School of Computer Science
Carnegie Mellon

Probabilistic Graphical Models

Learning one-node GM

Eric Xing

Lecture 5, September 23, 2009

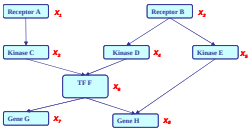

Reading:

Learning Graphical Models

The goal:

Given set of independent samples (***assignments*** of random variables), find the ***best*** (the most likely?) Bayesian Network (both DAG and CPDs)

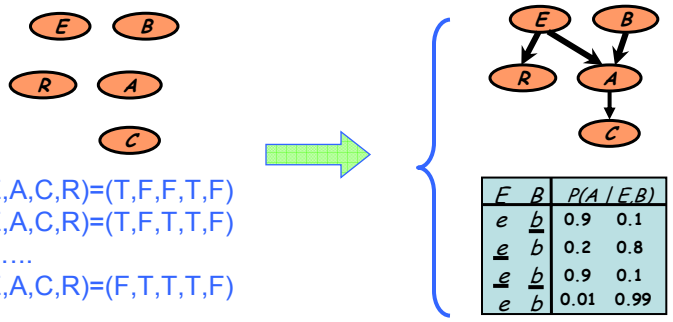

E	B	P(A | E,B)	
e	b	0.9	0.1
e	b	0.2	0.8
e	b	0.9	0.1
e	b	0.01	0.99

Learning Graphical Models

- Scenarios:
  - completely observed GMs
    - directed
    - undirected
  - partially or unobserved GMs
    - directed
    - undirected (an open research topic)
- Estimation principles:
  - Maximal likelihood estimation (MLE)
  - Bayesian estimation
  - Maximal conditional likelihood
  - Maximal "Margin"
  - Maximum entropy
- We use **learning** as a name for the process of estimating the parameters, and in some cases, the topology of the network, from data.

Score-based approach

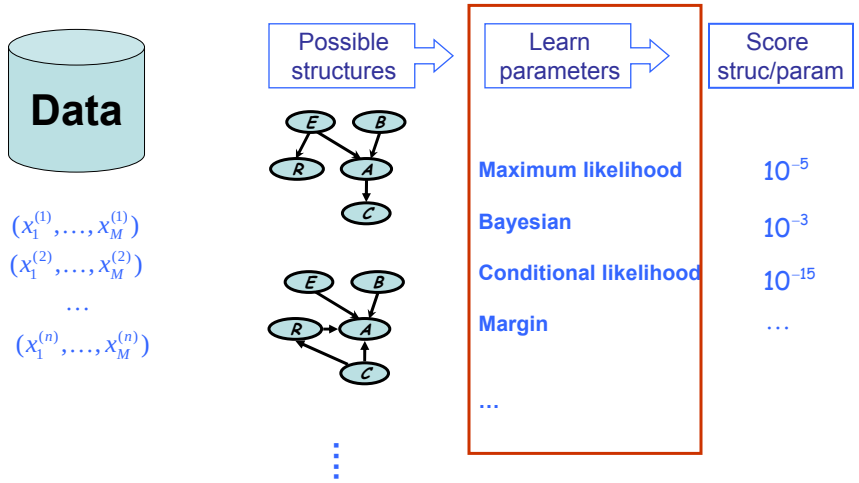

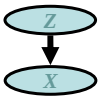

ML Parameter Est. for completely observed GMs of given structure

- The data:

  $\{(z^{(1)},x^{(1)}), (z^{(2)},x^{(2)}), (z^{(3)},x^{(3)}), \ldots (z^{(N)},x^{(N)})\}$

Parameter Learning

- Assume G is known and fixed,
  - from expert design
  - from an intermediate outcome of iterative structure learning
- Goal: estimate from a dataset of N independent, identically distributed (*iid*) training cases $D = \{x^{(1)}, \ldots, x^{(N)}\}$.
- In general, each training case $\mathbf{x}^{(n)} = (x_1^{(n)}, \ldots, x_M^{(n)})$ is a vector of M values, one per node,
  - the model can be completely observable, i.e., every element in x_n is known (no missing values, no hidden variables),
  - or, partially observable, i.e., $\exists i$, s.t. $x_i^{(n)}$ is not observed.
- **In this lecture we consider learning parameters for a single node.**

  -- Often known as "density estimation"

Density Estimation

- A Density Estimator learns a mapping from a set of attributes to a Probability

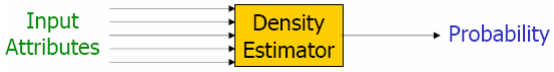

- Often know as *parameter estimation* if the distribution form is specified
  - Binomial, Gaussian …
- Four important issues:
  - Nature of the data (iid, correlated, …)
  - Objective function (MLE, MAP, Margin …)
  - Algorithm (simple algebra, gradient methods, EM, …)
  - Evaluation scheme (likelihood on test data, predictability, consistency, …)

Discrete Distributions

- Bernoulli distribution: Ber(p)

$$P(x) = \begin{cases} 1-p & \text{for } x = 0 \\ p & \text{for } x = 1 \end{cases} \quad \Rightarrow \quad P(x) = p^x (1-p)^{1-x}$$

