}_U$
  - compute the joint distribution over *query variables* $\mathbf{x}_Q$

$p(x_Q|x_E) = \frac{1}{Z'}\prod \phi'_c(x_c)$

- …often by transforming Markov network into one with fewer or simpler factors
  - Conditioning is easy
  - Marginalization is hard!

Conditioning: Single Factor

Suppose we have a single-factor MRF $p(x_1, x_2) = \frac{1}{Z}\phi(x_1, x_2)$ for two binary variables. We are given a fixed value for x_2, and want an MRF for $p(x_1|x_2)$, i.e.:

$$p(x_1|x_2) = \frac{1}{Z'}\phi'(x_1)$$

Observe

$$p(x_1|x_2) = \frac{p(x_1, x_2)}{p(x_2)} = \underbrace{\frac{1}{p(x_2)}\frac{1}{Z}}_{\frac{1}{Z'}}\underbrace{\phi(x_1, x_2)}_{\phi'(x_1)}$$

For fixed x_2, the conditional $p(x_1|x_2)$ is *proportional* to the joint $p(x_1, x_2)$. We can use the same factor, but hard-code x_2 so that only x_1 is a free variable:

$$\phi'(x_1) = \phi(x_1, x_2), \quad Z' = p(x_2)Z$$

x_1	x_2	ϕ	$p(x_1,x_2)$
0	0	1	.1
0	1	4	.4
1	0	3	.3
1	1	2	.2

$Z=10$

Observe $x_2=0$

x_1	ϕ'	$p(x_1\|x_2=0)$
0	1	.25
1	3	.75

$Z'=4$

Conditioning: General Case

For a general MRF, we can apply the same reasoning to *reduce* every factor by hard-coding the evidence variables

Factor Reduction: Example

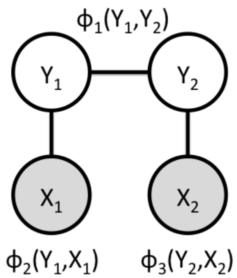

$\phi_1(Y_1,Y_2)$	$Y_2=0$	$Y_2=1$
$Y_1=0$	1	2
$Y_1=1$	7	2

$\phi_2(Y_1,X_1)$	$X_1=0$	$X_1=1$
$Y_1=0$	3	9
$Y_1=1$	4	1

$\phi_3(Y_2,X_2)$	$X_2=0$	$X_2=1$
$Y_2=0$	6	2
$Y_2=1$	2	7

Query: $P(Y_1,Y_2 \mid X_1=0, X_2=1)$

Factor Reduction: Step 1

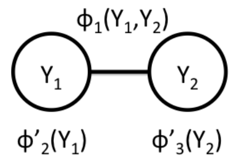

$\phi_1(Y_1,Y_2)$	$Y_2=0$	$Y_2=1$
$Y_1=0$	1	2
$Y_1=1$	7	2

$\phi_2(Y_1,X_1)$	$X_1=0$	$X_1=1$
$Y_1=0$	3	9
$Y_1=1$	4	1

$\phi_3(Y_2,X_2)$	$X_2=0$	$X_2=1$
$Y_2=0$	6	2
$Y_2=1$	2	7

Query: $P(Y_1,Y_2 \mid X_1=0, X_2=1)$

Factor Reduction: Step 2

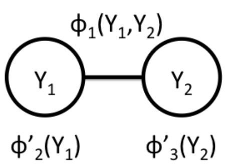

$\phi_1(Y_1,Y_2)$	$Y_2=0$	$Y_2=1$
$Y_1=0$	1	2
$Y_1=1$	7	2

$\phi_2(Y_1,X_1)$	$X_1=0$	$X_1=1$
$Y_1=0$	3	9
$Y_1=1$	4	1

$\phi_3(Y_2,X_2)$	$X_2=0$	$X_2=1$
$Y_2=0$	6	2
$Y_2=1$	2	7

Query: $P(Y_1,Y_2 \mid X_1=0, X_2=1)$

Factor Reduction: Step 2

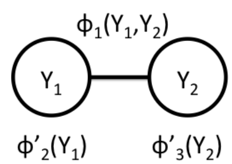

$\phi_1(Y_1,Y_2)$	$Y_2=0$	$Y_2=1$
$Y_1=0$	1	2
$Y_1=1$	7	2

	$\phi'_2(Y_1)$
$Y_1=0$	3
$Y_1=1$	4

	$\phi'_3(Y_2)$
$Y_2=0$	2
$Y_2=1$	7

Query: $P(Y_1,Y_2 \mid X_1=0, X_2=1) \propto \phi_1(Y_1,Y_2)\, \phi'_2(Y_1)\, \phi'_3(Y_2)$

Factor Reduction: General Algorithm

Suppose $p(\mathbf{x}) = \frac{1}{Z}\prod_{c\in\mathcal{C}}\phi_c(\mathbf{x}_c)$ and we observe $X_i = x_i$ for a single node i

