Markov Chains

Lecture #21

11/13/08

Announcements

- Programming Assignment #3 is due Dec. 11th.
  - – not as much time as it seems, given Thanksgiving break.
- Any questions?

Review: Inference on a chain

Summary:

$$\begin{aligned} P(D) &= \sum_{A,B,C} P(A,B,C,D) \\ &= \sum_{A,B,C} P(D|C)P(C|B)P(B|A)P(A) \\ &= \sum_{C} P(D|C) \sum_{B} P(C|B) \sum_{A} P(B|A)P(A) \end{aligned}$$

Review: multiple orders work

$$P(D) = \sum_{A,B,C} P(A)P(B \mid A)P(C \mid B)P(D \mid C)$$
$$P(D) = \sum_{A} P(A) \sum_{B} P(B \mid A) \sum_{C} P(C \mid B)P(D \mid C)$$
$$P(D) = \sum_{A} P(A) \sum_{B} P(B \mid A) f_c(B,D)$$
$$P(D) = \sum_{A} P(A) f_B(A,D)$$

Variable Elimination

- Suppose we're interested in $P(X_k)$
- Write query in the form

$$P(X_k) = \sum_{X_n,\ldots,X_{k+1},X_{k-1},\ldots,X_1} \prod_i P(X_i|Pa(X_i))$$

- Iteratively
  - – Move all irrelevant terms outside of innermost sum
  - – Perform innermost sum, getting a new term
  - – Insert the new term into the product

Dealing with Evidence

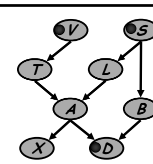

- How do we deal with evidence?
- Suppose get evidence $V = t, S = f, D = t$ and want to compute $P(L|V = t, S = f, D = t)$

$$P(L \mid V=t, S=f, D=t) = \frac{P(L, V=t, S=f, D=t)}{P(V=t, S=f, D=t)}$$

Dealing with Evidence

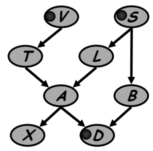

- We start by writing the factors:

$$\boxed{P(v)P(s)P(t \mid v)P(l \mid s)P(b \mid s)P(a \mid t,l)P(x \mid a)P(d \mid a,b)}$$

- Since we know that $V = t$, we don't need to eliminate V
- Instead, we can replace the factors $P(V)$ and $P(T|V)$ with

$$\boxed{f_{P(V)} = P(V = t) \quad f_{p(T|V)}(T) = P(T \mid V = t)}$$

- These "select" the appropriate parts of the original factors given the evidence

Markov chain Monte Carlo sampling

- Generates events by making random changes to the state variable.
- The next state is generated by sampling a value for one of the nonevidence variables conditioned on the current values.

Markov chains

- A Markov chain is a random process (infinite sequence of random variables)

  (X(0), X(1),…,X(t),…) that satisfies:

  P(X(t) | X(0),…X(t-1)) = P(X(t)|X(t-1))
- The probability of a particular state at time *t* depends only on the state at time *t*-1
- If the transition probabilities are fixed for all *t*, the chain is called *homogeneous* and is characterized by a transition matrix T.

$$T = \begin{pmatrix} 0.7 & 0.3 & 0 \\ 0.3 & 0.4 & 0.3 \\ 0 & 0.3 & 0.7 \end{pmatrix}$$

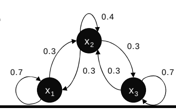

Markov chains

- In order for a Markov chain to be useful for sampling from *P(x)*, we require that for any starting state *x(0):*

$$\lim_{t \to \infty} P_t(x) = P(x)$$

- Equivalently, the stationary distribution of the Markov chain must be *P(x):*

$$P_{t+1}(\mathbf{x}') = \sum_{\mathbf{x}} P_t(\mathbf{x}) Q(\mathbf{x} \to \mathbf{x}')$$

the transition probability to *x'*

$$P(\mathbf{x}') = \sum_{\mathbf{x}} P(\mathbf{x}) Q(\mathbf{x} \to \mathbf{x}')$$

Using a Markov chain to sample

- If the Markov chain indeed converges to the desired distribution from, we can start in an arbitrary state, use the Markov chain to do a random walk for a while, and output the *x(t).*
- The resulting state will be sampled from *P(x)*.

