# SF2930 - Regression analysis <br> KTH Royal Institute of Technology, Stockholm 

Lecture 3 - Multiple linear regression (MPV 3, Iz 5.2)

February 14, 2022

## Todays lecture

- Multiple linear regression models
- Matrix notation for MLRM
- Least squares for MLRM
- The hat matrix
- Properties of the LS estimators
- Estimation of $\sigma^{2}$
- Random regression variables


## Multiple linear regression models

In many situations we need a more complicated model. Some examples of such models are given by, e.g.,

$$
y=\beta_{0}+\beta_{1} x_{1}+\beta_{2} x_{2}+\varepsilon
$$

or more generally

$$
y=\beta_{0}+\beta_{1} x_{1}+\beta_{2} x_{2}+\ldots+\beta_{k} x_{k}+\varepsilon
$$

These are said to be multiple linear regression models because they are linear in the regression coefficients $\beta_{0}, \beta_{1}, \ldots, \beta_{k} . x_{1}, x_{2}, \ldots, x_{k}$ are called regressors or prediction variables.

## Multiple linear regression models

More complex models, such as e.g.

$$
y=\beta_{0}+\beta_{1} x_{1}+\beta_{2} x_{1}^{2}+\beta_{3} x_{1}^{3}+\varepsilon
$$

can be treated by methods for multiple linear regression models by letting $x_{2}:=x_{1}^{2}$ and $x_{3}:=x_{1}^{3}$. Similarly,

$$
y=\beta_{0}+\beta_{1} x_{1}+\beta_{2} x_{2}+\beta_{3} x_{1} x_{2}+\beta_{4} e^{x_{3}}+\varepsilon
$$

can be treated by methods for multiple linear regression models by letting $x_{3}:=x_{1} x_{2}$ and $x_{4}:=e^{x_{3}}$.

## Matrix notation

Let $\mathbf{x}_{i}:=\left(1, x_{i 1}, x_{i 2}, \ldots, x_{i k}\right)$ denote the $i$ th observation of $\left(1, x_{1}, x_{2}, \ldots, x_{k}\right)$. Given $n$ observations, we define

$$
X:=\underbrace{\left(\begin{array}{ccccc}
1 & x_{11} & x_{12} & \ldots & x_{1 k} \\
1 & x_{21} & x_{22} & \ldots & x_{2 k} \\
1 & x_{31} & x_{32} & \ldots & x_{3 k} \\
\vdots & \vdots & \ddots & & \vdots \\
1 & x_{n 1} & x_{n 2} & \ldots & x_{n k}
\end{array}\right)}_{\text {The model matrix }} \quad \boldsymbol{\beta}:=\left(\begin{array}{c}
\beta_{0} \\
\beta_{1} \\
\vdots \\
\beta_{k}
\end{array}\right) \quad \mathbf{y}:=\left(\begin{array}{c}
y_{1} \\
y_{2} \\
\vdots \\
y_{n}
\end{array}\right) \quad \boldsymbol{\varepsilon}:=\left(\begin{array}{c}
\varepsilon_{1} \\
\varepsilon_{2} \\
\vdots \\
\varepsilon_{n}
\end{array}\right)
$$

Using this notation, the multiple linear regression model can be written as

$$
\underbrace{\mathbf{y}=X \boldsymbol{\beta}+\boldsymbol{\varepsilon}}_{\text {the regression function }}
$$

## Least squares

As in the 1-dimensional case, we want to choose $\beta_{0}, \beta_{1}, \ldots, \beta_{k}$ which minimizes the least-squares/loss function
$\underbrace{S\left(\beta_{1}, \beta_{2}, \ldots, \beta_{k}\right)}:=\sum \varepsilon_{i}^{2}=\sum\left(y_{i}-\left(\beta_{0}+\beta_{1} x_{i 1}+\ldots+\beta_{k} x_{i k}\right)\right)^{2}=\|\mathbf{y}-X \boldsymbol{\beta}\|_{2}^{2}$.
the loss function

## The least-squares normal equations

The equations $\frac{d S}{\beta_{1}}=0, \frac{d S}{\beta_{2}}=0, \ldots, \frac{d S}{\beta_{k}}=0$ are called the least-squares normal equations, and their solution the least squares estimators $\hat{\beta}_{0}, \hat{\beta}_{1}, \ldots, \hat{\beta}_{k}$.

