

## The Central Limit Theorem

- If $\left(X_{1}, X_{2}, \ldots X_{n}\right)$ are i.i.d. continuous random variables
- Then define $z=f\left(x_{1}, x_{2}, \ldots x_{n}\right)=\frac{1}{n} \sum_{i=1}^{n} x_{i}$
- As n-->infinity, $\mathrm{p}(\mathrm{z})--->$ Gaussian with mean $E\left[X_{i}\right]$ and variance $\operatorname{Var}\left[X_{i}\right]$

Somewhat of a justification for assuming Gaussian noise is common



## Bivariate Gaussians

Write r.v. $\mathbf{X}=\binom{X}{Y} \quad$ Then define $\quad X \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ to mean

$$
p(\mathbf{x})=\frac{1}{2 \pi\|\boldsymbol{\Sigma}\|^{1 / 2}} \exp \left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^{T} \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})\right)
$$

Where the Gaussian's parameters are...

$$
\boldsymbol{\mu}=\binom{\mu_{x}}{\mu_{y}} \quad \boldsymbol{\Sigma}=\left(\begin{array}{ll}
\sigma_{x}^{2} & \sigma_{x y} \\
\sigma_{x y} & \sigma_{y}^{2}
\end{array}\right)
$$

Where we insist that $\Sigma$ is symmetric non-negative definite
It turns out that $\mathrm{E}[\mathrm{X}]=\mu$ and $\operatorname{Cov}[\mathrm{X}]=\boldsymbol{\Sigma}$. (Note that this is a resulting property of Gaussians, not a definition)*

| *This note rates 7.4 on the pedanticness scale |
| :--- |
| Copyright $\odot 2001$, Andrew W. More |



## Multivariate Gaussians

Write r.v. $\mathbf{X}=\left(\begin{array}{c}X_{1} \\ X_{2} \\ \vdots \\ X_{m}\end{array}\right) \quad$ Then define $\quad X \sim N(\boldsymbol{\mu}, \mathbf{\Sigma})$ to mean

$$
p(\mathbf{x})=\frac{1}{(2 \pi)^{m / 2}\|\boldsymbol{\Sigma}\|^{1 / 2}} \exp \left(-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^{T} \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})\right)
$$

Where the Gaussian's
parameters have..

$$
\boldsymbol{\mu}=\left(\begin{array}{c}
\mu_{1} \\
\mu_{2} \\
\vdots \\
\mu_{m}
\end{array}\right) \quad \boldsymbol{\Sigma}=\left(\begin{array}{cccc}
\sigma_{1}^{2} & \sigma_{12} & \cdots & \sigma_{1 m} \\
\sigma_{12} & \sigma^{2} & \cdots & \sigma_{2 m} \\
\vdots & \vdots & \ddots & \vdots \\
\sigma_{1 m} & \sigma_{2 m} & \cdots & \sigma_{m}^{2}
\end{array}\right)
$$

Where we insist that $\Sigma$ is symmetric non-negative definite
Again, $E[X]=\mu$ and $\operatorname{Cov}[\mathrm{X}]=\Sigma$. (Note that this is a resulting property of Gaussians, not a definition)
Copyright © 2001, Andrew W. Moore Gaussians: Slide 13


## Spherical Gaussians



## Where are we now?

- We've seen the formulae for Gaussians
- We have an intuition of how they behave
- We have some experience of "reading" a Gaussian's covariance matrix
- Coming next:

Some useful tricks with Gaussians

Copyright © 2001, Andrew W. Moore
Gaussians: Slide 18

## Subsets of variables

Write $\mathbf{X}=\left(\begin{array}{c}X_{1} \\ X_{2} \\ \vdots \\ X_{m}\end{array}\right)$ as $\mathbf{X}=\binom{\mathbf{U}}{\mathbf{V}}$ where $\begin{array}{r}\mathbf{U}=\left(\begin{array}{c}X_{1} \\ \vdots \\ X_{m(u)}\end{array}\right) \\ \mathbf{V}=\left(\begin{array}{c}X_{m(u)+1} \\ \vdots \\ X_{m}\end{array}\right)\end{array}$

This will be our standard notation for breaking an m dimensional distribution into subsets of variables

## 

Write $\mathbf{X}=\left(\begin{array}{c}X_{1} \\ X_{2} \\ \vdots \\ X_{m}\end{array}\right)$ as $\mathbf{X}=\binom{\mathbf{U}}{\mathbf{V}}$ where $\mathbf{U}=\left(\begin{array}{c}X_{1} \\ \vdots \\ X_{m(u)}\end{array}\right), \mathbf{V}=\left(\begin{array}{c}X_{m(u)+1} \\ \vdots \\ X_{m}\end{array}\right)$
$\operatorname{IF}\binom{\mathbf{U}}{\mathbf{V}} \sim \mathrm{N}\left(\binom{\boldsymbol{\mu}_{u}}{\boldsymbol{\mu}_{v}},\left(\begin{array}{ll}\boldsymbol{\Sigma}_{u u} & \boldsymbol{\Sigma}_{u v} \\ \boldsymbol{\Sigma}_{u v}^{T} & \boldsymbol{\Sigma}_{v v}\end{array}\right)\right)$
THEN $U$ is also distributed as a Gaussian
$\mathbf{U} \sim N\left(\boldsymbol{\mu}_{u}, \boldsymbol{\Sigma}_{u u}\right)$
Copyright © 2001, Andrew W. Moore

