# Learning with Maximum Likelihood 

Note to other teachers and users of these slides. Andrew would be delighted if you found this source material useful in giving your own lectures. Feel free to use these slides verbatim, or to modify them to fit your own needs. PowerPoint originals are available. If you make use of a significant portion of these slides in your own lecture, please include this message, or the following link to the source repository of Andrew's tutorials: http://www.cs.cmu.edu/~awm/tutorials Comments and corrections gratefully received.

Andrew W. Moore Professor
School of Computer Science
Carnegie Mellon University
www.cs.cmu.edu/~awm awm@cs.cmu.edu 412-268-7599

## Maximum Likelihood learning of Gaussians for Data Mining

- Why we should care
- Learning Univariate Gaussians
- Learning Multivariate Gaussians
- What's a biased estimator?
- Bayesian Learning of Gaussians


## Why we should care

- Maximum Likelihood Estimation is a very very very very fundamental part of data analysis.
- "MLE for Gaussians" is training wheels for our future techniques
- Learning Gaussians is more useful than you might guess...


## Learning Gaussians from Data

- Suppose you have $x_{1}, x_{2}, \ldots x_{R} \sim$ (i.i.d) $N\left(\mu, \sigma^{2}\right)$
- But you don't know $\mu$

$$
\text { (you do know } \sigma^{2} \text { ) }
$$

MLE: For which $\mu$ is $x_{1}, x_{2}, \ldots x_{R}$ most likely?
MAP: Which $\mu$ maximizes $\mathrm{p}\left(\mu \mid x_{1}, x_{2}, \ldots x_{R}, \sigma^{2}\right)$ ?

## Learning Gaussians from Data

- Suppose you have $x_{1}, x_{2}, \ldots x_{R} \sim(i . i . d) N\left(\mu, \sigma^{2}\right)$
- But you don't know $\mu$

Sneer
(you do know $\sigma^{2}$ )
MLE: For which $\mu$ is $x_{1}, x_{2}, \ldots x_{R}$ most likely?
MAP: Which $\mu$ maximizes $\mathrm{p}\left(\mu \mid x_{1}, x_{2}, \ldots x_{R}, \sigma^{2}\right)$ ?

## Learning Gaussians from Data

- Suppose you have $x_{1}, x_{2}, \ldots x_{R} \sim(i . i . d) N\left(\mu, \sigma^{2}\right)$
- But you don't know $\mu$

Sneer
(you do know $\sigma^{2}$ )

MLE: For which $\mu$ is $x_{1}, x_{2}, \ldots x_{R}$ most likely?
MAP: Which $\mu$ maximizes $\mathrm{p}\left(\mu \mid x_{1}, x_{2}, \ldots x_{R}, \sigma^{2}\right)$ ?

Despite this, we'll spend $95 \%$ of our time on MLE. Why? Wait and see...

## MLE for univariate Gaussian

- Suppose you have $x_{1}, x_{2}, \ldots x_{R} \sim\left(\right.$ i.i.d) $N\left(\mu, \sigma^{2}\right)$
- But you don't know $\mu$ (you do know $\sigma^{2}$ )
- MLE: For which $\mu$ is $x_{1}, x_{2}, \ldots x_{R}$ most likely?

$$
\mu^{m l e}=\underset{\mu}{\arg \max } p\left(x_{1}, x_{2}, \ldots x_{R} \mid \mu, \sigma^{2}\right)
$$

$$
\begin{array}{ll}
\quad \text { Algebra Euphoria } \\
\mu^{m l e}=\underset{\mu}{\arg \max } p\left(x_{1}, x_{2}, \ldots x_{R} \mid \mu, \sigma^{2}\right) & \\
= & \text { (by i.i.d) } \\
= & \begin{array}{l}
\text { (monotonicity of } \\
\text { log) }
\end{array} \\
= & \begin{array}{l}
\text { (plug in formula } \\
\text { for Gaussian) }
\end{array} \\
= & \begin{array}{l}
\text { (after } \\
\text { simplification) }
\end{array}
\end{array}
$$

$$
\begin{array}{ll}
\quad \text { Algebra Euphoria } \\
\mu^{m l e}=\underset{\mu}{\arg \max } p\left(x_{1}, x_{2}, \ldots x_{R} \mid \mu, \sigma^{2}\right) \\
= & \underset{\mu}{\arg \max } \prod_{i=1}^{R} p\left(x_{i} \mid \mu, \sigma^{2}\right) \\
=\underset{\mu}{\arg \max } \sum_{i=1}^{R} \log p\left(x_{i} \mid \mu, \sigma^{2}\right) & \text { (by i.i.d) } \\
=\underset{\mu}{\arg \max } \frac{1}{\sqrt{2 \pi} \sigma} \sum_{i=1}^{R}-\frac{\left(x_{i}-\mu\right)^{2}}{2 \sigma^{2}} & \begin{array}{l}
\text { (monotonicity of } \\
= \\
\arg \min \sum^{R}\left(x_{i}-\mu\right)^{2}
\end{array} \\
\text { (plug in formula } \\
= & \begin{array}{l}
\text { (after } \\
\text { simplification) }
\end{array}
\end{array}
$$

