Methods of Applied Statistics I

STA2101H F LEC9101

Week 1

September 14 2022

Tweet

The Guy Medal in Gold is awarded to Nancy Reid @reid_nancy for her pioneering work on higher-order approximate inference

5:54 PM · Sep 13, 2022 · Twitter Web App

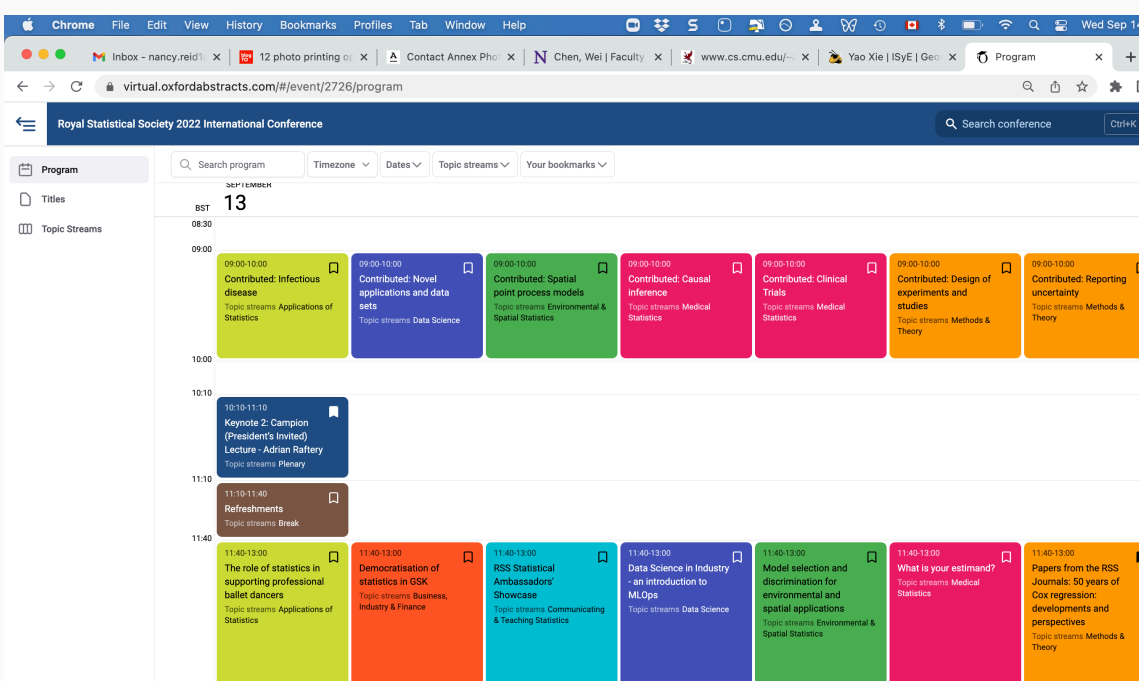
Chrome File Edit View History Bookmarks Profiles Tab Window Help
Wed Sep 14
Inbox - nancy.reid1
12 photo printing o
Contact Annex Phot
Chen, Wei | Faculty
www.cs.cmu.edu/~
Yao Xie | ISyE | Geo
Program
virtual.oxfordabstracts.com/#/event/2726/program
Royal Statistical Society 2022 International Conference
Search conference
Ctrl+K
Program
Titles
Topic Streams
Search program
Timezone
Dates
Topic streams
Your bookmarks
SEPTEMBER
BST
13
08:30
09:00
10:00
10:10
11:10
11:40
09:00-10:00
Contributed: Infectious disease
Topic streams Applications of Statistics
09:00-10:00
Contributed: Novel applications and data sets
Topic streams Data Science
09:00-10:00
Contributed: Spatial point process models
Topic streams Environmental & Spatial Statistics
09:00-10:00
Contributed: Causal inference
Topic streams Medical Statistics
09:00-10:00
Contributed: Clinical Trials
Topic streams Medical Statistics
09:00-10:00
Contributed: Design of experiments and studies
Topic streams Methods & Theory
09:00-10:00
Contributed: Reporting uncertainty
Topic streams Methods & Theory
10:10-11:10
Keynote 2: Campion (President's Invited) Lecture - Adrian Raftery
Topic streams Plenary
11:10-11:40
Refreshments
Topic streams Break
11:40-13:00
The role of statistics in supporting professional ballet dancers
Topic streams Applications of Statistics
11:40-13:00
Democratisation of statistics in GSK
Topic streams Business, Industry & Finance
11:40-13:00
RSS Statistical Ambassadors' Showcase
Topic streams Communicating & Teaching Statistics
11:40-13:00
Data Science in Industry - an introduction to MLOps
Topic streams Data Science
11:40-13:00
Model selection and discrimination for environmental and spatial applications
Topic streams Environmental & Spatial Statistics
11:40-13:00
What is your estimand?
Topic streams Medical Statistics
11:40-13:00
Papers from the RSS Journals: 50 years of Cox regression: developments and perspectives
Topic streams Methods & Theory

