## Lecture 11:

Bayesian Parameter Learning

Kevin Murphy

October 25, 2004

## MLE for general Bayes nets

- If we assume the parameters for each CPD are globally independent, and all nodes are fully observed, then the log-likelihood function decomposes into a sum of local terms, one per node:

$$
\log p(\mathcal{D} \mid \theta)=\log \prod_{m} \prod_{i} p\left(\mathbf{x}_{i}^{m} \mid \mathbf{x}_{\pi_{i}}, \theta_{i}\right)=\sum_{i} \sum_{m} \log p\left(\mathbf{x}_{i}^{m} \mid \mathbf{x}_{\pi_{i}}, \theta_{i}\right)
$$



## Example: A Directed Model

- Consider the distribution defined by the DAGM:

$$
p(\mathbf{x} \mid \theta)=p\left(\mathbf{x}_{1} \mid \theta_{1}\right) p\left(\mathbf{x}_{2} \mid \mathbf{x}_{1}, \theta_{2}\right) p\left(\mathbf{x}_{3} \mid \mathbf{x}_{1}, \theta_{3}\right) p\left(\mathbf{x}_{4} \mid \mathbf{x}_{2}, \mathbf{x}_{3}, \theta_{4}\right)
$$

- This is exactly like learning four separate small DAGMs, each of which consists of a node and its parents.



## MLE for Bayes nets with tabular CPDs

- Assume each CPD is represented as a table (multinomial) where

$$
\theta_{i j k} \stackrel{\text { def }}{=} P\left(X_{i}=j \mid X_{\pi_{i}}=k\right)
$$

- The sufficient statistics are just counts of family configurations

$$
N_{i j k} \stackrel{\text { def }}{=} \sum_{m} I\left(X_{i}^{m}=j, X_{\pi_{i}}^{m}=k\right)
$$

- The log-likelihood is

$$
\begin{aligned}
\ell & =\log \prod_{m} \prod_{i j k} \theta_{i j k}^{N_{i j k}} \\
& =\sum_{m} \sum_{i j k} N_{i j k} \log \theta_{i j k}
\end{aligned}
$$

- Using a Lagrange multiplier to enforce so $\sum_{j} \theta_{i j k}=1$ we get

$$
\hat{\theta}_{i j k}^{M L}=\frac{N_{i j k}}{\sum_{j^{\prime}} N_{i j^{\prime} k}}
$$

## Tied parameters

- Consider a time-invariant hidden Markov model (HMM)
- State transition matrix $A(i, j) \stackrel{\text { def }}{=} P\left(X_{t}=j \mid X_{t-1}=i\right)$,
- Discrete observation matrix $B(i, j) \stackrel{\text { def }}{=} P\left(Y_{t}=j \mid X_{t}=i\right)$
- State prior $\pi(i) \stackrel{\text { def }}{=} P\left(X_{1}=i\right)$.

The joint is

$$
P\left(X_{1: T}, Y_{1: T} \mid \theta\right)=P\left(X_{1} \mid \pi\right) \prod_{t=2}^{T} P\left(X_{t} \mid X_{t-1}, A\right) \prod_{t=1}^{T} P\left(Y_{t} \mid X_{t} ; B\right)
$$

## Learning a fully observed HMM

- The log-likelihood is

$$
\begin{aligned}
& \ell(\theta ; D)=\sum_{m} \log P\left(X_{1}=x_{1}^{m} \mid \pi\right) \\
& \quad+\sum_{t=2}^{T} P\left(X_{t}=x_{t}^{m} \mid X_{t-1}=x_{t-1}^{m}, A\right)+\sum_{t=1}^{T} P\left(Y_{t}=y_{t}^{m} \mid X_{t}=x_{t}^{m}, B\right)
\end{aligned}
$$

