## Maximum Likelihood Part Two ${ }^{1}$ STA 312 Fall 2023

${ }^{1}$ See last slide for copyright information.

## Background Reading

Maximum likelihood handout (see course home page)

## Overview

(1) No Formula for the MLE
(2) Multiple Parameters
(3) Numerical MLEs
(4) Hypothesis Tests
(5) Nonlinear functions

## Two more issues

- Maximum likelihood estimates are often not available in closed form.
- Multiple parameters.

Most real-world problems have both these features.

## No formula for the MLE

## All we need is one example to see the problem.

Let $X_{1}, \ldots, X_{n}$ be independent observations from a distribution with density

$$
f(x \mid \theta)= \begin{cases}\frac{1}{\Gamma(\theta)} e^{-x} x^{\theta-1} & \text { for } x \geq 0 \\ 0 & \text { for } x<0\end{cases}
$$

Where the parameter $\theta>0$. This is a gamma with $\alpha=\theta$ and $\lambda=1$.

$$
\begin{aligned}
\frac{\partial}{\partial \theta} \ell(\theta) & =\frac{\partial}{\partial \theta} \log \left(\prod_{i=1}^{n} \frac{1}{\Gamma(\theta)} e^{-x_{i}} x_{i}^{\theta-1}\right) \\
& =\frac{\partial}{\partial \theta} \log \left(\Gamma(\theta)^{-n} e^{-\sum_{i=1}^{n} x_{i}}\left(\prod_{i=1}^{n} x_{i}\right)^{\theta-1}\right) \\
& =\frac{\partial}{\partial \theta}\left(-n \log \Gamma(\theta)-\sum_{i=1}^{n} x_{i}+(\theta-1) \sum_{i=1}^{n} \log x_{i}\right) \\
& =-\frac{n \Gamma^{\prime}(\theta)}{\Gamma(\theta)}-0+\sum_{i=1}^{n} \log x_{i} \stackrel{\text { set }}{=} 0
\end{aligned}
$$

## Numerical MLE

## By computer

- The log likelihood defines a surface sitting over the parameter space.
- It could have hills and valleys and mountains.
- The value of the log likelihood is easy to compute for any given set of parameter values.
- This tells you the height of the surface at that point.
- Take a step uphill (blindfolded).
- Are you at the top? Compute the slopes of some secant lines.
- Take another step uphill.
- How big a step? Good question.
- Most numerical routines minimize a function of several variables.
- So minimize the minus log likelihood.


## Multiple parameters

Most real-world problems have a vector of parameters.

- Let $X_{1}, \ldots, X_{n}$ be a random sample from a normal distribution with expected value $\mu$ and variance $\sigma^{2}$.
The parameters $\mu$ and $\sigma^{2}$ are unknown.
- For $i=1, \ldots, n$, let $y_{i}=\beta_{0}+\beta_{1} x_{i, 1}+\cdots+\beta_{p-1} x_{i, p-1}+\epsilon_{i}$, where $\beta_{0}, \ldots, \beta_{p-1}$ are unknown constants.
$x_{i, j}$ are known constants.
$\epsilon_{1}, \ldots, \epsilon_{n}$ are independent $N\left(0, \sigma^{2}\right)$ random variables.
$\sigma^{2}$ is an unknown constant.
$y_{1}, \ldots, y_{n}$ are observable random variables.
The parameters $\beta_{0}, \ldots, \beta_{p-1}, \sigma^{2}$ are unknown.


## Multi-parameter MLE <br> You know most of this.

- Suppose there are $k$ parameters.
- The plane tangent to the log likelihood should be horizontal at the MLE.
- Partially differentiate the log likelihood (or minus log likelihood) with respect to each of the parameters.
- Set the partial derivatives to zero, obtaining $k$ equations in $k$ unknowns.
- Solve for the parameters, if you can.
- Is it really a maximum?
- There is a multivariate second derivative test.


## The Hessian matrix

$$
\mathbf{H}=\left[\frac{\partial^{2}(-\ell)}{\partial \theta_{i} \partial \theta_{j}}\right]
$$

- If there are $k$ parameters, the Hessian is a $k \times k$ matrix whose $(i, j)$ element is $\frac{\partial^{2}}{\partial \theta_{i} \partial \theta_{j}}(-\ell(\boldsymbol{\theta}))$.
- If the second derivatives are continuous, $\mathbf{H}$ is symmetric.
- If the gradient is zero at a point and $|\mathbf{H}| \neq 0$, then
- If all eigenvalues are positive at the point, local minimum.
- If all eigenvalues are negative at the point, local maximum.
- If there are both positive and negative eigenvalues at the point, saddle point.


