Lecture 4:

Probability Models

Sam Roweis

January 16, 2002

What's Inside the Nodes?

- We've focused a lot on the strucutre of the graphs in directed and undirected models. Today we'll look at specific functions that can live inside the nodes.
- For directed models we need functions $p(\mathbf{x}_i)$ for root nodes and $p(\mathbf{x}_i|\mathbf{x}_{\pi_i})$ for interior nodes.
- For undirected models we need clique potentials $\psi_C(\mathbf{x}_C)$ or energies $H_C(\mathbf{x}_C)$.
- We'll consider various types of nodes: binary/discrete (categorical), continuous, interval, and integer counts.
- For each type we'll see some basic *probability models* which are parametrized families of distributions.

Outline

- Inside Nodes
- (Conditional) Probability Tables
- Exponential Family Distributions
- Generalized Linear Models and Response Functions
- Basic Statistical Operations

Probability Tables & CPTs

- For discrete (categorical) nodes, the most basic parametrization is the probability table which lists $p(\mathbf{x}_i = k^{th} \text{ value})$.
- Since PTs must be nonnegative and sum to 1, for k-ary nodes there are $k-1$ free parameters.
- If a discrete node has discrete parent(s) we make one table for each setting of the parents: this is a *conditional probability table* or CPT.

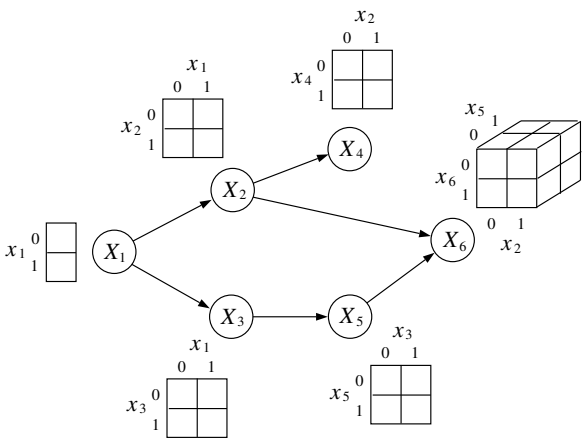

Exponential Family

- For a numeric random variable $\mathbf{x}$

$$p(\mathbf{x}|\eta) = h(\mathbf{x}) \exp\{\eta^\top T(\mathbf{x}) - A(\eta)\}$$
$$= \frac{1}{Z(\eta)} h(\mathbf{x}) \exp\{\eta^\top T(\mathbf{x})\}$$

  is an exponential family distribution with *natural parameter* η.

- Function $T(\mathbf{x})$ is a *sufficient statistic*.
- Function $A(\eta) = \log Z(\eta)$ is the log normalizer.
- Key idea: all you need to know about the data is captured in the summarizing function $T(\mathbf{x})$.

Poisson

- For an integer count variable with rate λ:

$$p(x|\lambda) = \frac{\lambda^x e^{-\lambda}}{x!}$$
$$= \frac{1}{x!} \exp\{x \log \lambda - \lambda\}$$

- Exponential family with:

$$\eta = \log \lambda$$
$$T(x) = x$$
$$A(\eta) = \lambda = e^\eta$$
$$h(x) = \frac{1}{x!}$$

- e.g. number of photons $\mathbf{x}$ that arrive at a pixel during a fixed interval given mean intensity λ
- Other count densities: binomial, exponential.

Bernoulli

- For a binary random variable with p(heads)=π:

$$p(x|\pi) = \pi^x (1-\pi)^{1-x}$$
$$= \exp\left\{ \log\left(\frac{\pi}{1-\pi}\right) x + \log(1-\pi) \right\}$$

- Exponential family with:

$$\eta = \log \frac{\pi}{1-\pi}$$
$$T(x) = x$$
$$A(\eta) = -\log(1-\pi) = \log(1+e^\eta)$$
$$h(x) = 1$$

- The logistic function relates the natural parameter and the chance of heads

$$\pi = \frac{1}{1+e^{-\eta}}$$

Multinomial

- For a set of integer counts on k trials

$$p(\mathbf{x}|\pi) = \frac{k!}{x_1! x_2! \cdots x_n!} \pi_1^{x_1} \pi_2^{x_2} \cdots \pi_n^{x_n}$$
$$= h(\mathbf{x}) \exp\left\{ \sum_i x_i \log \pi_i \right\}$$

- But the parameters are constrained: $\sum_i \pi_i = 1$. So we define the last one $\pi_n = 1 - \sum_{i=1}^{n-1} \pi_i$.

$$p(\mathbf{x}|\pi) = h(\mathbf{x}) \exp\left\{ \sum_{i=1}^{n-1} \log\left(\frac{\pi_i}{\pi_n}\right) x_i + k \log \pi_n \right\}$$

- Exponential family with:

$$\eta_i = \log \pi_i - \log \pi_n$$
$$T(x_i) = x_i$$
$$A(\eta) = -k \log \pi_n = k \log \sum_i e^{\eta_i}$$
$$h(\mathbf{x}) = k! / x_1! x_2! \cdots x_n!$$

- The *softmax* function relates the basic parameters to the natural ones:

$$\pi_i = \frac{e^{\eta_i}}{\sum_j e^{\eta_j}}$$

Gaussian (normal)

- For a continuous univariate random variable:

$$p(x|\mu,\sigma^2) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left\{-\frac{1}{2\sigma^2}(x-\mu)^2\right\}$$
$$= \frac{1}{\sqrt{2\pi}\sigma}\exp\left\{\frac{\mu x}{\sigma^2} - \frac{x^2}{2\sigma^2} - \frac{\mu^2}{2\sigma^2} - \log\sigma\right\}$$

- Exponential family with:

$$\eta = [\mu/\sigma^2\ ;\ -1/2\sigma^2]$$
$$T(x) = [x\ ;\ x^2]$$
$$A(\eta) = \log\sigma + \mu/2\sigma^2$$
$$h(x) = 1/\sqrt{2\pi}$$

- Note: a univariate Gaussian is a two-parameter distribution with a two-component vector of sufficient statistis.
- Other densities: Student-t, Laplacian.
- Sometimes we want only positive values: exponential, Gamma, log-normal.

