Chapter 1: Review

$$\underset{\sim}{f}: \mathbb{R}^p \longrightarrow \mathbb{R}^k \qquad \underset{\sim}{f}(\underset{\sim}{x}) = \begin{bmatrix} f_1(\underset{\sim}{x}) \\ \vdots \\ f_k(\underset{\sim}{x}) \end{bmatrix}$$

$$\underset{\sim}{x} \longmapsto \underset{\sim}{y}$$

Jacobian $\underset{\sim}{f}'(\underset{\sim}{x}) = \left[\frac{\partial f_i}{\partial x_j}(\underset{\sim}{x}) \right]$

$k \times p$

also denoted $J(\underset{\sim}{x}) = J_{\underset{\sim}{f}}(\underset{\sim}{x})$

Why? Linear approximation to diff. fcn. at $\underset{\sim}{x}_0$

$$\underset{\sim}{\hat{f}}_{\underset{\sim}{x}_0}(\underset{\sim}{x}) = \underset{\sim}{f}(\underset{\sim}{x}_0) + \underset{\sim}{f}'(\underset{\sim}{x}_0)(\underset{\sim}{x} - \underset{\sim}{x}_0)$$

$k \times 1$ $\quad p \times 1$ $\qquad k \times 1$ $\qquad k \times p$ $\quad p \times 1$

$p = 2 \quad k = 1$

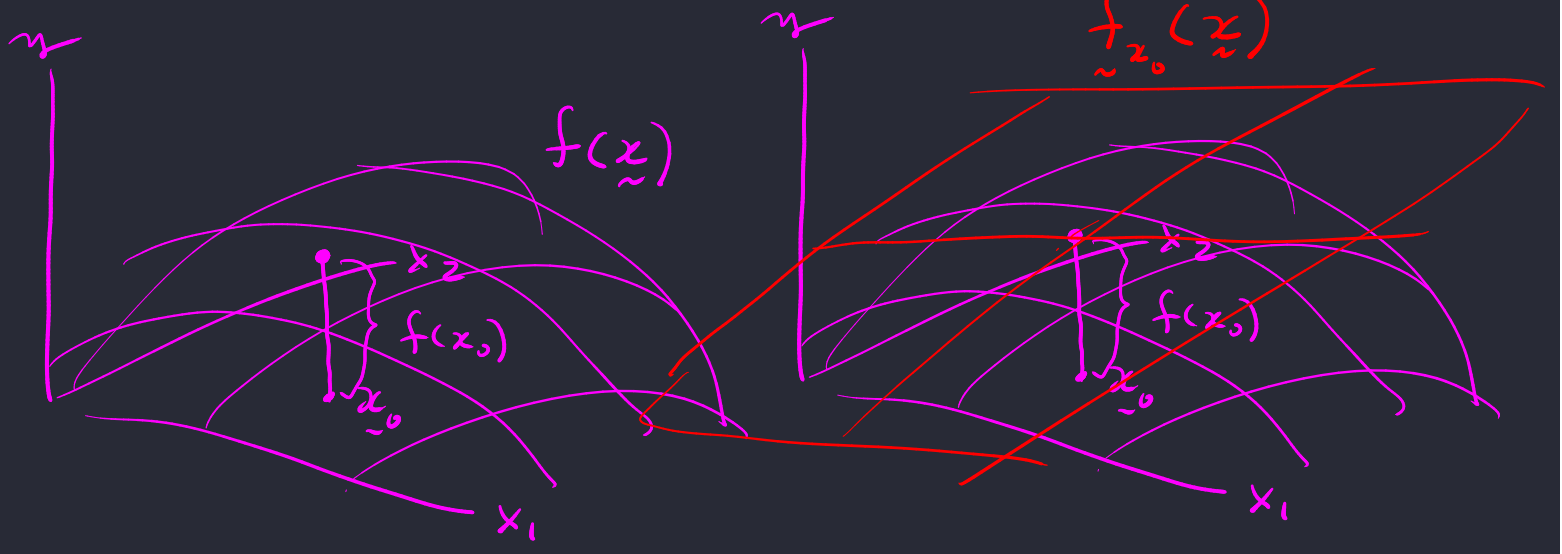

Special case: $k = p$

Then J is square and

$|J(\underset{\sim}{x})|$ is Jacobian determinant.

If $\underset{\sim}{f}$ is 1-1, differentiable $\mathbb{R}^p \to \mathbb{R}^p$

Under conditions: $\Omega \subseteq \mathbb{R}^p$:

$$\int_{\Omega} h(\underset{\sim}{x})\, d\underset{\sim}{x} = \int_{f(\Omega)} \frac{h(f^{-1}(\underset{\sim}{y}))}{\left\| \underbrace{J(f^{-1}(\underset{\sim}{y}))}_{\text{determinant}} \right\|}\, d\underset{\sim}{y}$$

← absolute value

$$\int_{\Omega} d\underset{\sim}{x}$$

Area (Volume) of Ω

$$\int_{f(\Omega)} d\underset{\sim}{y} = \int_{\Omega} \left\| J_f(\underset{\sim}{x}) \right\| d\underset{\sim}{x}$$

Area (Volume) of $f(\Omega)$

Jacobian of a linear (affine) transformation

Let A be a square matrix

$$\underset{\sim}{f}(\underset{\sim}{x}) = \underset{\sim}{a} + A\,\underset{\sim}{x}$$

Then $\underset{\sim}{f}'(x) = A$ (constant for all $\underset{\sim}{x}$)

Is $\underset{\sim}{f}'(\underset{\sim}{x}) = A\,\underset{\sim}{x}$? ☐ Yes ☐ No Why?

Let A be a square matrix $\ni |A| \neq 0$

$$\underset{\sim}{f}(\underset{\sim}{x}) = \underset{\sim}{a} + A\,\underset{\sim}{x}$$