## Intermission: A General Scalar MLE strategy

Task: Find MLE $\theta$ assuming known form for p(Data| $\theta$,stuff)

1. Write $L L=\log P($ Data| $\theta$,stuff)
2. Work out $\partial L L / \partial \theta$ using high-school calculus
3. Set $\partial L L / \partial \theta=0$ for a maximum, creating an equation in terms of $\theta$
4. Solve it*
5. Check that you've found a maximum rather than a minimum or saddle-point, and be careful if $\theta$ is constrained

> *This is a perfect example of something that works perfectly in all textbook examples and usually involves surprising pain if you need it for something new.

$$
\begin{aligned}
& \quad \quad \text { The MLE } \mu \\
& \mu^{\mu^{m l e}}=\underset{\mu}{\arg \max } p\left(x_{1}, x_{2}, \ldots x_{R} \mid \mu, \sigma^{2}\right) \\
& =\underset{\mu}{\arg \min } \sum_{i=1}^{R}\left(x_{i}-\mu\right)^{2} \\
& =\mu \text { s.t. } \quad 0=\frac{\partial \mathrm{LL}}{\partial \mu}=
\end{aligned}
$$

$$
=\text { (what?) }
$$

$$
\begin{gathered}
\text { The MLE } \mu \\
\mu^{m l e}=\underset{\mu}{\arg \max } p\left(x_{1}, x_{2}, \ldots x_{R} \mid \mu, \sigma^{2}\right) \\
=\underset{\mu}{\arg \min } \sum_{i=1}^{R}\left(x_{i}-\mu\right)^{2} \\
=\mu \text { s.t. } \quad 0=\frac{\partial \mathrm{LL}}{\partial \mu}=\frac{\partial}{\partial \mu} \sum_{i=1}^{R}\left(x_{i}-\mu\right)^{2} \\
-\sum_{i=1}^{R} 2\left(x_{i}-\mu\right) \\
\text { Thus } \quad \mu=\frac{1}{R} \sum_{i=1}^{R} x_{i}
\end{gathered}
$$

## Lawks-a-lawdy! $\mu^{m l e}=\frac{1}{R} \sum_{i=1}^{R} x_{i}$

- The best estimate of the mean of a distribution is the mean of the sample!

At first sight:
This kind of pedantic, algebra-filled and ultimately unsurprising fact is exactly the reason people throw down their "Statistics" book and pick up their "Agent Based Evolutionary Data Mining Using The Neuro-Fuzz Transform" book.

## A General MLE strategy

Suppose $\theta=\left(\theta_{1}, \theta_{2}, \ldots, \theta_{n}\right)^{\top}$ is a vector of parameters.
Task: Find MLE $\theta$ assuming known form for p(Data| $\theta$,stuff)

1. Write $L L=\log P($ Data| $\theta$,stuff)
2. Work out $\partial \mathrm{LL} / \partial \theta$ using high-school calculus

$$
\frac{\partial L L}{\partial \boldsymbol{\theta}}=\left(\begin{array}{c}
\frac{\partial L L}{\partial \theta_{1}} \\
\frac{\partial L L}{\partial \theta_{2}} \\
\vdots \\
\frac{\partial L L}{\partial \theta_{n}}
\end{array}\right)
$$

## A General MLE strategy

Suppose $\theta=\left(\theta_{1}, \theta_{2}, \ldots, \theta_{n}\right)^{\top}$ is a vector of parameters.
Task: Find MLE $\theta$ assuming known form for p(Data| $\theta$,stuff)

1. Write $L L=\log P($ Data| $\theta$,stuff)
2. Work out $\partial \mathrm{LL} / \partial \theta$ using high-school calculus
3. Solve the set of simultaneous equations

$$
\begin{gathered}
\frac{\partial L L}{\partial \theta_{1}}=0 \\
\frac{\partial L L}{\partial \theta_{2}}=0 \\
\vdots \\
\frac{\partial L L}{\partial \theta_{n}}=0
\end{gathered}
$$

## A General MLE strategy

Suppose $\theta=\left(\theta_{1}, \theta_{2}, \ldots, \theta_{n}\right)^{\top}$ is a vector of parameters.
Task: Find MLE $\theta$ assuming known form for p(Data| $\theta$,stuff)

1. Write $L L=\log P($ Data| $\theta$,stuff $)$
2. Work out $\partial \mathrm{LL} / \partial \theta$ using high-school calculus
3. Solve the set of simultaneous equations

$$
\begin{array}{ll}
\frac{\partial L L}{\partial \theta_{1}} & =0 \\
\frac{\partial L L}{\partial \theta_{2}} & =0
\end{array} \quad \begin{aligned}
& \text { 4. } \\
& \vdots \\
& \frac{\partial L L}{\partial \theta_{n}}
\end{aligned}=0 \begin{aligned}
& \text { Check that you're at } \\
& \text { a maximum }
\end{aligned}
$$