Today

1. Course introduction: course details, evaluation, syllabus, people
2. Upcoming events of interest
3. Review of linear regression
4. In the news: $\longrightarrow$ at the conference

STA 2101F: Methods of Applied Statistics I
Wednesday, 10am – 1 pm Eastern September 14 – December 7 2022
SF 3201

From the calendar:

This course will focus on principles and methods of applied statistical science. It is designed for MSc and PhD students in Statistics, and is required for the Applied Paper of the PhD comprehensive exams. The topics covered include: planning of studies, review of linear models, analysis of random and mixed effects models, model building and model selection, theory and methods for generalized linear models, and an introduction to nonparametric regression. Additional topics will be introduced as needed in the context of case studies in data analysis.
Prerequisites: ECO374H1/ECO375H1/STA302H1 (regression); STA305H1 (design of studies)

Course Delivery:

On September 14, the class will be delivered online at the scheduled

STA 2101F: Methods of Applied Statistics I
Wednesday, 10am – 1 pm Eastern September 14 – December 7 2022
SF 3201

From the calendar:

This course will focus on principles and methods of applied statistical science. It is designed for MSc and PhD students in Statistics, and is required for the Applied Paper of the PhD comprehensive exams. The topics covered include: planning of studies, review of linear models, analysis of random and mixed effects models, model building and model selection, theory and methods for generalized linear models, and an introduction to nonparametric regression. Additional topics will be introduced as needed in the context of case studies in data analysis.
Prerequisites: ECO374H1/ECO375H1/STA302H1 (regression); STA305H1 (design of studies)

Course Delivery:

Course Description

- Grading
- Academic Integrity
- Computing
- References Modules
- Contact

Use **Piazza** for course questions; email for personal questions

Today

1. Course introduction: technical issues, course details, evaluation, syllabus
2. Upcoming events of interest
3. Review of linear regression
4. In the news: $\longrightarrow$ at the conference

Upcoming events

- Thursdays 3.30 — Departmental Seminar link
- Mondays 3.30 — Data Science and Applied Research Seminar link
- Fridays 12.00 — Toronto Data Workshop
- Special:
  - September 29: CANSSI Ontario Research Day link
  - September 29, 30 3.30: Distinguished Lecture Series in Statistical Sciences link

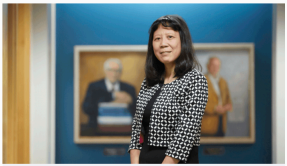

2022 DLSS: Xihong Lin

Professor, Department of BiostatisticsCoordinating Director, Program in Quantitative Genomics; Harvard T.H. Chan School of Public Health; Professor of Statistics, Department of Statistics, Harvard University

Sep 29 (3:30-4:30 pm): Lessons Learned from the COVID-19 Pandemic: A Statistician's Reflection
Sep 30 (3:30-4:30 pm): Ensemble Methods for Testing a Global Null Hypothesis