- Multinomial distribution: Mult(1, θ)
  - Multinomial (indicator) variable:

$$X = \begin{bmatrix} X_1 \\ X_2 \\ X_3 \\ X_4 \\ X_5 \\ X_6 \end{bmatrix}, \quad \text{where} \quad X_j = [0,1], \text{ and } \sum_{j\in\{1,\ldots,6\}} X_j = 1$$

$$X_j = 1 \text{ w.p. } \theta_j, \quad \sum_{j\in\{1,\ldots,6\}} \theta_j = 1 \ .$$

$$p(x(j)) = P(\{X_j = 1, \text{where } j \text{ index the dice-face}\})$$
$$= \theta_j = \theta_A^{x_A} \times \theta_C^{x_C} \times \theta_G^{x_G} \times \theta_T^{x_T} = \prod_k \theta_k^{x_k} = \theta^x$$

Continuous Distributions

- Uniform Probability Density Function

$$p(x) = 1/(b-a) \quad \text{for } a \le x \le b$$
$$= 0 \quad \text{elsewhere}$$

- Normal (Gaussian) Probability Density Function

$$p(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/2\sigma^2}$$

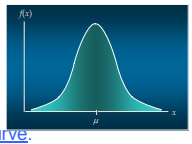

  - The distribution is symmetric, and is often illustrated as a bell-shaped curve.
  - Two parameters, μ (mean) and σ (standard deviation), determine the location and shape of the distribution.
  - The highest point on the normal curve is at the mean, which is also
  - The mean can be any numerical value: negative, zero, or positive.

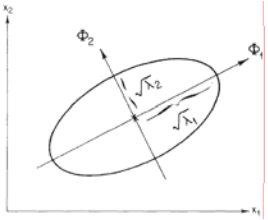

- Multivariate Gaussian

$$p(X;\vec{\mu},\Sigma) = \frac{1}{\left(\sqrt{2\pi}\right)^{n/2}|\Sigma|^{1/2}} \exp\left\{-\frac{1}{2}(X-\vec{\mu})^T\Sigma^{-1}(X-\vec{\mu})\right\}$$

Density Estimation Schemes

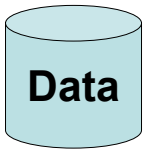

$(x_1^{(1)},\ldots,x_M^{(1)})$
$(x_1^{(2)},\ldots,x_M^{(2)})$
…
$(x_1^{(n)},\ldots,x_M^{(n)})$

Learn parameters → Algorithm → Score param

Maximum likelihood	**Analytical**	10^{-5}
Bayesian	**Gradient**	10^{-3}
Conditional likelihood	**EM**	10^{-15}
Margin	**Sampling**	…
Entropy	…	
…		

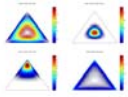

Parameter Learning from *iid* Data

- Goal: estimate distribution parameters θ from a dataset of N independent, identically distributed (*iid*), fully observed, training cases

$$\mathcal{D} = \{x^{(1)}, \ldots, x^{(n)}\}$$

- Maximum likelihood estimation (MLE)
  1. One of the most common estimators
  2. With iid and full-observability assumption, write $L(\theta)$ as the likelihood of the data:

$$\begin{aligned} L(\theta) &= P(x^{(1)}, x^{(2)}, \ldots, x^{(N)}; \theta) \\ &= P(x^{(1)}; \theta) P(x^{(2)}; \theta), \ldots, P(x^{(N)}; \theta) \\ &= \prod_{i=1}^{N} P(x^{(i)}; \theta) \end{aligned}$$

  3. pick the setting of parameters most likely to have generated the data we saw:

$$\theta^* = \arg\max_{\theta} L(\theta) \quad = \arg\max_{\theta} \log L(\theta)$$

Example: Bernoulli model

- Data:
  - We observed N ***iid*** coin tossing: D={1, 0, 1, …, 0}

- Representation:

  Binary r.v: $x^{(n)} \in \{0,1\}$

- Model: $P(x) = \begin{cases} 1-\theta & \text{for } x = 0 \\ \theta & \text{for } x = 1 \end{cases} \quad \Rightarrow \quad P(x) = \theta^x (1-\theta)^{1-x}$

- How to write the likelihood of a single observation $x^{(i)}$?

$$P(x^{(i)}) = \theta^{x^{(i)}} (1-\theta)^{1-x^{(i)}}$$

- The likelihood of dataset D={x_1, …,x_N}:

$$P(x^{(1)}, x^{(2)}, \ldots, x^{(N)} \mid \theta) = \prod_{i=1}^{N} P(x^{(i)} \mid \theta) = \prod_{i=1}^{N} \left(\theta^{x^{(i)}} (1-\theta)^{1-x^{(i)}}\right) = \theta^{\sum_{i=1}^{N} x^{(i)}} (1-\theta)^{\sum_{i=1}^{N} 1-x^{(i)}} = \theta^{\#\text{head}} (1-\theta)^{\#\text{tails}}$$

Maximum Likelihood Estimation

- Objective function:

$$\ell(\theta; D) = \log P(D \mid \theta) = \log \theta^{n_h} (1-\theta)^{n_t} = n_h \log \theta + (N - n_h) \log(1-\theta)$$

- We need to maximize this w.r.t. θ
- Take derivatives wrt θ

$$\frac{\partial \ell}{\partial \theta} = \frac{n_h}{\theta} - \frac{N - n_h}{1-\theta} = 0 \quad \Rightarrow \quad \hat{\theta}_{MLE} = \frac{n_h}{N} \quad \text{or} \quad \hat{\theta}_{MLE} = \frac{1}{N} \sum_i x^{(i)}$$

Frequency as sample mean

- Sufficient statistics
  - The counts, n_h, where $n_k = \sum_i x^{(i)}$, are **sufficient statistics** of data D

Overfitting

- Recall that for Bernoulli Distribution, we have

$$\hat{\theta}_{ML}^{head} = \frac{n^{head}}{n^{head} + n^{tail}}$$

- What if we tossed too few times so that we saw zero head?