We obtain a new MRF for $p(\mathbf{x}_{-i}|x_i)$ by the following procedure:

For each factor ϕ_c such that $i \in c$

- Replace $\phi_c(\mathbf{x}_c)$ by $\phi'_{c\backslash i}(\mathbf{x}_{c\backslash i}) := \phi_c(\overset{\text{free}}{\mathbf{x}_{c\backslash i}}, \overset{\text{hard-coded}}{x_i})$
- The $\mathbf{x}_{c\backslash i}$ variables remain "free", and x_i is hard-coded

To condition on many variables, we can repeat this procedure. Since order doesn't matter, we can hard-code all evidence variables in each factor at the same time.

Inference: Marginalization

Marginalization

Marginalization is the process of summing over some of the variables to get the marginal distribution of the remaining variables, or the partition function.

For example, the partition function is

$$Z = \overset{2\times 2\times 2\cdots 2 \;=\; 2^n}{\sum_{x_1}\sum_{x_2}\cdots\sum_{x_n}}\prod_{c\in\mathcal{C}}\phi_c(x_c)$$

Naively, this takes exponential time, but we can sometimes use the factorization structure to speed it up.

Example: Variable Elimination on a Chain

Consider the following MRF on a four-node "chain" graph:

$$p(x_1,x_2,x_3,x_4) = \frac{1}{Z}\phi_1(x_1)\phi_2(x_2)\phi_3(x_3)\phi_4(x_4)\phi_{12}(x_1,x_2)\phi_{23}(x_2,x_3)\phi_{34}(x_3,x_4)$$

x_i	$\phi_i(x_i)$
0	1
1	2

x_i	x_j	$\phi_{ij}(x_i,x_j)$
0	0	2
0	1	1
1	0	1
1	1	2

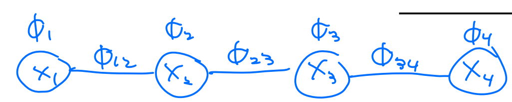

Let's compute Z:

$$= \sum_{x_1}\sum_{x_2}\sum_{x_3}\sum_{x_4} \phi_1(x_1)\phi_2(x_2)\phi_3(x_3)\phi_4(x_4)\phi_{12}(x_1,x_2)\phi_{23}(x_2,x_3)\phi_{34}(x_3,x_4)$$

$$= \sum_{x_1}\sum_{x_2}\sum_{x_3} \phi_1(x_1)\phi_2(x_2)\phi_3(x_3)\phi_{12}(x_1,x_2)\phi_{23}(x_2,x_3)\underbrace{\sum_{x_4}\overbrace{\phi_4(x_4)\phi_{34}(x_3,x_4)}^{2\times 2}}_{m_{4\to3}(x_3)}$$

$$= \sum_{x_1}\sum_{x_2} \phi_1(x_1)\phi_2(x_2)\phi_{12}(x_1,x_2)\underbrace{\sum_{x_3}\overbrace{\phi_3(x_3)\phi_{23}(x_2,x_3)m_{4\to3}(x_3)}^{2\times 2}}_{m_{3\to2}(x_2)}$$

$$= \sum_{x_1}\phi_1(x_1)\underbrace{\sum_{x_2}\overbrace{\phi_2(x_2)\phi_{12}(x_1,x_2)m_{3\to2}(x_2)}^{2\times 2}}_{m_{2\to1}(x_1)}$$

$$= \sum_{x_1}\phi_1(x_1)\,m_{2\to1}(x_1)$$

$$= \sum$$

Pictorially, this is how we changed the MRF

We eliminated x_4, x_3, x_2, x_1

x_2	x_3	$\phi_3 \cdot \phi_{23} \cdot m_{4\to3}$
0	0	$1\cdot2\cdot4$
0	1	$2\cdot1\cdot5$
1	0	$1\cdot1\cdot4$
1	1	$3\cdot2\cdot5$

x_3	x_4	$\phi_4 \cdot \phi_{34}$
0	0	$1\cdot2=2$
0	1	$2\cdot1=2$
1	0	$1\cdot1=1$
1	1	$2\cdot2=4$

x_3	$m_{4\to3}(x_3) = \sum_{x_4}\phi_4\cdot\phi_{34}$
0	4
1	5

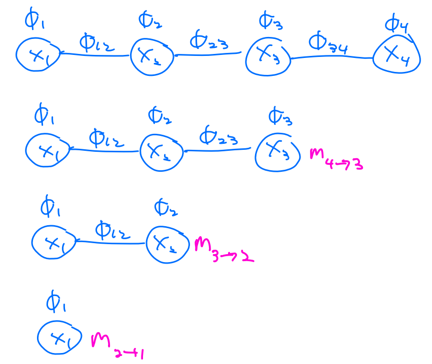

What if we want to compute the *unnormalized* marginal $\hat{p}(x_1)$?

$$\hat{p}(x_1) = \sum_{x_2}\sum_{x_3}\sum_{x_4} \phi_1(x_1)\phi_2(x_2)\phi_3(x_3)\phi_4(x_4)\phi_{12}(x_1,x_2)\phi_{23}(x_2,x_3)\phi_{34}(x_3,x_4)$$

$$= \cdots$$

$$= \phi_1(x_1)\,m_{2\to1}(x_1)$$

$$Z = \sum_{x_1}\hat{p}(x_1)$$

What if we want to compute the *actual* marginal $p(x_1)$?