Moments

- For numeric nodes, moment calculations are important.
- We can easily compute moments of any exponential family distribution by taking the derivatives of the log normalizer $A(\eta)$.
- The q^{th} derivative gives the q^{th} centred moment.

$$\frac{dA(\eta)}{d\eta} = \text{mean}$$
$$\frac{d^2A(\eta)}{d\eta^2} = \text{variance}$$
$$\ldots$$

- When the sufficient statistic is a vector, partial derivatives need to be considered.

Multivariate Gaussian Distribution

- For a continuous vector random variable:

$$p(x|\mu,\Sigma) = |2\pi\Sigma|^{-1/2}\exp\left\{-\frac{1}{2}(\mathbf{x}-\mu)^\top\Sigma^{-1}(\mathbf{x}-\mu)\right\}$$

- Exponential family with:

$$\eta = [\Sigma^{-1}\mu\ ;\ -1/2\Sigma^{-1}]$$
$$T(x) = [\mathbf{x}\ ;\ \mathbf{x}\mathbf{x}^\top]$$
$$A(\eta) = \log|\Sigma|/2 + \mu^\top\Sigma^{-1}\mu/2$$
$$h(x) = (2\pi)^{-n/2}$$

- Note: a univariate Gaussian is a two-parameter distribution with a two-component vector of sufficient statistis.

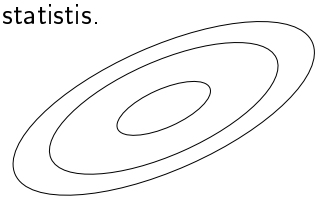

- You should know how to manipulate these, and condition on subsets of variables given others. Mostly linear algebra.

Nodes with Parents

- When the parent is discrete , we just have one density for each setting of the parent.
  e.g. table of natural parameters (exponential models) or table of tables (CPTs).
- When the parent is numeric, its value sets some of the parameters for the child node.
- A very common instance of this for regression is the "linear-Gaussian": $p(\mathbf{y}|\mathbf{x}) = \text{gauss}(\theta^\top\mathbf{x};\Sigma)$.
- For classification, often use bernoulli/multinomial whose paramters are some function $f(\theta^\top\mathbf{x})$.
- Generalized Linear Models: $p(\mathbf{y}|\mathbf{x})$ is exponential family with conditional mean $\mu = f(\theta^\top\mathbf{x})$.
- The function f is called the *response function*.
- If we chose f to be the inverse of the mapping b/w conditional mean and natural parameters then it is called the *canonical response function*.

$$\eta = \psi(\mu)$$
$$f(\cdot) = \psi^{-1}(\cdot)$$

Potential Functions

- We are much less constrained with potential functions, since they can be any positive function of the values of the clique nodes.
- Recall $\psi_C(\mathbf{x}_C) = \exp\{-H_C(\mathbf{x}_C)\}$
- A common choice for cliques which are pairs is:

$$H(\mathbf{x}) = \sum_i a_i \mathbf{x}_i + \sum_{\text{pairs } ij} w_{ij} \mathbf{x}_i \mathbf{x}_j$$

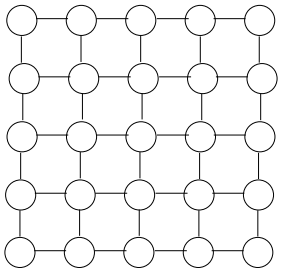

Fundamental Operations

- What can we do with a graphical model distribution?
- *Generate data.*
  For this you need to know how to sample from local models (directed) or how to do Gibbs or other sampling (undirected).
- *Compute log probabilities.*
  When all nodes are either observed or marginalized the result is a single number which is the log prob of the configuration.
- *Inference.*
  Compute expectations of some nodes given others which are observed or marginalized.
- *Learning.*
  Set the parameters of the local functions given some (partially) observed data.

Basic Statistical Problems

- Let's remind ourselves of the basic problems we discussed on the first day: *density estimation*, *clustering* *classification* and *regression*.
- Can always do joint density estimation and then condition to get what we want:
  Regression: $p(\mathbf{y}|\mathbf{x}) = p(\mathbf{y}, \mathbf{x})/p(\mathbf{x})$
  Classification: $p(c|\mathbf{x}) = p(c, \mathbf{x})/p(\mathbf{x})$
  Clustering: $p(c|\mathbf{x}) = p(c, \mathbf{x})/p(\mathbf{x})$ c unobserved

In general, if certain nodes are *always* observed we may not want to model their density:

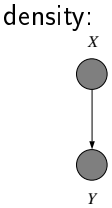

Regression
Classification

If certain nodes are *always* unobserved they are called *hidden* or *latent* variables (more later):

Clustering