## A General MLE strategy

Suppose $\theta=\left(\theta_{1}, \theta_{2}, \ldots, \theta_{n}\right)^{\top}$ is a vector of parameters.
Task: Find MLE $\theta$ assuming known form for p(Data| $\theta$,stuff)

1. Write $L L=\log P($ Data| $\theta$,stuff)
2. Work out $\partial \mathrm{LL} / \partial \theta$ using high-school calculus
3. Solve the set of simultaneous equations


## MLE for univariate Gaussian

- Suppose you have $x_{1}, x_{2}, \ldots x_{R} \sim\left(\right.$ i.i.d) $\mathrm{N}\left(\mu, \sigma^{2}\right)$
- But you don't know $\mu$ or $\sigma^{2}$
- MLE: For which $\theta=\left(\mu, \sigma^{2}\right)$ is $x_{1}, x_{2}, \ldots x_{R}$ most likely?
$\log p\left(x_{1}, x_{2}, \ldots x_{R} \mid \mu, \sigma^{2}\right)=-R\left(\log \pi+\frac{1}{2} \log \sigma^{2}\right)-\frac{1}{2 \sigma^{2}} \sum_{i=1}^{R}\left(x_{i}-\mu\right)^{2}$
$\frac{\partial L L}{\partial \mu}=\frac{1}{\sigma^{2}} \sum_{i=1}^{R}\left(x_{i}-\mu\right)$
$\frac{\partial L L}{\partial \sigma^{2}}=-\frac{R}{2 \sigma^{2}}+\frac{1}{2 \sigma^{4}} \sum_{i=1}^{R}\left(x_{i}-\mu\right)^{2}$


## MLE for univariate Gaussian

- Suppose you have $x_{1}, x_{2}, \ldots, x_{R} \sim\left(\right.$ i.i.d) $\mathrm{N}\left(\mu, \sigma^{2}\right)$
- But you don't know $\mu$ or $\sigma^{2}$
- MLE: For which $\theta=\left(\mu, \sigma^{2}\right)$ is $x_{1}, x_{2}, \ldots x_{R}$ most likely?

$$
\begin{aligned}
& \log p\left(x_{1}, x_{2}, \ldots x_{R} \mid \mu, \sigma^{2}\right)=-R\left(\log \pi+\frac{1}{2} \log \sigma^{2}\right)-\frac{1}{2 \sigma^{2}} \sum_{i=1}^{R}\left(x_{i}-\mu\right)^{2} \\
& 0=\frac{1}{\sigma^{2}} \sum_{i=1}^{R}\left(x_{i}-\mu\right) \\
& 0=-\frac{R}{2 \sigma^{2}}+\frac{1}{2 \sigma^{4}} \sum_{i=1}^{R}\left(x_{i}-\mu\right)^{2}
\end{aligned}
$$

## MLE for univariate Gaussian

- Suppose you have $x_{1}, x_{2}, \ldots, x_{R} \sim\left(\right.$ i.i.d) $\mathrm{N}\left(\mu, \sigma^{2}\right)$
- But you don't know $\mu$ or $\sigma^{2}$
- MLE: For which $\theta=\left(\mu, \sigma^{2}\right)$ is $x_{1}, x_{2}, \ldots x_{R}$ most likely? $\log p\left(x_{1}, x_{2}, \ldots x_{R} \mid \mu, \sigma^{2}\right)=-R\left(\log \pi+\frac{1}{2} \log \sigma^{2}\right)-\frac{1}{2 \sigma^{2}} \sum_{i=1}^{R}\left(x_{i}-\mu\right)^{2}$
$0=\frac{1}{\sigma^{2}} \sum_{i=1}^{R}\left(x_{i}-\mu\right) \Rightarrow \mu=\frac{1}{R} \sum_{i=1}^{R} x_{i}$
$0=-\frac{R}{2 \sigma^{2}}+\frac{1}{2 \sigma^{4}} \sum_{i=1}^{R}\left(x_{i}-\mu\right)^{2} \Rightarrow$ what?


## MLE for univariate Gaussian

- Suppose you have $x_{1}, x_{2}, \ldots x_{R} \sim(i . i . d) \mathrm{N}\left(\mu, \sigma^{2}\right)$
- But you don't know $\mu$ or $\sigma^{2}$
- MLE: For which $\theta=\left(\mu, \sigma^{2}\right)$ is $x_{1}, x_{2}, \ldots x_{R}$ most likely?

$$
\begin{gathered}
\mu^{m l e}=\frac{1}{R} \sum_{i=1}^{R} x_{i} \\
\sigma_{\text {mle }}^{2}=\frac{1}{R} \sum_{i=1}^{R}\left(x_{i}-\mu^{m l e}\right)^{2}
\end{gathered}
$$

