# CSEP 590A Summer 2006 

## Lecture 4

MLE, EM, RE, Expression

## FYI, re HW \#2: Hemoglobin History



Figure 8-76 An evolutionary scheme for the globin chains that carry oxygen in the blood of animals. The scheme emphasizes the $\beta$-like globin gene family. A relatively recent gene duplication of the $\gamma$-chain gene produced $\gamma^{\mathrm{G}}$ and $\gamma^{\mathrm{A}}$, which are fetal $\beta$-like chains of identical function. The location of the globin genes in the human genome is shown at the top of the figure. Alberts et al., 3rd ed..pg389

## Tonight

- MLE: Maximum Likelihood Estimators
- EM: the Expectation Maximization Algorithm
- Bio: Gene expression and regulation
- Next week: Motif description \& discovery


## MLE

## Maximum Likelihood Estimators

## Probability Basics, I

Ex.
Sample Space

$$
\{1,2, \ldots, 6\}
$$

Distribution

$$
p_{1}, \ldots, p_{6} \geq 0 ; \sum_{1 \leq i \leq 6} p_{i}=1 \quad f(x)>=0 ; \int_{\mathbb{R}} f(x) d x=1
$$

e.g.

$$
p_{1}=\cdots=p_{6}=1 / 6
$$


pdf, not probability

## Probability Basics, II

Ex.
Ex.

$$
E(g)=\int_{\mathbb{R}} g(x) f(x) d x
$$

Population

$$
\begin{array}{lll}
\text { mean } & \mu=\sum_{1 \leq i \leq 6} i p_{i} & \mu=\int_{\mathbb{R}} x f(x) d x \\
\text { variance } & \sigma^{2}=\sum_{1 \leq i \leq 6}(i-\mu)^{2} p_{i} & \sigma^{2}=\int_{\mathbb{R}}(x-\mu)^{2} f(x) d x
\end{array}
$$

Sample
mean
variance

$$
\begin{gathered}
\bar{x}=\sum_{1 \leq i \leq n} x_{i} / n \\
\bar{s}^{2}=\sum_{1 \leq i \leq n}\left(x_{i}-\bar{x}\right)^{2} / n
\end{gathered}
$$

## Parameter Estimation

- Assuming sample $x_{1}, x_{2}, \ldots, x_{n}$ is from a parametric distribution $\mathrm{f}(\mathrm{x} \mid \theta)$, estimate $\theta$.
- E.g.:

$$
\begin{aligned}
f(x) & =\frac{1}{\sqrt{2 \pi \sigma^{2}}} e^{-(x-\mu)^{2} /\left(2 \sigma^{2}\right)} \\
\theta & =\left(\mu, \sigma^{2}\right)
\end{aligned}
$$

## Maximum Likelihood Parameter Estimation

- One (of many) approaches to param. est.
- Likelihood of (indp) observations $x_{1}, x_{2}, \ldots, x_{n}$

$$
L\left(x_{1}, x_{2}, \ldots, x_{n}\right)=\prod_{i=1}^{n} f\left(x_{i} \mid \theta\right)
$$

- As a function of $\theta$, what $\theta$ maximizes the likelihood of the data actually observed
- Typical approach: $\frac{\partial}{\partial \theta} L(\vec{x} \mid \theta)=0$ or $\frac{\partial}{\partial \theta} \log L(\vec{x} \mid \theta)=0$


## Example I

$n$ coin flips, $x_{1}, x_{2}, \ldots, x_{n} ; n_{0}$ tails, $n_{1}$ heads, $n_{0}+n_{1}=n$; $\theta=$ probability of heads

$$
\begin{aligned}
L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \theta\right) & =(1-\theta)^{n_{0}} \theta^{n_{1}} \\
\log L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \theta\right) & =n_{0} \log (1-\theta)+n_{1} \log \theta \\
\frac{\partial}{\partial \theta} \log L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \theta\right) & =\frac{-n_{0}}{1-\theta}+\frac{n_{1}}{\theta}
\end{aligned}
$$



Setting to zero and solving:

$$
\theta=\frac{n_{1}}{n}
$$

(Also verify it's max, not min, \& not better on boundary)

Ex. 2: $x_{i} \sim N\left(\mu, \sigma^{2}\right), \sigma^{2}=1, \mu$ unknown

$$
\begin{aligned}
L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \theta\right) & =\prod_{1 \leq i \leq n} \frac{1}{\sqrt{2 \pi}} e^{-\left(x_{i}-\theta\right)^{2} / 2} \\
\ln L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \theta\right) & =\sum_{1 \leq i \leq n}-\frac{1}{2} \ln 2 \pi-\frac{\left(x_{i}-\theta\right)^{2}}{2} \\
\frac{d}{d \theta} \ln L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \theta\right) & =\sum_{1 \leq i \leq n}\left(x_{i}-\theta\right)
\end{aligned}
$$