Today

Start Recording

1. Course introduction: course details, evaluation, syllabus
2. Upcoming events of interest
3. Review of linear regression
4. Steps in analysis
5. In the news: $\longrightarrow$ at the conference

- Model:

$$Y = X \beta + \epsilon$$

- Model:

$$Y_{n\times 1} = X_{n\times p}\beta_{p\times 1} + \epsilon_{n\times 1}$$

- Model:

$$Y_{n\times 1} = X_{n\times p}\beta_{p\times 1} + \epsilon_{n\times 1}$$

- Equivalently:

$$y_i =$$

- Model:

$$Y_{n\times 1} = X_{n\times p}\beta_{p\times 1} + \epsilon_{n\times 1}$$

- Equivalently:

$$y_i =$$

- Standard Assumptions
  - y_i independent — equivalently ϵ_i independent — y is often called response
  - $\mathbb{E}(\epsilon_i) = 0$ — why?
  - $\text{var}(\epsilon_i) = \sigma^2$ — constant
  - x_i known, β to be estimated — x_i often called explanatory variables

- Model:

$$Y_{n\times 1} = X_{n\times p}\beta_{p\times 1} + \epsilon_{n\times 1}$$

- Equivalently:

$$y_i =$$

- Standard Assumptions
  - y_i independent — equivalently ϵ_i independent — y is often called response
  - $\mathbb{E}(\epsilon_i) = 0$ — why?
  - $\text{var}(\epsilon_i) = \sigma^2$ — constant
  - x_i known, β to be estimated — x_i often called explanatory variables
- More concisely:

$$\mathbb{E}(Y \mid X) = \quad , \quad \text{var}(Y \mid X) =$$

I ??

Nice big equation:

$$\begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} = \begin{pmatrix} x_{11} & \dots & x_{1p} \\ \vdots & \vdots & \vdots \\ x_{n1} & \dots & x_{np} \end{pmatrix} \begin{pmatrix} \\ \vdots \\ \\ \end{pmatrix} + \begin{pmatrix} \\ \vdots \\ \\ \end{pmatrix}$$

Nice big equation:

$$\begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} = \begin{pmatrix} x_{11} & \dots & x_{1p} \\ \vdots & \vdots & \vdots \\ x_{n1} & \dots & x_{np} \end{pmatrix} \begin{pmatrix} \\ \vdots \\ \\ \end{pmatrix} + \begin{pmatrix} \\ \vdots \\ \\ \end{pmatrix}$$

Or, if you prefer:

$$y_i = x_{i1}\beta_1 + x_{i2}\beta_2 + \cdots + x_{ip}\beta_p + \epsilon_i, \quad \epsilon_i \sim \qquad\qquad i = 1, \dots, n$$

Nice big equation:

$$\begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} = \begin{pmatrix} x_{11} & \dots & x_{1p} \\ \vdots & \vdots & \vdots \\ x_{n1} & \dots & x_{np} \end{pmatrix} \begin{pmatrix} \\ \vdots \\ \\ \end{pmatrix} + \begin{pmatrix} \\ \vdots \\ \\ \end{pmatrix}$$

Or, if you prefer:

$$y_i = x_{i1}\beta_1 + x_{i2}\beta_2 + \cdots + x_{ip}\beta_p + \epsilon_i, \quad \epsilon_i \sim \qquad\qquad i = 1, \dots, n$$

Or, if you prefer:

$$\mathbb{E}(y_i \mid x_i) = x_i^{\mathrm{T}}\beta, \qquad \mathrm{var}(y_i \mid x_i) = \sigma^2, \qquad i = 1, \dots, n$$

y_i independent

- often not completely clear: X might be fixed by design, or measured on each individual

e.g.?

- often not completely clear: X might be fixed by design, or measured on each individual

  e.g.?