## Unbiased Estimators

- An estimator of a parameter is unbiased if the expected value of the estimate is the same as the true value of the parameters.
- If $x_{1}, x_{2}, \ldots x_{R} \sim\left(\right.$ i.i.d) $N\left(\mu, \sigma^{2}\right)$ then

$$
E\left[\mu^{m l e}\right]=E\left[\frac{1}{R} \sum_{i=1}^{R} x_{i}\right]=\mu
$$

$\mu^{m / e}$ is unbiased

## Biased Estimators

- An estimator of a parameter is biased if the expected value of the estimate is different from the true value of the parameters.
- If $x_{1}, x_{2}, \ldots x_{R} \sim\left(\right.$ i.i.d) $\mathrm{N}\left(\mu, \sigma^{2}\right)$ then

$$
\begin{gathered}
E\left[\sigma_{m l e}^{2}\right]=E\left[\frac{1}{R} \sum_{i=1}^{R}\left(x_{i}-\mu^{\text {mle }}\right)^{2}\right]=E\left[\frac{1}{R}\left(\sum_{i=1}^{R} x_{i}-\frac{1}{R} \sum_{j=1}^{R} x_{j}\right)^{2}\right] \neq \sigma^{2} \\
\sigma^{2}{ }_{\text {m/e }} \text { is biased }
\end{gathered}
$$

## MLE Variance Bias

- If $x_{1}, x_{2}, \ldots x_{R} \sim\left(\right.$ i.i.d) $N\left(\mu, \sigma^{2}\right)$ then

$$
\begin{aligned}
& E\left[\sigma_{\text {mle }}^{2}\right]=E\left[\frac{1}{R}\left(\sum_{i=1}^{R} x_{i}-\frac{1}{R} \sum_{j=1}^{R} x_{j}\right)^{2}\right]=\left(1-\frac{1}{R}\right) \sigma^{2} \neq \sigma^{2} \\
& \text { Intuition check: consider the case of } \mathrm{R}=1 \\
& \text { Why should our guts expect that } \sigma^{2} \text { m/e would be an } \\
& \text { underestimate of true } \sigma^{2 ?} \\
& \text { How could you prove that? }
\end{aligned}
$$

## Unbiased estimate of Variance

- If $x_{1}, x_{2}, \ldots x_{R} \sim\left(\right.$ i.i.d) $N\left(\mu, \sigma^{2}\right)$ then
$E\left[\sigma_{\text {mle }}^{2}\right]=E\left[\frac{1}{R}\left(\sum_{i=1}^{R} x_{i}-\frac{1}{R} \sum_{j=1}^{R} x_{j}\right)^{2}\right]=\left(1-\frac{1}{R}\right) \sigma^{2} \neq \sigma^{2}$
So define $\quad \sigma_{\text {unbiased }}^{2}=\frac{\sigma_{\text {mle }}^{2}}{\left(1-\frac{1}{R}\right)}$ So $E\left[\sigma_{\text {unbiased }}^{2}\right]=\sigma^{2}$


## Unbiased estimate of Variance

- If $x_{1}, x_{2}, \ldots x_{R} \sim\left(\right.$ i.i.d) $\mathrm{N}\left(\mu, \sigma^{2}\right)$ then
$E\left[\sigma_{\text {mle }}^{2}\right]=E\left[\frac{1}{R}\left(\sum_{i=1}^{R} x_{i}-\frac{1}{R} \sum_{j=1}^{R} x_{j}\right)^{2}\right]=\left(1-\frac{1}{R}\right) \sigma^{2} \neq \sigma^{2}$
So define $\quad \sigma_{\text {unbiased }}^{2}=\frac{\sigma_{m l e}^{2}}{\left(1-\frac{1}{R}\right)} \quad$ So $E\left[\sigma_{\text {unbiased }}^{2}\right]=\sigma^{2}$

$$
\sigma_{\text {unbised }}^{2}=\frac{1}{R-1} \sum_{i=1}^{R}\left(x_{i}-\mu^{m l e}\right)^{2}
$$

## Unbiaseditude discussion

- Which is best?

$$
\begin{aligned}
& \sigma_{\text {mle }}^{2}=\frac{1}{R} \sum_{i=1}^{R}\left(x_{i}-\mu^{m l e}\right)^{2} \\
& \sigma_{\text {unbiased }}^{2}=\frac{1}{R-1} \sum_{i=1}^{R}\left(x_{i}-\mu^{m l e}\right)^{2}
\end{aligned}
$$

Answer:

- It depends on the task
-And doesn't make much difference once R--> large


## Don't get too excited about being unbiased

- Assume $x_{1}, x_{2}, \ldots x_{R} \sim(i . i . d) \mathrm{N}\left(\mu, \sigma^{2}\right)$
- Suppose we had these estimators for the mean

$$
\mu^{\text {suboptinal }}=\frac{1}{R+7 \sqrt{R}} \sum_{i=1}^{R} x_{i}
$$

Are either of these unbiased?

$$
\mu^{c r a p}=x_{1}
$$

Will either of them asymptote to the correct value as R gets large?
Which is more useful?