## Large-sample Theory

## Earlier results generalize to the multivariate case

The vector of MLEs is asymptotically normal. That is, multivariate normal.


## The Multivariate Normal

The multivariate normal distribution has many nice features. For us, the important ones are:

- It is characterized by a $k \times 1$ vector of expected values and a $k \times k$ variance-covariance matrix.
- Write $\mathbf{y} \sim N_{k}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.
- $\boldsymbol{\Sigma}=\left[\sigma_{i, j}\right]$ is a symmetric matrix with variances on the main diagonal and covariances on the off-diagonals.
- All the marginals are normal. $y_{j} \sim N\left(\mu_{j}, \sigma_{j, j}\right)$.


## The vector of MLEs is asymptotically multivariate

 normal. (Thank you, Mr. Wald)$$
\widehat{\boldsymbol{\theta}}_{n} \dot{\sim} N_{k}\left(\boldsymbol{\theta}, \frac{1}{n} \boldsymbol{\mathcal { I }}(\boldsymbol{\theta})^{-1}\right)
$$

- Compare $\widehat{\theta}_{n} \dot{\sim} N\left(\theta, \frac{1}{n I(\theta)}\right)$.
- $\boldsymbol{I}(\boldsymbol{\theta})$ is the Fisher information matrix.
- Specifically, the Fisher information in one observation.
- A $k \times k$ matrix

$$
\mathcal{I}(\boldsymbol{\theta})=\left[-E\left(\frac{\partial^{2}}{\partial \theta_{i} \partial \theta_{j}} \log f(Y ; \boldsymbol{\theta})\right)\right]
$$

- The Fisher Information in the whole sample is $n \boldsymbol{I}(\boldsymbol{\theta})$.


## $\widehat{\boldsymbol{\theta}}_{n}$ is asymptotically $N_{k}\left(\boldsymbol{\theta}, \frac{1}{n} \mathcal{I}(\boldsymbol{\theta})^{-1}\right)$

- Asymptotic covariance matrix of $\widehat{\boldsymbol{\theta}}_{n}$ is $\frac{1}{n} \boldsymbol{\mathcal { I }}(\boldsymbol{\theta})^{-1}$, and of course we don't know $\boldsymbol{\theta}$.
- For tests and confidence intervals, we need a good approximate asymptotic covariance matrix,
- Based on a good estimate of the Fisher information matrix.
- $\boldsymbol{I}\left(\widehat{\boldsymbol{\theta}}_{n}\right)$ would do.
- But it's inconvenient: Need to compute partial derivatives and expected values in

$$
\mathcal{I}(\boldsymbol{\theta})=\left[E\left[-\frac{\partial^{2}}{\partial \theta_{i} \partial \theta_{j}} \log f(Y ; \boldsymbol{\theta})\right]\right]
$$

and then substitute $\widehat{\boldsymbol{\theta}}_{n}$ for $\boldsymbol{\theta}$.

## The observed Fisher information

Approximate

$$
\frac{1}{n} \mathcal{I}(\boldsymbol{\theta})^{-1}=\left[n E\left[-\frac{\partial^{2}}{\partial \theta_{i} \partial \theta_{j}} \log f(Y ; \boldsymbol{\theta})\right]\right]^{-1}
$$

with

$$
\widehat{\mathbf{V}}_{n}=\left(\left[-\frac{\partial^{2}}{\partial \theta_{i} \partial \theta_{j}} \ell(\boldsymbol{\theta}, \mathbf{Y})\right]_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}_{n}}\right)^{-1}
$$

As in the univariate case, substitute the MLE for the parameter instead of taking the expected value.

## Compare the Hessian and (Estimated) Asymptotic Covariance Matrix

- $\widehat{\mathbf{V}}_{n}=\left(\left[-\frac{\partial^{2}}{\partial \theta_{i} \partial \theta_{j}} \ell(\boldsymbol{\theta}, \mathbf{Y})\right]_{\theta=\hat{\boldsymbol{\theta}}_{n}}\right)^{-1}$
- Hessian at MLE is $\mathbf{H}=\left[-\frac{\partial^{2}}{\partial \theta_{i} \partial \theta_{j}} \ell(\boldsymbol{\theta}, \mathbf{Y})\right]_{\theta=\widehat{\boldsymbol{\theta}}_{n}}$
- So to estimate the asymptotic covariance matrix of $\boldsymbol{\theta}$, just invert the Hessian.
- The Hessian is usually available as a by-product of a numerical search for the MLE.
- Because it's needed for the second derivative test.