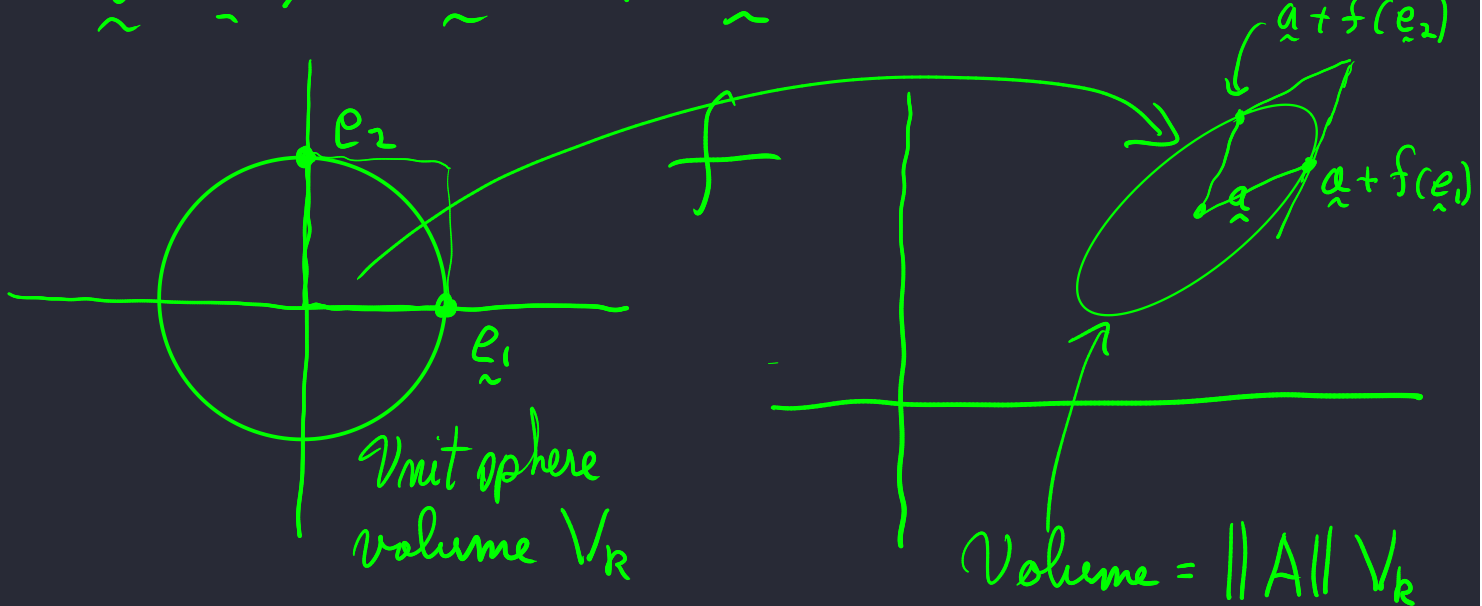

Transforming a random vector:

Random vector: $\underset{\sim}{X}$ in $\mathbb{R}^p$, density h_X

$\underset{\sim}{Y} = \underset{\sim}{f}(\underset{\sim}{X})$, f 1-1 + differentiable

then density for $\underset{\sim}{Y}$ is

$$h_{\underset{\sim}{Y}}(\underset{\sim}{y}) = \frac{h_{\underset{\sim}{X}}(\underset{\sim}{f}^{-1}(\underset{\sim}{y}))}{\|J(f^{-1}(\underset{\sim}{y}))\|}$$

Special case: $k = 1$

$f'(\underset{\sim}{x})$ is called the [redacted]

of $P = 1$: [redacted]

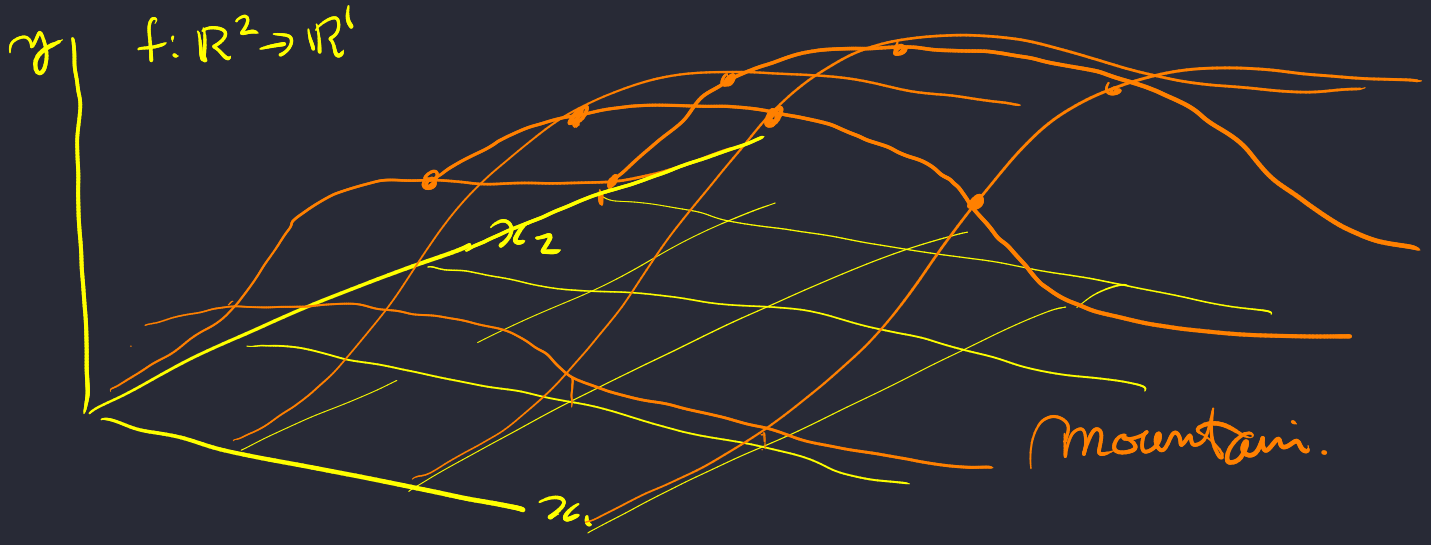

From above:

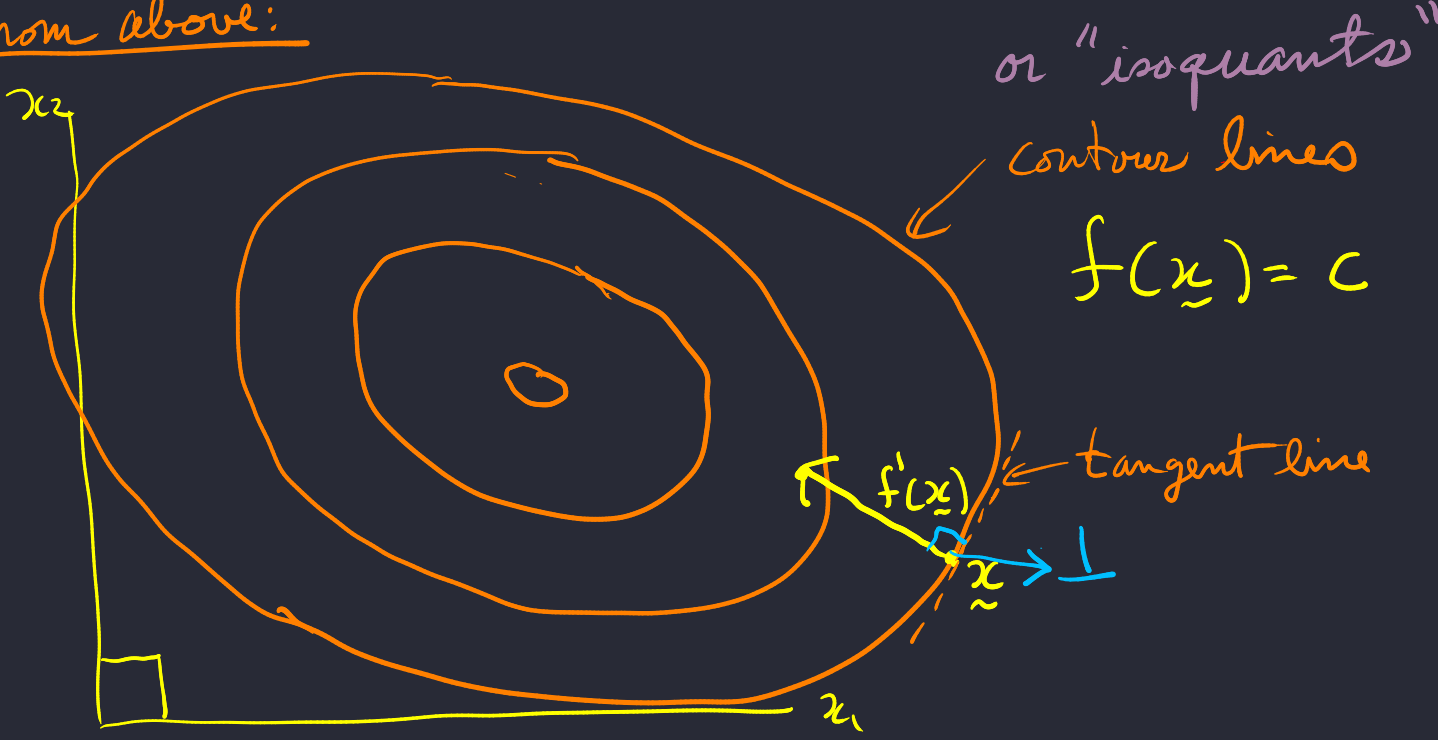

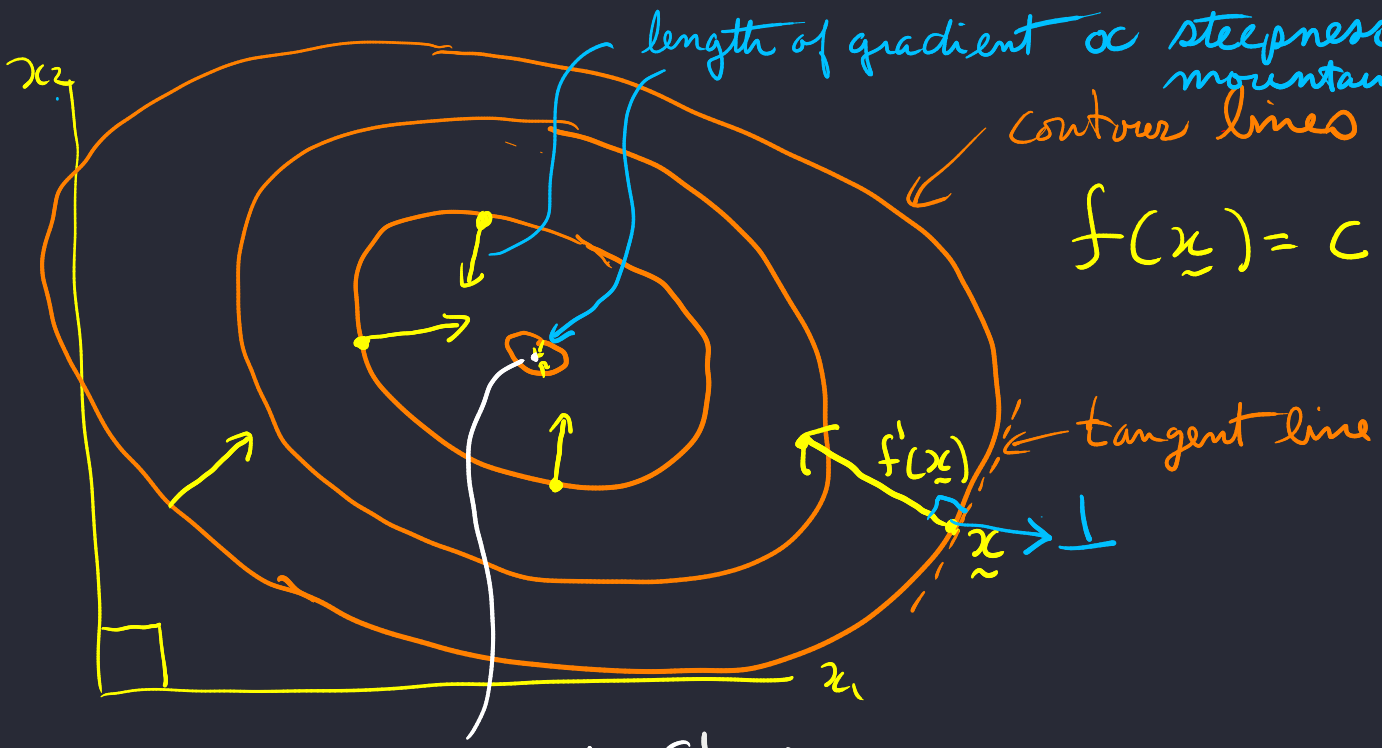

At the summit $f'(\underset{\sim}{x}) = \underset{\sim}{0}$

Note: $f(\underset{\sim}{x}) \in \mathbb{R}^1$

$f'(\underset{\sim}{x}) \in \mathbb{R}^2$

Hessian — 2nd derivative

$f : \mathbb{R}^p \to \mathbb{R}$

$$f''(\underset{\sim}{x}) = \left[\frac{\partial f_i}{\partial x_j}(\underset{\sim}{x}) \right] \text{ is a } p \times p \text{ matrix}$$

Quadratic approximation at $\underset{\sim}{x}_0$

$$\hat{f}_{\underset{\sim}{x}_0}(\underset{\sim}{x}) = f(\underset{\sim}{x}_0) + f'(\underset{\sim}{x}_0)(\underset{\sim}{x} - \underset{\sim}{x}_0) + \frac{1}{2}(\underset{\sim}{x} - \underset{\sim}{x}_0)' f''(\underset{\sim}{x}_0)(\underset{\sim}{x} - \underset{\sim}{x}_0)$$

... and beyond ... Taylor's theorem.

Inference:

- Get a sample from an unknown distribution (population).
- What does the sample tell you about the distribution

Example: Distribution: $N(\underbrace{\mu, \sigma^2}_{\text{unknown}})$

Sample $n = 10$

0.85887292	0.05047837	1.18787131	1.18794773	-0.15461645	-0.75116726	0.10407667	-0.26744312	0.10795396	1.44423848

mean = 0.3768 sd = 0.7386

So ...?

(log) likelihood function:

$\ln f(\underset{\sim}{x}; \mu, \sigma)$ as a function of μ & σ

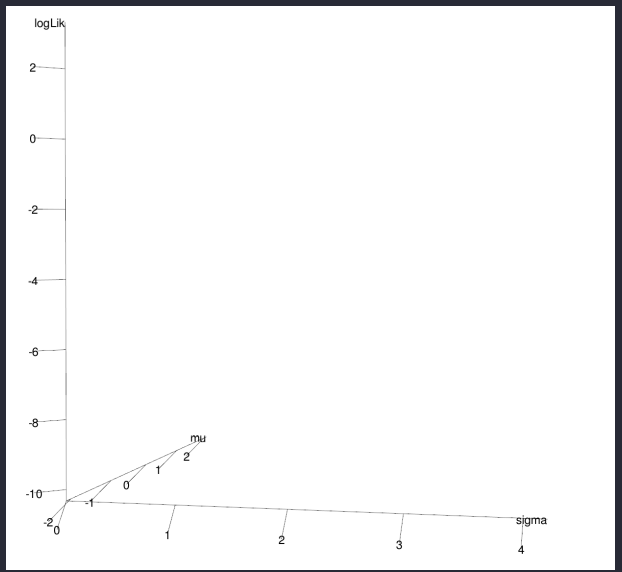

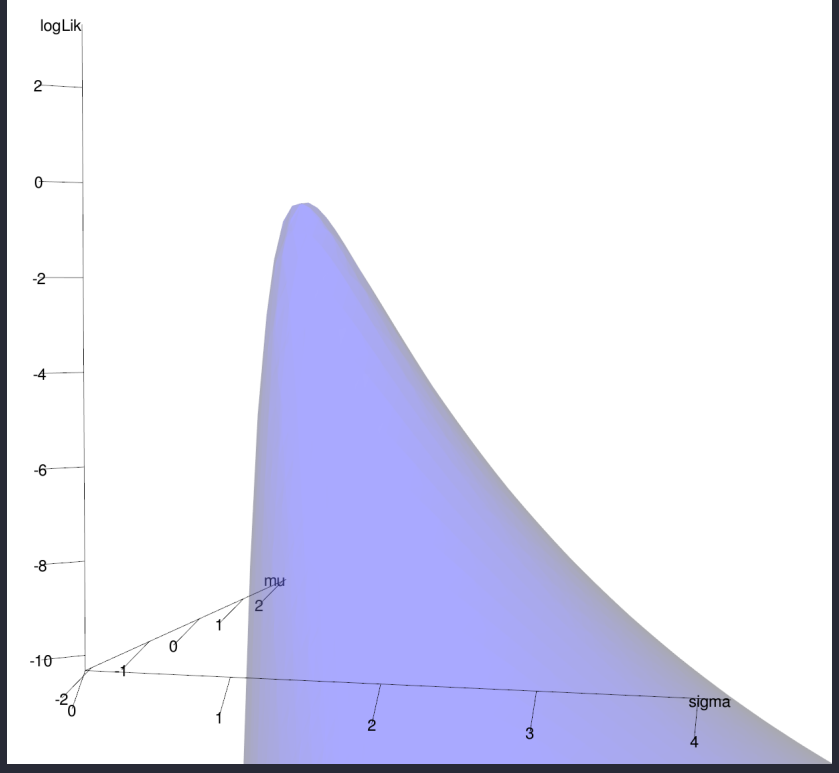

What do we do with this???

- Find the summit?
- Slice it?
- Do a quadratic approximation?
- Consider a gradient at some "H_0"?
- Multiply it by a prior density?
  - and use Bayes Formula to get a posterior.
- Multiply it by a prior and use
  - MCMC
  - HMC

} to simulate from the posterior.

Answer:

The great divide

Everyone uses the likelihood.

Bayesian

Combine it with a prior to create a posterior probability for parameters

Frequentist

Priors are "subjective"

- Posit one or more hypothetical "true" value.
- Consider the sampling variability of the likelihood under these hypothetical values.

Bayesian: Bayes formula

$$\pi(\underset{\sim}{\theta} \mid \underset{\sim}{x}_{obs}) = \frac{\overbrace{f(\underset{\sim}{x}_{obs} \mid \underset{\sim}{\theta})}^{\text{likelihood}} \overbrace{\pi(\underset{\sim}{\theta})}^{\text{prior}}}{\underbrace{\int f(\underset{\sim}{x}_{obs} \mid \underset{\sim}{\theta}) \pi(\underset{\sim}{\theta}) \, d\underset{\sim}{\theta}}_{\text{norming "constant"}}}$$

FINDING THIS CAN BE HARD

Some modern solutions:

Use only $f(x_{obs} \mid \underset{\sim}{\theta})\,\pi(\underset{\sim}{\theta})$

and MCMC or HMC

Frequentist methods

All based on sampling variation of likelihood.

i.e. What would happen if we were to resample over and over.

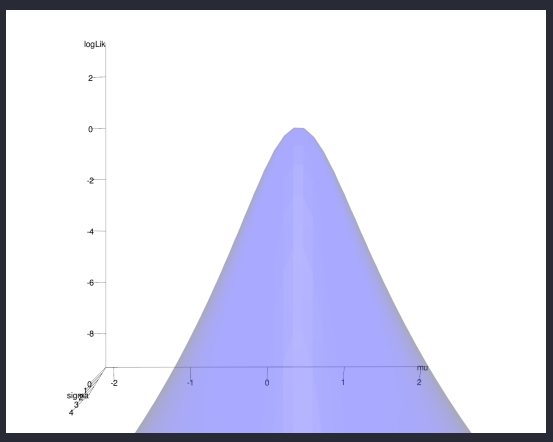

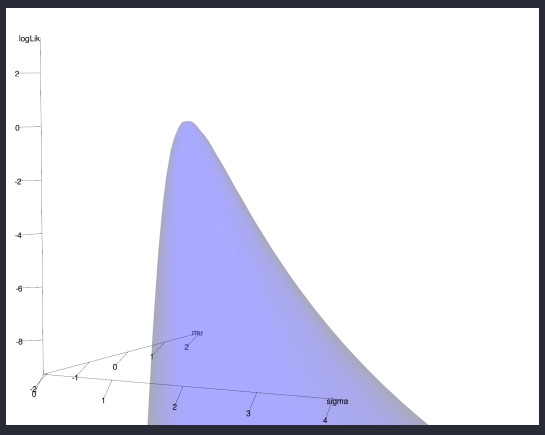

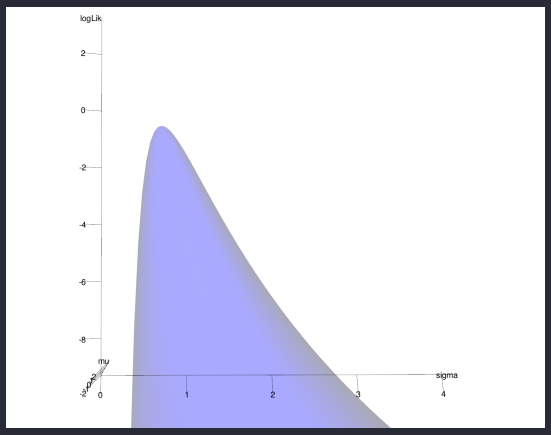

Here, we know from normal theory that max occurs at

$\mu = \bar{x}$, $\sigma = \hat{\sigma}$

$= 0.3768$ $= 0.7387$

How close to "true" value?

Three standard approaches

1) Likelihood Ratio Tests (Wilk's)
   Use random variability in height of log lik.
2) Wald Tests: Use quadratic approx. at MLE.
3) Fisher score (aka "Rao") test
   Use gradient at hypothetical value

> Wouldn't it be better to use the gradient at the MLE?

why? or Why not?

Likelihood Ratio = log Likelihood difference

Take a hypothetical value, e.g. $\mu=0, \sigma=1$ and consider random variability in height of log Lik. under repeated sampling.

Issue: height is arbitrary for a continuous distribution. So need to normalize by specifying height (doesn't matter which) at a choice of μ, σ.

Common choice: height = 0 when μ, σ at hypothetical values.

Let's take a number of samples from $N(0,1)$

```
> set.seed(4567)
> sams <- lapply(rep(10,5), rnorm, mean = 0, sd = 1)
> sams
[[1]]
 [1] -0.7358250 -0.9025461  0.2524151  0.6150259  1.3545344  1.6023185  0.4432036  0.2114059 -0.3523372  1.5262018

[[2]]
 [1]  1.5186540  1.2311758 -0.3413057  0.3480193  0.1033780  0.3460563  0.1758544  0.9994781 -1.6544060 -1.4433098

[[3]]
 [1]  0.18435328 -0.18692255 -0.05300318  0.67699652  1.72128013  1.15726733 -0.02079913  0.30526505  0.01921551 -0.34847418

[[4]]
 [1]  1.72528724 -0.11417979 -0.47668452  1.17492289  0.87798073 -1.10289564 -0.94173811 -0.01198302  0.04414184  0.19493947

[[5]]
 [1] -0.7108700  0.3049154  0.6592707 -1.0953988  0.6362406  0.5217352  0.4086028 -0.5748284  1.3277637 -0.5658254
```