- We can optimize each parameter $(A, B, \pi)$ separately.
- Define $A(i, j)=P\left(X_{t}=j \mid X_{t-1}=i\right)$.
- $A$ is a stochastic matrix: $\sum_{j} A(i, j)=1$
- Each row of $A$ is multinomial distribution.
- So MLE is the fraction of transitions from $i$ to $j$

$$
\hat{A}_{M L}(i, j)=\frac{\# i \rightarrow j}{\sum_{k} \# i \rightarrow k}=\frac{\sum_{m} \sum_{t=2}^{T} I\left(X_{t-1}^{m}=i, X_{t}^{m}=j\right)}{\sum_{m} \sum_{t=2}^{T} I\left(X_{t-1}^{m}=i\right)}
$$

- If the states $X_{t}$ represent words, this is called a bigram language model.
- Note that $\hat{A}_{M L}(i, j)=0$ if the particular $i, j$ pair did not occur in the training data; this is called the sparse data problem.
- We will solve this using a prior.


## DIRICHLET PRIORS

- Let $X \in\{1, \ldots, K\}$ have a multinomial distribution

$$
P(X \mid \theta)=\theta_{1}^{I(X=1)} \theta_{2}^{I(X=2)} \cdots \theta_{K}^{I(X=k)}
$$

- For a set of data $X^{1}, \ldots, X^{N}$, the sufficient statistics are the counts $N_{i}=\sum_{n} I\left(X_{n}=i\right)$.
- Consider a Dirichlet prior with hyperparameters $\alpha$

$$
p(\theta \mid \alpha)=\mathcal{D}(\theta \mid \alpha)=\frac{1}{Z(\alpha)} \cdot \theta_{1}^{\alpha_{1}-1} \cdot \theta_{2}^{\alpha_{2}-1} \cdots \theta_{K}^{\alpha_{K}-1}
$$

where $Z(\alpha)$ is the normalizing constant

- The Dirichlet prior is conjugate to (has the same form as) the multinomial likelihood.
- The $\alpha_{k}$ act like pseudo (virtual) counts.


## Normalization constant

- $Z(\alpha)$ is the normalizing constant

$$
\begin{aligned}
Z(\alpha) & =\int \cdots \int \theta_{1}^{\alpha_{1}-1} \cdots \theta_{K}^{\alpha_{K}-1} d \theta_{1} \cdots d \theta_{K} \\
& =\frac{\Gamma\left(\sum_{i=1}^{K} \alpha_{i}\right)}{\prod_{i=1}^{K} \Gamma\left(\alpha_{i}\right)}
\end{aligned}
$$

- $\Gamma(\alpha)$ is the gamma function:

$$
\Gamma(\alpha)=\int_{0}^{\infty} t^{\alpha-1} e^{-t} d t
$$

- For integers, $\Gamma(n+1)=n$ !


## DIRICHLET POSTERIOR

- Likelihood, prior, posterior:

$$
\begin{aligned}
P(\vec{N} \mid \vec{\theta}) & =\prod_{i=1}^{K} \theta_{i}^{N_{i}} \\
p(\theta \mid \alpha) & =\mathcal{D}(\theta \mid \alpha)=\frac{1}{Z(\alpha)} \cdot \theta_{1}^{\alpha_{1}-1} \cdot \theta_{2}^{\alpha_{2}-1} \cdots \theta_{K}^{\alpha_{K}-1} \\
p(\theta \mid \vec{N}, \vec{\alpha}) & =\frac{1}{Z(\alpha) p(\vec{N} \mid \alpha)} \theta_{1}^{\alpha_{1}+N_{1}} \cdots \theta_{K}^{\alpha_{K}+N_{k}} \\
& =\mathcal{D}\left(\alpha_{1}+N_{1}, \ldots, \alpha_{K}+N_{K}\right)
\end{aligned}
$$

- Marginal likelihood (evidence):

$$
P(\vec{N} \mid \vec{\alpha})=\int p(\vec{N} \mid \vec{\alpha}) p(\vec{\theta} \mid \vec{\alpha}) d^{K} \theta=\frac{Z(\vec{N}+\vec{\alpha})}{Z(\vec{\alpha})}
$$