## Least squares

We have

$$
\begin{aligned}
S(\boldsymbol{\beta}) & =\sum \varepsilon_{i}^{2}=\|\mathbf{y}-X \boldsymbol{\beta}\|_{2}^{2}=(\mathbf{y}-X \boldsymbol{\beta})^{T}(\mathbf{y}-X \boldsymbol{\beta}) \\
& =\mathbf{y}^{T} \mathbf{y}-2 \mathbf{y}^{T} X \boldsymbol{\beta}+(X \boldsymbol{\beta})^{T} X \boldsymbol{\beta} \\
& =\mathbf{y}^{T} \mathbf{y}-2 \mathbf{y}^{T} X \boldsymbol{\beta}+\boldsymbol{\beta}^{T} X^{T} X \boldsymbol{\beta} .
\end{aligned}
$$

Consequently, the least squares solution $\hat{\boldsymbol{\beta}}$ must satisfy

$$
\underbrace{0=\left.\frac{d S}{d \boldsymbol{\beta}}\right|_{\hat{\boldsymbol{\beta}}}=-2 X^{T} \mathbf{y}+2 X^{T} X \hat{\boldsymbol{\beta}}}_{\text {least-squares normal equations }}
$$

Solving for $\hat{\boldsymbol{\beta}}$, we obtain

$$
\hat{\boldsymbol{\beta}}=\left(X^{T} X\right)^{-1} X^{T} \mathbf{y} .
$$

From linear algebra, we know that the matrix $X^{T} X$ in invertible exactly if the columns of $X=\left(\mathbf{x}_{1}, \mathbf{x}_{2}, \ldots, \mathbf{x}_{n}\right)^{T}$ are linearly independent.

## Geometric interpretation



## Example

1 pairs (df00[,3:13], col $=" \# 703457 ", \operatorname{pch}=16, \quad l w d=0, \quad \operatorname{cex}=.7$, xaxt='n', yaxt='n')







 $\square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square \square$

## Example

${ }_{1}$ df00.model <- lm(people_fully_vaccinated_per_hundred~gdp_per _capita, data $=$ df00)
summary (df00.model)

Call:
 capita, data $=$ df00)

Residuals:

| Min | $1 Q$ | Median | $3 Q$ | Max |
| ---: | ---: | ---: | ---: | ---: |
| -16.428 | -6.176 | -0.675 | 7.997 | 14.445 |

Coefficients:

|  | Estimate | Std. Error | t value | $\operatorname{Pr}(>\|t\|)$ |  |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: |
| (Intercept) | $1.929 e+01$ | $6.075 e+00$ | 3.175 | 0.00588 | $* *$ |
| gdp_per_capita | $1.194 e-03$ | $1.957 e-04$ | 6.100 | $1.53 e-05$ | $* * *$ |


Residual standard error: 9.665 on 16 degrees of freedom
(33 observations deleted due to missingness)
Multiple R-squared: 0.6993, Adjusted R-squared: 0.6805
F-statistic: 37.21 on 1 and 16 DF , p-value: $1.534 \mathrm{e}-05$

## Example



## Example

```
df00.model2 <- lm(people_fully_vaccinated_per_hundred~gdp_
```

    per_capita+hospital_beds_per_thousand, data = df00)
    summary (df00.model2)

Call:
lm(formula = people_fully_vaccinated_per_hundred ~ gdp_per_ capita + hospital_beds_per_thousand, data $=d f 00)$

Residuals:

| Min | $1 Q$ | Median | $3 Q$ | Max |
| ---: | ---: | ---: | ---: | ---: |
| -13.7639 | -4.4811 | 0.0485 | 5.8690 | 12.7837 |

Coefficients:
Estimate Std. Error t value $\operatorname{Pr}(>|t|)$
(Intercept) $33.7305585 \quad 10.3210623 \quad 3.268 \quad 0.00519$ **
$\begin{array}{lllll}\text { gdp_per_capita } & 0.0011229 & 0.0001901 & 5.908 & 2.88 e-05\end{array}$ *** hospital_beds_ -2.1185866 1.2567801 -1.686 0.11253

