Lecture 9 - Apr 25, 2011
CSE 515, Statistical Methods, Spring 2011

Instructor: Su-In Lee
University of Washington, Seattle

## Parameter estimation

- Maximum likelihood estimation (MLE)
- Parameter estimation based on observations
- Bayesian approach
- Incorporate our prior knowledge


Bayesian network


## Maximum Likelihood Estimator

- The Coin example - general case
- X: result of a coin toss (head or tail)
- Training data (instances) $D=\left\langle x[1], \ldots x[m]>\right.$ ( $M_{H}$ heads and $M_{T}$ tails)
- Parameters: $\mathrm{P}(\mathrm{X}=\mathrm{h})=\theta$
- Goal: find $\theta \in[0,1]$ that predicts the data well
- "Predicts the data well" = likelihood of the data given $\theta$

$$
L(\theta: D)=P(D: \theta)=P(x[1], \ldots, x[m]: \theta)
$$

- MLE: Find $\theta$ maximizing likelihood

$$
L(\theta: D)=\prod_{i=1}^{m} P(x[i] \mid x[1], \ldots, x[i-1], \theta)=\prod_{i=1}^{m} P(x[i] \mid \theta)=\theta^{M_{H}}(1-\theta)^{M_{T}}
$$

- Equivalent to maximizing log-likelihood $l(\theta: D)=\log P(D: \theta)=M_{H} \log \theta+M_{T} \log (1-\theta)$
- Differentiating the log-likelihood and solving for $\theta$, we get that the maximum likelihood parameter:

$$
\theta_{\text {mle }}=\operatorname{argmaxl}(\theta: D)=\frac{M_{H}}{M_{H}+M_{T}}
$$

## Sufficient Statistics

- For computing the parameter $\theta$ of the coin toss example, we only needed $M_{H}$ and $M_{T}$ since

$$
L(\theta: D)=P(D: \theta)=\theta^{M_{H}}(1-\theta)^{M_{T}}
$$

$\rightarrow M_{H}$ and $M_{T}$ are sufficient statistics


## Sufficient Statistics

- A function $s(D)$ is a sufficient statistic from instances to a vector in $\mathfrak{R}^{k}$ if, for any two datasets $D$ and $D^{\prime}$ and any $\theta \in \Theta$, we have

$$
\sum_{x[i] \in D} s(x[i])=\sum_{x[i] \in D^{\prime}} s(x[i]) \Rightarrow L(D: \theta)=L\left(D^{\prime}: \theta\right)
$$

- We often refer to the tuple $\sum_{x(i j \in D} s(x[i])$ as the sufficient statistics
of the data set D .
- In coin toss experiment, $M_{H}$ and $M_{T}$ are sufficient statistics



## Sufficient Statistics for Multinomial

- Y: multinomial, k values (e.g. result of a dice throw)
- A sufficient statistics for a dataset $D$ over $Y$ is the tuple of counts $<M_{1}, \ldots M_{k}>$ such that $M_{i}$ is the number of times that the $Y=y^{i}$ in $D$
- Likelihood function: $L(D: \theta)=\prod_{i=1}^{k} \theta_{i}^{M_{i}}$ where $\theta_{i}=P\left(Y=y^{i}\right)$
- MLE Principle: Choose $\Theta$ that maximize $L(D: \Theta)$
- Multinomial MLE: $\quad \theta^{i}=\frac{M_{i}}{\sum_{i=1}^{m} M_{i}}$


## Sufficient Statistic for Gaussian

- Gaussian distribution: $X \sim N\left(\mu, \sigma^{2}\right)$
- Probability density function (pdf): $\quad p(X)=\frac{1}{\sqrt{2 \pi} \sigma} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^{2}}$
- Rewrite as $p(X)=\frac{1}{\sqrt{2 \pi} \sigma} \exp \left(-x^{2} \frac{1}{2 \sigma^{2}}+x \frac{\mu}{\sigma^{2}}-\frac{\mu^{2}}{\sigma^{2}}\right)$
$\rightarrow$ sufficient statistics for Gaussian: $\left\langle M, \Sigma_{m} x[m], \Sigma_{m} x[m]^{2}\right\rangle$
- MLE Principle: Choose $\Theta$ that maximize $L(\mathrm{D}: \Theta)$
- Multinomial MLE: $\quad \mu=\frac{1}{M} \sum_{m} x[m]$

$$
\sigma=\sqrt{\frac{1}{M} \sum_{m}(x[m]-\mu)^{2}}
$$

## MLE for Bayesian Networks

- Parameters
- $\theta_{x} 0, \theta_{x^{1}}$
- $\theta_{y} 00_{x} 0, \theta_{y} 1\left|{ }_{x} 0, \theta_{y} 0_{x} 1, \theta_{y}\right|_{x}{ }^{1}$
- Training data:

| $\boldsymbol{X}$ |  |
| :--- | :--- |
| $\mathrm{x}^{0}$ | $\mathrm{x}^{1}$ |
| 0.7 | 0.3 |