## Hierarchical Bayesian models

- $\theta$ are the parameters for the likelihood $p(X \mid \theta)$
- $\alpha$ are the parameters for the prior $p(\theta \mid \alpha)$.
- We can have hyper-hyper-parameters, etc.
- We stop when the choice of hyper ${ }^{n}$-parameters makes no difference to the marginal likelihood; typically make hyper-parameters constants.
- Type-II maximum likelihood (empirical Bayes) = computing point estimates of $\alpha$ :

$$
\hat{\alpha}_{M L}=\arg \max _{\alpha} p(\vec{\alpha} \mid \vec{N})=\arg \max _{\alpha} p(\vec{N} \mid \vec{\alpha}) p(\vec{\alpha})
$$



## Beta priors

- Consider a coin toss $X \in\{h, t\}$.
- The Dirichlet distribution becomes the beta distribution:

$$
p(\theta)=\frac{1}{Z(\alpha)} \theta^{\alpha_{h}-1}(1-\theta)^{\alpha_{t}-1}
$$

- If $\alpha_{h}=\alpha_{t}=1$, we have a uniform (Laplace) prior.
- The posterior mean (predicted probability of heads) is

$$
\begin{aligned}
P\left(X=h \mid \alpha_{h}, \alpha_{t}\right) & =\int_{0}^{1} d \theta \quad P(X=1 \mid \theta) p(\theta) \\
& =\int_{0}^{1} d \theta \quad \theta p(\theta)=\frac{\alpha_{h}}{\alpha_{h}+\alpha_{t}}
\end{aligned}
$$

- Hence $\alpha_{h}$ is the number of virtual heads we have seen in our prior "database"; similarly for $\alpha_{t}$.
- The strength of the prior is measured by the equivalent sample size $\alpha_{h}+\alpha_{t}$.


## SEqUential Bayesian updating

- Start with beta prior $p\left(\theta \mid \alpha_{h}, \alpha_{t}\right)=\mathcal{B}\left(\theta ; \alpha_{h}, \alpha_{t}\right)$.
- Observe $N$ trials with $N_{h}$ heads and $N_{t}$ tails. Posterior becomes

$$
p\left(\theta \mid \alpha_{h}, \alpha_{t}, N_{h}, N_{t}\right)=\mathcal{B}\left(\theta ; \alpha_{h}+N_{h}, \alpha_{t}+N_{t}\right)=\mathcal{B}\left(\theta ; \alpha_{h}^{\prime}, \alpha_{t}^{\prime}\right)
$$

- Observe another $N^{\prime}$ trials with $N_{h}^{\prime}$ heads and $N_{t}^{\prime}$ tails. Posterior becomes

$$
\begin{aligned}
p\left(\theta \mid \alpha_{h}^{\prime}, \alpha_{t}^{\prime}, N_{h}^{\prime}, N_{t}^{\prime}\right) & =\mathcal{B}\left(\theta ; \alpha_{h}^{\prime}+N_{h}^{\prime}, \alpha_{t}^{\prime}+N_{t}^{\prime}\right) \\
& =\mathcal{B}\left(\theta ; \alpha_{h}+N_{h}+N_{h}^{\prime}, \alpha_{t}+N_{t}+N_{t}^{\prime}\right)
\end{aligned}
$$

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


## Effect of prior strength

- Let $N=N_{h}+N_{t}$ be number of samples (observations).
- Let $N^{\prime}$ be the number of pseudo observations (strength of prior) and define the prior means

$$
\alpha_{h}=N^{\prime} \alpha_{h}^{\prime}, \quad \alpha_{t}=N^{\prime} \alpha_{t}^{\prime}, \quad \alpha_{h}^{\prime}+\alpha_{t}^{\prime}=1
$$