## Connection to Numerical Optimization

- Suppose we are minimizing the minus log likelihood by a direct search.
- We have reached a point where the gradient is close to zero. Is this point a minimum?
- The Hessian is a matrix of mixed partial derivatives. If all its eigenvalues are positive at a point, the function is concave up there.
- Partial derivatives are usually approximated by the slopes of secant lines - no need to calculate them symbolically.
- It's the multivariable second derivative test.


## So to find the estimated asymptotic covariance matrix

- Minimize the minus log likelihood numerically.
- The Hessian at the place where the search stops is usually available.
- Invert it to get $\widehat{\mathbf{V}}_{n}$.
- This is so handy that sometimes we do it even when a closed-form expression for the MLE is available.


## Estimated Asymptotic Covariance Matrix $\widehat{\mathbf{V}}_{n}$ is Useful

- Asymptotic standard error of $\widehat{\theta}_{j}$ is the square root of the $j$ th diagonal element.
- Denote the asymptotic standard error of $\widehat{\theta}_{j}$ by $S_{\widehat{\theta}_{j}}$.
- Thus

$$
Z_{j}=\frac{\widehat{\theta}_{j}-\theta_{j}}{S_{\widehat{\theta}_{j}}}
$$

is approximately standard normal.

## Confidence Intervals and $Z$-tests

Have $Z_{j}=\frac{\widehat{\theta}_{j}-\theta_{j}}{S_{\widehat{\theta}_{j}}}$ approximately standard normal, yielding

- Confidence intervals: $\widehat{\theta}_{j} \pm S_{\widehat{\theta}_{j}} z_{\alpha / 2}$
- Test $H_{0}: \theta_{j}=\theta_{0}$ using

$$
Z=\frac{\widehat{\theta}_{j}-\theta_{0}}{S_{\widehat{\theta}_{j}}}
$$

Some null hypotheses involve multiple parameters For example,

$$
\begin{array}{ll}
H_{0}: & \beta_{1}=\beta_{2}=\beta_{3}=0 \\
H_{0}: & \frac{1}{3}\left(\theta_{1}+\theta_{2}+\theta_{3}\right)=\frac{1}{3}\left(\theta_{4}+\theta_{5}+\theta_{6}\right)=\frac{1}{2}\left(\theta_{7}+\theta_{8}\right)
\end{array}
$$

## Two hypothesis tests for multi-parameter problems

## They also apply to single-parameter models

- Wald tests and likelihood ratio tests.
- They both apply to linear null hypotheses of the form $H_{0}: \mathbf{L} \boldsymbol{\theta}=\mathbf{h}$
- Where $\mathbf{L}$ is an $r$ by $k$ matrix with linearly independent rows.
- This kind of null hypothesis is familiar from linear regression (STA302).


## Example

Linear regression with 4 explanatory variables

- $\boldsymbol{\theta}=\left(\beta_{0}, \beta_{1}, \beta_{2}, \beta_{3}, \beta_{4}, \sigma^{2}\right)$
- $H_{0}: \beta_{1}=\beta_{2}=\beta_{3}=0$
- $H_{0}: \mathbf{L} \boldsymbol{\theta}=\mathbf{0}$

$$
\left(\begin{array}{llllll}
0 & 1 & 0 & 0 & 0 & 0 \\
0 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & 0 & 0
\end{array}\right)\left(\begin{array}{l}
\beta_{0} \\
\beta_{1} \\
\beta_{2} \\
\beta_{3} \\
\beta_{4} \\
\sigma^{2}
\end{array}\right)=\left(\begin{array}{l}
0 \\
0 \\
0
\end{array}\right)
$$

## Another example of $H_{0}: \mathbf{L} \boldsymbol{\theta}=\mathbf{h}$

## A collection of linear constraints on the parameter $\boldsymbol{\theta}$

Example with $k=7$ parameters: $H_{0}$ has three parts

- $\theta_{1}=\theta_{2}$ and
- $\theta_{6}=\theta_{7}$ and
- $\frac{1}{3}\left(\theta_{1}+\theta_{2}+\theta_{3}\right)=\frac{1}{3}\left(\theta_{4}+\theta_{5}+\theta_{6}\right)$

$$
\left(\begin{array}{rrrrrrr}
1 & -1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & -1 \\
1 & 1 & 1 & -1 & -1 & -1 & 0
\end{array}\right)\left(\begin{array}{c}
\theta_{1} \\
\theta_{2} \\
\theta_{3} \\
\theta_{4} \\
\theta_{5} \\
\theta_{6} \\
\theta_{7}
\end{array}\right)=\left(\begin{array}{l}
0 \\
0 \\
0
\end{array}\right)
$$

Notice the number of rows in $\mathbf{L}$ is the number of $=$ signs in $H_{0}$.