- tuples <x[m],y[m]> m=1,...,M
- Likelihood function:

$$
\begin{aligned}
L(D: \theta) & =\prod_{m=1}^{M} P(x[m], y[m]: \theta) \\
& =\prod_{m=1}^{M} P\left(x[m]: \theta_{X}\right) P\left(y[m] \mid x[m]: \theta_{Y \mid X}\right) \\
& =\left(\prod_{m=1}^{M} P\left(x[m]: \theta_{X}\right)\right)\left(\prod_{m=1}^{M} P\left(y[m] \mid x[m]: \theta_{Y \mid X}\right)\right)
\end{aligned}
$$


$\rightarrow$ Likelihood decomposes into two separate terms, one for each variable ("decomposability of the likelihood function") ${ }_{8}$

## MLE for Bayesian Networks

- Terms further decompose by CPDs:

$$
\begin{aligned}
& \prod_{m=1}^{M} P(y[m] \mid x[m]: \theta)=\prod_{m \times x \mid m \times x^{0}} P\left(y[m] \mid x[m]: \theta_{Y \mid X}\right) \prod_{\substack{m \times(m)=x^{2}}} P\left(y[m] \mid x[m]: \theta_{Y \mid X}\right)
\end{aligned}
$$

- By sufficient statistics
where $M\left[x^{1}, y^{1}\right]$ is the number of data instances in which $X$ takes the value $x^{1}$ and $Y$ takes the value $y^{1}$
- MLE

$$
\theta_{y^{0} \times x^{1}}=\frac{M\left[x^{1}, y^{0}\right]}{M\left[x^{1}, y^{0}\right]+M\left[x^{1}, y^{1}\right]}=\frac{M\left[x^{1}, y^{0}\right]}{M\left[x^{1}\right]}
$$

## MLE for Bayesian Networks

- Likelihood for Bayesian network

$$
\begin{aligned}
L(\Theta: D) & =\prod_{m}^{m} P(x[m]: \Theta) \\
& =\prod_{m} \prod_{i} P\left(x_{i}[m] \mid P a_{i}[m]: \Theta_{i}\right) \\
& =\prod_{i}\left[\prod_{m} P\left(x_{i}[m] \mid P a_{i}[m]: \Theta_{i}\right)\right] \\
& =\prod_{i} L_{i}\left(\boldsymbol{\theta}_{x_{i} \mid P x_{i}}: X_{i}, P a_{i}\right) \quad \begin{array}{l}
\text { Conditional likelihood } \\
\text { or "Local likelihood" }
\end{array}
\end{aligned}
$$

$\rightarrow$ if $\theta_{\mathrm{X}_{\mathrm{i}} \mid \mathrm{Pa}\left(\mathrm{X}_{\mathrm{i}}\right)}$ are disjoint then MLE can be computed by maximizing each local likelihood separately

## MLE for Table CPD BayesNets

- Multinomial CPD

$$
\begin{aligned}
L_{Y}\left(D: \theta_{Y \mid \mathbf{X}}\right) & =\prod_{m} \theta_{y[m] \mathbf{X}[m]} \\
& =\prod_{\mathbf{x} \in \operatorname{Val}(\mathbf{X})}\left[\prod_{y \in \operatorname{Val}(\mathrm{Y})} \theta_{y \mid \mathbf{X}}{ }^{M[\mathbf{X}, y]}\right]
\end{aligned}
$$

- For each value $\mathbf{x} \in \mathbf{X}$ we get an independent multinomial problem where the MLE is

$$
\theta_{y^{i} \mid x}=\frac{M\left[x, y^{i}\right]}{M[x]}
$$

## MLE for Tree CPDs

- Assume tree CPD with known tree structure
$L\left(D: \theta_{X \mid Y, Z}\right)$
$=\prod \prod^{P\left(x \mid y, z, \theta_{x \mid Y, z}\right)^{M[y, z, x]}}$
$=\prod^{y}\left(\theta_{x \mid y^{2}}{ }^{M^{0}\left[y^{0}, x^{0}, x\right]}\right) \cdot\left(\theta_{x \mid y^{0}}{ }^{M\left[y^{0}, z^{2}, x\right]}\right) \cdot\left(\theta_{x \mid y^{1}, z^{1}}{ }^{M\left[y^{1}, x^{2}, x\right]}\right) \cdot\left(\theta_{x \mid y^{1}, z^{1}}{ }^{M\left[y^{1}, z^{1}, x\right]}\right)$