## MLE for m-dimensional Gaussian

- Suppose you have $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R} \sim($ i.i.d) $\mathrm{N}(\mu, \Sigma)$
- But you don't know $\mu$ or $\Sigma$
- MLE: For which $\theta=(\mu, \Sigma)$ is $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R}$ most likely?

$$
\begin{aligned}
\boldsymbol{\mu}^{m l e} & =\frac{1}{R} \sum_{k=1}^{R} \mathbf{x}_{k} \\
\boldsymbol{\Sigma}^{m l e} & =\frac{1}{R} \sum_{k=1}^{R}\left(\mathbf{x}_{k}-\boldsymbol{\mu}^{m l e}\right)\left(\mathbf{x}_{k}-\boldsymbol{\mu}^{m l e}\right)^{T}
\end{aligned}
$$

## MLE for m-dimensional Gaussian

- Suppose you have $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R} \sim($ i.i.d) $N(\mu, \Sigma)$
- But you don't know $\mu$ or $\Sigma$
- MLE: For which $\theta=(\mu, \Sigma)$ is $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R}$ most likely?

$$
\begin{aligned}
& \boldsymbol{\mu}^{m l e}=\frac{1}{R} \sum_{k=1}^{R} \mathbf{x}_{k} \quad \mu_{i}^{m l e}=\frac{1}{R} \sum_{k=1}^{R} \mathbf{x}_{k i}\left\{\begin{array}{l}
\text { Where } 1 \leq \mathrm{i} \leq \mathrm{m} \\
\boldsymbol{\Sigma}^{m l e}=\frac{1}{R} \sum_{k=1}^{R}\left(\mathbf{x}_{k}-\mu^{m l e}\right)\left(\mathbf{x}_{k}-\mu^{m l e}\right)^{T} \\
\text { And } x_{\mathrm{k}} \text { is value of the } \\
\text { ith component of } \mathbf{x}_{\mathrm{k}} \\
\text { (the } \mathrm{i}^{\text {th }} \text { attribute of } \\
\text { the } \mathrm{k}^{\text {th }} \text { record) }
\end{array}\right. \\
& \begin{array}{l}
\text { And } \mu_{i}{ }^{m / e} \text { is the ith } \\
\text { component of } \mu^{m / e}
\end{array}
\end{aligned}
$$

## MLE for m-dimensional Gaussian

- Suppose you have $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R} \sim($ i.i.d) $\mathrm{N}(\mu, \Sigma)$
- But you don't know $\mu$ or $\Sigma$
- MLE: For which $\theta=(\mu, \Sigma)$ is $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R}$ most likely?

$$
\begin{aligned}
& \boldsymbol{\mu}^{m l e}=\frac{1}{R} \sum_{k=1}^{R} \mathbf{x}_{k} \quad \begin{array}{l}
\text { Where } 1 \leq \mathrm{i} \leq \mathrm{m}, 1 \leq \mathrm{j} \leq \mathrm{m} \\
\text { And } x_{\mathrm{ki}} \text { is value of the } \mathrm{i}^{\text {th }} \\
\text { component of } \mathbf{x}_{\mathrm{k}} \text { (the } \mathrm{i}^{\text {th }} \\
\text { attribute of the } \mathrm{k}^{\text {th }} \text { record) } \\
\boldsymbol{\Sigma}^{m l e}=\frac{1}{R} \sum_{k=1}^{R}\left(\mathbf{x}_{k}-\mu^{m l e}\right)\left(\mathbf{x}_{k}-\mu^{m l e}\right)^{T} \\
\begin{array}{l}
\text { And } \sigma_{\mathrm{ij}} \text { mle is the }(\mathrm{i}, \mathrm{j})^{\text {th }} \\
\text { component of } \Sigma^{m / e}
\end{array} \\
\sigma_{i j}^{m l e}=\frac{1}{R} \sum_{k=1}^{R}\left(\mathbf{x}_{k i}-\mu_{i}^{m l e}\right)\left(\mathbf{x}_{k j}-\mu_{j}^{m l e}\right)
\end{array}
\end{aligned}
$$

## MLE for m-dimensinnal Rauccian

- Suppose you have $\mathbf{x}_{1}, \boldsymbol{X}_{2,}$ • A: Just plug through the MLE
- But you don't know/ $\mu$ or $\Sigma$ recipe.
- MLE: For which $\theta=(\mu, \Sigma)$ is $\begin{aligned} & \text { Note how } \Sigma^{m / e} \text { is forced to be } \\ & \text { symmetric non-negative definite }\end{aligned}$
$\boldsymbol{\mu}^{m l e}=\frac{1}{R} \sum_{k=1}^{R} \mathbf{x}_{k} \quad \boldsymbol{\Sigma}^{m l e}=\frac{1}{R} \sum_{k=1}^{R}\left(\mathbf{x}_{k}-\mu^{m l e}\right)\left(\mathbf{x}_{k}-\mu^{m l e}\right)^{T} \begin{aligned} & \text { Note the unbiased case } \\ & \text { How many datapoints would you } \\ & \text { need before the Gaussian has a }\end{aligned}$
$\boldsymbol{\Sigma}^{\text {unbiased }}=\frac{\boldsymbol{\Sigma}^{m l e}}{1-\frac{1}{R}}=\frac{1}{R-1} \sum_{k=1}^{R}\left(\mathbf{x}_{k}-\mu^{m l e}\right)\left(\mathbf{x}_{k}-\mu^{m l e}\right)^{T}$