\section*{| Gaussian Marginals |
| :--- |
| are Gaussian |
| $\binom{\mathbf{y}}{\mathbf{v}} \rightarrow \xrightarrow{\substack{\text { Margin. } \\ \text { alize }}} \rightarrow \mathbf{U}$ |}

Write $\mathbf{X}=\left(\begin{array}{c}X_{1} \\ X_{2} \\ \vdots \\ X_{m}\end{array}\right)$ as $\mathbf{X}=\binom{\mathbf{U}}{\mathbf{V}}$ where $\left.\mathbf{U}=\binom{X_{1}}{\vdots}, \mathbf{V}=\binom{X_{m(u)+1}}{\vdots}, ~ \begin{array}{c}\text { How would you prove } \\ \text { this? }\end{array}\right)$
$\operatorname{IF}\binom{\mathbf{U}}{\mathbf{V}} \sim \mathrm{N}\left(\binom{\boldsymbol{\mu}_{u}}{\boldsymbol{\mu}_{v}},\left(\begin{array}{ll}\boldsymbol{\Sigma}_{u u} & \boldsymbol{\Sigma}_{u v} \\ \boldsymbol{\Sigma}_{u v}^{T} & \boldsymbol{\Sigma}_{v v}\end{array}\right)\right)$
THEN $U$ is also distributed as a Gaussian this?
$\operatorname{IF}\binom{\mathbf{U}}{\mathbf{V}} \sim \mathrm{N}\left(\binom{\boldsymbol{\mu}_{u}}{\boldsymbol{\mu}_{v}},\left(\begin{array}{ll}\boldsymbol{\Sigma}_{u u} & \boldsymbol{\Sigma}_{u v} \\ \boldsymbol{\Sigma}_{u v}^{T} & \boldsymbol{\Sigma}_{v v}\end{array}\right)\right)$
This fact is not immediately obvious

Obvious, once we know it's a Gaussian (why?)
(
Gaussians: Slide 21
$\mathbf{U} \sim \mathrm{N}\left(\boldsymbol{\mu}_{u}, \boldsymbol{\Sigma}_{u u}\right)$
$=\int p(\mathbf{u}, \mathbf{v}) d \mathbf{v}$
$=$ (snore...)
Copyright © 2001, Andrew W. Moore
Gaussians: Slide 22
Conditional of Gaussian is Gaussian

$\operatorname{IF}\binom{\mathbf{U}}{\mathbf{V}} \sim \mathrm{N}\left(\binom{\boldsymbol{\mu}_{u}}{\boldsymbol{\mu}_{v}},\left(\begin{array}{cc}\boldsymbol{\Sigma}_{u u} & \boldsymbol{\Sigma}_{u v} \\ \boldsymbol{\Sigma}_{u v}^{T} & \boldsymbol{\Sigma}_{v v}\end{array}\right)\right)$ THEN $\mathbf{U} \mid \mathbf{V} \sim \mathrm{N}\left(\boldsymbol{\mu}_{u \mid v}, \boldsymbol{\Sigma}_{u \mid v}\right)$ where

$$
\begin{gathered}
\boldsymbol{\mu}_{u \mid v}=\boldsymbol{\mu}_{u}+\boldsymbol{\Sigma}_{u v}^{T} \boldsymbol{\Sigma}_{v v}^{-1}\left(\mathbf{V}-\boldsymbol{\mu}_{v}\right) \\
\boldsymbol{\Sigma}_{u \mid v}=\boldsymbol{\Sigma}_{u u}-\boldsymbol{\Sigma}_{u v}^{T} \boldsymbol{\Sigma}_{v v}^{-1} \boldsymbol{\Sigma}_{u v}
\end{gathered}
$$



Copyright © 2001, Andrew W. Moore





## Assume...

- You are an intellectual snob
- You have a child


## Intellectual snobs with children

- ...are obsessed with IQ
- In the world as a whole, IQs are drawn from a Gaussian N(100,15²)



## IQ tests

- If you take an IQ test you'll get a score that, on average (over many tests) will be your IQ
- But because of noise on any one test the score will often be a few points lower or higher than your true IQ.
SCORE | IQ ~N(IQ,10²)


## Assume...

- You drag your kid off to get tested
- She gets a score of 130
- "Yippee" you screech and start deciding how to casually refer to her membership of the top 2\% of IQs in your Christmas newsletter.