  We have $\hat{\theta}_{ML}^{head} = 0,$ and we will predict that the probability of seeing a head next is zero!!!

- Frequentist vs. Bayesian estimate

Bayesian Parameter Estimation

- Treat the distribution parameters θ also as a *random variable*
- The *a posteriori* distribution of θ after seem the data is:

$$p(\theta \mid D) = \frac{p(D \mid \theta)\, p(\theta)}{p(D)} = \frac{p(D \mid \theta)\, p(\theta)}{\int p(D \mid \theta)\, p(\theta) d\theta}$$

This is Bayes Rule

$$\text{posterior} = \frac{\text{likelihood} \times \text{prior}}{\text{marginal likelihood}}$$

Bayes, Thomas (1763) An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London,* **53:370-418**

The prior p(.) encodes our prior knowledge about the domain

Frequentist Parameter Estimation

Two people with different priors $p(\theta)$ will end up with different estimates $p(\theta|D)$.

- Frequentists dislike this "subjectivity".
- Frequentists think of the parameter as a fixed, unknown constant, not a random variable.
- Hence they have to come up with different "objective" **estimators** (ways of computing from data), instead of using Bayes' rule.
  - These estimators have different properties, such as being "unbiased", "minimum variance", etc.
  - The maximum likelihood estimator, is one such estimator, which is simple and has good statistical properties.

Discussion

θ or $p(\theta)$, this is the problem!

Maximum Likelihood Estimation

- The log-likelihood is monotonically related to the likelihood:

$$\ell(\theta; D) = \log p(D \mid \theta) = \sum_{n=1}^{N} \log p(x^{(n)} \mid \theta)$$

- The Idea underlying maximum likelihood estimation (MLE): pick the setting of parameters most likely to have generated the data we saw:

$$\hat{\theta}_{ML} = \arg\max_{\theta} \ell(\theta; D)$$

- Problem of MLE:
  - Overfitting: means that "some of the relationships that appear statistically significant are actually just noise. It occurs when the complexity of the statistical model is too great for the amount of data that you have"
  - Often the MLE **overfits** the training data, so it is common to maximize a **regularized** log-likelihood instead:

$$\hat{\theta}' = \arg\max_{\theta} \ell(\theta; D) - c(\theta)$$

  - Insufficient training data can lead to spurious estimator (e.g., certain possible values are not observed due to data sparsity), so it is common to **smooth** the estimated parameter

Being a pragmatic frequentist

- Maximum *a posteriori* (MAP) estimation:

$$\hat{\theta}_{MAP} = \arg\max_{\theta} p(\theta \mid D) = \arg\max_{\theta} \ell(\theta; D) + \log p(\theta)$$

- Smoothing with pseudo-counts
  - Recall that for Binomial Distribution, we have

$$\hat{\theta}_{MLE}^{head} = \frac{n^{head}}{n^{head} + n^{tail}}$$

  - What if we tossed too few times so that we saw zero head?
    We have $\hat{\theta}_{ML}^{head} = 0$, and we will predict that the probability of seeing a head next is zero!!!
  - The rescue: $\hat{\theta}_{MLE}^{head} = \frac{n^{head} + n'}{n^{head} + n^{tail} + n'}$ **But are we still objective?**
    - Where n' is know as the pseudo- (imaginary) count

Bayesian estimation for Bernoulli

- Beta distribution:

$$P(\theta; \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \theta^{\alpha-1}(1-\theta)^{\beta-1} = B(\alpha, \beta)\theta^{\alpha-1}(1-\theta)^{\beta-1}$$

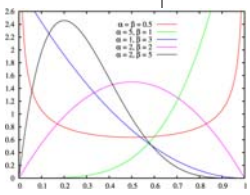

  - When x is discrete $\Gamma(x+1) = x\Gamma(x) = x!$

- Posterior distribution of θ:

$$P(\theta \mid x^{(1)}, \ldots, x^{(N)}) = \frac{p(x^{(1)}, \ldots, x^{(N)} \mid \theta)\, p(\theta)}{p(x^{(1)}, \ldots, x^{(N)})} \propto \theta^{n_h}(1-\theta)^{n_t} \times \theta^{\alpha-1}(1-\theta)^{\beta-1} = \theta^{n_h+\alpha-1}(1-\theta)^{n_t+\beta-1}$$

  - Notice the isomorphism of the posterior to the prior,
  - such a prior is called a **conjugate prior**
  - α *and* β are hyperparameters (parameters of the prior) and correspond to the number of "virtual" heads/tails (pseudo counts)

Bayesian estimation for Bernoulli, con'd

- Posterior distribution of θ:

$$P(\theta \mid x^{(1)},\ldots,x^{(N)}) = \frac{p(x^{(1)},\ldots,x^{(N)} \mid \theta)\,p(\theta)}{p(x^{(1)},\ldots,x^{(N)})} \propto \theta^{n_h}(1-\theta)^{n_t} \times \theta^{\alpha-1}(1-\theta)^{\beta-1} = \theta^{n_h+\alpha-1}(1-\theta)^{n_t+\beta-1}$$