- If measured, then should we consider its distribution? E.g. should our model be $(y_i, x_i^{\mathrm{T}}) \sim$??

  some $(p+1)$-dimensional distribution

- often not completely clear: X might be fixed by design, or measured on each individual

  e.g.?

- If measured, then should we consider its distribution? E.g. should our model be $(y_i, x_i^{\mathrm{T}}) \sim$??

  some $(p+1)$-dimensional distribution

- Almost always in regression settings we condition on X, as on previous slide

  ancillary statistic

- often not completely clear: X might be fixed by design, or measured on each individual

  e.g.?

- If measured, then should we consider its distribution? E.g. should our model be $(y_i, x_i^{\mathrm{T}}) \sim$??

  some $(p+1)$-dimensional distribution

- Almost always in regression settings we condition on X, as on previous slide

  ancillary statistic

- often not emphasized: interpretation of β_j

- often not completely clear: X might be fixed by design, or measured on each individual

  e.g.?

- If measured, then should we consider its distribution? E.g. should our model be $(y_i, x_i^{\mathrm{T}}) \sim$??

  some $(p+1)$-dimensional distribution

- Almost always in regression settings we condition on X, as on previous slide

  ancillary statistic

- often not emphasized: interpretation of β_j
  - version 1: effect on the expected response of a unit change in jth explanatory variable, all other variables held fixed

- often not completely clear: X might be fixed by design, or measured on each individual

  e.g.?

- If measured, then should we consider its distribution? E.g. should our model be $(y_i, x_i^{\mathrm{T}}) \sim$??

  some $(p+1)$-dimensional distribution

- Almost always in regression settings we condition on X, as on previous slide

  ancillary statistic

- often not emphasized: interpretation of β_j
  - version 1: effect on the expected response of a unit change in jth explanatory variable, all other variables held fixed
  - version 2:

$$\beta_j = \frac{\partial \mathbb{E}(y_i \mid x_{ij})}{\partial x_{ij}} \qquad \frac{\partial \mathbb{E}(y \mid x_j)}{\partial x_j}$$

notation ambiguous, see CD §6.5.2

Least squares estimation

- Definition

$$\hat{\beta}_{LS} := \min_{\beta} \sum_{i=1}^{n} (y_i - x_i^T \beta)^2$$

Least squares estimation

- Definition

$$\hat{\beta}_{LS} := \min_{\beta} \sum_{i=1}^{n} (y_i - x_i^T \beta)^2$$

- Equivalently,

Least squares estimation

- Definition

$$\hat{\beta}_{LS} := \min_{\beta} \sum_{i=1}^{n} (y_i - x_i^T \beta)^2$$

- Equivalently,

- Equivalently,

$$\hat{\beta}_{LS} :=$$

$L2$ distance

Least squares estimation

- Definition

$$\hat{\beta}_{LS} := \min_{\beta} \sum_{i=1}^{n} (y_i - x_i^T \beta)^2$$

- Equivalently,

- Equivalently,

$$\hat{\beta}_{LS} :=$$

$L2$ distance

- Equivalently, $\hat{\beta}_{LS}$ is the solution of the score equation

$$X^T(y - X\beta) = 0$$

?how?

Least squares estimation

- Definition

$$\hat{\beta}_{LS} := \min_{\beta} \sum_{i=1}^{n} (y_i - x_i^{\mathrm{T}} \beta)^2$$

- Equivalently,

- Equivalently,

$$\hat{\beta}_{LS} :=$$

$L2$ distance

- Equivalently, $\hat{\beta}_{LS}$ is the solution of the score equation

$$X^{\mathrm{T}}(y - X\beta) = 0$$

?how?

- Solution

$$\hat{\beta}_{LS} =$$

check dimensions

... least squares estimation

- Solution

$$\hat{\beta}_{LS} = (X^{\mathrm{T}}X)^{-1}(X^{\mathrm{T}}y)$$

check dimensions

ASIDE: here and following all assume X is fixed

... least squares estimation

- Solution

$$\hat{\beta}_{LS} = (X^{\mathrm{T}}X)^{-1}(X^{\mathrm{T}}y)$$

check dimensions

- Expected value

$$\mathbb{E}(\hat{\beta}_{LS}) =$$

why?