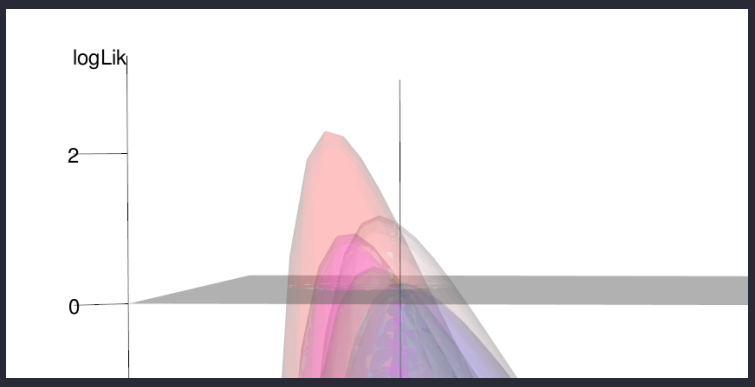

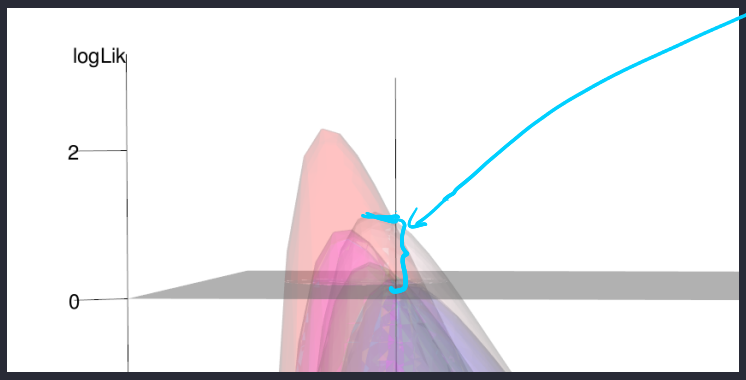

5 random log Likelihood mountains

Summit altitude above 0 has asymptotic $\frac{1}{2}\chi^2_d$ distribution

d = # of parameters

So back to original data.
What to do with:

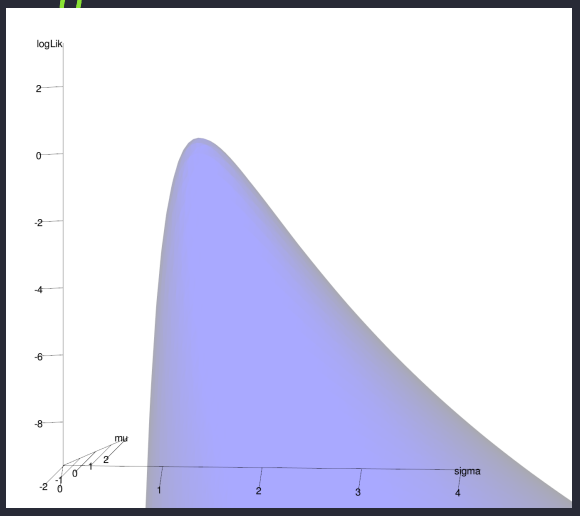

Answer: Invert sampling distribution

Take all points (μ, σ) such that the height of the summit is within $\frac{1}{2}\chi^{2,.95}_2 = 2.9957$ of height of log Likelihood.

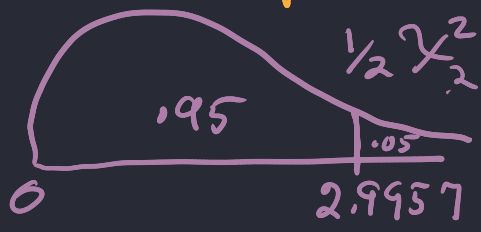

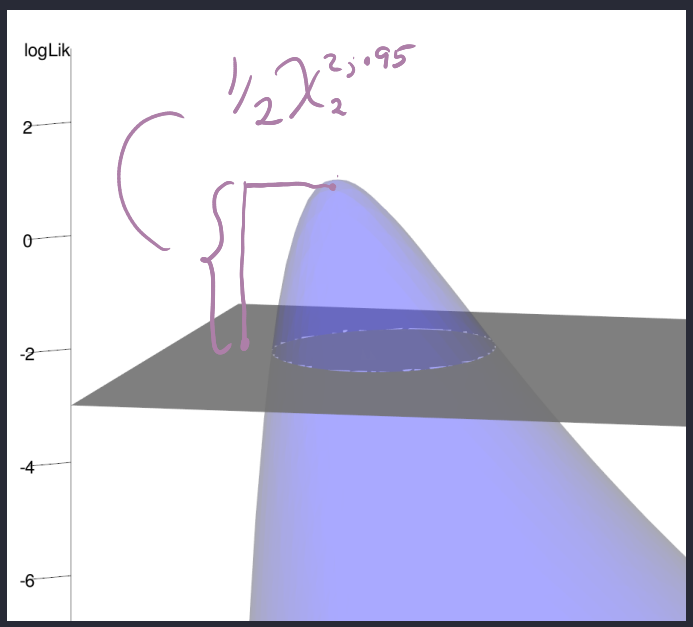

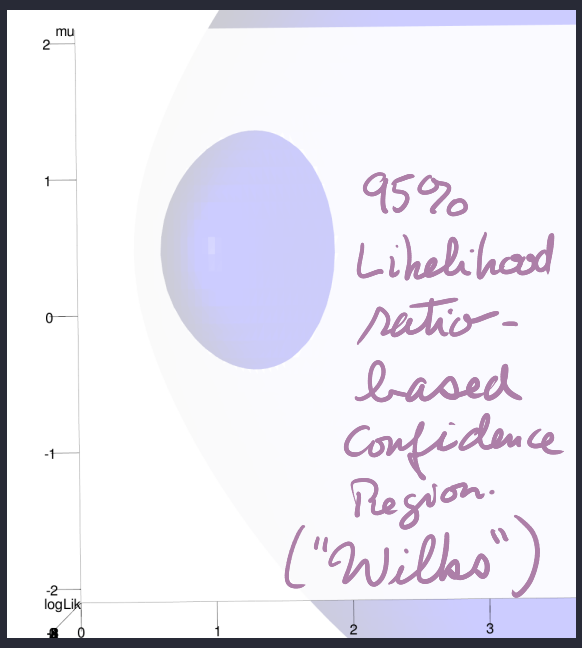

Wald: Same idea with quadratic approx at MLE (shape?)

Gradient = Score :