## Confidence intervals

We need to talk
We need to discuss how accurate we expect $\mu^{m / e}$ and $\Sigma^{m / e}$ to be as a function of $R$

And we need to consider how to estimate these accuracies from data...
-Analytically *
-Non-parametrically (using randomization and bootstrapping) *
But we won't. Not yet.
*Will be discussed in future Andrew lectures...just before we need this technology.

## Structural error

Actually, we need to talk about something else too..
What if we do all this analysis when the true distribution is in fact not Gaussian?
How can we tell? *
How can we survive? *
*Will be discussed in future Andrew lectures...just before we need this technology.

## Gaussian MLE in action

Using R=392 cars from the "MPG" UCI dataset supplied by Ross Quinlan


## Data-starved Gaussian MLE

Using three subsets of MPG.
Each subset has 6 randomly-chosen cars.




## Multivariate MLE

|  | mean | cov |  |  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| mpg | 23.4459 | 60.9181 | -10.3529 | -657.585 | -233.858 | -5517.44 | 9.11551 | 16.6915 |
| cylinders | 5.47194 | -10.3529 | 2.9097 | 169.722 | 55.3482 | 1300.42 | -2.37505 | -2.17193 |
| displacement | 194.412 | -657.585 | 169.722 | 10950.4 | 3614.03 | 82929.1 | -156.994 | -142.572 |
| horsepower | 104.469 | -233.858 | 55.3482 | 3614.03 | 1481.57 | 28265.6 | -73.187 | -59.0364 |
| weight | 2977.58 | -5517.44 | 1300.42 | 82929.1 | 28265.6 | 721485 | -976.815 | -967.228 |
| acceleration | 15.5413 | 9.11551 | -2.37505 | -156.994 | -73.187 | -976.815 | 7.61133 | 2.95046 |
| modelyear | 75.9796 | 16.6915 | -2.17193 | -142.572 | -59.0364 | -967.228 | 2.95046 | 13.5699 |

## Covariance matrices are not exciting to look at

## Being Bayesian: MAP estimates for Gaussians

- Suppose you have $\mathbf{x}_{1,}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R} \sim(i . i . d) N(\mu, \Sigma)$
- But you don't know $\mu$ or $\Sigma$
- MAP: Which $(\mu, \Sigma)$ maximizes $p\left(\mu, \Sigma \mid \mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R}\right)$ ?

Step 1: Put a prior on ( $\mu, \Sigma$ )

Being Bayesian: MAP estimates for Gaussians

- Suppose you have $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R} \sim($ i.i.d) $N(\mu, \Sigma)$
- But you don't know $\mu$ or $\Sigma$
- MAP: Which $(\mu, \Sigma)$ maximizes $\mathrm{p}\left(\mu, \Sigma \mid \mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R}\right)$ ?

Step 1: Put a prior on ( $\mu, \Sigma$ )
Step 1a: Put a prior on $\Sigma$

$$
\left(v_{0}-m-1\right) \Sigma \sim \operatorname{IW}\left(v_{0},\left(v_{0}-m-1\right) \Sigma_{0}\right)
$$

This thing is called the Inverse-Wishart distribution.

A PDF over SPD matrices!

## Dain~ D...ñinn. NA^D estimates for Gaussians

- $v_{0}$ small: "I am not sure
 guess of $\Sigma$
- $v_{0}$ large: "I'm pretty sure imiz about my guess of $\Sigma_{0}{ }^{\prime \prime} \quad \mathrm{E}[\Sigma]=\Sigma_{0}$

Step 1: Pu Ior on ( $\mu$,
Step 1a: Put , rior on $\Sigma$

$$
\left(v_{0}-m-1\right) \Sigma \sim \mathcal{W}\left(v_{0},\left(v_{0}-m-1\right) \Sigma_{0}\right)
$$

This thing is called the Inverse-Wishart distribution.

A PDF over SPD matrices!

Being Bayesian: MAP estimates for Gaussians

- Suppose you have $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R} \sim($ i.i.d) $N(\mu, \Sigma)$
- But you don't know $\mu$ or $\Sigma$
- MAP: Which $(\mu, \Sigma)$ maximizes $\mathrm{p}\left(\mu, \Sigma \mid \mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R}\right)$ ?