ASIDE: here and following all assume X is fixed

... least squares estimation

- Solution

$$\hat{\beta}_{LS} = (X^{\mathrm{T}}X)^{-1}(X^{\mathrm{T}}y)$$

check dimensions

- Expected value

$$\mathbb{E}(\hat{\beta}_{LS}) =$$

why?

- Least squares estimates are unbiased

ASIDE: here and following all assume X is fixed

... least squares estimation

- Solution

$$\hat{\beta}_{LS} = (X^{\mathrm{T}}X)^{-1}(X^{\mathrm{T}}y)$$

check dimensions

- Expected value

$$\mathbb{E}(\hat{\beta}_{LS}) =$$

why?

- Least squares estimates are unbiased
- Variance

really variance-covariance matrix

$$\mathrm{var}(\hat{\beta}_{LS}) = (X^{\mathrm{T}}X)^{-1}X^{\mathrm{T}}\mathrm{var}(y)X(X^{\mathrm{T}}X)^{-1} = (X^{\mathrm{T}}X)^{-1}X^{\mathrm{T}}\sigma^2 IX(X^{\mathrm{T}}X)^{-1} = \sigma^2(X^{\mathrm{T}}X)^{-1}$$

ASIDE: here and following all assume X is fixed

What about the normal distribution?

SM §8.2

- If we further assume $\epsilon_i \sim N(0, \sigma^2)$ (and independent across i), then

What about the normal distribution?

SM §8.2

- If we further assume $\epsilon_i \sim N(0, \sigma^2)$ (and independent across i), then
- $y \mid X \sim N(X\beta, \sigma^2 I)$, and

What about the normal distribution?

- If we further assume $\epsilon_i \sim N(0, \sigma^2)$ (and independent across i), then
- $y \mid X \sim N(X\beta, \sigma^2 I)$, and
- the likelihood function is

$$L(\beta, \sigma^2; y) = \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left\{-\frac{1}{2\sigma^2}(y - X\beta)^T(y - X\beta)\right\},$$

What about the normal distribution?

SM §8.2

- If we further assume $\epsilon_i \sim N(0, \sigma^2)$ (and independent across i), then
- $y \mid X \sim N(X\beta, \sigma^2 I)$, and
- the likelihood function is

$$L(\beta, \sigma^2; y) = \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left\{-\frac{1}{2\sigma^2}(y - X\beta)^T(y - X\beta)\right\},$$

- the log-likelihood function is

$$\ell(\beta, \sigma^2; y) = -\frac{n}{2}\log(\sigma^2) - \frac{1}{2\sigma^2}(y - X\beta)^{\mathrm{T}}(y - X\beta),$$

constants in params don't matter

What about the normal distribution?

SM §8.2

- If we further assume $\epsilon_i \sim N(0, \sigma^2)$ (and independent across i), then
- $y \mid X \sim N(X\beta, \sigma^2 I)$, and
- the likelihood function is

$$L(\beta, \sigma^2; y) = \frac{1}{(2\pi\sigma^2)^{n/2}} \exp\left\{-\frac{1}{2\sigma^2}(y - X\beta)^T(y - X\beta)\right\},$$

- the log-likelihood function is

$$\ell(\beta, \sigma^2; y) = -\frac{n}{2}\log(\sigma^2) - \frac{1}{2\sigma^2}(y - X\beta)^{\mathrm{T}}(y - X\beta),$$

constants in params don't matter

- the maximum likelihood estimate of β is

$$\hat{\beta}_{ML} = (X^{\mathrm{T}}X)^{-1}X^{\mathrm{T}}y = \hat{\beta}_{LS}$$

... what about the normal distribution?