- Maximum *a posteriori* (MAP) estimation:

$$\theta_{MAP} = \arg\max_{\theta} \log P(\theta \mid x^{(1)},\ldots,x^{(N)})$$

- Posterior mean estimation:

Bata parameters can be understood as pseudo-counts

$$\theta_{Bayes} = \int \theta p(\theta \mid D)d\theta = C\int \theta \times \theta^{n_h+\alpha-1}(1-\theta)^{n_t+\beta-1}d\theta = \frac{n_h+\alpha}{N+\alpha+\beta}$$

- Prior strength: A=α+β
  - A can be interoperated as the size of an imaginary data set from which we obtain the **pseudo-counts**

Effect of Prior Strength

- Suppose we have a uniform prior (α=β=1/2), and we observe $\vec{n} = (n_h = 2, n_t = 8)$
- Weak prior A = 2. Posterior prediction:

$$p(x = h \mid n_h = 2, n_t = 8, \vec{\alpha} = \vec{\alpha}' \times 2) = \frac{1+2}{2+10} = 0.25$$

- Strong prior A = 20. Posterior prediction:

$$p(x = h \mid n_h = 2, n_t = 8, \vec{\alpha} = \vec{\alpha}' \times 20) = \frac{10+2}{20+10} = 0.40$$

- However, if we have enough data, it washes away the prior. e.g., $\vec{n} = (n_h = 200, n_t = 800)$. Then the estimates under weak and strong prior are $\frac{1+200}{2+1000}$ and $\frac{10+200}{20+1000}$, respectively, both of which are close to 0.2

Being a "subjective" Bayesian

- The prior p(.) encodes our prior knowledge about the domain
  - therefore Bayesian estimation has been criticized for being "subjective"
- Empirical Bayes – fit prior from "training" data

How estimators should be used?

- $\widehat{\theta}_{MAP}$ is not Bayesian (even though it uses a prior) since it is a point estimate.
- Consider predicting the future. A sensible way is to combine predictions based on all possible values of θ, weighted by their posterior probability, this is what a Bayesian will do:

$$p(x_{new} \mid \mathbf{x}) = \int p(x_{new}, \theta \mid \mathbf{x}) d\theta$$
$$= \int p(x_{new} \mid \theta, \mathbf{x}) p(\theta \mid \mathbf{x}) d\theta$$
$$= \int p(x_{new} \mid \theta) p(\theta \mid \mathbf{x}) d\theta$$

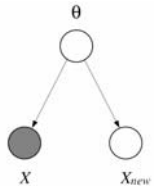

- A frequentist will typically use a "plug-in" estimator such as ML/MAP:

$$p(x_{new} \mid \mathbf{x}) = p(x_{new} \mid \widehat{\theta}_{ML}), \quad \text{or,} \quad p(x_{new} \mid \mathbf{x}) = p(x_{new} \mid \widehat{\theta}_{MAP})$$

  - The Bayesian estimate will collapse to MAP for concentrated posterior

Frequentist vs. Beyesian

- This is a “theological” war.
- Advantages of Bayesian approach:
  - Mathematically elegant.
  - Works well when amount of data is much less than number of parameters (e.g., one-shot learning).
  - Easy to do incremental (sequential) learning.
  - Can be used for model selection (max likelihood will always pick the most complex model).
- Advantages of frequentist approach:
  - Mathematically/ computationally simpler.
  - "objective", unbiased, invariant to reparameterization
- As $|\mathcal{D}| \to \infty$, the two approaches become the same:

$$p(\theta \mid D) \to \delta(\theta, \hat{\theta}_{ML})$$

Simplest GMs: the building blocks

Density estimation

Parametric and nonparametric methods

Regression

Linear, conditional mixture, nonparametric

Classification

Generative and discriminative approach

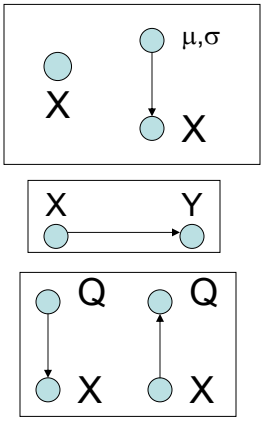

Plates

- A plate is a "macro" that allows subgraphs to be replicated

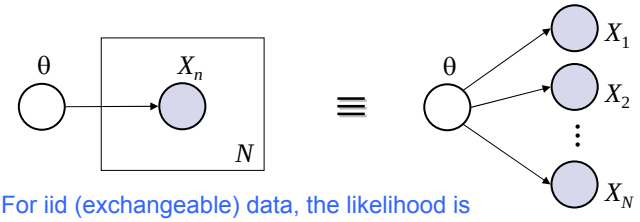

  - For iid (exchangeable) data, the likelihood is

$$p(\mathcal{D} \mid \theta) = \prod_n p(x_n \mid \theta)$$

  - We can represent this as a Bayes net with N nodes.
    - The rules of plates are simple: repeat every structure in a box a number of times given by the integer in the corner of the box (e.g. N), updating the plate index variable (e.g. n) as you go.
    - Duplicate every arrow going into the plate and every arrow leaving the plate by connecting the arrows to each copy of the structure.