$=\prod_{x}^{x}\left(\theta_{x \mid y^{0}}^{M\left[y^{0}, x\right]}\right) \cdot\left(\theta_{x \mid y^{1}, z^{2}}^{M\left[y^{1}, z^{0}, x\right]}\right) \cdot\left(\theta_{x y^{1}, z^{1}}{ }^{M\left[y^{1}, z^{1}, x\right]}\right)$


Terms for $\left\langle\mathrm{y}^{0}, \mathrm{z}^{0}\right\rangle$ and $\left.<\mathrm{y}^{0}, \mathrm{z}^{1}\right\rangle$ can be combined

Optimization can be done by leaves


## MLE for Tree CPD BayesNets

- Tree CPD T, leaves I

$$
\begin{aligned}
L_{Y}\left(D: \theta_{Y \mid \mathbf{X}}\right) & =\prod_{m} P\left(y[m] \mid \mathbf{x}[m]: \theta_{Y \mid \mathbf{X}}\right) \\
& =\prod_{m}^{m} \theta_{y[m] \mid l(x[m])} \\
& =\prod_{l \in \operatorname{Leaves}(T)}\left[\prod_{y \in V \mid(Y)} \theta_{y \mid l}^{M[c, y]}\right]
\end{aligned}
$$

- For each value $\mathrm{I} \in \mathrm{Leaves}(\mathrm{T})$ we get an independent multinomial problem where the MLE is

$$
\theta_{y^{\prime} \mid}=\frac{M\left[c_{1}, y^{\prime}\right]}{M\left[c_{1}\right]} \quad M\left[c_{l}\right]=\sum_{x:(x)=1} M\left[x, y^{\prime}\right]
$$

## Limitations of MLE

- A thumbtack is tossed 10 times, and comes out 'head' 3 of the 10 tosses $\rightarrow$ Probability of head $=0.3$
- A coin is tossed 10 times, and comes out 'head' 3 of the 10 tosses $\rightarrow$ Probability of head $=0.3$
- A coin is tossed 1,000,000 times, and comes out 'head' 300,000 of the 1,000,000 tosses $\rightarrow$ Probability of head $=0.3$
- Would you place the same bet on the next thumbtack toss as you would on the next coin toss?
- We need to incorporate prior knowledge
- Prior knowledge should only be used as a guide


## Bayesian Inference

- Assumptions
- Given a fixed $\theta$ tosses are independent
- If $\theta$ is unknown tosses are not marginally independent - each toss tells us something about $\theta$
- The following network captures our assumptions



## Bayesian Inference

- Joint probabilistic model

$$
\begin{aligned}
P(x[1], \ldots, x[M], \theta) & =P(x[1], \ldots, x[M] \mid \theta) P(\theta) \\
& =P(\theta) \prod_{i=1}^{M} P(x[i] \mid \theta) \\
& =P(\theta) \theta^{M_{H}}(1-\theta)^{M_{T}}
\end{aligned}
$$



- Posterior probability over $\theta$

$$
P(\theta \mid x[1], \ldots, x[M])=\frac{\overbrace{P(x[1], \ldots, x[M] \mid \theta)}^{\text {Likelihood }} \overbrace{P(\theta)}^{\text {Prior }}}{\underbrace{P(x[1], \ldots, x[M])}_{\text {Normalizing factor }}}
$$

For a uniform prior, posterior is the normalized likelihood

## Bayesian Prediction

- Predict the data instance from the previous ones

$$
\begin{aligned}
& P(x[M+1] \mid x[1], \ldots, x[M]) \\
& \quad=\int_{\theta} P(x[M+1], \theta \mid x[1], \ldots, x[M]) d \theta \\
& \quad=\int_{\theta} P(x[M+1] \mid x[1], \ldots, x[M], \theta) P(\theta \mid x[1], \ldots, x[M]) d \theta \\
& \quad=\int_{\theta} P(x[M+1] \mid \theta) P(\theta \mid x[1], \ldots, x[M]) d \theta
\end{aligned}
$$