- maximum likelihood estimate of β is

$$\hat{\beta}_{ML} = (X^{\mathrm{T}}X)^{-1}X^{\mathrm{T}}y = \hat{\beta}_{LS}$$

… what about the normal distribution?

- maximum likelihood estimate of β is

$$\hat{\beta}_{ML} = (X^{\mathrm{T}}X)^{-1}X^{\mathrm{T}}y = \hat{\beta}_{LS}$$

- distribution of $\hat{\beta}$ is normal

why?

$$\hat{\beta} \sim N_p(\beta, \sigma^2(X^{\mathrm{T}}X)^{-1})$$

... what about the normal distribution?

- maximum likelihood estimate of β is

$$\hat{\beta}_{ML} = (X^{\mathrm{T}}X)^{-1}X^{\mathrm{T}}y = \hat{\beta}_{LS}$$

- distribution of $\hat{\beta}$ is normal

why?

$$\hat{\beta} \sim N_p(\beta, \sigma^2(X^{\mathrm{T}}X)^{-1})$$

- distribution of $\hat{\beta}_j$ is

$$N(\beta_j, \sigma^2(X^{\mathrm{T}}X)^{-1}_{jj}), \quad j = 1, \ldots, p$$

... what about the normal distribution?

- maximum likelihood estimate of β is

$$\hat{\beta}_{ML} = (X^{\mathrm{T}}X)^{-1}X^{\mathrm{T}}y = \hat{\beta}_{LS}$$

- distribution of $\hat{\beta}$ is normal

why?

$$\hat{\beta} \sim N_p(\beta, \sigma^2(X^{\mathrm{T}}X)^{-1})$$

- distribution of $\hat{\beta}_j$ is

$$N(\beta_j, \sigma^2(X^{\mathrm{T}}X)^{-1}_{jj}), \quad j = 1, \ldots, p$$

- maximum likelihood estimate of σ^2 is $\frac{1}{n}(y - X\hat{\beta})^{\mathrm{T}}(y - X\hat{\beta})$

... what about the normal distribution?

- maximum likelihood estimate of β is

$$\hat{\beta}_{ML} = (X^{\mathrm{T}}X)^{-1}X^{\mathrm{T}}y = \hat{\beta}_{LS}$$

- distribution of $\hat{\beta}$ is normal

why?

$$\hat{\beta} \sim N_p(\beta, \sigma^2(X^{\mathrm{T}}X)^{-1})$$

- distribution of $\hat{\beta}_j$ is

$$N(\beta_j, \sigma^2(X^{\mathrm{T}}X)^{-1}_{jj}), \quad j = 1, \ldots, p$$

- maximum likelihood estimate of σ^2 is $\frac{1}{n}(y - X\hat{\beta})^{\mathrm{T}}(y - X\hat{\beta})$
- but we use

$$\tilde{\sigma}^2 = \frac{1}{n-p}(y - X\hat{\beta})^{\mathrm{T}}(y - X\hat{\beta})$$

Pause

(1) I'm lost

(2) I'm good

(3) I'm bored

HW Question Week 1

STA2101F 2022

Due September 21 2022 11.59 pm

Homework to be submitted through Quercus

You can submit this HW in Word, Latex, or R Markdown, but in future please use R Markdown. If you are using Word or Latex with a R script for the computational work, then this R script should be provided as an Appendix. In the document itself you would just include properly formatted output.

You are welcome to discuss questions with others, but the solutions and code must be written independently. Any R output that is included in a solution should be formatted as part of the discussion (i.e. not cut and pasted from the Console).

The dataset `wafer` concerns a study on semiconductors. You can get more information about the data with `?wafer`; you will first need `library(faraway);data(wafer)`, and possibly `install.packages("faraway")`. The questions below are adapted from LM Ch.3.