Step 1: Put a prior on ( $\mu, \Sigma$ )
Step 1a: Put a prior on $\Sigma$

$$
\left.\begin{array}{cl}
\left(v_{0}-\mathrm{m}-1\right) \Sigma \sim \operatorname{IW}\left(v_{0}\left(v_{0}-\mathrm{m}-1\right) \Sigma_{0}\right) \\
\text { Step 1b: Put a prior on } \mu \mid \Sigma \\
\mu \mid \Sigma \sim N\left(\mu_{0}, \Sigma / \kappa_{0}\right)
\end{array}\right\} \begin{aligned}
& \text { Together, " } \Sigma \text { " and } \\
& \text { " } \mu \mid \Sigma \text { " define a } \\
& \text { joint distribution } \\
& \text { on }(\mu, \Sigma)
\end{aligned}
$$

## Being Bayesian: MAP estimates for Gaussians

- Suppose you have $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{n} \sim(i . i . d) N(u . \Sigma)$
- But you don't know $\mu$ or ${ }^{2} \kappa_{0}$ small: "I am not sure
- MAP: Which ( $\mu, \Sigma$ ) maximi about my guess of $\mu_{0}{ }^{\prime \prime}$ )?

$$
\begin{aligned}
& \mu_{0} \text { : My best guess of } \mu \mu, \Sigma \text { ) } \\
& \mathrm{E}[\mu]=\mu_{0} \\
& \Sigma \\
& \left(v_{0}-m-1\right) \Sigma \sim \quad v_{0},\left(v_{0}-m-1\right) \Sigma \quad \text { Together, " } \Sigma \text { " and } \\
& \text { Step 1b: Put a p or on } \mu \mid \Sigma \\
& \mu \mid \Sigma \sim N\left(\mu_{0}, \Sigma / \kappa_{0}\right) \\
& \kappa_{0} \text { large: "I'm pretty sure } \\
& \text { about my guess of } \mu_{0} \text { " } \\
& \text { Together, " } \Sigma \text { " and } \\
& \text { " } \mu \mid \Sigma \text { " define a } \\
& \text { joint distribution } \\
& \text { on ( } \mu, \Sigma \text { ) }
\end{aligned}
$$

Notice how we are forced to express our ignorance of $\mu$ proportionally to $\Sigma$

## Being Bayesian: MAP estimates for Gaussians

- Suppose you have $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R} \sim($ i.i.d) $N(\mu, \Sigma)$
- But you don't know $\mu$ or $\Sigma$
- MAP: Which $(\mu, \Sigma)$ maximizes $\mathrm{p}\left(\mu, \Sigma \mid \mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R}\right)$ ?

Step 1: Put a prior on ( $\mu, \Sigma$ )
Why do we use this form of prior?
Step la: Put a prior on $\Sigma$

$$
\left(v_{0}-m-1\right) \Sigma \sim \operatorname{IW}\left(v_{0},\left(v_{0}-m-1\right) \Sigma_{0}\right)
$$

Step 1b: Put a prior on $\mu \mid \Sigma$

$$
\mu \mid \Sigma \sim N\left(\mu_{0}, \Sigma / \kappa_{0}\right)
$$

## Being Bayesian: MAP estimates for Gaussians

- Suppose you have $\mathbf{x}_{1,}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R} \sim(i . i . d) N(\mu, \Sigma)$
- But you don't know $\mu$ or $\Sigma$
- MAP: Which ( $\mu, \Sigma$ ) maximizes $\mathrm{p}\left(\mu, \Sigma \mid \mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R}\right)$ ?

Step 1: Put a prior on ( $\mu, \Sigma$ )
Step 1a: Put a prior on $\Sigma$
$\left(v_{0}-m-1\right) \Sigma \sim \operatorname{IW}\left(v_{0},\left(v_{0}-m-1\right) \Sigma_{0}\right)$
Step 1b: Put a prior on $\mu \mid \Sigma$
Why do we use this form of prior?
Actually, we don't have to But it is computationally and algebraically convenient...
...it's a conjugate prior.

$$
\mu \mid \Sigma \sim N\left(\mu_{0}, \Sigma / \kappa_{0}\right)
$$

## Being Bayesian: MAP estimates for Gaussians

- Suppose you have $\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R} \sim($ i.i.d) $\mathrm{N}(\mu, \Sigma)$
- MAP: Which $(\mu, \Sigma)$ maximizes $\mathrm{p}\left(\mu, \Sigma \mid \mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \mathbf{x}_{R}\right)$ ?

Step 1: Prior: $\left(v_{0}-m-1\right) \Sigma \sim \operatorname{IW}\left(v_{0},\left(v_{0}-m-1\right) \Sigma_{0}\right), \mu \mid \Sigma \sim N\left(\mu_{0}, \Sigma / \kappa_{0}\right)$
Step 2:

$$
\overline{\mathbf{x}}=\frac{1}{R} \sum_{k=1}^{R} \mathbf{x}_{k} \quad \boldsymbol{\mu}_{R}=\frac{\kappa_{0} \boldsymbol{\mu}_{0}+R \overline{\mathbf{x}}}{\kappa_{0}+R} \quad \begin{aligned}
& v_{R}=v_{0}+R \\
& \kappa_{R}=\kappa_{0}+R
\end{aligned}
$$