- Solve for uniform prior $P(\theta)=1$ (for $0 \leq \theta \leq 1$ ) and binomial variable
$P\left(x[M+1]=x^{1} \mid x[1], \ldots, x[M]\right)=\frac{1}{P(x[1], \ldots, x[M])} \int_{\theta} \theta \cdot \theta^{M_{H}} \cdot(1-\theta)^{M_{T}}$
"Bayesian estimate" $\left.\longrightarrow=\frac{M_{H}+1}{M_{H}+M_{T}+2} \longleftrightarrow \begin{array}{l}\text { "Imaginary } \\ \text { counts" } \\ 17\end{array}\right]$


## Example: Binomial Data

- Prior: uniform for $\theta$ in $[0,1]$
- $P(\theta)=1$
$\rightarrow P(\theta \mid \mathrm{D})$ is proportional to the likelihood $\mathrm{L}(\mathrm{D}: \theta)$

$$
P(\theta \mid x[1], \ldots x[M]) \propto P(x[1], \ldots x[M] \mid \theta)
$$

$$
\left(M_{H}, M_{T}\right)=(4,1)
$$

- MLE for $\mathrm{P}(\mathrm{X}=\mathrm{H})$ is $4 / 5=0.8$
- Bayesian prediction is $5 / 7=0.71$

$$
P(x[M+1]=H \mid D)=\int \theta \cdot P(\theta \mid D) d \theta=\frac{5}{7}=0.7142 \ldots
$$

## Dirichlet Priors

- A Dirichlet prior is specified by a set of (non-negative) hyper-parameters $\alpha_{1}, \ldots \alpha_{k}$ so that
$\theta=\left[\theta_{1}, \ldots, \theta_{k}\right] \sim \operatorname{Dirichlet}\left(\alpha_{1}, \ldots \alpha_{k}\right)$ if
- $p(\theta)=\frac{1}{Z} \prod_{k} \theta_{k}^{\alpha_{k}-1}$ where $\quad \sum_{k} \theta_{k}=1, \quad \Gamma(x)=\int_{0}^{\infty} t^{x-1} e^{-t} d t$

$$
\text { and } Z=\frac{\prod_{i=1}^{k} \Gamma\left(\alpha_{i}\right)}{\Gamma\left(\sum_{i=1}^{k} \alpha_{i}\right)} .
$$

- Intuitively, hyper-parameters correspond to the number of imaginary counts before starting the coin toss experiment


## Dirichlet Priors - Example



## Dirichlet Priors

- Dirichlet priors have the property that the posterior is also Dirichlet
- Prior is $\operatorname{Dir}\left(\alpha_{1}, \ldots \alpha_{k}\right) \quad p(\theta)=\frac{1}{Z} \prod_{k} \theta_{k}^{\alpha_{k}-1}$
- Data counts are $M_{1}, \ldots, M_{k}$
- Posterior is $\operatorname{Dir}\left(\alpha_{1}+M_{1}, \ldots \alpha_{k}+M_{k}\right) \quad p(\theta \mid D)=\frac{1}{Z^{\prime}} \prod_{k} \theta_{k}^{\alpha_{k}+M_{k}-1}$
- The hyperparameters $\alpha_{1}, \ldots, \alpha_{k}$ can be thought of as "imaginary" counts from our prior experience
- Equivalent sample size $=\alpha_{1}+\ldots+\alpha_{k}$
- The larger the equivalent sample size the more confident we are in our prior


## Effect of Priors

- Prediction of $P(X=H)$ after seeing data with $M_{H}=0.2 M$, $M_{T}=0.8 \mathrm{M}$ as a function of the sample size



## Effect of Priors (cont.)

- In real data, Bayesian estimates are less sensitive to noise in the data



## General Formulation

- Joint distribution over D, $\theta$

$$
P(D, \theta)=P(D \mid \theta) P(\theta)
$$

- Posterior distribution over parameters

$$
P(\theta \mid D)=\frac{P(D \mid \theta) P(\theta)}{P(D)}
$$

- $P(D)$ is the marginal likelihood of the data

$$
P(D)=\int_{\theta} P(D \mid \theta) P(\theta) d \theta
$$

- As we saw, likelihood can be described compactly using sufficient statistics
- We want conditions in which posterior is also compact - E.g. Dirichlet priors


## Conjugate Families

- A family of priors $P(\theta: \alpha)$ is conjugate to a model $P(\xi \mid \theta)$ if for any possible dataset $D$ of i.i.d samples from $P(\xi \mid \theta)$ and choice of hyperparameters $\alpha$ for the prior over $\theta$, there are hyperparameters $\alpha^{\prime}$ that describe the posterior, i.e.,
$\mathrm{P}\left(\theta: \alpha^{\prime}\right) \propto \mathrm{P}(\mathrm{D} \mid \theta) \mathrm{P}(\theta: \alpha)$
- Posterior has the same parametric form as the prior
- Dirichlet prior is a conjugate family for the multinomial likelihood
- Conjugate families are useful since:
- Many distributions can be represented with hyperparameters
- They allow for sequential update within the same representation
- In many cases we have closed-form solutions for prediction