Looking at likelihood from above

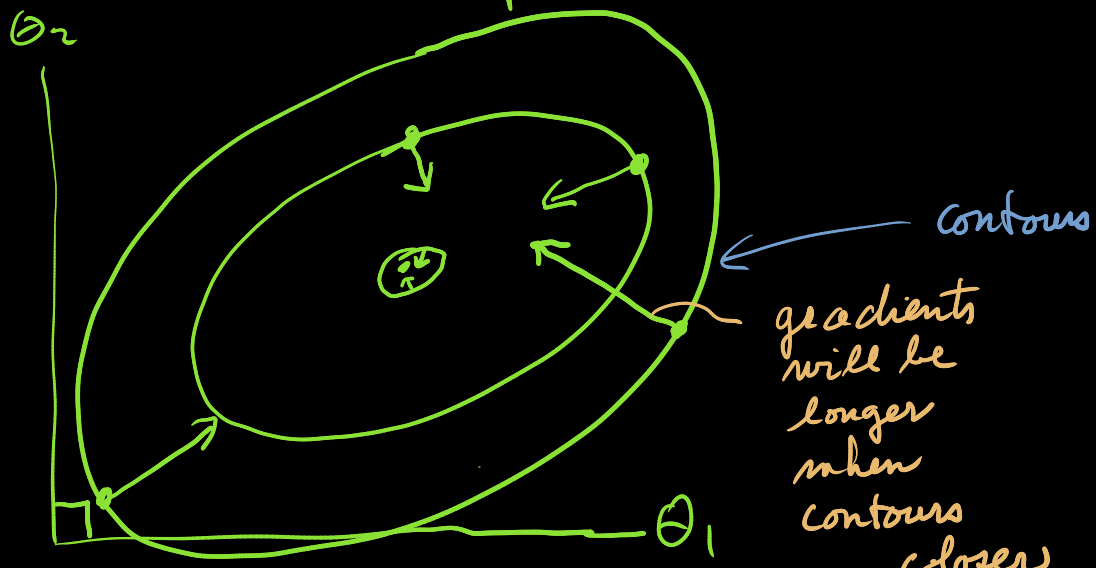

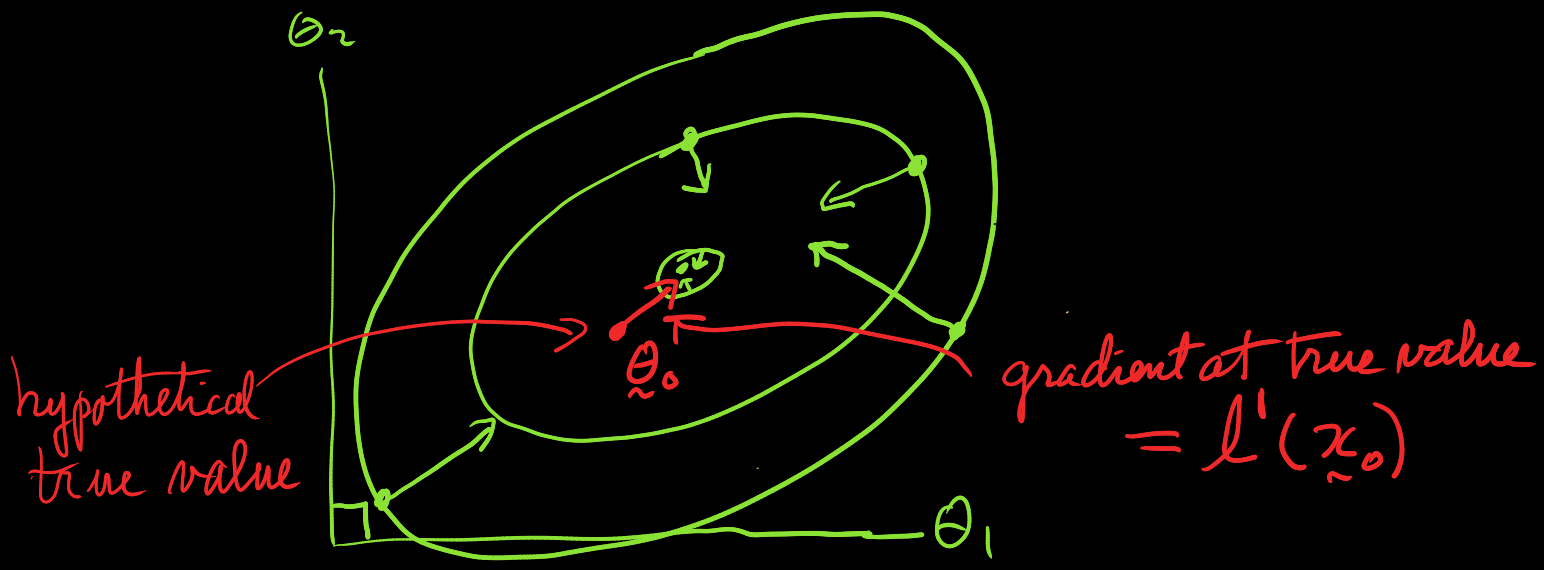

Score Moments :

① $E_{\underset{\sim}{\theta}_0}(l'(\underset{\sim}{\theta}_0)) = \underset{\sim}{0}$

② $Var_{\underset{\sim}{\theta}_0}(l'(\underset{\sim}{\theta}_0)) = E_{\underset{\sim}{\theta}_0}(-l''(\underset{\sim}{\theta}_0))$

$= I(\theta_0)$ "Fisher information"

Proof: Start with $\int f(\underset{\sim}{y} \mid \theta)\, d\mu(\underset{\sim}{y}) = 1$

Regularity assumptions:

- Fixed support
- $\frac{d}{d\underset{\sim}{\theta}}$ & $\int$ commute many times.

$$\int \underbrace{f(\underset{\sim}{y} \mid \theta)}_{\text{function of } \underset{\sim}{y} \text{ \& } \underset{\sim}{\theta}}\, d\mu(\underset{\sim}{y}) = 1$$

function of $\square$

$$\frac{d}{d\theta} \int f(\underset{\sim}{y} \mid \underset{\sim}{\theta})\, d\mu(\underset{\sim}{y}) = \square$$

$$= \int \frac{\partial}{\partial \theta} f(\underset{\sim}{y} \mid \underset{\sim}{\theta})\, d\mu(\underset{\sim}{y})$$

$$= \int \qquad\qquad d\mu(\underset{\sim}{y}) \qquad (*)$$

$$= E(\quad) =$$

Fisher Information: From (*) above:

$$0 = \int \ell'(\underset{\sim}{y}|\theta)\, f(\underset{\sim}{y}|\theta)\, d\mu(\underset{\sim}{y})$$

Diff w.r.t. θ:

$$0 = \frac{d}{d\theta} \int \ell'(\underset{\sim}{y}|\theta)\, f(\underset{\sim}{y}|\theta)\, d\mu(\underset{\sim}{y})$$

o o o

Finish - it's fun

$$Var_{\underset{\sim}{\theta}}(\ell'(\underset{\sim}{\theta})) = E_{\underset{\sim}{\theta}}(-\ell''(\underset{\sim}{\theta}))$$

Define: Fisher Information: $I(\theta) = E_\theta(-\ell''(\underset{\sim}{\theta}))$

Distribution of gradients around $\underset{\sim}{\theta}_0$

given $\underset{\sim}{\theta}_0$ true value

SD ellipse

$\{\underset{\sim}{\theta} : \underset{\sim}{\theta}^T I(\underset{\sim}{\theta}_0)^{-1} \underset{\sim}{\theta} = 1\}$

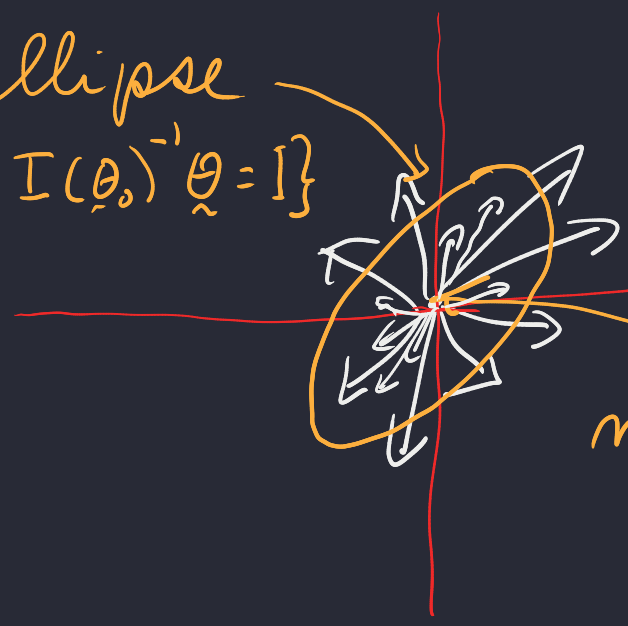

"Shape" $\sqrt{I(\theta_0)}$

not asymptotic

mean = $\underset{\sim}{0}$ ∴ Very Important

How do we get distribution of $\hat{\theta}$ from the distribution of $\underset{\sim}{l}'(\underset{\sim}{\theta})$?