Take $\hat{p}(x_1)$ and normalize it

$$Z = \sum_{x_1}\hat{p}(x_1), \quad p(x_1) = \frac{1}{Z}\hat{p}(x_1)$$

Lesson: always normalize at the end

x_1	$\hat{p}(x_1)$	$p(x_1)$
0	3	3/20
1	17	17/20
	20	

What if we eliminate x_3 first?

$$= \sum_{x_1}\sum_{x_2}\sum_{x_4} \phi_1(x_1)\phi_2(x_2)\phi_4(x_4)\phi_{12}(x_1,x_2)\underbrace{\overbrace{\sum_{x_3}\phi_3(x_3)\phi_{23}(x_2,x_3)\phi_{34}(x_3,x_4)}^{2\times2\times2}}_{\tau_{24}(x_2,x_4)}$$

$$= \sum_{x_1}\sum_{x_2}\sum_{x_4} \phi_1(x_1)\phi_2(x_2)\phi_4(x_4)\phi_{12}(x_1,x_2)\,\tau_{24}(x_2,x_4)$$

$$= \dots$$

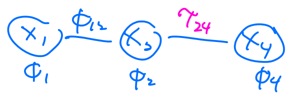

Correct, but less efficient due to larger intermediate factors

What if our graph is a star graph? $Z = \sum_{x_0}\sum_{x_1}\dots\sum_{x_5}\phi_{01}(x_0,x_1)\cdot\dots\cdot\phi_{05}(x_0,x_5)$

$$= \sum_{x_1}\sum_{x_2}\dots\sum_{x_5}\underbrace{\overbrace{\sum_{x_0}\phi_{01}(x_0,x_1)\cdot\dots\cdot\phi_{05}(x_0,x_5)}^{2\times2\times2\times2\times2\times2}}_{\tau(x_1,x_2,x_3,x_4,x_5)}$$

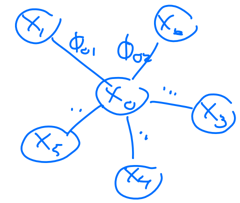

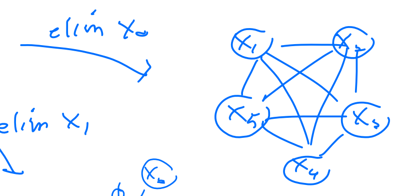

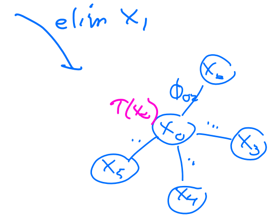

The Variable Elimination Algorithm

Variable elimination is an algorithm to compute any marginal distribution in any MRF

In words: pick a variable x_i to eliminate, multiply together all factors containing x_i to get an intermediate factor, then sum out x_i

- Let $F = \{\phi_c : c \in \mathcal{C}\}$ be the set of factors
- For each variable i in **some elimination order** (may not include all variables)
  - Let $A = \{\phi_c \in F : i \in c\}$ be the set of factors whose scope contains i
  - Let $\phi_a(\mathbf{x}_a) = \prod_{\phi_c \in A} \phi_c(\mathbf{x}_c)$ be the product of factors in A, with scope a equal to the union of the scopes of the individual factors
  - Let $\psi_i(\mathbf{x}_{a\setminus i}) = \sum_{x_i} \phi_a(\mathbf{x}_{a\setminus i}, x_i)$ be the result of summing out x_i
  - Let $F = F \setminus A \cup \{\psi_i\}$

The final set of factors forms an MRF for the marginal distribution of the variables that were not eliminated.

Variable Elimination Discussion

- The efficiency of variable elimination depends on the maximum size of the intermediate factors created, which depends on the elimination ordering
  - Inference in MRFs is NP-hard, so we can't always find a good elimination ordering.
  - Finding the best elimination ordering for a given MRF is also NP-hard!
- It's always efficient to eliminate leaves if present (intermediate factors are no larger than original ones)
  - $\implies$ for trees, we can find an efficient elimination ordering
  - In fact, because the elimination ordering is predictable in trees, we can realize extra efficiencies when answering multiple queries through a dynamic programming approach known as **message passing**