- Then posterior mean is a convex combination of the prior mean and the MLE:

$$
\begin{aligned}
P\left(X=h \mid \alpha_{h}, \alpha_{t}, N_{h}, N_{t}\right) & =\frac{\alpha_{h}+N_{h}}{\alpha_{h}+N_{h}+\alpha_{t}+N_{t}} \\
& =\frac{N^{\prime} \alpha_{h}^{\prime}+N_{h}}{N+N^{\prime}} \\
& =\frac{N^{\prime}}{N+N^{\prime}} \alpha_{h}^{\prime}+\frac{N}{N+N^{\prime}} \frac{N_{h}}{N} \\
& =\lambda \alpha_{h}^{\prime}+(1-\lambda) \frac{N_{h}}{N}
\end{aligned}
$$

where $\lambda=N^{\prime} /\left(N+N^{\prime}\right)$.

## Effect of PRIOR STRENGTH

- Suppose we have a uniform prior $\alpha_{h}^{\prime}=\alpha_{t}^{\prime}=0.5$, and we observe $N_{h}=3, N_{t}=7$.
- Weak prior $N^{\prime}=2$. Posterior prediction:

$$
P\left(X=h \mid \alpha_{h}=1, \alpha_{t}=1, N_{h}=3, N_{t}=7\right)=\frac{3+1}{3+1+7+1}=\frac{1}{3} \approx 0.33
$$

- Strong prior $N^{\prime}=20$. Posterior prediction:

$$
\frac{3+10}{3+10+7+10}=\frac{13}{30} \approx 0.43
$$

- However, if we have enough data, it washes away the prior. e.g., $N_{h}=300, N_{t}=700$. Estimates are $\frac{300+1}{1000+2}$ and $\frac{300+10}{1000+20}$, both of which are close to 0.3


## PRIOR SMOOTHS PARAMETER ESTIMATES

- The MLE can change dramatically with small sample sizes.
- The MAP estimate changes much more smoothly (depending on the strength of the prior).
- This is called regularization.
- Lower blue=MLE, red = beta(1,1), pink = beta(5,5), upper blue $=$ beta $(10,10)$



## Bayesian parameter estimation for general BNs

- Defn 13.4.1: global parameter independence means $p(\theta)=\prod_{i} p\left(\theta_{i}\right)$, where $\theta_{i}$ are the parameters for CPD for $X_{i}$.
- If we assume global parameter independence, and have fully observed data, then the parameter posterior decomposes into a sum of local terms, one per node:

$$
\log p(\theta \mid \mathcal{D})=\sum_{i} \sum_{m} \log p\left(\mathbf{x}_{i}^{m} \mid \mathbf{x}_{\pi_{i}}, \theta_{i}\right)+\log p\left(\theta_{i}\right)
$$



## BAYESIAN PARAMETER ESTIMATION FOR BNs

 with tabular CPDs- Defn 13.4.4: local parameter independence means $p\left(\theta_{i}\right)=\prod_{k} p\left(\theta_{i, \cdot, k}\right)$, where $\theta_{i, j, k}=P\left(X_{i}=j \mid X_{\pi_{i}}=k\right)$ is the row of the CPT corresponding to conditioning case $k$.
- If we assume global and local parameter independence, and have fully observed data, then the parameter posteriors are

$$
P\left(\theta_{i, \cdot, k} \mid D\right)=\mathcal{D}\left(\alpha_{i, 1, k}+N_{i, 1, k}, \ldots, \alpha_{i, S, k}+N_{i, S, k}\right)
$$

- Posterior for $\theta_{y \mid x^{0}}$ and $\theta_{y \mid x^{1}}$ is factorized despite v-structure on $y_{m}$ because CPT acts like a multiplexer.



## Where do The priors come from?

- We can define $\alpha_{i j k}^{\prime}=N^{\prime} P^{\prime}\left(X_{i}=j \mid X_{\pi_{i}}=k\right)$, where $N^{\prime}$ is the strength of our prior and $P^{\prime}$ is some Bayes net that summarizes our virtual database of pseudo counts.
- This is called the BDe (Bayesian-Dirichlet likelihood equivalent) prior.
- Type-II ML $=$ learning $P^{\prime}$ from data.