Residual standard error: 9.153 on 15 degrees of freedom
(33 observations deleted due to missingness)
Multiple R-squared: 0.7472, Adjusted R-squared: 0.7135 F-statistic: 22.17 on 2 and 15 DF , p-value: 3.318e-05

## Example



## Example

df00.model3 <- lm(people_fully_vaccinated_per_hundred~gdp_ per_capita+I (gdp_per_capita~2), data $=d f 00)$
summary (df00.model3)
Call:
 capita + I (gdp_per_capita~2), data $=d f 00)$

Residuals:

| Min | $1 Q$ | Median | $3 Q$ | Max |
| ---: | ---: | ---: | ---: | ---: |
| -17.6333 | -3.6830 | 0.4621 | 6.3833 | 11.8959 |

Coefficients:

| (Intercept) | $7.418 \mathrm{e}+00$ | $1.007 \mathrm{e}+01$ | 0.737 | 0.47262 |  |
| :--- | ---: | ---: | ---: | ---: | ---: |
| gdp_per_capita | $2.261 \mathrm{e}-03$ | $7.593 \mathrm{e}-04$ | 2.978 | 0.00939 |  |
| I (gdp_per_capita-2) | $-1.956 \mathrm{e}-08$ | $1.347 \mathrm{e}-08$ | -1.451 | 0.16729 |  |$*$


Residual standard error: 9.347 on 15 degrees of freedom
(33 observations deleted due to missingness)
Multiple R-squared: 0.7363 , Adjusted R-squared: 0.7012
F-statistic: 20.94 on 2 and 15 DF , p-value: $4.55 \mathrm{e}-05$

## Example



## The hat matrix

## The hat matrix

If $\mathbf{x}=\left(1, x_{0}, x_{1}, \ldots, x_{k}\right)^{T}$, then the corresponding fitted model will be

$$
y=\mathbf{x}^{T} \hat{\boldsymbol{\beta}}=\hat{\beta}_{0}+\hat{\beta}_{1} x_{1}+\ldots+\hat{\beta}_{k} x_{k}
$$

and the vector $\hat{\mathbf{y}}=\left(\hat{y}_{1}, \hat{y}_{2}, \ldots, \hat{y}_{n}\right)$ will be given by

$$
\hat{\mathbf{y}}=X \hat{\beta}=\underbrace{X\left(X^{T} X\right)^{-1} X^{T}}_{\mathbf{H}} \mathbf{y} .
$$

The matrix $\mathbf{H}:=X\left(X^{T} X\right)^{-1} X^{T}$ is called the hat matrix.

## Leverage

Since $\hat{\mathbf{y}}=\mathbf{H y}$, we can think of $\mathbf{H}_{i j}$ as expressing how much leverage the variable $y_{j}$ exerts on the fitted value $\hat{y}_{i}$.

## Residuals

Let $\mathbf{e}:=\mathbf{y}-\hat{\mathbf{y}}$ be the vector of residuals. Then $\mathbf{e}=(I-H) \mathbf{y}$, and hence $\operatorname{Cov}(\mathbf{e})=\sigma^{2}(I-H)$.

## The hat matrix

## Leverage

Since $\hat{\mathbf{y}}=\mathbf{H y}$, we can think of $\mathbf{H}_{i j}$ as expressing how much leverage the variable $y_{j}$ exerts on the fitted value $\hat{y}_{i}$.

## Residuals

Let $\mathbf{e}:=\mathbf{y}-\hat{\mathbf{y}}$ be the vector of residuals. Then $\mathbf{e}=(I-H) \mathbf{y}$, and hence $\operatorname{Cov}(\mathbf{e})=\sigma^{2}(I-H)$.



## The hat matrix

## Model testing

If the model assumption is correct, we have

$$
\operatorname{Cov}(\mathbf{e}, \hat{\mathbf{y}})=\operatorname{Cov}((I-H) \mathbf{y}, H \mathbf{y})=0
$$

Consequently, a scatterplot of $\mathbf{e}$ vs $\hat{\mathbf{y}}$ should have no apparent slope or other pattern.