Inference

- If you really like likelihood theory, the expected Fisher information is SM §8.2.3

$$\mathcal{I}(\beta, \sigma^2) = \begin{pmatrix} \sigma^{-2} X^T X & 0 \\ 0 & \frac{1}{2} n \sigma^{-4} \end{pmatrix}$$

$\mathcal{I}^{-1}$ gives (asymptotic) variance of MLE

Inference

- If you really like likelihood theory, the expected Fisher information is SM §8.2.3

$$\mathcal{I}(\beta, \sigma^2) = \begin{pmatrix} \sigma^{-2} X^T X & 0 \\ 0 & \frac{1}{2} n \sigma^{-4} \end{pmatrix}$$

$\mathcal{I}^{-1}$ gives (asymptotic) variance of MLE

- but just using previous slide we have

$$\frac{\hat{\beta}_j - \beta_j}{\sigma[\{(X^T X)^{-1}\}_{jj}]^{1/2}} \sim N(0, 1)$$

Inference

- If you really like likelihood theory, the expected Fisher information is SM §8.2.3

$$\mathcal{I}(\beta, \sigma^2) = \begin{pmatrix} \sigma^{-2} X^T X & 0 \\ 0 & \frac{1}{2} n \sigma^{-4} \end{pmatrix}$$

$\mathcal{I}^{-1}$ gives (asymptotic) variance of MLE

- but just using previous slide we have

$$\frac{\hat{\beta}_j - \beta_j}{\sigma [\{(X^T X)^{-1}\}_{jj}]^{1/2}} \sim N(0, 1)$$

- and

$$\frac{\hat{\beta}_j - \beta_j}{\tilde{\sigma} [\{(X^T X)^{-1}\}_{jj}]^{1/2}} \sim t_{n-p}$$

Example

```
install.packages("faraway")
library(faraway)
data(prostate)
head(prostate)
```

Example

```
install.packages("faraway")
library(faraway)
data(prostate)
head(prostate)


model1 <- lm(lpsa ~ ., data = prostate)

summary(model1)
Coefficients:
             Estimate Std. Error t value Pr(>|t|)
(Intercept)  0.669337   1.296387   0.516  0.60693
lcavol       0.587022   0.087920   6.677 2.11e-09 ***
lweight      0.454467   0.170012   2.673  0.00896 **
age         -0.019637   0.011173  -1.758  0.08229 .
lbph         0.107054   0.058449   1.832  0.07040
```

Example

```
summary(model1)
Coefficients:
             Estimate Std. Error t value Pr(>|t|)
(Intercept)  0.669337   1.296387   0.516  0.60693
lcavol       0.587022   0.087920   6.677 2.11e-09 ***
lweight      0.454467   0.170012   2.673  0.00896 **
age         -0.019637   0.011173  -1.758  0.08229 .
lbph         0.107054   0.058449   1.832  0.07040 .
svi          0.766157   0.244309   3.136  0.00233 **
lcp         -0.105474   0.091013  -1.159  0.24964
gleason      0.045142   0.157465   0.287  0.77503
pgg45        0.004525   0.004421   1.024  0.30886
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1
```

322 views

Women in Statistics and Data Science

Follow @WomeninStat 18.7K followers

Jul 23rd 2020, 14 tweets, 4 min read

Bookmark Save as PDF + My Authors

Today, we're going to play a game I'm calling "IT'S JUST A LINEAR MODEL" (IJALM).

It works like this: I name a model for a quantitative response Y, and then you guess whether or not IJALM.

Many special cases $\mathbb{E}(Y \mid X) = X\beta$

- $y_i = \beta_0 + \beta_1 x_i + \epsilon_i, \quad i = 1, \ldots, n$

1st column of X?

Many special cases $\mathbb{E}(Y \mid X) = X\beta$

- $y_i = \beta_0 + \beta_1 x_i + \epsilon_i, \quad i = 1, \ldots, n$

1st column of X?

- $y_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + \beta_3 x_i^3 + \beta_4 x_i^4 + \beta_5 x_i^5 \epsilon_i$

Many special cases $\mathbb{E}(Y \mid X) = X\beta$

- $y_i = \beta_0 + \beta_1 x_i + \epsilon_i, \quad i = 1, \ldots, n$

1st column of X?