## Example of BayEsian parameter learning

- Suppose we draw $X_{1: 37}^{1: N} \sim P\left(X_{1: 37} \mid \theta^{*}\right)$ from the ICU-Alarm BN.
- Then we estimate

$$
\hat{\theta}=\arg \max _{\theta} P\left(X^{1: N} \mid \theta\right) P\left(\theta \mid \alpha^{\prime}, N^{\prime}\right)
$$

for different sample sizes $N$ and prior strengths $N^{\prime}$ (with uniform prior $\left.\alpha_{i j k}^{\prime}=1 /\left|X_{i}\right|\right)$.

- We compare answers using the Kullback-Leibler divergence

$$
K L\left(P\left(X \mid \theta^{*}\right) \| P(X \mid \hat{\theta})\right)=\sum_{x} P\left(x \mid \theta^{*}\right) \log \frac{P\left(x \mid \theta^{*}\right)}{P(x \mid \hat{\theta})}
$$

where $K L(P \| Q) \geq 0$ measures the "distance" of the approximation $Q$ from truth $P$.

## Example of Bayesian parameter learning




- If $N_{i j k}=0$ in training but $P\left(X_{i}=j \mid X_{\pi}=k, \theta^{*}\right)>0$, then $K L\left(P^{*} \| \hat{P}\right)=\infty$, since

$$
K L\left(P\left(X \mid \theta^{*}\right) \| P(X \mid \hat{\theta})\right)=\sum_{x} P\left(x \mid \theta^{*}\right) \log \frac{P\left(x \mid \theta^{*}\right)}{P(x \mid \hat{\theta})}
$$

- Dirichlet smoothing helps a lot!
- Optimal prior strength $=5$.


## Application: LANGUAGE MODELING

- A bigram model predicts $P\left(X_{t}=j \mid X_{t-1}=i, \theta\right)=\theta_{i j}$.
- Often the data is sparse so $N_{i j}=0$ so $\theta_{i j}=0$.
- A standard hack is to use backoff smoothing or deleted interpolation:

$$
\hat{P}\left(X_{t} \mid X_{t-1}\right)=\lambda f_{x_{t}}+(1-\lambda) f_{x_{t} \mid x_{t-1}}
$$

where $\lambda$ is set by cross valdiation and $f_{i}$ and $f_{j \mid i}$ are empirical frequencies.

- A similar effect can be gotten using a hierarchical prior.


## Application: LANGUAGE modeling



- Assign the same Dirichlet prior $\alpha m_{i}$ to each row of the transition matrix.
- So the prediction is

$$
P(i \mid j, D, \alpha m)=\frac{f_{i \mid j}+\alpha m_{i}}{\sum_{i^{\prime}} f_{i^{\prime} \mid j}+\alpha m_{i^{\prime}}}=\lambda_{j} m_{i}+\left(1-\lambda_{j}\right) f_{i \mid j}
$$

where $\lambda_{j}=\frac{\alpha}{f_{j}+\alpha}$.

- This is like adaptive deleted interpolation.


## Application: LANGUAGE MODELING



- We can optimize the hyperparameters using numerical methods (e.g., conjugate gradient), which is faster than cross validation.

$$
(\alpha m)^{M A P}=\arg \max P(D \mid \alpha m)
$$

- We could consider more realistic priors, e.g., mixtures of Dirichlets to account for types of words (adjectives, verbs, etc.)