## Properties of $\hat{\boldsymbol{\beta}}$

## The expected value

$$
\begin{aligned}
\mathbb{E}[\hat{\boldsymbol{\beta}}] & =\mathbb{E}\left[\left(X^{T} X\right)^{-1} X^{T} \mathbf{y}\right]=\left(X^{T} X\right)^{-1} X^{T} \mathbb{E}[\mathbf{y}]=\left(X^{T} X\right)^{-1} X^{T} \mathbb{E}[X \boldsymbol{\beta}+\boldsymbol{\varepsilon}] \\
& =\left(X^{T} X\right)^{-1} X^{T} X \boldsymbol{\beta}=\boldsymbol{\beta}
\end{aligned}
$$

In other words, if the model is correct, then $\hat{\boldsymbol{\beta}}$ is an unbiased estimator of $\boldsymbol{\beta}$.

## The covariance

$$
\begin{aligned}
& \operatorname{Cov}(\hat{\boldsymbol{\beta}})=\mathbb{E}[(\hat{\boldsymbol{\beta}}-\mathbb{E}[\hat{\boldsymbol{\beta}}])(\hat{\boldsymbol{\beta}}-\mathbb{E}[\hat{\boldsymbol{\beta}}])]=\operatorname{Var}(\hat{\boldsymbol{\beta}}-\mathbb{E}[\hat{\boldsymbol{\beta}}]) \\
& \quad=\operatorname{Var}\left(\left(X^{T} X\right)^{-1} X^{T} \mathbf{y}\right)=\left(X^{T} X\right)^{-1} X^{T} \operatorname{Var}(\mathbf{y})\left(\left(X^{T} X\right)^{-1} X^{T}\right)^{T} \\
& \quad=\left(X^{T} X\right)^{-1} X^{T} \sigma^{2} I\left(\left(X^{T} X\right)^{-1} X^{T}\right)^{T}=\sigma^{2}\left(X^{T} X\right)^{-1} X^{T}\left(\left(X^{T} X\right)^{-1} X^{T}\right)^{T} \\
& \quad=\sigma^{2}\left(X^{T} X\right)^{-1} X^{T} X\left(X^{T} X\right)^{-1}=\sigma^{2}\left(X^{T} X\right)^{-1} .
\end{aligned}
$$

## Gauss-Markov theorem

Theorem
$\hat{\boldsymbol{\beta}}$ is the best linear unbiased estimator of $\boldsymbol{\beta}$.

## Estimation of $\sigma^{2}$

## Residuals

We let $\mathbf{e}:=\mathbf{y}-\hat{\mathbf{y}}$ be the vector of residuals.

## The residual sum of squares

We have

$$
\begin{aligned}
S S_{R e s} & =\sum\left(y_{i}-\hat{y}_{i}\right)^{2}=(\mathbf{y}-\hat{\mathbf{y}})^{T}(\overbrace{\mathbf{y}-\hat{\mathbf{y}}}^{\mathrm{e}})=(\mathbf{y}-X \hat{\boldsymbol{\beta}})^{T}(\mathbf{y}-X \hat{\boldsymbol{\beta}}) \\
& =\mathbf{y}^{T} \mathbf{y}-\left(X \hat{\boldsymbol{\beta}}^{T} \mathbf{y}-\mathbf{y}^{T}(X \hat{\boldsymbol{\beta}})+(X \hat{\boldsymbol{\beta}})^{T} X \hat{\boldsymbol{\beta}}\right. \\
& =\mathbf{y}^{T} \mathbf{y}-\hat{\boldsymbol{\beta}}^{T} X^{T} \mathbf{y}-\mathbf{y}^{T} X \hat{\boldsymbol{\beta}}+\hat{\boldsymbol{\beta}}^{T} \underbrace{X^{T} X \hat{\boldsymbol{\beta}}}_{=X^{T} \mathbf{y}}=\mathbf{y}^{T} \mathbf{y}-\mathbf{y}^{T} X \hat{\boldsymbol{\beta}}
\end{aligned}
$$

$S S_{\text {Res }}$ has $n-(k+1)$ degrees of freedom. Also, one can show that $\mathbb{E}\left[S S_{\text {Res }}\right]=\sigma^{2}(n-(k+1))$.
The residual mean squared

$$
\hat{\sigma}^{2}=M S_{R e s}=\frac{S S_{R e s}}{n-(k+1)}
$$

Note that $\hat{\sigma}^{2}$ is an unbiased estimate of $\sigma^{2}$. As in the simple linear regression case, this estimate depends on the model.