And verify it's max, not min \& not better on boundary


$$
=\left(\sum_{1 \leq i \leq n} x_{i}\right)-n \theta=0
$$

$$
\hat{\theta}=\left(\sum_{1 \leq i \leq n} x_{i}\right) / n=\bar{x}
$$

## Ex 3: $x_{i} \sim N\left(\mu, \sigma^{2}\right), \mu, \sigma^{2}$ both unknown

$$
\begin{aligned}
\ln L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \theta_{1}, \theta_{2}\right) & =\sum_{1 \leq i \leq n}-\frac{1}{2} \ln 2 \pi \theta_{2}-\frac{\left(x_{i}-\theta_{1}\right)^{2}}{2 \theta_{2}} \\
\frac{\partial}{\partial \theta_{1}} \ln L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \theta_{1}, \theta_{2}\right) & =\sum_{1 \leq i \leq n} \frac{\left(x_{i}-\theta_{1}\right)}{\theta_{2}}=0 \\
\hat{\theta}_{1} & =\left(\sum_{1 \leq i \leq n} x_{i}\right) / n=\bar{x}
\end{aligned}
$$

## Ex. 3, (cont.)

$$
\begin{aligned}
\ln L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \theta_{1}, \theta_{2}\right) & =\sum_{1 \leq i \leq n}-\frac{1}{2} \ln 2 \pi \theta_{2}-\frac{\left(x_{i}-\theta_{1}\right)^{2}}{2 \theta_{2}} \\
\frac{\partial}{\partial \theta_{2}} \ln L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \theta_{1}, \theta_{2}\right) & =\sum_{1 \leq i \leq n}-\frac{1}{2} \frac{2 \pi}{2 \pi \theta_{2}}+\frac{\left(x_{i}-\theta_{1}\right)^{2}}{2 \theta_{2}^{2}}=0 \\
\hat{\theta}_{2} & =\left(\sum_{1 \leq i \leq n}\left(x_{i}-\hat{\theta}_{1}\right)^{2}\right) / n=\bar{s}^{2}
\end{aligned}
$$

A consistent, but biased estimate of population variance. (An example of overfitting.) Unbiased estimate is:

$$
\hat{\theta}_{2}=\sum_{1 \leq i \leq n} \frac{\left(x_{i}-\hat{\theta}_{1}\right)^{2}}{n-1}
$$

Moral: MLE is a great idea, but not a magic bullet

## EM

The Expectation-Maximization Algorithm

## More Complex Example

This?


## Gaussian Mixture Models / Model-based Clustering



Parameters $\theta$

| means | $\mu_{1}$ | $\mu_{2}$ |
| :--- | :--- | :--- |
| variances | $\sigma_{1}^{2}$ | $\sigma_{2}^{2}$ |
| mixing parameters | $\tau_{1}$ | $\tau_{2}=1-\tau_{1}$ |

P.D.F.
$f\left(x \mid \mu_{1}, \sigma_{1}^{2}\right) \quad f\left(x \mid \mu_{2}, \sigma_{2}^{2}\right)$
Likelihood

$$
\begin{aligned}
& L\left(x_{1}, x_{2}, \ldots, x_{n} \mid \mu_{1}, \mu_{2}, \sigma_{1}^{2}, \sigma_{2}^{2}, \tau_{1}, \tau_{2}\right) \\
& \quad=\prod_{i=1}^{n} \sum_{j=1}^{2} \tau_{j} f\left(x_{i} \mid \mu_{j}, \sigma_{j}^{2}\right)
\end{aligned}
$$

No closedform max




## A What-lf Puzzle

- Likelihood $\theta$
- $L(x_{1}, x_{2}, \ldots, x_{n} \mid \overbrace{\mu_{1}, \mu_{2}, \sigma_{1}^{2}, \sigma_{2}^{2}, \tau_{1}, \tau_{2}})$

$$
=\prod_{i=1}^{n} \sum_{j=1}^{2} \tau_{j} f\left(x_{i} \mid \mu_{j}, \sigma_{j}^{2}\right)
$$

- Messy: no closed form solution known for finding $\theta$ maximizing $L$
- But what if we knew the hidden data?

$$
z_{i j}= \begin{cases}1 & \text { if } x_{i} \text { drawn from } f_{j} \\ 0 & \text { otherwise }\end{cases}
$$