## Bayesian Estimation in BayesNets

Bayesian network for parameter estimation

Bayesian network


- Instances are independent given the parameters
- (x[m' y [ $\left.\mathrm{m}^{\prime}\right]$ ) are d-separated from (x[m],y[m]) given $\theta$
- Priors for individual variables are a priori independent
- Global independence of parameters $P(\theta)=\prod_{i} P\left(\theta_{X_{i} \mid P\left(X_{i}\right)}\right)$


## Bayesian Estimation in BayesNets



- Posteriors of $\theta$ are independent given complete data
- Complete data d-separates parameters for different CPDs
- $P\left(\theta_{X}, \theta_{Y \mid X} \mid D\right)=P\left(\theta_{X} \mid D\right) P\left(\theta_{Y \mid X} \mid D\right)$
- As in MLE, we can solve each estimation problem separately


## Bayesian Estimation in BayesNets

Bayesian network for parameter estimation
Bayesian network


## Bayesian Estimation in BayesNets



- Posteriors of $\theta$ can be computed independently
- For multinomial $\theta_{x_{i} \mid p a_{j}}$ posterior is Dirichlet with parameters $\left(\alpha_{x_{i}=1 \mid p a_{i}}+M\left[\mathrm{X}_{\mathrm{i}}=1 \mid \mathrm{pa}_{\mathrm{i}}\right]\right), \ldots,\left(\alpha_{\mathrm{x}_{\mathrm{i}}=k \mid p \mathrm{a}_{\mathrm{i}}}+\mathrm{M}\left[\mathrm{X}_{\mathrm{i}}=\mathrm{k} \mid \mathrm{pa}_{\mathrm{i}}\right]\right)$
- $P\left(X_{i}[M+1]=x_{i} \mid P q_{i}[M+1]=p q_{i}, D\right)=\frac{\alpha_{x_{i} \mid p q}+M\left[x_{i}, p q_{i}\right]}{\sum \alpha_{x_{i} \mid p q}+M\left[x_{i}, p q_{i}\right]}$


## Assessing Priors for BayesNets

- We need the $\alpha\left(\mathrm{x}_{\mathrm{i}}, \mathrm{pa}_{\mathrm{i}}\right)$ for each node $\mathrm{x}_{\mathrm{i}}$
- We can use initial parameters $\Theta_{0}$ as prior information
- Need also an equivalent sample size parameter $\mathrm{M}^{\prime}$
- Then, we let $\alpha\left(x_{i}, \mathrm{pa}_{\mathrm{i}}\right)=\mathrm{M}^{\prime} \cdot \mathrm{P}\left(\mathrm{x}_{\mathrm{i}}, \mathrm{pa}_{\mathrm{i}} \mid \Theta_{0}\right)$
- This allows to update a network using new data
- Example network for priors
- $P(X=0)=P(X=1)=0.5$
- $P(Y=0)=P(Y=1)=0.5$
- $\mathrm{M}^{\prime}=1$
- Note: $\alpha\left(x_{0}\right)=0.5 \alpha\left(x_{0}, y_{0}\right)=0.25$


## Case Study: ICU Alarm Network

- The "Alarm" network
- 37 variables
- Experiment
- Sample instances
- Learn parameters
- MLE
- Bayesian



## Case Study: ICU Alarm Network



- MLE performs worst
- Prior $M^{\prime}=5$ provides best smoothing


## Parameter Estimation Summary

- Estimation relies on sufficient statistics
- For multinomials these are of the form $M\left[x_{i}, \mathrm{pa}_{\mathrm{i}}\right]$
- Parameter estimation

$$
\begin{array}{cc}
\hat{\theta}_{x_{i} \mid p a_{i}}=\frac{M\left[x_{i}, p a_{i}\right]}{M\left[p a_{i}\right]} & P\left(x_{i} \mid p a_{i}, D\right)=\frac{\alpha_{x_{i}, p a_{i}}+M\left[x_{i}, p a_{i}\right]}{\alpha_{p a_{i}}+M\left[p a_{i}\right]} \\
\text { MLE } & \text { Bayesian (Dirichlet) }
\end{array}
$$

- Bayesian methods also require choice of priors
- MLE and Bayesian are asymptotically equivalent
- Both can be implemented in an online manner by accumulating sufficient statistics


## Acknowledgement

- These lecture notes were generated based on the slides from Prof Eran Segal.