Discrete Distributions

- Bernoulli distribution: Ber(p)

$$P(x) = \begin{cases} 1-p & \text{for } x = 0 \\ p & \text{for } x = 1 \end{cases} \quad \Rightarrow \quad P(x) = p^x (1-p)^{1-x}$$

- Multinomial distribution: Mult(1, θ)
  - Multinomial (indicator) variable:

$$X = \begin{bmatrix} X_1 \\ X_2 \\ X_3 \\ X_4 \\ X_5 \\ X_6 \end{bmatrix}, \quad \text{where} \quad X_j = [0,1], \text{ and } \sum_{j\in\{1,\ldots,6\}} X_j = 1$$

$$X_j = 1 \text{ w.p. } \theta_j, \quad \sum_{j\in\{1,\ldots,6\}} \theta_j = 1 \ .$$

$$p(x(j)) = P(\{X_j = 1, \text{where } j \text{ index the dice-face}\})$$
$$= \theta_j = \theta_A^{x_A} \times \theta_C^{x_C} \times \theta_G^{x_G} \times \theta_T^{x_T} = \prod_k \theta_k^{x_k} = \theta^x$$

Discrete Distributions

- Multinomial distribution: Mult(n,θ)
  - Count variable:

$$n = \begin{bmatrix} n_1 \\ \vdots \\ n_K \end{bmatrix}, \qquad \text{where } \sum_j n_j = N$$

$$p(n) = \frac{N!}{n_1!n_2!\cdots n_K!}\theta_1^{n_1}\theta_2^{n_2}\cdots\theta_K^{n_K} = \frac{N!}{n_1!n_2!\cdots n_K!}\theta^n$$

Example: multinomial model

- Data:
  - We observed N ***iid*** die rolls (K-sided): D={5, 1, K, …, 3}
- Representation:

Unit basis vectors: $x_n = \begin{pmatrix} x_{n,1} \\ x_{n,2} \\ \vdots \\ x_{n,K} \end{pmatrix}$, where $x_{n,k} = \{0,1\}$, and $\sum_{k=1}^{K} x_{n,k} = 1$

- Model:

$$X_{n,k} = 1 \text{ w.p. } \theta_k, \text{ and } \sum_{k\in\{1,\ldots K\}} \theta_k = 1$$

GM:

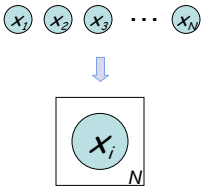

- How to write the likelihood of a single observation x_n?

$$P(x_i) = P(\{x_{n,k} = 1, \text{where } k \text{ index the die-side of the } n\text{th roll}\})$$
$$= \theta_k = \theta_1^{x_{n,1}} \times \theta_2^{x_{n,2}} \times \cdots \times \theta_K^{x_{n,K}} = \prod_{k=1}^{K} \theta_k^{x_{n,k}}$$

- The likelihood of dataset D={x_1, …,x_N}:

$$P(x_1, x_2, \ldots, x_N \mid \theta) = \prod_{n=1}^{N} P(x_n \mid \theta) = \prod_{n=1}^{N}\left(\prod_k \theta_k^{x_{n,k}}\right) = \prod_k \theta_k^{\sum_{n=1}^{N} x_{n,k}} = \prod_k \theta_k^{n_k}$$

MLE: constrained optimization with Lagrange multipliers

- Objective function:

$$\ell(\theta;D) = \log P(D \mid \theta) = \log \prod_k \theta_k^{n_k} = \sum_k n_k \log \theta_k$$

- We need to maximize this subject to the constrain $\sum_{k=1}^{K} \theta_k = 1$
- Constrained cost function with a Lagrange multiplier

$$\bar{\ell} = \sum_k n_k \log \theta_k + \lambda\left(1 - \sum_{k=1}^{K} \theta_k\right)$$

- Take derivatives wrt θ_k

$$\frac{\partial \bar{\ell}}{\partial \theta_k} = \frac{n_k}{\theta_k} - \lambda = 0$$

$$n_k = \lambda \theta_k \Rightarrow \sum_k n_k = N = \lambda \sum_k \theta_k = \lambda \quad \Rightarrow \quad \hat{\theta}_{k,MLE} = \frac{n_k}{N} \quad \text{or} \quad \hat{\theta}_{k,MLE} = \frac{1}{N}\sum_n x_{n,k}$$

Frequency as sample mean

- Sufficient statistics
  - The counts, $\vec{n} = (n_1, \cdots, n_K)$, $n_k = \sum_n x_{n,k}$, are **sufficient statistics** of data D

Bayesian estimation:

- Dirichlet distribution:

$$P(\theta) = \frac{\Gamma(\sum_k \alpha_k)}{\prod_k \Gamma(\alpha_k)} \prod_k \theta_k^{\alpha_k - 1} = C(\alpha) \prod_k \theta_k^{\alpha_k - 1}$$

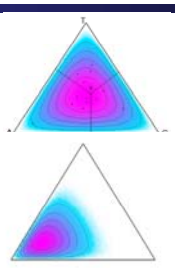

GM:

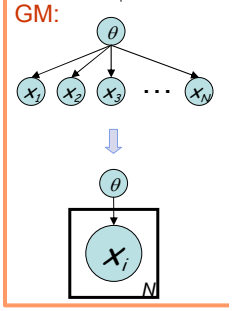

- Posterior distribution of θ:

$$P(\theta \mid x_1, \ldots, x_N) = \frac{p(x_1, \ldots, x_N \mid \theta) p(\theta)}{p(x_1, \ldots, x_N)} \propto \prod_k \theta_k^{n_k} \prod_k \theta_k^{\alpha_k - 1} = \prod_k \theta_k^{\alpha_k + n_k - 1}$$

  - Notice the isomorphism of the posterior to the prior,
  - such a prior is called a **conjugate prior**

Dirichlet parameters can be understood as pseudo-counts

- Posterior mean estimation:

$$\theta_k = \int \theta_k \, p(\theta \mid D) d\theta = C \int \theta_k \prod_k \theta_k^{\alpha_k + n_k - 1} d\theta = \frac{n_k + \alpha_k}{N + |\alpha|}$$

More on Dirichlet Prior:

- Where is the normalize constant $\mathcal{C}(\alpha)$ come from?

$$\frac{1}{C(\alpha)} = \int \cdots \int \theta_1^{\alpha_1 - 1} \cdots \theta_K^{\alpha_K - 1} d\theta_1 \cdots d\theta_K = \frac{\prod_k \Gamma(\alpha_k)}{\Gamma\left(\sum_k \alpha_k\right)}$$

  - Integration by parts
  - $\Gamma(\alpha)$ is the gamma function: $\Gamma(\alpha) = \int_0^\infty t^{\alpha-1} e^{-t} dt$
  - For inregers, $\Gamma(n+1) = n!$
- Marginal likelihood:

$$p(\{x_1, \ldots, x_N\} \mid \vec{\alpha}) = p(\vec{n} \mid \vec{\alpha}) = \int p(\vec{n} \mid \vec{\theta}) p(\vec{\theta} \mid \vec{\alpha}) d\vec{\theta} = \frac{C(\vec{\alpha})}{C(\vec{n} + \vec{\alpha})}$$

- Posterior in closed-form:

$$P(\vec{\theta} \mid \{x_1, \ldots, x_N\}, \vec{\alpha}) = \frac{p(\vec{n} \mid \theta) p(\theta \mid \vec{\alpha})}{p(\vec{n} \mid \vec{\alpha})} = C(\vec{n} + \vec{\alpha}) \prod_k \theta_k^{\alpha_k + n_k - 1} \quad = \mathrm{Dir}(\vec{n} + \vec{\alpha})$$

- Posterior predictive rate:

$$p(x_{N+1} = i \mid \{x_1, \ldots, x_N\}, \vec{\alpha}) = \int C(\vec{n} + \vec{\alpha}) \prod_k \theta_k^{\alpha_k + n_k - 1} \times \theta_i^{\alpha_k + n_k} d\vec{\theta} = \frac{C(\vec{n} + \vec{\alpha})}{C(\vec{n} + \vec{\alpha} + x_N)} \quad = \frac{n_i + \alpha_i}{|\vec{n}| + |\vec{\alpha}|}$$

Sequential Bayesian updating

- Start with Dirichlet prior $P(\vec{\theta} \mid \vec{\alpha}) = \mathrm{Dir}(\vec{\theta} : \vec{\alpha})$
- Observe N' samples with sufficient statistics $\vec{n}'$. Posterior becomes:

$$P(\vec{\theta} \mid \vec{\alpha}, \vec{n}') = \mathrm{Dir}(\vec{\theta} : \vec{\alpha} + \vec{n}')$$

- Observe another N'' samples with sufficient statistics $\vec{n}''$. Posterior becomes:

$$P(\vec{\theta} \mid \vec{\alpha}, \vec{n}', \vec{n}'') = \mathrm{Dir}(\vec{\theta} : \vec{\alpha} + \vec{n}' + \vec{n}'')$$

- So sequentially absorbing data in any order is equivalent to batch update.

Effect of Prior Strength

- Let $N = |\vec{n}| = \sum_k n_k$ be the number of observed samples
- Let $A = |\vec{\alpha}| = \sum_k \alpha_k$ be the number of "pseudo observations"
  ---- the strength of the prior
- Let $\vec{\alpha}' = |\vec{\alpha}| / A$ denote the prior means
- Then posterior mean is a convex combination of the prior mean and the MLE:

$$p(x_{N+1} = i \mid \{x_1, \ldots, x_N\}, \vec{\alpha}) = \frac{n_i + \alpha_i}{|\vec{n}| + |\vec{\alpha}|} = \frac{n_i + \alpha_i}{N + A}$$

$$= \frac{A}{N+A}\frac{\alpha_i}{A} + \frac{N}{N+A}\frac{n_i}{N}$$

$$= \lambda \alpha_i' + (1-\lambda)\hat{\theta}_{k,MLE}$$

where $\lambda = \frac{A}{N+A}$.