## EM as Egg vs Chicken

- IF $z_{i j}$ known, could estimate parameters $\theta$
- IF parameters $\theta$ known, could estimate $\mathrm{z}_{\mathrm{ij}}$
- But we know neither; (optimistically) iterate:
- E: calculate expected $\mathrm{z}_{\mathrm{ij}}$, given parameters
- M: calc "MLE" of parameters, given $E\left(z_{i j}\right)$


## The E-step

- Assume $\theta$ known \& fixed
- A (B): the event that $x_{i}$ was drawn from $f_{1}\left(f_{2}\right)$
- D: the observed datum $x_{i}$
- Expected value of $\mathrm{z}_{\mathrm{il}}$ is $\mathrm{P}(\mathrm{A} \mid \mathrm{D})$

$$
\left.\begin{array}{rl}
P(A \mid D) & =\frac{P(D \mid A) P(A)}{P(D)} \\
P(D) & =P(D \mid A) P(A)+P(D \mid B) P(B) \\
& =f_{1}\left(x_{i} \mid \theta_{1}\right) \tau_{1}+f_{2}\left(x_{i} \mid \theta_{2}\right) \tau_{2}
\end{array}\right\} \begin{gathered}
\text { Repeat } \\
\text { for } \\
\text { each } \\
\mathrm{x}_{\mathrm{i}}
\end{gathered}
$$

## The M-Step

Goal is to find MLE $\theta$ of:

$$
L\left(x_{1}, \ldots, x_{n}, z_{11}, z_{12}, \ldots, z_{n 2} \mid \theta\right)
$$

$x_{i}$ 's are known;
Would be easy if $z_{i j}$ 's also known, but they aren't.
Instead, maximize expected likelihood of visible data

$$
E\left(L\left(x_{1}, \ldots, x_{n} \mid \theta\right)\right)
$$

where expectation is over distribution of hidden data $\left(z_{i j}\right.$ 's)

## M-step Details

(For simplicity, assume $\sigma_{1}=\sigma_{2}=\sigma ; \tau_{1}=\tau_{2}=.5$ )

$$
\begin{aligned}
L(\vec{x}, \vec{z} \mid \theta) & =\prod_{1 \leq i \leq n} \frac{1}{\sqrt{2 \pi \sigma^{2}}} \exp \left(-\sum_{1 \leq j \leq 2} z_{i j} \frac{\left(x_{i}-\mu_{j}\right)^{2}}{2 \sigma^{2}}\right) \\
E[\log L(\vec{x}, \vec{z} \mid \theta)] & =E\left[\sum_{1 \leq i \leq n}\left(-\frac{1}{2} \log 2 \pi \sigma^{2}-\sum_{1 \leq j \leq 2} z_{i j} \frac{\left(x_{i}-\mu_{j}\right)^{2}}{2 \sigma^{2}}\right)\right] \\
& =\sum_{1 \leq i \leq n}\left(-\frac{1}{2} \log 2 \pi \sigma^{2}-\sum_{1 \leq j \leq 2} E\left[z_{i j}\right] \frac{\left(x_{i}-\mu_{j}\right)^{2}}{2 \sigma^{2}}\right)
\end{aligned}
$$

Find $\theta$ maximizing this as before, using $E\left[z_{i j}\right]$ found in E-step. Result:
$\mu_{j}=\sum_{i=1}^{n} E\left[z_{i j}\right] x_{i} / \sum_{i=1}^{n} E\left[z_{i j}\right]$ (intuit: avg, weighted by subpop prob)

## EM Summary

- Fundamentally a max likelihood parameter estimation problem
- Useful if analysis is more tractable when $0 / I$ hidden data $\mathbf{z}$ known
- Iterate:
- E-step: estimate $\mathrm{E}(\mathrm{z})$ given $\theta$
- M-step: estimate $\theta$ maximizing $E(l i k e l i h o o d)$ given $\mathrm{E}(\mathrm{z})$


## EM Issues

- Under mild assumptions (sect II.6), EM is guaranteed to increase likelihood with every E-M iteration, hence will converge.
- But may converge to local, not global, max. (Recall the 4-bump surface...)
- Issue is probably intrinsic, since EM is often applied to NP-hard problems (including clustering, above, and motif-discovery, soon)
- Nevertheless, widely used, often effective


## Relative entropy

## Relative Entropy

- AKA Kullback-Liebler Distance/Divergence, AKA Information Content
- Given distributions P, Q

$$
H(P \| Q)=\sum_{x \in \Omega} P(x) \log \frac{P(x)}{Q(x)}
$$

Notes:
Let $P(x) \log \frac{P(x)}{Q(x)}=0$ if $P(x)=0$ [since $\left.\lim _{y \rightarrow 0} y \log y=0\right]$
Undefined if $0=Q(x)<P(x)$