## CPDs FOR CONTINUOUS NODES

- So far we have considered the case where $p(y \mid x, \theta)$ can be represented as a multinomial (table).
- Now we consider the case where some nodes may be continuous.

| $X$ | $Y$ | $p(Y \mid X)$ |
| :---: | :---: | :---: |
| $\mathbb{R}^{n}$ | $\mathbb{R}^{m}$ | regression |
| $\mathbb{R}^{n}$ | $\{0,1\}$ | binary classification |
| $\{0,1\}^{n}$ | $\{0,1\}$ | binary classification |
| $\mathbb{R}^{n}$ | $\{1, \ldots, K\}$ | multiclass classification |
| $\{1, \ldots, K\}$ | $\mathbb{R}^{n}$ | conditional density modeling |

## Exponential Family

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

$$
\begin{aligned}
p(\mathbf{x} \mid \eta) & =h(\mathbf{x}) \exp \left\{\eta^{\top} T(\mathbf{x})-A(\eta)\right\} \\
& =\frac{1}{Z(\eta)} h(\mathbf{x}) \exp \left\{\eta^{\top} T(\mathbf{x})\right\}
\end{aligned}
$$

is an exponential family distribution with
natural (canonical) parameter $\eta$.

- Function $T(\mathbf{x})$ is a sufficient statistic.
- Function $A(\eta)=\log Z(\eta)$ is the log normalizer.
- Examples: Bernoulli, multinomial, Gaussian, Poisson, gamma,...
- A distribution $p(x)$ has finite sufficient statistics (independent of number of data cases) iff it is in the exponential family.


## MLE for Exponential Family

- For iid data, the log-likelihood is

$$
\begin{aligned}
\ell(\eta ; \mathcal{D}) & =\log \prod_{m} h\left(x^{m}\right) \exp \left(\eta^{T} T\left(x^{m}\right)-A(\eta)\right) \\
& =\left(\sum_{m} \log h\left(\mathbf{x}^{m}\right)\right)-M A(\eta)+\left(\eta^{\top} \sum_{m} T\left(\mathbf{x}^{m}\right)\right)
\end{aligned}
$$

- Take derivatives and set to zero:

$$
\begin{aligned}
\frac{\partial \ell}{\partial \eta} & =\sum_{m} T\left(\mathbf{x}^{m}\right)-M \frac{\partial A(\eta)}{\partial \eta}=0 \\
\Rightarrow \frac{\partial A(\eta)}{\partial \eta} & =\frac{1}{M} \sum_{m} T\left(\mathbf{x}^{m}\right) \\
\hat{\mu}_{\mathrm{ML}} & =\frac{1}{M} \sum_{m} T\left(\mathbf{x}^{m}\right)
\end{aligned}
$$

- This amounts to moment matching.
- We can infer the canonical parameters using $\hat{\eta}_{M L}=\psi\left(\hat{\mu}_{M L}\right)$


## Linear Regression




## Multivariate Linear Regression

- Consider vector-valued input $X \in R^{k}$ going to vector-valued output $Y \in R^{d}$ via regression matrix $A \in R^{k \times d}$.

$$
p(y \mid x)=(2 \pi)^{-d / 2}|\Sigma|^{-\frac{1}{2}} \exp \left[-\frac{1}{2}(y-A x)^{T} \Sigma^{-1}(y-A x)\right]
$$

- Log-likelihood

$$
\ell=-\frac{1}{2} \sum_{m} \log |\Sigma|-\frac{1}{2} \sum_{m}\left(y_{m}-A x_{m}\right)^{T} \Sigma^{-1}\left(y_{m}-A x_{m}\right)
$$

- To take derivatives wrt a matrix, we use the following identity

$$
\frac{\partial\left((M a+b)^{T} C(M a+b)\right)}{\partial M}=\left(C+C^{T}\right)(M a+b) a^{T}
$$

where $A=M, a=-x_{m}$ and $b=y_{m}$.