One of the most used results in statistics:

$$f(x) = a + bx + \frac{cx^2}{2} \quad \text{for } c < 0$$

Then max f at $\boxed{\hat{x} = -\frac{b}{c}}$

Multivariate version:

$$f(\underset{\sim}{x}) = a + \underset{\sim}{b}^T \underset{\sim}{x} + \underset{\sim}{x}^T C \underset{\sim}{x}/2$$

for C negative definite

Then max f at $\boxed{\hat{\underset{\sim}{x}} = -C^{-1}\underset{\sim}{b}}$

What if C is positive definite ? Neither ?

Approx. Distribution of $\hat{\underset{\sim}{\theta}}$

If $l(\underset{\sim}{\theta})$ close to quadratic

$$l(\underset{\sim}{\theta}) \approx l(\underset{\sim}{\theta}_0) + l'(\underset{\sim}{\theta}_0)(\underset{\sim}{\theta} - \underset{\sim}{\theta}_0) - (\underset{\sim}{\theta} - \underset{\sim}{\theta}_0)^T I(\underset{\sim}{\theta}_0)(\underset{\sim}{\theta} - \underset{\sim}{\theta}_0)/2$$

Max at $\hat{\underset{\sim}{\theta}} \approx \underset{\sim}{\theta}_0 + I^{-1}(\underset{\sim}{\theta}_0)\, l'(\underset{\sim}{\theta}_0)$

$$E_{\underset{\sim}{\theta}_0}(\hat{\underset{\sim}{\theta}}) \approx \underset{\sim}{\theta}_0 + \underset{\sim}{0}$$

$$Var_{\underset{\sim}{\theta}_0}(\hat{\theta}) \approx I^{-1}(\underset{\sim}{\theta}_0) I(\underset{\sim}{\theta}_0) I^{-1}(\underset{\sim}{\theta}_0) = I^{-1}(\underset{\sim}{\theta}_0)$$

Asymptotic distribution of MLE $\hat{\theta}$

$$E(\hat{\underset{\sim}{\theta}}) \longrightarrow \underset{\sim}{\theta}_0$$

$$Var(\hat{\underset{\sim}{\theta}}) \simeq I^{-1}(\underset{\sim}{\theta}_0)$$

$$\{\underset{\sim}{\theta} : (\underset{\sim}{\theta} - \underset{\sim}{\theta}_0)^T I(\theta_0)(\theta - \theta_0) = 1\}$$

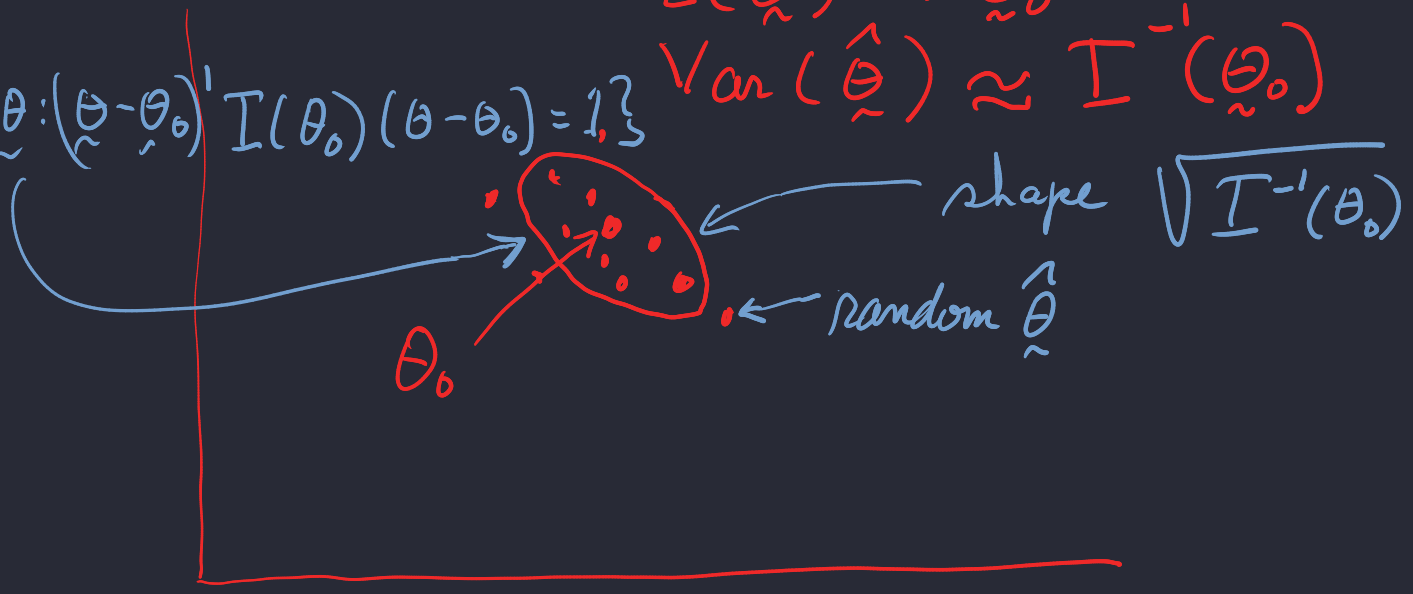

What if you are only interested in μ?

Bayes: "easy" use marginal posterior for μ

Frequentist: Not always clear

- e.g Behrens-Fisher problem
  Inference for $\mu_1 - \mu_2$ with 2 normal samples when you can't assume $\sigma_1 = \sigma_2$.
- Use additional principles to get e.g. t-test, etc.
- Use profile likelihood

Exponential families:

Special families that
- make things easy for
  - Bayesian: conjugate priors
  - Frequentist: finite-dimensional sufficient statistics as $n \uparrow$

Note: Likelihood is always sufficient.
(whole function, not necessarily MLE)

Skip

$n = 1$

$$f(\underset{\sim}{x} \mid \underset{\sim}{\psi}) = h(\underset{\sim}{x})\, g(\underset{\sim}{\psi}) \exp\left(\sum_{i=1}^{k} \underbrace{\theta_i(\underset{\sim}{\psi})}_{\text{canonical parameter}} T_i(\underset{\sim}{x}) \right)$$

n i.i.d

$$f(x_1, \ldots, x_n \mid \underset{\sim}{\psi}) = \prod_{i=1}^{n} h(x_i)\, g^n(\underset{\sim}{\psi}) \times \exp\left\{ \sum_{i=1}^{k} \theta_i(\underset{\sim}{\psi}) \underbrace{\left(\sum_{j=1}^{n} T_j(\underset{\sim}{x}) \right)}_{\text{suff. stat.}} \right\}$$

Redo using exponential form

Example: $N(\mu, \sigma)$

$$f(x \mid \mu, \sigma) = \frac{1}{\sqrt{2\pi}\,\sigma} \exp\left\{ -\frac{1}{2\sigma^2}(x-\mu)^2 \right\}$$

$$= \frac{1}{\sqrt{2\pi}\,\sigma} \exp\left\{ -\frac{1}{2\sigma^2} \cdot x^2 + \frac{\mu}{\sigma^2} x - \frac{1}{2\sigma^2}\mu^2 \right\}$$

$$= \underbrace{\frac{1}{\sqrt{2\pi}}}_{h} \underbrace{\frac{1}{\sigma}\exp\left\{-\frac{\mu^2}{2\sigma^2}\right\}}_{g} \times \exp\left\{\underbrace{\frac{1}{\sigma^2}\times\left(-\frac{x^2}{2}\right)}_{\theta_1 \times T_1} + \underbrace{\frac{\mu}{\sigma^2}x}_{\theta_2 \times T_2}\right\}$$