Stationary distribution

$$Q = \begin{pmatrix} 0.7 & 0.3 & 0 \\ 0.3 & 0.4 & 0.3 \\ 0 & 0.3 & 0.7 \end{pmatrix}$$

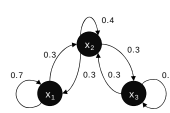

- The stationary distribution of this chain is (0.33,0.33,0.33)

Markov chains for sampling

- To ensure that the chain converges to a unique stationary distribution the following conditions are sufficient:
  - *Irreducibility*: every state is eventually reachable from any start state; for all x,y there exists a t such that $P_t(y) > 0$ when starting at x
  - *Aperiodicity*: the chain doesn't get caught in cycles.
- The process is *ergodic* if it is both irreducible and aperiodic

Detailed balance

- To ensure that the stationary distribution of the Markov chain is $P(x)$ it is sufficient for P and Q to satisfy the *detailed balance* (*reversibility*) condition:

$$P(\mathbf{x})Q(\mathbf{x} \to \mathbf{x}') = P(\mathbf{x}')Q(\mathbf{x}' \to \mathbf{x})$$

Given that detailed balance holds:

$$\begin{aligned}\sum_{\mathbf{x}} P(\mathbf{x})Q(\mathbf{x} \to \mathbf{x}') &= \sum_{\mathbf{x}} P(\mathbf{x}')Q(\mathbf{x}' \to \mathbf{x}) \\ &= P(\mathbf{x}) \sum_{\mathbf{x}} Q(\mathbf{x}' \to \mathbf{x}) \\ &= P(\mathbf{x}')\end{aligned}$$

Gibbs sampling

- Idea: To transition from one state (variable assignment) to another by:
  - Pick a variable X_j,
  - Sample its value from the conditional distribution

  $P(x_j \mid x_1,\ldots x_{j-1},x_{j+1},\ldots x_n)$
- In a Bayesian network x_j depends only on a subset of the variables.

Markov Blanket

- Variables are independent of their non-descendents given their parents
- Variables are independent of *everything else in the network* given their *Markov blanket*.
- So, to sample a node, only need to condition on its Markov blanket:

$P(x_j \mid x_1,\ldots x_{j-1},x_{j+1},\ldots x_n) =$
$P(x_j \mid MB(x_j))$

The Gibbs sampling algorithm

GIBBS$(X, \boldsymbol{e}, bn, N)$ returns estimate of $P(X|\boldsymbol{e})$

$\boldsymbol{N}[x]$ - counts the number of times each value of X was observed

$\boldsymbol{x}[j]$ - the current state of the network $\boldsymbol{x}[0]$ initialized with random values for the nonevidence variables

for $j = 1$ to N do

  for each nonevidence variable X_i

    sample X_i from $P(X_i|MB(X_i))$

    $\boldsymbol{N}[x] = \boldsymbol{N}[x] + 1$, where x is the value of X in $\boldsymbol{x}[j]$

Convergence of Gibbs sampling

- Gibbs sampling satisfies detailed balance:

$$\begin{aligned}P(\mathbf{x}|\mathbf{e})P(x_i'|\bar{\mathbf{x}}_i,\mathbf{e}) &= P(x_i,\bar{\mathbf{x}}_i|\mathbf{e})P(x_i'|\bar{\mathbf{x}}_i,\mathbf{e}) \\ &= P(x_i|\bar{\mathbf{x}}_i,\mathbf{e})P(\bar{\mathbf{x}}_i,\mathbf{e})P(x_i'|\bar{\mathbf{x}}_i,\mathbf{e}) \\ &= P(x_i|\bar{\mathbf{x}}_i,\mathbf{e})P(x_i',\bar{\mathbf{x}}_i|\mathbf{e}) \\ &= P(\mathbf{x}'|\mathbf{e})P(x_i|\bar{\mathbf{x}}_i',\mathbf{e})\end{aligned}$$

$\bar{\mathbf{x}}_i$ denotes the variables other than $\boldsymbol{x}_i$

Gibbs sampling example

- Consider a 2 variable network:

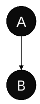

	a	-a
	0.5	0.5

	b	-b
a	0.8	0.2
-a	0.2	0.8

- Initialize randomly
- Sample variables alternately

Practical issues

- How many iterations?
- When to stop?