## Theorem: $H(P \| Q) \geq 0$

$$
\begin{aligned}
H(P \| Q) & =\sum_{x} P(x) \log \frac{P(x)}{Q(x)} \\
& \geq \sum_{x} P(x)\left(1-\frac{Q(x)}{P(x)}\right) \\
& =\sum_{x}(P(x)-Q(x)) \\
& =\sum_{x} P(x)-\sum_{x} Q(x) \\
& =1-1 \\
& =0
\end{aligned}
$$

Furthermore: $\mathrm{H}(\mathrm{P} \| \mathrm{Q})=0$ if and only if $\mathrm{P}=\mathrm{Q}$ Bottom line: "bigger" means "more different"

## Gene Expression \& Regulation

## Gene Expression

- Recall a gene is a DNA sequence
- To say a gene is expressed means that it
I. is transcribed from DNA to RNA

2. the mRNA is processed in various ways
3. is exported from the nucleus (eukaryotes)
4. is translated into protein

- A key point: not all genes are expressed all the time, in all cells, or at equal levels


## Transcription

- RNA polymerase complex
- E. coli: 5 proteins $\left(2 \alpha, \beta, \beta^{\prime}, \sigma\right)$ $\sigma$ is initiation factor; finds promoter, then released/replaced by elongation factors
- Eukaryotes: 3 pols, each > 10 subunits
- attaches to DNA, melts helix, makes RNA copy ( $5^{\prime} \rightarrow 3$ ') of template ( $3^{\prime} \rightarrow 5^{\prime}$ ) at $\sim 30 \mathrm{nt} / \mathrm{sec}$




## 5' Processing: Capping

- methylated G added to 5' end, and methyl added to ribose of Ist nucleotide of transcript
- probably helps distinguish protein-coding mRNAs from other RNA junk
- prevents degradation
- facilitates start of translation


## 3' Processing: Poly A

(Eukaryotes)

- Transcript cleaved after AAUAAA (roughly)
- pol keeps running (until it falls off) but no 5' cap added to strand downstream of poly A site, so it's rapidly degraded
- 10s - 100s of A's added to 3' end of transcript - its "poly A tail"


## More processing: Splicing

- Also in eukaryotes, most genes are spliced: protein coding exons are interrupted by non-coding introns, which are cut out \& degraded, exons spliced together
- More details about this when we get to gene finding



## Nuclear Export

- In eukaryotes, mature mRNAs are actively transported out of the nucleus \& ferried to specific destinations (e.g., mitochondria, ribosomes)


## Regulation

- In most cells, pro- or eukaryote, easily a I0,000-fold difference between least- and most-highly expressed genes
- Regulation happens at all steps. E.g., some transcripts can be sequestered then released, or rapidly degraded, some are weakly translated, some are very actively translated, some are highly transcribed, some are not transcribed at all
- Below, focus on Ist step only: transcriptional regulation


## DNA Binding Proteins

A variety of DNA binding proteins
("transcription factors"; a significant fraction, perhaps $10 \%$ ?, of all human proteins) modulate transcription of protein coding genes

## The Double Helix


(a) Computer-generated Image of DNA
(by Mel Prueitt)
(b) Uncoiled DNA Fragment


## As shown, the two strands coil

about each other in a fashion such that all the bases project inward toward the helix axis. The two strands are held together by hydrogen bonds (pink rods) linking each base projecting from one backbone to its so-called complementary base projecting from the other backbone. The base $A$ always bonds to $T$ ( A and T are comple-

Shown in (b)
is an uncoiled fragment of (a three complementary base pai chemist's viewpoint, each stra a polymer made up of four re called deoxyribonucleotides

## In the groove

## Different

 patterns of potential H bonds at edges of different base pairs, accessible esp. in major groove

## Helix-Turn-Helix DNA Binding Motif



## H-T-H Dimers



Alberts, et al.

> Bind 2 DNA patches, ~ I turn apart Increases both specificity and affinity


# Zinc Finger Motif 



## Leucine Zipper Motif



Homo-/hetero-dimers and combinatorial control

## Bacterial Met Repressor a beta-sheet DNA binding domain <br> Negative feedback loop:

 high Met level $\Rightarrow$ repress Met synthesis genes

## Summary

- Learning from data:
- MLE: Max Likelihood Estimators
- EM: Expectation Maximization (MLE w/hidden data)
- Expression \& regulation
- Expression: creation of gene products
- Regulation: when/where/how much of each gene product; complex and critical
- Next week: using MLE/EM to find regulatory motifs in biological sequence data