Hierarchical Bayesian Models

- θ are the parameters for the likelihood $p(x|\theta)$
- α are the parameters for the prior $p(\theta|\alpha)$.
- We can have hyper-hyper-parameters, etc.
- We stop when the choice of hyper-parameters makes no difference to the marginal likelihood; typically make hyper-parameters constants.
- Where do we get the prior?
  - Intelligent guesses
  - Empirical Bayes (Type-II maximum likelihood)
    → computing point estimates of α :

$$\hat{\vec{\alpha}}_{MLE} = \arg\max_{\vec{\alpha}} = p(\vec{n} \mid \vec{\alpha})$$

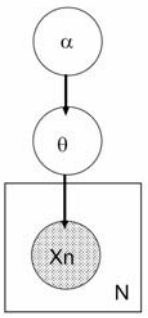

Limitation of Dirichlet Prior:

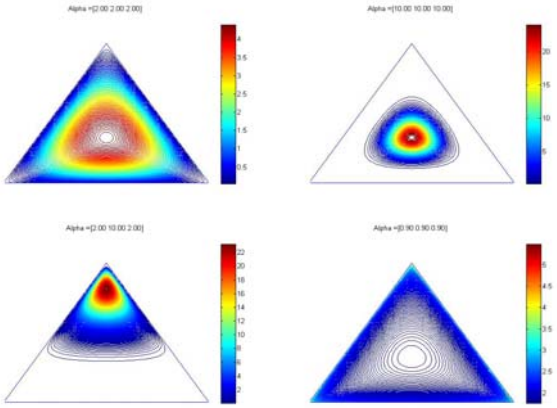

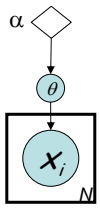

The Logistic Normal Prior

$$\theta \sim LN_K(\mu, \Sigma)$$

$$\gamma \sim N_{K-1}(\mu, \Sigma) \qquad \gamma_K = 0$$

$$\theta_i = \exp\left\{\gamma_i - \log\left(1 + \sum_{i=1}^{K-1} e^{\gamma_i}\right)\right\}$$

$$C(\gamma) = \log\left(1 + \sum_{i=1}^{K-1} e^{\gamma_i}\right)$$

Problem

- Log Partition Function
- Normalization Constant

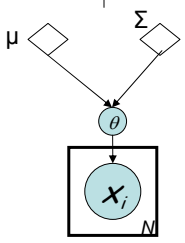

- Pro: co-variance structure
- Con: non-conjugate (we will discuss how to solve this later)

Logistic Normal Densities

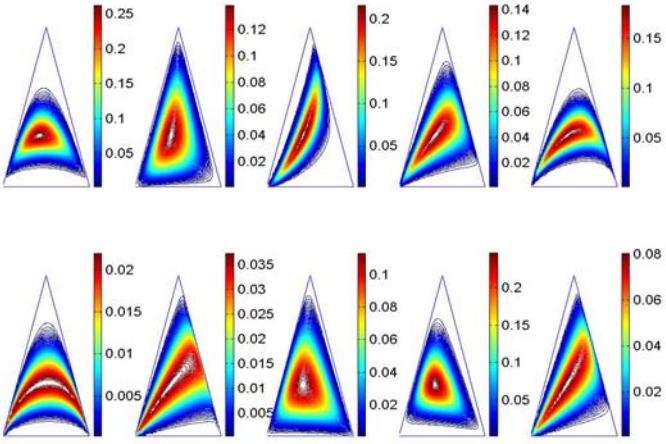

Logistic Normal

Example: univariate-Gaussian

- Data:
  - We observed N ***iid*** real samples:
    D={-0.1, 10, 1, -5.2, …, 3}
- Model: $P(x) = \left(2\pi\sigma^2\right)^{-1/2} \exp\left\{-(x-\mu)^2 / 2\sigma^2\right\}$
- Log likelihood:

$$\ell(\theta; D) = \log P(D \mid \theta) = -\frac{N}{2}\log(2\pi\sigma^2) - \frac{1}{2}\sum_{n=1}^{N}\frac{(x_n-\mu)^2}{\sigma^2}$$

- MLE: take derivative and set to zero:

$$\frac{\partial \ell}{\partial \mu} = (1/\sigma^2)\sum_n (x_n - \mu)$$

$$\frac{\partial \ell}{\partial \sigma^2} = -\frac{N}{2\sigma^2} + \frac{1}{2\sigma^4}\sum_n (x_n-\mu)^2$$

$$\Rightarrow \quad \mu_{MLE} = \frac{1}{N}\sum_n (x_n)$$

$$\sigma^2_{MLE} = \frac{1}{N}\sum_n (x_n - \mu_{ML})^2$$

GM:

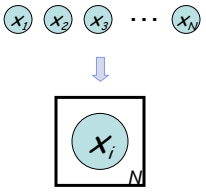

MLE for a multivariate-Gaussian

- It can be shown that the MLE for μ *and* Σ is

$$\mu_{MLE} = \frac{1}{N}\sum_n (x_n)$$

$$\Sigma_{MLE} = \frac{1}{N}\sum_n (x_n - \mu_{ML})(x_n - \mu_{ML})^T = \frac{1}{N} S$$

$$x_n = \begin{pmatrix} x_{n,1} \\ x_{n,2} \\ \vdots \\ x_{n,K} \end{pmatrix}$$

$$X = \begin{pmatrix} --- x_1^T --- \\ --- x_2^T --- \\ \vdots \\ --- x_N^T --- \end{pmatrix}$$

where the scatter matrix is

$$S = \sum_n (x_n - \mu_{ML})(x_n - \mu_{ML})^T = \left(\sum_n x_n x_n^T\right) - N\mu_{ML}\mu_{ML}^T$$