## Wald Test for $H_{0}: \mathbf{L} \boldsymbol{\theta}=\mathbf{h}$

Based on $(\mathbf{x}-\boldsymbol{\mu})^{\top} \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu}) \sim \chi^{2}(p)$

$$
W_{n}=\left(\mathbf{L} \widehat{\boldsymbol{\theta}}_{n}-\mathbf{h}\right)^{\top}\left(\mathbf{L} \widehat{\mathbf{V}}_{n} \mathbf{L}^{\top}\right)^{-1}\left(\mathbf{L} \widehat{\boldsymbol{\theta}}_{n}-\mathbf{h}\right)
$$

- Looks like the formula for the general linear $F$-test in regression.
- Asymptotically chi-squared under $H_{0}$.
- Reject for large values of $W_{n}$.
- $d f=$ number of rows in $\mathbf{L}$.
- Number of linear constraints specified by $H_{0}$.


## The Wtest Function

$$
W_{n}=\left(\mathbf{L} \widehat{\boldsymbol{\theta}}_{n}-\mathbf{h}\right)^{\top}\left(\mathbf{L} \widehat{\mathbf{V}}_{n} \mathbf{L}^{\top}\right)^{-1}\left(\mathbf{L} \widehat{\boldsymbol{\theta}}_{n}-\mathbf{h}\right)
$$

Wtest $=$ function(L,Tn, Vn,h=0) \# HO: L theta $=\mathrm{h}$
\# For Wald tests based on numerical MLEs, Tn = theta-hat, \# and Vn is the inverse of the Hessian.
\{

```
value = numeric(3)
names (value) = c("W","df","p-value")
\(r=\operatorname{dim}(L)[1]\)
\(\mathrm{W}=\mathrm{t}(\mathrm{L} \% * \% \mathrm{Tn}-\mathrm{h}) \% * \%\) solve \((\mathrm{L} \% * \% \operatorname{Vn} \% * \% \mathrm{t}(\mathrm{L})) \% * \%\)
    ( \(\mathrm{L} \% * \% \mathrm{Tn}-\mathrm{h}\) )
\(\mathrm{W}=\) as.numeric (W)
pval = 1-pchisq(W,r)
value[1] \(=W\); value[2] \(=r\); value[3] = pval
return(value)
\}
```


## Likelihood ratio tests

- $X_{1}, \ldots, X_{n} \stackrel{i . i . d .}{\sim} F_{\theta}, \theta \in \Theta$
- $H_{0}: \theta \in \Theta_{0}$ v.s. $H_{1}: \theta \in \Theta \cap \Theta_{0}^{c}$

$$
\begin{aligned}
G^{2} & =-2 \log \left(\frac{\max _{\theta \in \Theta_{0}} L(\theta)}{\max _{\theta \in \Theta} L(\theta)}\right)=-2 \log \frac{L\left(\widehat{\theta}_{0}\right)}{L(\widehat{\theta})} \\
& =2\left(\ell(\widehat{\theta})-\ell\left(\widehat{\theta}_{0}\right)\right)
\end{aligned}
$$

- Under $H_{0}, G^{2}$ has an approximate chi-squared distribution for large $n$.
- Degrees of freedom $=$ number of (non-redundant, linear) equalities specified by $H_{0}$.
- Reject when $G^{2}$ is large.


## Example: Multinomial with 3 categories

- Parameter space is 2-dimensional.
- Unrestricted MLE is $\left(p_{1}, p_{2}\right)$ : Sample proportions.
- $H_{0}: \theta_{1}=2 \theta_{2}$


## Parameter space for $H_{0}: \theta_{1}=2 \theta_{2}$

Red dot is unrestricted MLE, Black square is restricted MLE


## Comparing Likelihood Ratio and Wald tests

- Asymptotically equivalent under $H_{0}$, meaning $\left(W_{n}-G_{n}^{2}\right) \xrightarrow{p} 0$
- Under $H_{1}$,
- Both have the same approximate distribution (non-central chi-square).
- Both go to infinity as $n \rightarrow \infty$.
- But values are not necessarily close for the same data set.
- Likelihood ratio test tends to get closer to the right Type I error probability for small samples.
- Wald can be more convenient when testing lots of hypotheses, because you only need to fit the model once.
- Wald can be more convenient if it's a lot of work to write the restricted likelihood.