$\left(v_{R}+m-1\right) \boldsymbol{\Sigma}_{R}=\left(v_{0}+m-1\right) \boldsymbol{\Sigma}_{0}+\sum_{k=1}^{R}\left(\mathbf{x}_{k}-\overline{\mathbf{x}}\right)\left(\mathbf{x}_{k}-\overline{\mathbf{x}}\right)^{T}+\frac{\left(\overline{\mathbf{x}}-\boldsymbol{\mu}_{0}\right)\left(\overline{\mathbf{x}}-\boldsymbol{\mu}_{0}\right)^{T}}{1 / \kappa_{0}+1 / R}$
Step 3: Posterior: $\left(v_{R}+m-1\right) \Sigma \sim \operatorname{IW}\left(v_{R},\left(v_{R}+m-1\right) \Sigma_{R}\right)$,

$$
\mu \mid \Sigma \sim N\left(\mu_{R}, \Sigma / \kappa_{R}\right)
$$

Result: $\mu^{\text {map }}=\mu_{\mathrm{R}}, \mathrm{E}\left[\Sigma \mid \mathbf{X}_{1}, \mathbf{X}_{2}, \ldots \mathbf{X}_{R}\right]=\Sigma_{\mathrm{R}}$

## Being Bayesian Look carefully at what these formulae are doing. It's all very sensible.

- Suppose you hav ${ }^{\bullet}$ Conjugate priors mean prior form and posterior
- MAP: Which ( $\mu, \Sigma$ statistics" of the data.

Step 1: Prior: $\left(v_{0}-m-1\right) \Sigma \sim$ -The marginal distribution on $\mu$ is a student-t
-One point of view: it's pretty academic if $R>30$
Step 2:

Step 2: \begin{tabular}{|l|l|l|}

$\overline{\mathbf{x}}=\frac{1}{R} \sum_{k=1}^{R} \mathbf{x}_{k}$ \& $\boldsymbol{\mu}_{R}=\frac{\kappa_{0} \boldsymbol{\mu}_{0}+R \overline{\mathbf{x}}}{\kappa_{0}+R}$ \& | $v_{R}=v_{0}+R$ |
| :--- |
| $\kappa_{R}=\kappa_{0}+R$ | <br>

\hline
\end{tabular}

$\left(v_{R}+m-1\right) \boldsymbol{\Sigma}_{R}=\left(v_{0}+m-1\right) \boldsymbol{\Sigma}_{0}+\sum_{k=1}^{R}\left(\mathbf{x}_{k}-\overline{\mathbf{x}}\right)\left(\mathbf{x}_{k}-\overline{\mathbf{x}}\right)^{T}+\frac{\left(\overline{\mathbf{x}}-\boldsymbol{\mu}_{0}\right)\left(\overline{\mathbf{x}}-\boldsymbol{\mu}_{0}\right)^{T}}{1 / \kappa_{0}+1 / R}$
Step 3: Posterior: $\left(v_{\mathrm{R}}+\mathrm{m}-1\right) \Sigma \sim \operatorname{IW}\left(v_{\mathrm{R}},\left(v_{\mathrm{R}}+m-1\right) \Sigma_{\mathrm{R}}\right)$,
$\mu \mid \Sigma \sim N\left(\mu_{R}, \Sigma / \kappa_{R}\right)$
Result: $\mu^{\text {map }}=\mu_{\mathrm{R}}, \mathrm{E}\left[\Sigma \mid \mathbf{X}_{1}, \mathbf{X}_{2}, \ldots \mathbf{x}_{R}\right]=\Sigma_{\mathrm{R}}$

## Where we're at

|  | Categorical inputs only | Real-valued inputs only | Mixed Real / Cat okay |
| :---: | :---: | :---: | :---: |
| Classifier Predict | Joint BC Naïve BC |  | Dec Tree |
| $\xrightarrow{2}:$Density <br> Estimator$\rightarrow$Prob- <br> ability | Joint DE Naïve DE | Gauss DE |  |
| Regressor Predict |  |  |  |

## What you should know

- The Recipe for MLE
- What do we sometimes prefer MLE to MAP?
- Understand MLE estimation of Gaussian parameters
- Understand "biased estimator" versus "unbiased estimator"
- Appreciate the outline behind Bayesian estimation of Gaussian parameters


## Useful exercise

- We'd already done some MLE in this class without even telling you!
- Suppose categorical arity-n inputs $\mathrm{x}_{1}, \mathrm{x}_{2}, \ldots$ $x_{R} \sim$ (i.i.d.) from a multinomial

$$
M\left(p_{1}, p_{2}, \ldots p_{n}\right)
$$

where

$$
\mathrm{P}\left(\mathrm{x}_{k}=j \mid \mathbf{p}\right)=\mathrm{p}_{j}
$$

- What is the MLE $\mathbf{p}=\left(\mathrm{p}_{1}, \mathrm{p}_{2}, \ldots \mathrm{p}_{n}\right)$ ?