Express g as a function of $\theta_1 = \frac{1}{\sigma^2}$ $\theta_2 = \frac{\mu}{\sigma^2}$

Now: $\sigma^2 = \frac{1}{\theta_1}$ $\mu = \sigma^2\theta_2 = \theta_2/\theta_1$

So $k(\underset{\sim}{\theta}) = \frac{1}{\sigma}\exp\left\{-\frac{\mu^2}{2\sigma^2}\right\}$

$$= \sqrt{\theta_1}\exp\left\{-\frac{1}{2}\frac{\theta_2^2}{\theta_1^2}\times\theta_1\right\}$$

$$= \sqrt{\theta_1}\exp\left\{-\frac{1}{2}\theta_2^2/\theta_1\right\}$$

Then $E\left(\begin{pmatrix}T_1\\T_2\end{pmatrix}\middle|\,\underset{\sim}{\theta}\right) = k'(\theta)$

$$Var\left(\begin{pmatrix}T_1\\T_2\end{pmatrix}\middle|\,\underset{\sim}{\theta}\right) = k''(\theta)$$

$k^{(n)}(\theta)$ is cumulant generating function.

EXER: Check my math

Easier formulation of Exponential families

Density: Single observation $\underset{\sim}{x}$.

$$f_X(\underset{\sim}{x} \mid \underset{\sim}{\theta}) = \exp\left\{\underset{\sim}{\eta}(\underset{\sim}{\theta}) \cdot \underset{\sim}{T}(\underset{\sim}{x}) - A(\underset{\sim}{\theta}) + B(x)\right\}$$

So support of f must not depend on $\underset{\sim}{\theta}$

e.g. $U(\theta, \theta+1)$ is not an exponential family (why?)

IID Sample $\underset{\sim}{x}_1, \ldots, \underset{\sim}{x}_n$

$$\prod_{i=1}^{n} f_X(\underset{\sim}{x}_i \mid \underset{\sim}{\theta}) = \exp\left\{\underset{\sim}{\eta}(\underset{\sim}{\theta}) \cdot \sum_{i=1}^{n} \underset{\sim}{T}(\underset{\sim}{x}_i) - nA(\underset{\sim}{\theta}) + \sum_{i=1}^{n} B(\underset{\sim}{x}_i)\right\}$$

Same form as above

Sufficiency: $\sum_{i=1}^{n} \underset{\sim}{T}(\underset{\sim}{x}_i)$ is sufficient and the sufficient statistic has dimension equal to the dimension of $\underset{\sim}{T}$.

- $\underset{\sim}{\eta}$ is canonical parameter
- $K(\underset{\sim}{\eta}) = A(\underset{\sim}{\theta}(\underset{\sim}{\eta}))$ is the cumulant fcn

($\underset{\sim}{\theta}(\underset{\sim}{\eta})$: inverse of $\underset{\sim}{\eta}(\underset{\sim}{\theta})$)

$$K'(\underset{\sim}{\eta}) = E_{\underset{\sim}{\eta}}(\underset{\sim}{T}) \qquad K''(\underset{\sim}{\eta}) = \mathrm{Var}_{\underset{\sim}{\eta}}(\underset{\sim}{T})$$

$N(\mu, \sigma)$ (Painful but salubrious example)

$$f(x|\mu,\sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{1}{2\sigma^2}(x-\mu)^2\right\}$$

$$= \exp\left\{-\frac{1}{2\sigma^2}(x^2 - 2\mu x + \mu^2) - \frac{1}{2}\log(2\pi\sigma^2)\right\}$$

$$= \exp\left\{\underbrace{-\frac{1}{2\sigma^2}\cdot x^2 + \frac{\mu}{\sigma^2}\cdot x}_{\eta(\underset{\sim}{\theta})\cdot \underset{\sim}{T}(x)} \underbrace{- \frac{\mu^2}{2\sigma^2} - \frac{1}{2}\log\sigma^2}_{-A(\underset{\sim}{\theta})} \underbrace{- \frac{1}{2}\log(2\pi)}_{B(x) = \text{constant}}\right\}$$

$$\eta_1 = -\frac{1}{2\sigma^2} \quad \eta_2 = \frac{\mu}{\sigma^2}$$

$$T_1(x) = x^2 \quad T_2(x) = x$$

To get K, express $A(\theta)$ as a function of η

$$\sigma^2 = \frac{-1}{2\eta_1} \qquad (\text{note: } \eta_1 < 0)$$

$$\mu = \eta_2\sigma^2 = -\frac{\eta_2}{2\eta_1}$$

So $A = \frac{\mu^2}{2\sigma^2} + \frac{1}{2}\log\sigma^2$

$$= \frac{\eta_2^2/4\eta_1^2}{-1/\eta_1} + \frac{1}{2}\log\left(\frac{-1}{2\eta_1}\right)$$

$$= -\frac{\eta_2^2}{4\eta_1} - \frac{1}{2}\log(-2\eta_1)$$

$$A'(\underset{\sim}{\eta}) = \begin{pmatrix} \frac{\eta_2^2}{4\eta_1^2} - \frac{1}{2}\times\frac{1}{-2\eta_1}\times -2 \\ -\frac{2\eta_2}{4\eta_1} \end{pmatrix} = \begin{pmatrix} \mu^2+\sigma^2 \\ \mu \end{pmatrix}$$

Phew...

The proof is much easier than the example

Suppose A is expressed in terms of η

$$f(\underset{\sim}{x}\,|\,\underset{\sim}{\eta}) = \exp\{\underset{\sim}{\eta}\cdot T(\underset{\sim}{x}) - A(\underset{\sim}{\eta}) + B(\underset{\sim}{x})\}$$

Moment-generating function:

$$m_{\underset{\sim}{\eta}}(\underset{\sim}{t}) = E_{\underset{\sim}{\eta}}(e^{\underset{\sim}{t}\cdot\underset{\sim}{T}})$$

$$= \int e^{\underset{\sim}{t}\cdot\underset{\sim}{T}}\, e^{\underset{\sim}{\eta}\cdot\underset{\sim}{T} - A(\underset{\sim}{\eta}) + B(\underset{\sim}{x})}\, d\mu(\underset{\sim}{x})$$

density for $\underset{\sim}{\eta}+\underset{\sim}{t}$

$$= \int \exp\{(\underset{\sim}{t}+\underset{\sim}{\eta})\cdot T - A(\underset{\sim}{\eta}+\underset{\sim}{t}) + A(\underset{\sim}{\eta}+\underset{\sim}{t}) - A(\underset{\sim}{\eta}) + B(\underset{\sim}{x})\}\, d\mu$$

$$= \exp\{A(\eta+\underset{\sim}{t}) - A(\eta)\}$$

cumulant generating function

$$K_{\eta}(\underset{\sim}{t}) = \log m_{\eta}(\underset{\sim}{t})$$

$$= A(\eta+\underset{\sim}{t}) - A(\eta)$$

$$E_{\eta}(\underset{\sim}{T}) = K'_{\eta}(\underset{\sim}{t})_{\underset{\sim}{t}=\underset{\sim}{0}} = \frac{\partial}{\partial \underset{\sim}{t}} \left(A(\eta+t) - A(\eta)\right)\Big|_{\underset{\sim}{t}=\underset{\sim}{0}}$$

$$= A'(\eta)$$

$$\mathrm{Var}_{\eta}(T) = K''(\eta)$$