- $y_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + \beta_3 x_i^3 + \beta_4 x_i^4 + \beta_5 x_i^5 \epsilon_i$
- $y_i = \beta_0 \pm \beta_1 + \epsilon_i$

Many special cases $\mathbb{E}(Y \mid X) = X\beta$

- $y_i = \beta_0 + \beta_1 x_i + \epsilon_i, \quad i = 1, \ldots, n$

1st column of X?

- $y_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + \beta_3 x_i^3 + \beta_4 x_i^4 + \beta_5 x_i^5 \epsilon_i$
- $y_i = \beta_0 \pm \beta_1 + \epsilon_i$
- $y_i = \beta_0 + \beta_1 \sin(x_i) + \beta_2 \cos(x_i) + \epsilon_i$

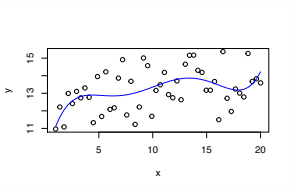

Many special cases $\mathbb{E}(Y \mid X) = X\beta$

- $y_i = \beta_0 + \beta_1 x_i + \epsilon_i, \quad i = 1, \ldots, n$

1st column of X?

- $y_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + \beta_3 x_i^3 + \beta_4 x_i^4 + \beta_5 x_i^5 \epsilon_i$
- $y_i = \beta_0 \pm \beta_1 + \epsilon_i$
- $y_i = \beta_0 + \beta_1 \sin(x_i) + \beta_2 \cos(x_i) + \epsilon_i$
- $y_i = \gamma_0 x_{1i}^{\gamma_1} x_{2i}^{\gamma_2} \eta_i, \quad \eta_i \sim$ positive r.v.

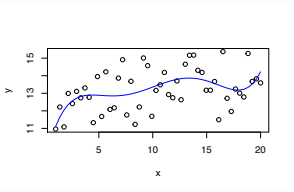

Many special cases $\mathbb{E}(Y \mid X) = X\beta$

- $y_i = \beta_0 + \beta_1 x_i + \epsilon_i, \quad i = 1, \ldots, n$

1st column of X?

- $y_i = \beta_0 + \beta_1 x_i + \beta_2 x_i^2 + \beta_3 x_i^3 + \beta_4 x_i^4 + \beta_5 x_i^5 \epsilon_i$
- $y_i = \beta_0 \pm \beta_1 + \epsilon_i$
- $y_i = \beta_0 + \beta_1 \sin(x_i) + \beta_2 \cos(x_i) + \epsilon_i$
- $y_i = \gamma_0 x_{1i}^{\gamma_1} x_{2i}^{\gamma_2} \eta_i, \quad \eta_i \sim$ positive r.v.
- $y_i = \varphi_0 + \sum_{k=1}^{K} \varphi_k s_k(x_i) + \epsilon_i$

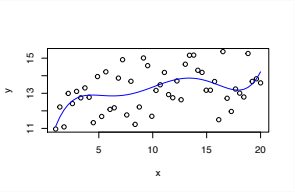

e.g smoothing splines

The linear model

- expected value $\mathbb{E}(y) =$ linear in β

$\longrightarrow$ `Sep142022.Rmd`

The linear model

- expected value $\mathbb{E}(y) =$ linear in β
- measured with additive error $y = \mathbb{E}(y) + \epsilon, \quad \epsilon \sim$

$\longrightarrow$ `Sep142022.Rmd`

The linear model

- expected value $\mathbb{E}(y) =$ linear in β
- measured with additive error $y = \mathbb{E}(y) + \epsilon, \quad \epsilon \sim$
- generalizations

$\epsilon \sim$

$\longrightarrow$ `Sep142022.Rmd`

Today

1. Course introduction: technical issues, course details, evaluation, syllabus, people
2. Upcoming events of interest
3. Review of linear regression
4. In the news: $\longrightarrow$ at the conference