## Multivariate Linear Regression

- Log-likelihood:

$$
\ell=-\frac{1}{2} \sum_{m} \log |\Sigma|-\frac{1}{2} \sum_{m}\left(y_{m}-A x_{m}\right)^{T} \Sigma^{-1}\left(y_{m}-A x_{m}\right)
$$

- Using

$$
\frac{\partial\left((M a+b)^{T} C(M a+b)\right)}{\partial M}=\left(C+C^{T}\right)(M a+b) a^{T}
$$

we have

$$
\begin{aligned}
\frac{\partial \ell}{\partial A} & =-\frac{1}{2} \sum_{m} 2 \Sigma^{-1}\left(y_{m}-A x_{m}\right) x_{m}^{T} \\
& =-\Sigma^{-1} \sum_{m} y_{m} x_{m}^{T}-A \sum_{m} x_{m} x_{m}^{T} \\
& \stackrel{\text { def }}{=}-\Sigma^{-1} S_{Y X^{\prime}}-A S_{X X^{\prime}}=0
\end{aligned}
$$

where $S_{Y X^{\prime}}$ and $S_{X X^{\prime}}$ are the sufficient statistics. Hence

$$
A=S_{Y X^{\prime}} S_{X X^{\prime}}^{-1}
$$

## 1D Linear Regression

- For the vector case,

$$
A=S_{Y X^{\prime}} S_{X X^{\prime}}^{-1}
$$

where $S_{Y X^{\prime}}=\sum_{m} y_{m} x_{m}^{T}$ and $S_{X X^{\prime}}=\sum_{m} x_{m} x_{m}^{T}$.

- In the special case of scalar outputs, let $A=\theta^{T}$, and the design matrix $X=\left[x_{m}^{T}\right]$ stacked as rows and $Y=\left[y_{m}\right]$ a column vector. Then we get the normal equations

$$
\theta=\left(X^{T} X\right)^{-1} X^{T} Y
$$



- For scalar (1D) output

$$
p\left(y_{n} \mid x_{n}, \theta, \sigma^{2}\right) p\left(\theta \mid \mu, \tau^{2}\right) p\left(\sigma^{2} \mid \alpha, \beta\right)
$$

Gaussian $\times$ Gaussian $\times$ Gamma

- For vector output

$$
p\left(y_{n} \mid x_{n}, A, \Sigma\right) p\left(A \mid \mu, \tau^{2}\right) p(\Sigma \mid \alpha, \beta)
$$

Gaussian $\times$ matrix-Gaussian $\times$ Wishart

## MLE for Generalized Linear Models

- GLIM with scale parameter $\phi$ and canonical parameter $\eta=\theta^{T} x$ :

$$
p(y \mid x, \theta, \phi)=h(y, \phi) \exp \left(\frac{\eta^{T} y-A(\eta)}{\phi}\right)
$$

- Log-likelihood

$$
\ell=\sum_{n} \log h\left(y_{n}\right)+\frac{1}{\phi} \sum_{n}\left(\theta^{T} x_{n} y_{n}-A\left(\eta_{n}\right)\right)
$$

- Derivative of Log-likelihood

$$
\begin{aligned}
\frac{d \ell}{d \theta} & =\frac{1}{\phi} \sum_{n}\left(x_{n} y_{n}-\frac{d A\left(\eta_{n}\right)}{d \eta_{n}} \frac{d \eta_{n}}{d \theta}\right) \\
& =\frac{1}{\phi} \sum_{n}\left(y_{n}-\mu_{n}\right) x_{n} \\
& =\frac{1}{\phi} X^{T}(y-\mu)
\end{aligned}
$$

- Derivative of Log-likelihood

$$
\frac{d \ell}{d \theta}=\frac{1}{\phi} \sum_{n}\left(y_{n}-\mu_{n}\right) x_{n}
$$

- Stochastic gradient ascent $=$ least mean squares (LMS) algorithm:

$$
\theta^{t+1}=\theta^{t}+\rho\left(y_{n}-\mu_{n}^{t}\right) x_{n}
$$

where $\mu_{n}^{t}=\theta^{(t) T} x_{n}$ and $\rho$ is a step size.