## Random regression variables

Recall that the regression variables $x_{1}, x_{2}, \ldots$ can be either

1. non-random (e.g. data from planned experiments), or
2. random (common when we use already collected data).



So far, we have only covered case 1. In this case, we found the LS estimates by minimizing $S(\boldsymbol{\beta})=\|\mathbf{y}-X \boldsymbol{\beta}\|_{2}^{2}$. Note that this means that we are minimizing the errors at the sampled points, and not for general points.
If the regression variables are random, we instead want to minimize

$$
S(\boldsymbol{\beta})=\mathbb{E}\left[\|\mathbf{y}-\mathbf{X} \boldsymbol{\beta}\|_{2}^{2}\right]=\mathbb{E}\left[(\mathbf{y}-X \boldsymbol{\beta})^{T}(\mathbf{y}-\mathbf{X} \boldsymbol{\beta})\right]
$$

## Random regression variables

The "best" estimates
Let

$$
S(\boldsymbol{\beta})=\mathbb{E}\left[\|\mathbf{y}-\mathbf{X} \boldsymbol{\beta}\|_{2}^{2}\right]=\mathbb{E}\left[(\mathbf{y}-X \boldsymbol{\beta})^{T}(\mathbf{y}-\mathbf{X} \boldsymbol{\beta})\right]
$$

Differentiating, we obtain

$$
\frac{d S(\boldsymbol{\beta})}{d \boldsymbol{\beta}}=-2 \mathbb{E}\left[\mathbf{X}^{T} \mathbf{y}-\mathbf{X}^{T} \mathbf{X} \boldsymbol{\beta}\right]=-2 \mathbb{E}\left[\mathbf{X}^{T} \mathbf{y}\right]+2 \mathbb{E}\left[\mathbf{X}^{T} \mathbf{X}\right] \boldsymbol{\beta}
$$

and hence the "best" estimates are given by

$$
\tilde{\boldsymbol{\beta}}:=\mathbb{E}\left[\mathbf{X}^{T} \mathbf{X}\right]^{-1} \mathbb{E}\left[\mathbf{X}^{T} \mathbf{y}\right]
$$

Equivalently,

$$
\left(\tilde{\beta}_{1}, \ldots, \hat{\beta}_{k}\right):=\Sigma_{X X}^{-1} \Sigma_{X Y} \quad \text { and } \quad \tilde{\beta}_{0}:=\mu_{\mathbf{y}}-\mu_{\mathbf{X}}\left(\hat{\beta}_{1}, \ldots, \hat{\beta}_{k}\right)^{T}
$$

## An approximation of the "best" estimates

Note that since the joint distribution of $(\mathbf{X}, \mathbf{y})$ is in general not known, we cannot really use the above "best" estimeates.
For this reason, we let $X_{*}:=\left(x_{i j}\right)$ (so that $X=\left(1, X_{*}\right)$ ), and approximate these estimators by

$$
\left\{\begin{array}{l}
\left(\hat{\beta}_{1}, \ldots, \hat{\beta}_{k}\right):=\left(\left(X_{*}-\bar{X}_{*}\right)^{T}\left(X_{*}-\bar{X}_{*}\right)\right)^{-1}\left(X_{*}-\bar{X}_{*}\right)^{T}(y-\bar{y}) \\
\hat{\beta}_{0}:=\bar{y}-\bar{X}\left(\hat{\beta}_{1}, \ldots, \hat{\beta}_{k}\right)^{T}
\end{array}\right.
$$

## Commentes

$\rightarrow$ The above estimators are identical to the LSE given for the non-random setting.
$\rightarrow$ All previous results are applicable if $y \mid x \sim N\left(\beta_{0}+\beta_{1} x, \sigma^{2}\right)$ and the sampled $x$ are independent (with some unknown distribution).