- The sufficient statistics are $\Sigma_n x_n$ and $\Sigma_n x_n x_n^T$.
- Note that $X^T X = \Sigma_n x_n x_n^T$ may not be full rank (eg. if $N < D$), in which case Σ_{ML} is not invertible

Bayesian parameter estimation for a Gaussian

- There are various reasons to pursue a Bayesian approach
  - We would like to update our estimates sequentially over time.
  - We may have prior knowledge about the expected magnitude of the parameters.
  - The MLE for Σ may not be full rank if we don't have enough data.
- We will restrict our attention to conjugate priors.
- We will consider various cases, in order of increasing complexity:
  - Known σ, unknown μ
  - Known μ, unknown σ
  - Unknown μ and σ

Bayesian estimation: unknown μ, known σ

- Normal Prior:

$$P(\mu) = (2\pi\tau^2)^{-1/2} \exp\{-(\mu-\mu_0)^2/2\tau^2\}$$

- Joint probability:

$$P(x,\mu) = (2\pi\sigma^2)^{-N/2} \exp\left\{-\frac{1}{2\sigma^2}\sum_{n=1}^{N}(x_n-\mu)^2\right\}$$
$$\times (2\pi\tau^2)^{-1/2} \exp\{-(\mu-\mu_0)^2/2\tau^2\}$$

- Posterior:

$$P(\mu \mid x) = (2\pi\tilde{\sigma}^2)^{-1/2} \exp\{-(\mu-\tilde{\mu})^2/2\tilde{\sigma}^2\}$$

where $\tilde{\mu} = \frac{N/\sigma^2}{N/\sigma^2+1/\tau^2}\bar{x} + \frac{1/\tau^2}{N/\sigma^2+1/\tau^2}\mu_0$, and $\tilde{\sigma}^2 = \left(\frac{N}{\sigma^2}+\frac{1}{\tau^2}\right)^{-1}$

Sample mean

GM:

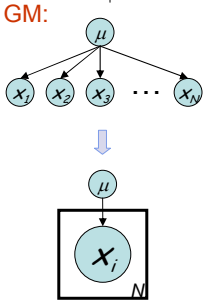

Bayesian estimation: unknown μ, known σ

$$\mu_N = \frac{N/\sigma^2}{N/\sigma^2+1/\sigma_0^2}\bar{x} + \frac{1/\sigma_0^2}{N/\sigma^2+1/\sigma_0^2}\mu_0, \qquad \tilde{\sigma}^2 = \left(\frac{N}{\sigma^2}+\frac{1}{\sigma_0^2}\right)^{-1}$$

- The posterior mean is a convex combination of the prior and the MLE, with weights proportional to the relative noise levels.
- The precision of the posterior $1/\sigma^2_N$ is the precision of the prior $1/\sigma^2_0$ plus one contribution of data precision $1/\sigma^2$ for each observed data point.

- Sequentially updating the mean
  - $\mu*$ = 0.8 (unknown), $(\sigma^2)*$ = 0.1 (known)
  - Effect of single data point

$$\mu_1 = \mu_0 + (x-\mu_0)\frac{\sigma_0^2}{\sigma^2+\sigma_0^2} = x - (x-\mu_0)\frac{\sigma_0^2}{\sigma^2+\sigma_0^2}$$

  - Uninformative (vague/ flat) prior, $\sigma^2_0 \to \infty$

$$\mu_N \to \mu_0$$

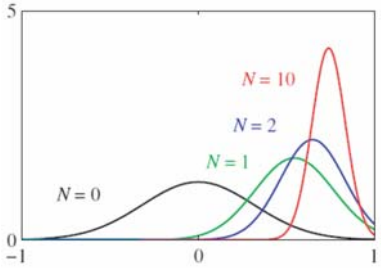

Other scenarios

- Known μ, unknown $\lambda = 1/\sigma_2$
  - The conjugate prior for λ is a Gamma with shape a_0 and rate (inverse scale) b_0

$$p(\lambda|a,b) = \frac{1}{\Gamma(a)} b^a \lambda^{a-1} \exp(-b\lambda)$$

  - The conjugate prior for σ^2 is Inverse-Gamma

$$IG(\sigma^2|a,b) = \frac{1}{\Gamma(a)} b^a (\sigma^2)^{-(a+1)} \exp(-b/(\sigma^2))$$

- Unknown μ and unknown σ_2
  - The conjugate prior is
    Normal-Inverse-Gamma

$$\begin{aligned} P(\mu,\sigma^2) &= P(\mu|\sigma^2)P(\sigma^2) \\ &= \mathcal{N}(\mu|m,\sigma^2 V)\ IG(\sigma^2|a,b) \end{aligned}$$

  - Semi conjugate prior

- Multivariate case:
  - The conjugate prior is
    Normal-Inverse-Wishart

$$\begin{aligned} P(\mu,\Sigma) &= P(\mu|\Sigma)P(\Sigma) \\ &= \mathcal{N}(\mu|\mu_0,\frac{1}{\kappa_0}\Sigma)\ \mathcal{IW}(\Sigma|\Lambda_0^{-1},\nu_0) \end{aligned}$$

Summary

- Learning scenarios:
  - Data
  - Objective function
  - Frequetist and Bayesian

- Learning single-node GM – density estimation
  - Typical discrete distribution
  - Typical continuous distribution
  - Conjugate priors