## BATCH LEARNING FOR CANONICAL GLIMS

- Hessian

$$
\begin{aligned}
H & =\frac{d^{2} \ell}{d \theta d \theta^{T}}=\frac{d}{d \theta^{T}} \frac{1}{\phi} \sum_{n} x_{n}\left(y_{n}-\mu_{n}\right)=-\frac{1}{\phi} \sum_{n} x_{n} \frac{d \mu_{n}}{d \theta^{T}} \\
& =-\frac{1}{\phi} \sum_{n} x_{n} \frac{d \mu_{n}}{d \eta_{n}} \frac{d \eta_{n}}{d \theta^{T}} \\
& =-\frac{1}{\phi} \sum_{n} x_{n} \frac{d \mu_{n}}{d \eta_{n}} x_{n}^{T} \text { since } \eta_{n}=\theta^{T} x_{n} \\
& =-\frac{1}{\phi} X^{T} W X
\end{aligned}
$$

where $X=\left[x_{n}^{T}\right]$ is the design matrix and

$$
W=\operatorname{diag}\left(\frac{d \mu_{1}}{d \eta_{1}}, \ldots, \frac{d \mu_{N}}{d \eta_{N}}\right)
$$

$$
\begin{aligned}
\nabla_{\theta} \ell & =\frac{1}{\phi} X^{T}(y-\mu) \\
H & =-\frac{1}{\phi} X^{T} W x \\
\theta^{t+1} & =\theta^{T}+H^{-1} \nabla_{\theta} \ell \\
& =\left(X^{T} W^{t} X\right)^{-1}\left[X^{T} W^{t} X \theta^{t}+X^{T}\left(y-\mu^{t}\right)\right] \\
& =\left(X^{T} W^{t} X\right)^{-1} X^{T} W^{t} z^{t}
\end{aligned}
$$

where the adjusted response is

$$
z^{t}=X \theta^{t}+\left(W^{t}\right)^{-1}\left(y-\mu^{t}\right)
$$

We iteratively reoptimize

$$
\theta^{t+1}=\arg \min _{\theta}(z-X \theta)^{T} W(z-X \theta)
$$

This Newton-Raphson procedure will (usually) find the global optimum starting from $\theta=0$.

$$
\begin{aligned}
\mu & =\sigma(\eta)=\frac{1}{1+e^{-\eta}}=\sigma\left(\theta^{T} x\right)=p(y=1 \mid x, \theta) \\
\frac{d \mu}{d \eta} & =\mu(1-\mu) \\
W & =\left(\begin{array}{lll}
\mu_{1}\left(1-\mu_{1}\right) & & \\
& \cdots & \\
& & \mu_{n}\left(1-\mu_{n}\right)
\end{array}\right)
\end{aligned}
$$

## LOGISTIC REGRESSION: PRACTICAL ISSUES

- It is very common to use penalized maximum likelihood.

$$
\begin{aligned}
p(y= \pm 1 \mid x, \theta) & =\sigma\left(y \theta^{T} x\right)=\frac{1}{1+\exp \left(-y \theta^{T} x\right)} \\
p(\theta) & \sim \mathcal{N}\left(0, \lambda^{-1} I\right) \\
\ell(\theta) & =\sum_{n} \log \sigma\left(y_{n} \theta^{T} x_{n}\right)-\frac{\lambda}{2} \theta^{T} \theta
\end{aligned}
$$

- IRLS takes $O\left(N d^{2}\right)$ per iteration, where $N=$ number of training cases and $d=$ size of input $x$.
- Quasi-Newton methods, that approximate the Hessian, work faster.
- Conjugate gradient takes $O(N d)$ per iteration, and usually works best in practice.
- Stochastic gradient descent can also be used if $N$ is large c.f. perceptron rule:

$$
\nabla_{\theta} \ell(\theta)=\left(1-\sigma\left(y_{n} \theta^{T} x_{n}\right)\right) y_{n} x_{n}-\lambda \theta
$$