$\mathrm{P}\left(\mathrm{X}<130 \mid \mu=100, \sigma^{2}=15^{2}\right)=$ $\mathrm{P}\left(\mathrm{X}<2 \mid \mu=0, \sigma^{2}=1\right)=$ $\operatorname{erf}(2)=0.977$


## Maximum Likelihood IQ

- IQ~N(100,15²)
- SIIQ ~ N(IQ, $\left.10^{2}\right)$
- $\mathrm{S}=130$

$$
I Q^{m l e}=\underset{i q}{\arg \max } p(s=130 \mid i q)
$$

- The MLE is the value of the hidden parameter that makes the observed data most likely
- In this case

$$
I Q^{m l e}=130
$$

Copyright © 2001, Andrew W. Moore
Gaussians: Slide 36


## What we really want:

- IQ~N(100,15²)
- SIIQ ~N(IQ, $\left.10^{2}\right)$
- $\mathrm{S}=130$
- Question: What is IQ | (S=130)?


## Called the Posterior

 Distribution of IQCopyright © 2001, Andrew W. Moore

## Which tool or tools?

- $\operatorname{IQ} \sim N\left(100,15^{2}\right)$
- SIIQ ~N(IQ, $\left.10^{2}\right)$
- $\mathrm{S}=130$
- Question: What is IQ | $(S=130)$ ?


$$
\begin{aligned}
& \mathbf{U} \left\lvert\, \mathbf{V} \rightarrow \begin{array}{c}
\text { Chain } \\
\text { Rule }
\end{array}\right. \rightarrow\binom{\mathbf{U}}{\mathbf{V}} \\
& \text { Gaussians: Slide } 39
\end{aligned}
$$

Copyright © 2001, Andrew W. Moore


Plan

- IQ~N(100,15²)
- SIIQ ~N(IQ, $\left.10^{2}\right)$
- $\mathrm{S}=130$
- Question: What is IQ | $(S=130)$ ?
$\underset{\sim}{S \mid I Q} \rightarrow \underset{\substack{\text { Chain } \\ \text { Rule }}}{ } \rightarrow\binom{S}{I Q} \rightarrow$ Swap $\left.\rightarrow\binom{I Q}{S} \rightarrow \underset{\substack{\text { Condition- } \\ \text { alize }}}{\text { a }} \rightarrow \mathrm{IQ} \right\rvert\, S$

Copyright © 2001, Andrew W. Moore
Gaussians: Slide 40

That was an important result! It explains how to combine noisy measurements (sensor fusion)

So I will do it again in 1D

> Combining Measurements: 1D
> - True value $x$
> - Measurements $m_{1}, m 2: E\left(m_{1}-x\right)=0, \operatorname{Var}\left(m_{1}\right)=\sigma_{1}{ }^{2}$, $E\left(m_{2}-x\right)=0, \operatorname{Var}\left(m_{2}\right)=\sigma_{2}{ }^{2}$, independent
> - Linear estimate $x=k_{1} m_{1}+k_{2} m_{2}$
> - Unbiased estimate means $\mathrm{k}_{2}=1-\mathrm{k}_{1}$ so $\mathrm{E}(x)=\mathrm{x}$
> - Minimize $\operatorname{Var}(x)=\mathrm{k}_{1}{ }^{2} \sigma_{1}{ }^{2}+\left(1-\mathrm{k}_{1}\right)^{2} \sigma_{2}{ }^{2}$
> - So $\partial \operatorname{Var}(x) / \partial \mathrm{k}_{1}=0 \rightarrow 2 \mathrm{k}_{1}\left(\sigma_{1}{ }^{2}+\sigma_{2}{ }^{2}\right)-2 \sigma_{2}{ }^{2}=0$
> - So $\mathrm{k}_{1}=\sigma_{2}{ }^{2} /\left(\sigma_{1}{ }^{2}+\sigma_{2}{ }^{2}\right), \mathrm{k}_{2}=\sigma_{1}{ }^{2} /\left(\sigma_{1}{ }^{2}+\sigma_{2}{ }^{2}\right)$
> So $\operatorname{Var}(x)=\sigma_{1}{ }^{2} \sigma_{2}{ }^{2} /\left(\sigma_{1}{ }^{2}+\sigma_{2}{ }^{2}\right)$
> What happens when $\sigma_{2}{ }^{2}=0$ ? $\sigma_{2}{ }^{2}=$ infinity?
> BLUE: Best Linear Unbiased Estimator
> Copyright © 2001, Andrew w. Moore Gaussians: Slide 43

Your pride and joy's posterior IQ

- If you did the working, you now have $\mathrm{p}(\mathrm{IQ} \mid \mathrm{S}=130)$
- This is a density, not a number!
- If you have to give the most likely IQ given the score you should give

$$
I Q^{\operatorname{map}}=\underset{i q}{\arg \max } p(i q \mid s=130)
$$

- This is the mean for a Gaussian
- MAP means "Maximum A-posteriori"

What you should know

- The Gaussian PDF formula off by heart
- Understand the workings of the formula for a Gaussian
- Be able to understand the Gaussian tools described so far
- Have a rough idea of how you could prove them
- Be happy with how you could use them
- Understand the Bayesian approach to combining information Gaussians: Slide 45

Copyright © 2001, Andrew W. Moore