## Non-linear functions of the parameter vector

- Most tests are about linear combinations of the model parameters.
- Sometimes we want tests and confidence intervals for non-linear functions of $\boldsymbol{\theta} \in \mathbb{R}^{k}$.
- Like $\frac{\alpha}{\lambda^{2}}$ (variance of a gamma).
- Fortunately, smooth functions of an asymptotically multivariate normal random vector are asymptotically normal.


## Theorem based on the delta method of Cramér

## The delta method is more general than this.

Let $\boldsymbol{\theta} \in \mathbb{R}^{k}$. Under the conditions for which $\widehat{\boldsymbol{\theta}}_{n}$ is asymptotically $N_{k}\left(\boldsymbol{\theta}, \mathbf{V}_{n}\right)$ with $\mathbf{V}_{n}=\frac{1}{n} \mathcal{I}(\boldsymbol{\theta})^{-1}$, let the function $g: \mathbb{R}^{k} \rightarrow \mathbb{R}$ be such that the elements of $\dot{\mathrm{g}}(\boldsymbol{\theta})=\left(\frac{\partial g}{\partial \theta_{1}}, \ldots, \frac{\partial g}{\partial \theta_{k}}\right)$ are continuous in a neighbourhood of the true parameter vector $\boldsymbol{\theta}$. Then

$$
g(\widehat{\boldsymbol{\theta}}) \dot{\sim} N\left(g(\boldsymbol{\theta}), \dot{\mathrm{g}}(\boldsymbol{\theta}) \mathbf{V}_{n} \dot{\mathrm{~g}}(\boldsymbol{\theta})^{\top}\right)
$$

Note that the asymptotic variance $\dot{\mathrm{g}}(\boldsymbol{\theta}) \mathbf{V}_{n} \dot{\mathrm{~g}}(\boldsymbol{\theta})^{\top}$ is a matrix product: $(1 \times k)$ times $(k \times k)$ times $(k \times 1)$.

The standard error of $g(\widehat{\boldsymbol{\theta}})$ is $\sqrt{\dot{\mathrm{g}}(\widehat{\boldsymbol{\theta}}) \widehat{\mathbf{V}}_{n} \dot{\mathrm{~g}}(\widehat{\boldsymbol{\theta}})^{\top}}$.

## Example of $\dot{\mathrm{g}}(\boldsymbol{\theta})=\left(\frac{\partial g}{\partial \theta_{1}}, \ldots, \frac{\partial g}{\partial \theta_{k}}\right)$

- Variance of gamma is $g(\alpha, \lambda)=\frac{\alpha}{\lambda^{2}}$.
- $\theta_{1}=\alpha, \theta_{2}=\lambda, k=2$,
- $\operatorname{So} \dot{\mathrm{g}}(\boldsymbol{\theta})$ is $1 \times 2$.

$$
\begin{aligned}
\dot{\mathrm{g}} & =\left(\frac{\partial g}{\partial \alpha}, \frac{\partial g}{\partial \lambda}\right) \\
& =\left(\frac{1}{\lambda^{2}}, \alpha(-2) \lambda^{-3}\right) \\
& =\left(\frac{1}{\lambda^{2}}, \frac{-2 \alpha}{\lambda^{3}}\right)
\end{aligned}
$$

Then, $\dot{\mathrm{g}}(\widehat{\boldsymbol{\theta}}) \widehat{\mathbf{V}}_{n} \dot{\mathrm{~g}}(\widehat{\boldsymbol{\theta}})^{\top}$ is easy if you have $\widehat{\mathbf{V}}_{n}$.

## Copyright Information

This slide show was prepared by Jerry Brunner, Department of Statistics, University of Toronto. It is licensed under a Creative Commons Attribution - ShareAlike 3.0 Unported License. Use any part of it as you like and share the result freely. The $\mathrm{IAT}_{\mathrm{E}} \mathrm{X}$ source code is available from the course website:
http://www.utstat.toronto.edu/brunner/oldclass/312f23

